electronics

MDPI

Article

A Simple Introduction to Quantum Annealing in Antennas and Propagation

Marco Donald Migliore [1,2,3]

[1] DIEI (Dipartimento di Ingegeria Elettrica e dell'Informazione "Maurizio Scarano"), University of Cassino and Southern Lazio, Via G. Di Biasio 43, 03043 Cassino, Italy; mdmiglio@unicas.it
[2] European University of Technology EUT+, 03043 Cassino, Italy
[3] CNIT (Consorzio Nazionale Interuniversitario per le Telecomunicazioni), Viale G.P. Usberti, 43124 Parma, Italy

Abstract

The objective of this paper is to introduce the application of quantum annealing (QA) to electromagnetic (EM) engineering. We demonstrate that numerous EM design and inverse problems naturally admit a mathematical formulation in terms of binary quadratic interactions. By utilizing didactic, simplified examples, we illustrate how this underlying physical structure allows for a direct mapping of EM problems onto Quadratic Unconstrained Binary Optimization (QUBO) models and equivalent Ising Hamiltonians. All numerical experiments are conducted using Simulated Annealing as a classical proxy, since the primary contribution of this work is the QUBO formulation itself rather than its execution on quantum hardware. Our results suggest that a broad class of EM problems—including array thinning, reconfigurable intelligent surface optimization, subarray partitioning, and electromagnetic inverse scattering—are suitable candidates for optimization using near-future quantum annealing architectures, once hardware connectivity and noise floors reach the required specifications.

Keywords: quantum annealing; quadratic unconstrained binary optimization (QUBO); array thinning; inverse scattering; reconfigurable intelligent surfaces; subarray partitioning

Check for updates

Academic Editor: Wiesław Leonski

Received: 28 April 2026
Revised: 26 May 2026
Accepted: 28 May 2026
Published: 1 June 2026

1. Introduction

The objective of this paper is to introduce the application of quantum annealing (QA) within the field of electromagnetics (EM). We present examples of EM design and inverse problems that naturally lend themselves to a formulation based on binary quadratic interactions, facilitating a direct mapping onto current quantum optimization architectures. Quantum computing leverages fundamental physical phenomena, such as superposition and entanglement, to process information in a manner distinct from classical computation [1]. While classical computers manipulate bits restricted to discrete values $\{0, 1\}$, quantum computers operate on qubits, where the state space of the system scales exponentially with the number of subsystems. This property, which provides the theoretical basis for a potential quantum advantage in specific problem classes, has garnered significant interest within the electromagnetics community, particularly in areas regarding antennas and propagation [2–11].

Currently, two major hardware paradigms dominate the field: gate-based quantum computers [12] and quantum annealers [13,14]. Although both are rooted in quantum mechanics, they differ substantially in architecture, control mechanisms, and target applications. Gate-based quantum computers are conceptually analogous to classical digital

 https://doi.org/10.3390/electronics15112390

processors, implementing algorithms through sequences of discrete, elementary quantum gates. This universal approach is designed to eventually realize any quantum algorithm, such as Shor's factoring or Grover's search. In contrast, quantum annealing is a specialized paradigm primarily tailored for combinatorial optimization. Rather than executing a sequence of logical gates, a quantum annealer evolves a physical system toward a final minimum-energy configuration (the ground state) that is programmed to correspond to the solution of the encoded problem.

Quantum annealing is often regarded as being more accessible for near-term practical exploration than universal gate-based systems. This is primarily because its hardware requirements—specifically regarding coherence times and connectivity—are less restrictive for certain optimization tasks [15]. While a definitive, general quantum speedup over the most advanced classical heuristics remains a subject of ongoing research, the technology has already been explored for experimental applications in traffic-flow optimization, scheduling, and logistics.

In this paper, we focus exclusively on the annealing approach. Readers interested in a comprehensive introduction to gate-based quantum computing and its specific applications to antennas and propagation are referred to [3]; similarly, introductory works on quantum annealing specifically tailored for the antennas and propagation community can be found in [4,5]. With respect to these works, the present manuscript follows the same QUBO-mapping perspective but extends it in two directions. First, it addresses two configurations—subarray partitioning and electromagnetic inverse scattering—which, to the best of our knowledge, have not been previously cast as QUBO problems in the antennas and propagation literature. Second, it surfaces an operational point that is easy to overlook when moving from simulated-annealing proxies to physical QPU hardware: the two approaches enforce the cardinality constraint in fundamentally different ways. In SA, we control the move set, so we can use swap moves that preserve cardinality by construction—every configuration the search visits has the same number of selected assets. Feasibility is therefore guaranteed by the moves themselves, and the cardinality penalty is redundant: it adds the same constant to every state and never influences the search. A quantum annealer does not restrict the search to feasible state; hence, the constraints should be taken into account in a different way. The constraint can only be expressed through the problem Hamiltonian, which makes the penalty term mandatory rather than optional. This shifts the burden onto penalty-weight tuning: the weight must be large enough to suppress infeasible solutions, yet a large weight competes with the objective coefficients and, given the limited precision of the hardware, can degrade solution quality.

It is important to stress that all numerical experiments presented in this paper are solved using Simulated Annealing (SA), a well-known classical stochastic heuristic for QUBO models that mimics thermodynamic cooling. In this work, SA is used as a classical verification proxy for the optimization process, while the term Quantum Annealing (QA) is reserved for the corresponding physical paradigm and future hardware-based implementations. The primary contribution of this work is therefore the formulation of electromagnetic problems as QUBO models, demonstrating that their mathematical structure is naturally compatible with quantum annealing hardware. The actual execution on a quantum processing unit (QPU), such as the D-Wave Advantage system, is left as a direct follow-up step: once the QUBO matrix Q is assembled, it can be submitted to a QPU without any modification to the problem formulation.

2. Quantum Annealing and Hardware Implementation

This section provides a brief introduction to quantum annealing (QA). Quantum annealing is a meta-heuristic optimization that seeks the low-energy configurations of a physical system. In the ideal adiabatic framework, the system evolves according to a time-dependent Hamiltonian [3]:

$$\hat{H}(t) = A(t)\hat{H}_0 + B(t)\hat{H}_P \,, \tag{1}$$

where $\hat{H}_0$ is the initial Hamiltonian with a known, easily prepared ground state (typically representing a uniform superposition of all possible states), $\hat{H}_P$ is the problem Hamiltonian whose ground state encodes the target solution, and $A(t), B(t)$ are functions defining the annealing schedule. Initially, the system is prepared in the ground state of $\hat{H}_0$ by setting $A(0) \gg B(0)$. Over the annealing period $t \in [0, t_f]$, $A(t)$ is decreased to zero while $B(t)$ increases to its maximum value. According to the adiabatic theorem, if the evolution is sufficiently slow and the system remains isolated, it stays close to the instantaneous ground state of $\hat{H}(t)$, eventually reaching the ground state of $\hat{H}_P$. Modern commercial quantum annealers typically operate with annealing times on the order of microseconds.

The physical mechanism by which QA explores these cost landscapes differs fundamentally from classical counterparts like Simulated Annealing (SA). In SA, the probability of escaping a local minimum across an energy barrier of height ΔE scales as $\exp(-\Delta E/k_B T)$, meaning that thermal hopping is primarily governed by barrier height. Conversely, QA leverages a transverse magnetic field to initiate quantum tunneling through energy barriers. The tunneling rate depends heavily on both the height and the spatial width of the barrier. Consequently, QA is expected to provide substantial acceleration on optimization landscapes characterized by tall but narrow potential barriers, which are notoriously difficult for classical thermal fluctuations to surmount [16,17], while offering a more modest advantage over SA on wide, shallow configurations.

It must also be emphasized that even in an ideal, perfectly noiseless quantum device, finding the absolute ground state is not mathematically guaranteed for all problem instances. According to the adiabatic theorem, if the minimum energy gap between the ground state and the first excited state becomes extremely narrow during the evolution, the required annealing time increases exponentially, inducing non-adiabatic transitions into excited states. In practical near-term hardware, this intrinsic constraint is compounded by extrinsic environmental noise, thermal fluctuations, and finite coherence times that prevent the system from being truly isolated. Consequently, a single annealing cycle may return a low-energy excited state rather than the absolute global minimum. Because the execution time for a single cycle is extremely short, the process is typically repeated many times to sample a large distribution of candidate solutions. By selecting the lowest-energy configuration from these samples, one can obtain the global optimum, or a sufficiently high-quality sub-optimal solution, with a certain success probability. In this sense, current hardware is best viewed as a probabilistic sampler of low-energy states rather than a deterministic global optimizer.

In hardware implementations, the problem Hamiltonian is typically expressed in the form of an Ising model [18]. Originating from statistical physics, the Ising model describes a system of interacting binary spins. The energy function associated with the problem Hamiltonian is:

$$E(z_1, \ldots, z_N) = \sum_i h_i z_i + \sum_{i<j} J_{ij} z_i z_j \,, \tag{2}$$

where $z_i \in \{-1, +1\}$ are spin variables, h_i denotes local longitudinal fields acting on individual spins, and J_{ij} denotes pairwise coupling coefficients. The configuration of spins that minimizes this energy corresponds to the ground state. To utilize a quantum annealer, optimization problems must be mapped into this Ising form so that the coefficients h_i and J_{ij} can be programmed as physical biases and couplings in the device.

Equivalently, the same class of problems can be formulated as Quadratic Unconstrained Binary Optimization (QUBO) models:

$$\min_{\mathbf{x}\in\{0,1\}^N} \mathbf{x}^T Q\mathbf{x}\,, \tag{3}$$

where $\mathbf{x}$ is a vector of binary decision variables and $Q \in \mathbb{R}^{N\times N}$ is the QUBO matrix. It is useful to emphasize that while the QUBO formulation mathematically matches the native Ising structure of commercial quantum annealers, QUBO itself is a well-established paradigm in classical combinatorial optimization [19]. Consequently, translating EM problems into a QUBO framework does not merely prepare them for future quantum acceleration; it immediately unlocks a wide array of powerful classical heuristics digital solvers. Prominent examples include comprehensive open-source algorithmic suites [20], as well as simulated bifurcation algorithms and machines rooted in classical and nonlinear mechanics [21,22], including very recent variants exhibiting quantum-like features [23]. Therefore, the QUBO models presented herein serve as a unified mathematical bridge to both quantum processing units and state-of-the-art classical optimization suites, significantly maximizing their potential engineering impact.

The transition between binary variables x_i and Ising spins z_i is performed through the transformation $x_i = \frac{1-z_i}{2}$. For an upper-triangular QUBO matrix, the mapping yields:

$$J_{ij} = \frac{1}{4}Q_{ij} \quad (i<j), \qquad h_i = -\frac{1}{2}Q_{ii} - \frac{1}{4}\sum_{j\neq i} Q_{ij} \tag{4}$$

resulting in $x^T Qx = \sum_{i<j} J_{ij}z_iz_j + \sum_i h_i z_i + E_{\text{offset}}$, where the constant scalar energy offset is given by:

$$E_{\text{offset}} = \frac{1}{2}\sum_i Q_{ii} + \frac{1}{4}\sum_{i<j} Q_{ij} \tag{5}$$

Accounting for E_{offset} ensures strict numerical equivalence between the absolute energies computed in the binary and spin domains. This mapping identifies two primary classes of physically programmable parameters in the quantum annealer:

- *Local fields (h_i):* In hardware architectures such as D-Wave, these correspond to qubit biases. These parameters represent the "individual preference" of a qubit to settle in a specific state, regardless of its neighbors.
- *Couplings (J_{ij}):* These represent the "relationships" or interference between pairs of elements. A negative coupling ($J_{ij} < 0$) favors aligned spin states (ferromagnetic interaction), which in EM models typically represents constructive interference. Conversely, a positive coupling ($J_{ij} > 0$) favors opposite states (antiferromagnetic interaction), representing destructive interference. In hardware, these correspond to inter-qubit couplers.

Implementing these models on physical hardware involves addressing the structural constraints of the Quantum Processing Unit (QPU). For instance, the D-Wave Advantage system [24] utilizes the Pegasus topology [25], which restricts each physical qubit to a maximum of 15 neighbors. When the logical connectivity of a problem exceeds this hardware limit, minor embedding is required. This technique maps a single logical variable to a "chain" of multiple physical qubits, effectively embedding the problem's logical graph into the sparser physical graph of the QPU. While minor embedding enables the solution of highly connected problems, it consumes extra physical qubits and introduces additional constraints, such as chain strength, which can reduce the overall success probability of finding the global optimum.

3. Applications in Antennas and Propagation Problems

The suitability of many electromagnetic design and inverse problems for quantum annealing (QA) stems from the fact that several relevant performance metrics are quadratic functions of the electromagnetic field. Quantities such as radiated power, received power, and least-squares data misfit can all be expressed in a quadratic form. In many antenna and propagation problems, the field F at a given observation point can be represented as a linear superposition of contributions from discrete sources or scattering elements:

$$F = \sum_i a_i x_i \,, \tag{6}$$

where $a_i \in \mathbb{C}$ represents the complex contribution of the i-th element and x_i is a design variable controlling its state (e.g., activation, phase, or presence). The corresponding power, proportional to $|F|^2$, expands as:

$$|F|^2 = \sum_i |a_i|^2 x_i^2 + \sum_{i \neq j} a_i a_j^* x_i x_j \,. \tag{7}$$

For binary variables $x_i \in \{0, 1\}$, where $x_i^2 = x_i$, this simplifies to:

$$|F|^2 = \sum_i |a_i|^2 x_i + \sum_{i<j} 2\,\mathrm{Re}\{a_i a_j^*\} x_i x_j \,. \tag{8}$$

This expression demonstrates that electromagnetic interference between different field contributions naturally produces the structure required for a Quadratic Unconstrained Binary Optimization (QUBO) formulation. A similar situation arises in electromagnetic inverse problems. If the measurement model is expressed as $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}$, where $\mathbf{n}$ is noise, the least-squares data misfit can be written in quadratic form:

$$\|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2 = \mathbf{x}^T \mathbf{A}^H \mathbf{A}\mathbf{x} - 2\mathrm{Re}\{\mathbf{y}^H \mathbf{A}\mathbf{x}\} + \|\mathbf{y}\|^2 \,. \tag{9}$$

Quadratic formulations of this type are well established in the antennas and propagation community and are widely utilized within the framework of convex optimization [26,27]. Convex minimization provides efficient tools for solving such problems when the optimization variables are continuous. However, many practical problems in antenna engineering involve intrinsically discrete design variables. Examples include element activation in array thinning, quantized phase states in reconfigurable intelligent surfaces, subarray assignment in phased arrays, and binary occupancy in sparse inverse reconstruction.

When variables are restricted to discrete sets, the problem becomes combinatorial, and the convexity properties that enable efficient continuous optimization no longer apply. In these cases, the quadratic structure of the underlying electromagnetic model is preserved while the variables are restricted to binary or discrete values. This leads naturally to QUBO models, which can be interpreted as the search for the ground state of an interacting spin system. Such problems are suitable for solution using quantum annealing hardware or classical algorithms inspired by annealing dynamics. Several characteristics make electromagnetic problems particularly compatible with this approach:

- Discrete variables: Many design choices correspond to on/off states, quantized phases, or binary structural decisions.
- Quadratic objectives: Performance metrics often derive from power, energy, or least-squares fitting criteria.
- Structured interactions: The resulting matrices often exhibit Toeplitz or block-Toeplitz structures induced by array geometry and wave propagation physics.

- Physical mapping: The quadratic models map directly onto interacting spin systems, enabling implementation on quantum hardware.

It is also important to contextualize quantum annealing with respect to popular evolutionary metaheuristics widely adopted in electromagnetic optimization, such as Genetic Algorithms (GA) and Particle Swarm Optimization (PSO). While GA and PSO are versatile, population-based frameworks primarily tailored for continuous or mixed-integer search spaces, QA targets a rigid, specific algebraic structure (QUBO/Ising) restricted to binary decision variables. For problems that naturally admit a compact QUBO mapping, QA explores the landscape via quantum mechanical evolution rather than via stochastic operators in parameter space, which can be advantageous on very complex combinatorial landscapes. However, a universal 'quantum advantage' over well-tuned classical GA/PSO heuristics remains unproven and highly problem-dependent due to hardware embedding constraints and noise thresholds; hence, QA should be viewed as a complementary tool specialized for highly discrete, combinatorial EM challenges.

4. Examples of Electromagnetic Problems Implemented Using QUBO

In the following subsections, four representative electromagnetic problems exhibiting the aforementioned characteristics are presented, and their corresponding QUBO formulations are derived. It should be noted that the application of quantum annealing to antennas and propagation problems has been investigated by several research groups, covering topics such as array sparsification and RIS control (see, for example, [7–11]). To the best of the authors' knowledge, applications to subarray partitioning and electromagnetic inverse scattering have not yet been extensively reported in the open literature.

The objective of this section is to demonstrate how these diverse problems can be analyzed within a unified framework, highlighting the common underlying principles that characterize different electromagnetic applications. It is important to clarify that the following examples are primarily pedagogical in nature; they are designed to clarify how such problems can be represented by QUBO models. While these cases are simplified for clarity, the underlying methodology can be scaled to manage more complex and realistic industrial applications.

4.1. Array Thinning

Array thinning consists of selecting a subset of M active elements from N candidate antenna positions to reduce sidelobe levels while preserving specific main-beam characteristics. In this example, the optimization is formulated to suppress the integrated sidelobe energy while maintaining radiation in the broadside direction ($\theta = 0°$). Let $x_n \in \{0, 1\}$ denote the activation state of the n-th antenna element. For a linear array with inter-element spacing d (see Figure 1) , the array factor is defined as:

$$F(\theta) = \sum_{n=0}^{N-1} x_n a'_n e^{jk_0 dn \sin\theta}, \tag{10}$$

where k_0 is the free-space wavenumber and a'_n represents a fixed complex excitation coefficient. The corresponding radiated power pattern is proportional to $|F(\theta)|^2$. The integrated sidelobe level is defined as:

$$\mathrm{ISL} = 10\log_{10}\left(\frac{\int_{\Theta_{SL}} |F(\theta)|^2\, d\theta}{\max_\theta\ |F(\theta)|^2}\right), \tag{11}$$

where $\Theta_{SL} = \{\theta : \theta' < |\theta| \leq 90^\circ\}$ is the sidelobe region within the visible space, and θ' is the half-angular width of the main lobe, defining the boundary between the main beam and the sidelobe region. The cost function to be minimized is the discrete approximation of this integral:

$$C_{SL}(\mathbf{x}) = \sum_{\theta_k \in \Theta_{SL}} w_k |F(\theta_k)|^2, \tag{12}$$

where w_k denotes quadrature weights. Expanding the squared magnitude and extracting the real part to ensure a symmetric interaction matrix, the sidelobe QUBO matrix Q_{SL} is defined as:

$$Q_{SL,nm} = \mathrm{Re}\left\{ \sum_{\theta_k \in \Theta_{SL}} w_k \, a'_n a'^*_m \, e^{jk_0 d(n-m)\sin\theta_k} \right\}. \tag{13}$$

In the case of uniform excitation ($a'_n = 1$), Q_{SL} exhibits a Toeplitz structure where $Q_{nm} = q_{n-m}$.

To promote the formation of a directive beam at broadside, a constraint matrix $Q_{BS} = -\mathrm{Re}\{\mathbf{v}_0 \mathbf{v}_0^H\}$ is introduced, where $\mathbf{v}_0$ is the steering vector at $\theta = 0^\circ$. The total objective function takes the form:

$$C(\mathbf{x}) = \lambda_{SL} \, \mathbf{x}^T Q_{SL} \mathbf{x} + \lambda_{BS} \, \mathbf{x}^T Q_{BS} \mathbf{x}, \tag{14}$$

where λ_{SL} and λ_{BS} are weighting coefficients. Both matrices are normalized to the range $[-1, 1]$ before weighting, making the hyperparameters scale-invariant.

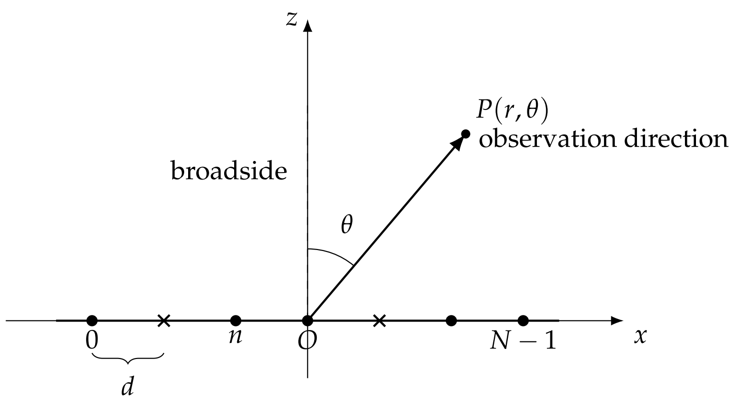

Figure 1. Geometry of a thinned linear array. Filled circles denote active elements, while crossed positions denote inactive ones. Broadside is along the z-axis.

The QUBO model was applied to a linear array of $N = 128$ elements with inter-element spacing $d = \lambda/2$, thinned to $M = 32$ active elements (thinning ratio 25%). The angular grid consists of 4000 uniformly spaced points over the visible region $\theta \in [-90^\circ, 90^\circ]$, with main lobe boundary $\theta' = 3^\circ$, chosen to encompass the main lobe of the array factor while excluding the first transition sidelobe, so that the minimization is applied only to directions sufficiently far from broadside. The weighting coefficients are set to $\lambda_{SL} = 1.0$ and $\lambda_{BS} = 5.0$. The QUBO problem is solved via simulated annealing with 400 independent restarts.

Figure 2 shows the normalized radiation pattern of the optimized array. Table 1 summarizes the performance metrics for three reference configurations.

The QUBO-thinned array achieves ISL $= -12.29$ dB, representing a reduction of approximately 8 dB with respect to the fully populated uniform array and approximately 5 dB with respect to a random selection of $M = 32$ elements. It should be noted that the cost function directly penalizes the integrated sidelobe energy rather than the peak sidelobe level; consequently, the SLL of -8.81 dB is higher than that of the uniform array, as the optimizer distributes sidelobe energy broadly rather than suppressing individual peaks. This trade-off is inherent to ISL minimization and can be addressed by introducing angle-dependent weights w_k that emphasize directions close to the main lobe, where the first and highest sidelobes typically occur. This example demonstrates that electromagnetic domain knowledge—specifically the analytic form of the array factor and the quadratic structure of radiated power—can be naturally incorporated into the QUBO energy function, allowing the optimization landscape to directly reflect the underlying physics of wave interference.

Table 1. Upper figure: Array thinning performance comparison ($N = 128$, $d = \lambda/2$, visible region $\theta \in [-90^\circ, 90^\circ]$); lower figure: array excitation.

Configuration	Active Elements	ISL [dB]	SLL [dB]
Uniform array ($N = 128$)	128	≈ -4	-13.26
Random thinning ($M = 32$)	32	≈ -7	≈ -10
QUBO thinning ($M = 32$)	32	-12.29	-8.81

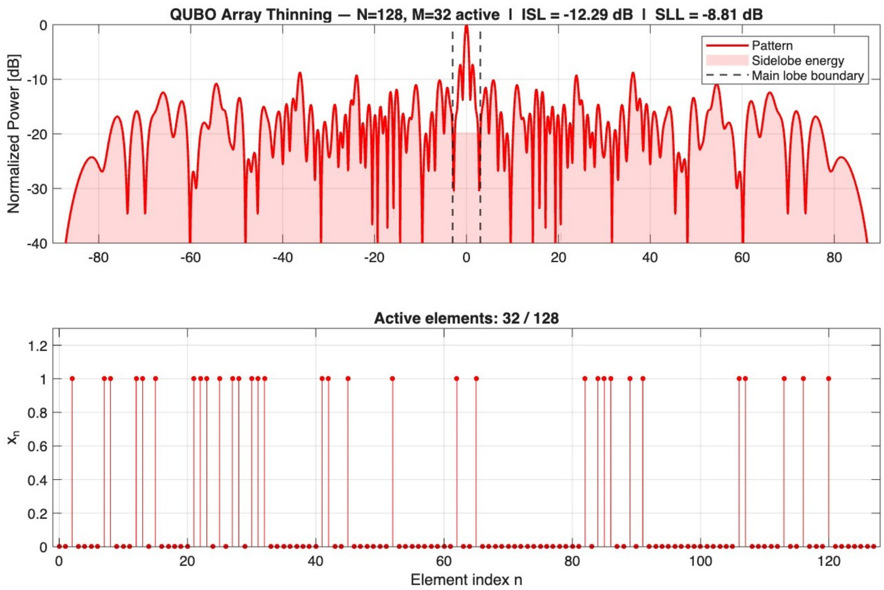

Figure 2. Normalized radiation pattern of the QUBO-thinned array ($N = 128$, $M = 32$, $d = \lambda/2$). The shaded region indicates the sidelobe area used in the ISL computation ($10^\circ < |\theta| \leq 90^\circ$).

4.2. Reconfigurable Intelligent Surfaces

A reconfigurable intelligent surface (RIS) is composed of a set of passive reflecting elements whose local phase response can be programmed to shape the propagation channel between a transmitter and a receiver [28]. Figure 3 illustrates a simplified one-dimensional geometry used to introduce the corresponding QUBO formulation. Consider an RIS with N passive elements. In this pedagogical binary case, each element can impose one of two phase states:

$$\phi_i \in \{0, \pi\}\,. \tag{15}$$

These two states are represented by spin variables $s_i \in \{-1, +1\}$, where $s_i = +1$ corresponds to phase 0 and $s_i = -1$ corresponds to phase π. Let g_i denote the complex channel coefficient associated with the TX $\to$ RIS$_i$ $\to$ RX path. Under the narrowband assumption and neglecting multiple scattering:

$$g_i = \alpha_i\, e^{-jk_0(d_{t,i}+d_{r,i})}\,, \tag{16}$$

where $d_{t,i}$ and $d_{r,i}$ are the propagation distances from the transmitter and to the receiver respectively, $k_0 = 2\pi/\lambda$, and $\alpha_i = (d_{t,i}d_{r,i})^{-1}$ models free-space amplitude attenuation. The effective channel produced by the RIS is $H_{\text{eff}} = \sum_{i=1}^{N} g_i s_i$, and the received power P is proportional to:

$$P = |H_{\text{eff}}|^2 = \left|\sum_{i=1}^{N} g_i s_i\right|^2 = \sum_{i=1}^{N}\sum_{j=1}^{N} g_i g_j^* s_i s_j \,. \tag{17}$$

Although the double summation is guaranteed to be real-valued by the Hermitian symmetry of the form, the individual coefficients $g_i g_j^*$ are generally complex. Since quantum annealing hardware requires real-valued interactions, the objective function must be expressed using only real quantities. Exploiting the pairwise cancellation of imaginary parts ($\mathrm{Im}\{g_i g_j^*\} + \mathrm{Im}\{g_j g_i^*\} = 0$), the received power reduces to:

$$P = \sum_{i,j} \mathrm{Re}\{g_i g_j^*\}\, s_i s_j \,. \tag{18}$$

Maximizing P is therefore equivalent to minimizing the Ising energy:

$$H(\mathbf{s}) = -\sum_{i<j} J_{ij} s_i s_j - \sum_i h_i s_i \,, \tag{19}$$

where $J_{ij} = -\mathrm{Re}\{g_i g_j^*\}$ are the real-valued pairwise couplings encoding wave interference between elements, and $h_i = -|g_i|^2$ are the local fields encoding the individual element contributions. This demonstrates how the RIS phase optimization maps directly onto the native Ising representation of a quantum annealer, with programmable parameters determined entirely by the wave propagation geometry.

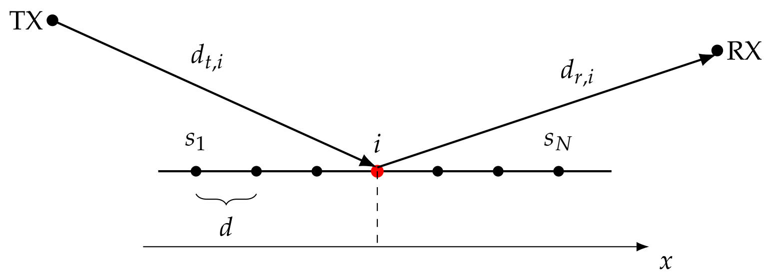

Figure 3. One-dimensional RIS model used to introduce the binary phase optimization problem. The contribution of the i-th RIS element is associated with the propagation distances $d_{t,i}$ and $d_{r,i}$.

The method was tested on an 8×8 planar RIS ($N = 64$ elements) at $f = 28$ GHz ($\lambda \approx 10.7$ mm), with half-wavelength element spacing. The transmitter is located at $\mathbf{r}_{TX} = (-30, 0, 5)$ m and the receiver at $\mathbf{r}_{RX} = (20, 10, 1.5)$ m, with the RIS centered at the origin. The receiver is therefore off-axis with respect to the TX–RIS line, at an elevation angle $\theta \approx -7°$ and azimuth $\phi \approx 27°$ from the RIS boresight. The optimization is solved via simulated annealing with $n = 30$ independent restarts of 2×10^5 iterations each.

Figure 4 shows the optimized phase configurations considering quantized ideal (global binary optimum) and QUBO solution. Figure 5 shows the corresponding far-field power maps as a function of elevation θ and azimuth ϕ; the white cross marks the direction of the receiver. The figures show that the main beam is correctly steered toward the receiver, and the QUBO solution achieves the global minimum, confirming the potentiality of QUBO approach.

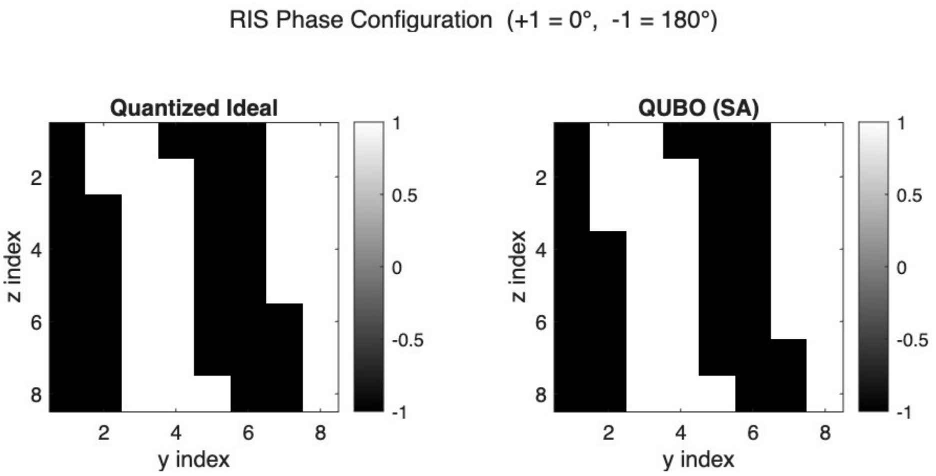

Figure 4. Optimized phase configurations of the 8 × 8 RIS ($N = 64$ elements, $f = 28$ GHz). Each cell represents one element: white = $s_i = +1$ (phase 0); black = $s_i = -1$ (phase π). From left to right: quantized ideal (global binary optimum) and QUBO solution obtained via simulated annealing.

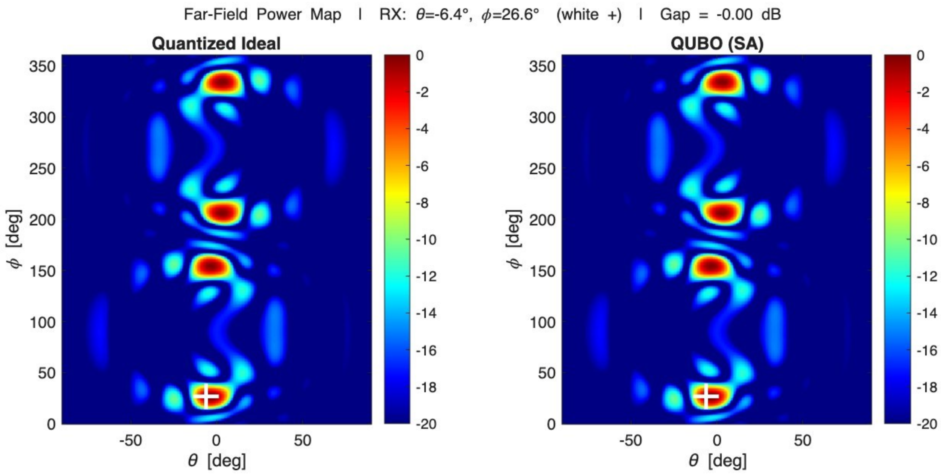

Figure 5. Far-field power maps of the RIS as a function of elevation angle θ and azimuth angle ϕ, normalized to the peak of the continuous-phase reference. From left to right: quantized ideal and QUBO solution. The white cross marks the direction of the receiver ($\theta \approx -7°$, $\phi \approx 27°$). In all three cases the main beam is correctly steered toward the receiver.

4.3. Subarray Partitioning

In many phased-array systems, antenna elements are grouped into subarrays to reduce hardware complexity [29]. The objective is to determine an assignment of N elements to K subarrays that minimizes the integrated sidelobe level (ISL). Consider a linear array with inter-element spacing d (see Figure 6). Let $x_{nk} \in \{0, 1\}$ indicate whether element n belongs to subarray k. Each element must belong to exactly one subarray, and in this pedagogical example, subarrays contain an equal number of elements. The effective weight is $w_n = \sum_{k=1}^{K} v_k x_{nk}$, where v_k is the subarray excitation. The cost function $C = \sum_{\theta \in \Theta_{SL}} |F(\theta)|^2$ expands to:

$$C = \sum_{n,m} \sum_{k,l} Q_{(nk)(ml)} x_{nk} x_{ml} \, , \tag{20}$$

where $Q_{(nk)(ml)} = \mathrm{Re}\{v_k v_l^* \sum_{\theta \in \Theta_{SL}} e^{jk_0(r_n - r_m)\sin\theta}\}$. This partitioning problem naturally maps to a QUBO of form $\min \mathbf{x} \in \{0,1\}^{NK} \mathbf{x}^T Q \mathbf{x}$. Numerical experiments (Figures 7 and 8) on an $N = 32$ array with $K = 4$ subarrays show that the QUBO-optimized interleaved configuration reduces the ISL to 0.33 dB, compared to 6.6 dB for a conventional contiguous partition.

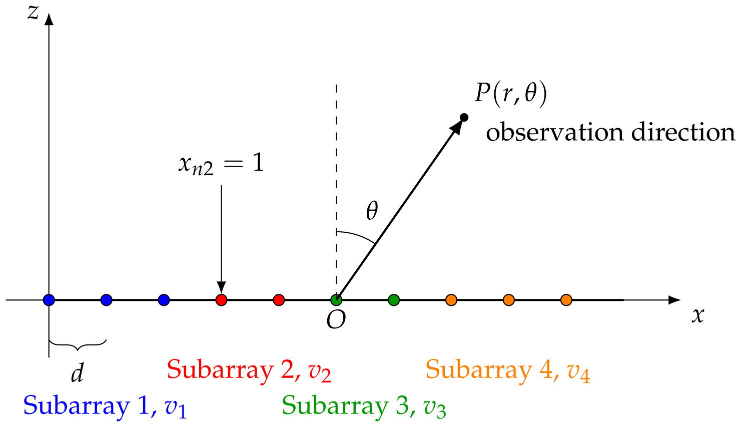

Figure 6. Geometry of the subarray partitioning problem. Binary variables x_{nk} describe the assignment of element n to subarray k.

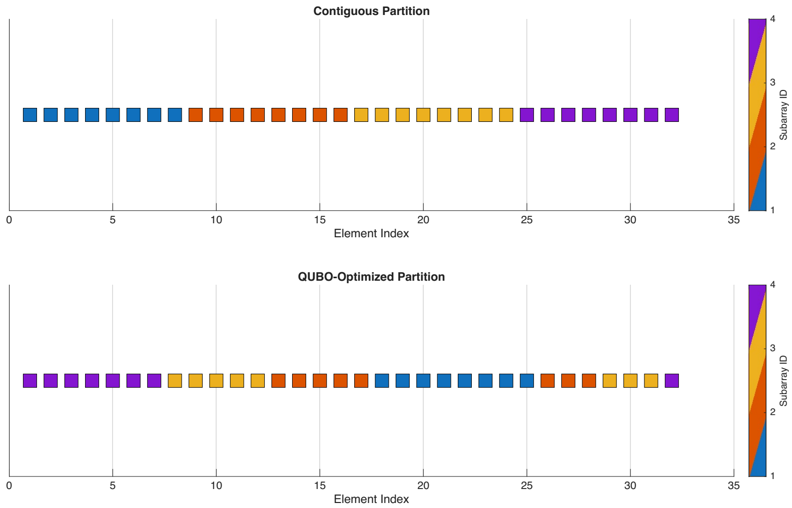

Figure 7. Element-to-subarray assignment for the considered array. (**Top**) conventional contiguous partition, where adjacent elements belong to the same subarray. (**Bottom**) optimized partition obtained through the QUBO solver.

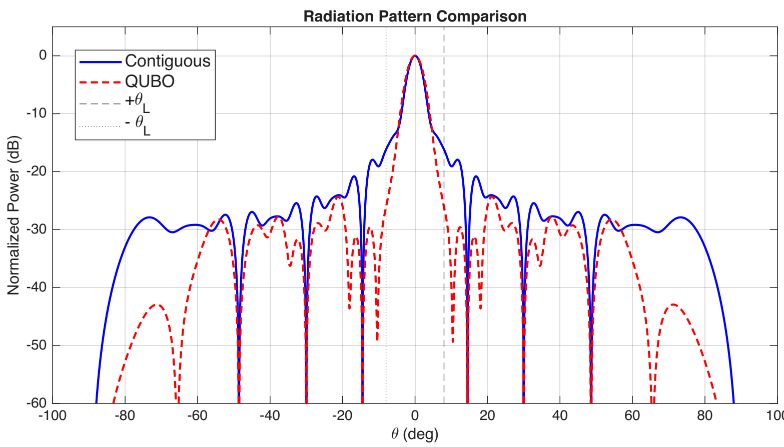

Figure 8. Comparison of the normalized radiation patterns obtained with a conventional contiguous subarray partition (solid blue line) and with the partition obtained through the QUBO optimization (red dashed line). The vertical dashed lines indicate the main-lobe region $|\theta| \leq \theta_L$.

4.4. Electromagnetic Inverse Scattering

Electromagnetic inverse scattering aims to reconstruct the unknown spatial distribution of scattering objects from measurements of the scattered field. Each cell is characterized by a binary variable $x_i \in \{0,1\}$, where $x_i = 1$ indicates the presence of a scatterer. The investigation domain spans a square region of physical aperture $D \times D = 2.5\lambda \times 2.5\lambda$, discretized into a 10×10 grid where each cell has a linear dimension of $\Delta x = \Delta y = 0.25\lambda$. The $M = 64$ measurement points are distributed along a circular probing trajectory enclosing the scattering domain, providing full angular spectral support. Under the Born approximation, the forward model is the linear system

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}, \tag{21}$$

where $\mathbf{y} \in \mathbb{C}^M$ collects the $M = 64$ receiver measurements, $\mathbf{n} \sim \mathcal{CN}(\mathbf{0}, \sigma^2\mathbf{I})$ is complex Gaussian noise with $\sigma = 0.5$ (SNR = 22.5 dB), and $\mathbf{A} \in \mathbb{C}^{M\times N}$ is the sensing matrix whose entries are given by the scalar 2D Green's function [30] modulated by the incident plane wave:

$$A_{mn} = \frac{e^{-jk_0 r_{mn}}}{\sqrt{r_{mn}}}\, e^{-jk_0 y_n}. \tag{22}$$

Since $\mathrm{rank}(\mathbf{A}) = 64 < N = 100$, the problem is severely underdetermined. The regularized cost function is:

$$C(\mathbf{x}) = \|\mathbf{A}\mathbf{x} - \mathbf{y}\|^2 + \alpha_{\text{tik}}\|\mathbf{x}\|^2 + \lambda_{\text{card}}\left(\sum_{i=1}^{N} x_i - K\right)^2 - \sum_{\langle i,j\rangle} \beta_{ij}\, x_i x_j, \tag{23}$$

where the four terms account for data fidelity, Tikhonov regularization ($\alpha_{\text{tik}} = 0.5$), enforcement of the known sparsity level $K = 14$ ($\lambda_{\text{card}} = 3.0$), and a spatial compactness prior ($\beta_{ij} = 2.0$ for direct neighbors, $\beta_{ij} = 1.0$ for diagonal neighbors). Exploiting the binary identity $x_i^2 = x_i$, all linear terms are absorbed into the diagonal of the quadratic matrix, and the problem is cast in standard QUBO form $\min_{\mathbf{x}\in\{0,1\}^N} \mathbf{x}^T Q\mathbf{x}$.

The QUBO is solved via simulated annealing with $n = 15$ independent restarts of 8×10^5 iterations each, using swap moves that preserve the cardinality constraint exactly at every step. An important operational distinction must be noted regarding this last constraint. Because the classical solver utilizes localized swap moves that strictly preserve the number of active scatterers at each step, $\sum_i x_i = K$ remains invariant. As a result, since $x_i^2 = x_i$ for binary variables, the Tikhonov term reduces to the constant $\alpha_{\text{tik}}K$ and the cardinality term $(\sum_i x_i - K)^2$ vanishes identically; neither contributes to the energy differences driving the classical optimization path, which is therefore guided solely by data fidelity and spatial compactness. Conversely, on physical quantum annealing hardware, native operations are based on independent qubit transitions that do not conserve particle number. For an actual QPU deployment, the cardinality penalty parameter (λ_{card}) is mandatory, as it establishes the physical energy barriers required to dynamically restrict the spin evolution to the feasible solution manifold.

This last constraint requires an important operational remark, because it is enforced very differently in the two settings. In the classical solver, the search uses localized swap moves that conserve the number of active scatterers at every step, so $\sum_i x_i = K$ stays fixed throughout. Two simplifications follow. Since $x_i^2 = x_i$ for binary variables, the Tikhonov term equals the constant $\alpha_{\text{tik}}K$, and the cardinality term $(\sum_i x_i - K)^2$ is identically zero. Neither term varies between the configurations the solver visits, so neither affects the energy differences that guide the search; the classical optimization is driven purely by data fidelity and spatial compactness. A physical quantum annealer behaves differently:

its native operations flip qubits independently and do not conserve particle number, so the search is not confined to fixed-cardinality configurations. In this case the cardinality penalty λ_{card} is no longer optional but mandatory, since it is the only mechanism that raises the energy of infeasible configurations and thereby places the ground state on the feasible solution manifold.

Figure 9 shows the ground truth ($K = 14$ active pixels in a non-convex arrangement) and the QUBO reconstruction. In this controlled test case, the solver achieves a perfect reconstruction (F1 = 1.00) at SNR = 22.5 dB. It should be noted that this result is obtained under favorable conditions: the sparsity level K is assumed known, and the clustering prior implicitly reflects the compact nature of the target. Performance under unknown sparsity, more complex geometries, or lower SNR regimes remains a subject for future investigation. This example illustrates how electromagnetic domain knowledge—such as sparsity, spatial compactness, and the structure of the Green's function—can be embedded directly into the QUBO energy landscape, without requiring any post-processing to enforce discrete constraints.

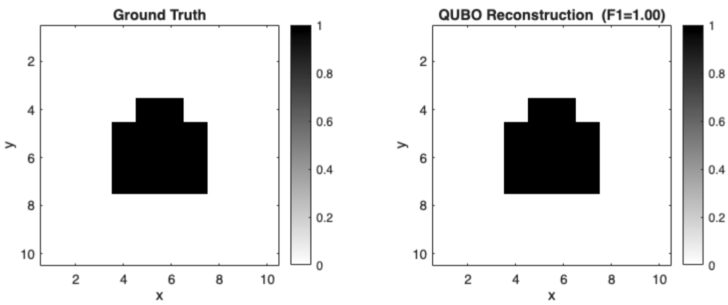

Figure 9. Reconstruction of a source distribution using QUBO.

It must be stressed that the current numerical experiment represents a pedagogical test case where the target aligns perfectly with the inversion grid. In real-world inverse scattering applications, such an assumption is far too restrictive. The choice of grid discretization directly dictates the ill-conditioning of the problem. Indeed, a central challenge in inverse configurations is properly matching the discretization layout with the true number of degrees of freedom of the problem to be solved [31], thereby mitigating severe ill-conditioning. Several advanced techniques have been proposed to handle this issue, most notably the Iterative Multi-Scaling Approach (IMSA) and adaptive multi-resolution frameworks, which initiate the inversion from a coarse grid and progressively refine the mesh in subsequent steps [32]. The implementation of these multi-scaling frameworks is beyond the scope of this paper, which limits its focus to foundational, pedagogical examples.

It also must be explicitly acknowledged that this numerical validation is conducted under highly managed ideal bounds, specifically featuring a well-defined noise threshold (SNR = 22.5 dB) and a known target sparsity constraint ($K = 14$). In highly practical engineering environments, the true sparsity profile is rarely available a priori, and lower signal-to-noise ratios can severely degrade data fidelity. Operating under unknown sparsity or compromised noise regimes will require the dynamic tuning of the penalty multipliers (λ_{card} and α_{tik}) to prevent the solver from collapsing into trivial or highly infeasible regions, representing an essential direction for true applications.

5. Conclusions

The unified QUBO formulation presented in this paper provides a systematic bridge between electromagnetic optimization problems and quantum annealing architectures. Conceptually, this mapping reveals a profound connection between the physics of electromagnetic wave interference and the dynamics of interacting spin systems. While the numerical experiments in this work leverage classical Simulated Annealing as a verification proxy, the exactness of the derived QUBO formulations means they can be submitted to a physical Quantum Processing Unit (QPU) without algebraic modification. Transitioning to real hardware will require careful management of physical constraints, such as minor-embedding overhead and chain strength selection under constrained topologies. Validating these exact formulations on a physical quantum annealer remains a direct and necessary follow-up step to benchmark performance against non-ideal hardware noise floors. In this respect, it should be emphasized that Simulated Annealing, employed here purely as a classical proxy, validates the correctness of the QUBO formulation but does not reproduce the distinctive mechanisms of quantum annealing: it relies on thermal hopping rather than quantum tunneling, and it does not capture hardware-specific challenges such as minor embedding, chain-strength selection, and finite-precision biases. Accordingly, the SA results reported here should be interpreted as a verification of the problem mapping rather than as an estimate of the true potential of QA, which can only be assessed on physical hardware. It is also worth stressing the critical role played by the cardinality constraint across the antenna applications considered in this work: array thinning, subarray partitioning, and sparse inverse reconstruction all rely on a prescribed number of active elements or scatterers to define a physically valid solution. As discussed above, on a quantum annealer this constraint cannot be guaranteed by the search dynamics and must be encoded as a penalty term in the problem Hamiltonian, so that its proper weighting becomes essential to keep the sampled low-energy states within the feasible solution manifold.

Although Simulated Annealing (SA) provides satisfactory results for the pedagogical examples discussed in this work, the primary advantage of Quantum Annealing (QA) lies in its unique ability to explore a configuration space that scales exponentially with the number of qubits. Such spaces become computationally prohibitive for classical heuristics when the energy landscape is characterized by high and narrow potential barriers. While these barriers are difficult to overcome through classical thermal fluctuations, they can be effectively bypassed via quantum tunneling. Despite this theoretical alignment, current hardware remains limited by environmental noise and sparse physical connectivity, often necessitating minor embedding techniques that consume additional qubit resources. Consequently, the near-term future suggests a hybrid paradigm where quantum annealers serve as specialized probabilistic samplers within broader classical workflows. As annealing hardware continues to evolve toward higher connectivity and reduced noise floors, electromagnetics is poised to become a significant application domain for quantum-assisted optimization.

Funding: This was supported in part by the PRIN2022 P20224NMCF—TERRAIN, and by the European Union—NextGenerationEU under the project NRRP RESTART, RESearch and innovation on future Telecommunications systems and networks, to make Italy more smART PE 00000001 (CUP D43C22003080001).

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Desurvire, E. *Classical and Quantum Information Theory: An Introduction for the Telecom Scientist*; Cambridge University Press: Cambridge, UK, 2009.
2. Tabi, Z.I. Evaluation of Quantum Annealer Performance via the Massive MIMO Problem. *IEEE Access* **2021**, *9*, 131658–131671. [CrossRef]
3. Migliore, M.D. Quantum computing for antennas and propagation problems: A gentle introduction. *IEEE Antennas Propag. Mag.* **2024**, *67*, 10–18. [CrossRef]

4. Lee, S.; Lim, Q.J.; Ross, C.; Lee, E.; Han, S.; Kim, Y.; Peng, Z.; Kim, S. Quantum annealing for electromagnetic engineers—Part I. *IEEE Antennas Propag. Mag.* **2024**, *67*, 29–38. [CrossRef]
5. Lee, S.; Lim, Q.J.; Ross, C.; Lee, E.; Han, S.; Kim, Y.; Peng, Z.; Kim, S. Quantum annealing for electromagnetic engineers—Part II. *IEEE Antennas Propag. Mag.* **2025**, *68*, 37–46. [CrossRef]
6. Migliore, M.D. Classical and quantum processing in the deep physical layer. *IEEE Access* **2023**, *11*, 52969–52982. [CrossRef]
7. Tosi, L.; Rocca, P.; Anselmi, N.; Massa, A. Array-antenna power-pattern analysis through quantum computing. *IEEE Trans. Antennas Propag.* **2023**, *71*, 3251–3259. [CrossRef]
8. Lim, Q.J.; Ross, C.; Ghosh, A.; Vook, F.W.; Gradoni, G.; Peng, Z. Quantum-assisted combinatorial optimization for reconfigurable intelligent surfaces. *IEEE Trans. Antennas Propag.* **2023**, *72*, 147–159. [CrossRef]
9. Rocca, P. Antenna array thinning through quantum Fourier transform. *IEEE Access* **2021**, *9*, 124313–124323. [CrossRef]
10. Krikidis, I.; Psomas, C.; Singh, A.K.; Jamieson, K. Optimizing configuration selection in reconfigurable-antenna MIMO systems. *IEEE Trans. Commun.* **2024**, *72*, 8010–8023. [CrossRef]
11. Ross, C. Engineering reflective metasurfaces with Ising Hamiltonian and quantum annealing. *IEEE Trans. Antennas Propag.* **2021**, *70*, 2841–2854. [CrossRef]
12. McMahon, D. *Quantum Computing Explained*; John Wiley & Sons: Hoboken, NJ, USA, 2008.
13. Farhi, E.; Goldstone, J.; Gutmann, S.; Lapan, J.; Lundgren, A.; Preda, D. A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. *Science* **2001**, *292*, 472–475. [CrossRef] [PubMed]
14. Johnson, M.W.; Amin, M.H.S.; Gildert, S.; Lanting, T.; Hamze, F.; Dickson, N.; Harris, R.; Berkley, A.J.; Johansson, J.; Bunyk, P.; et al. Quantum annealing with manufactured spins. *Nature* **2011**, *473*, 194–198. [CrossRef]
15. Yarkoni, S.; Raponi, E.; Bäck, T.; Schmitt, S. Quantum annealing for industry applications: Introduction and review. *Rep. Prog. Phys.* **2022**, *85*, 104001. [CrossRef]
16. Boixo, S.; Rønnow, T.F.; Isakov, S.V.; Wang, Z.; Wecker, D.; Lidar, D.A.; Martinis, J.M.; Troyer, M. Evidence for quantum annealing with more than one hundred qubits. *Nat. Phys.* **2014**, *10*, 218–224. [CrossRef]
17. Denchev, V.S.; Boixo, S.; Isakov, S.V.; Ding, N.; Babbush, R.; Smelyanskiy, V.; Martinis, J.; Neven, H. What is the computational value of finite-range tunneling? *Phys. Rev. X* **2016**, *6*, 031015. [CrossRef]
18. Kadowaki, T.; Nishimori, H. Quantum annealing in the transverse Ising model. *Phys. Rev. E* **1998**, *58*, 5355–5363. [CrossRef]
19. Punnen, A.P. (Ed.) *The Quadratic Unconstrained Binary Optimization Problem: Theory, Algorithms, and Applications*; Springer: Cham, Switzerland, 2022.
20. Dunning, I.; Gupta, S.; Silberholz, J. What Works Best When? A Systematic Evaluation of Heuristics for Max-Cut and QUBO. *Inf. J. Comput.* **2018**, *30*, 608–624. [CrossRef]
21. Goto, H.; Tatsumura, K.; Dixon, A. R. Combinatorial optimization by simulating adiabatic bifurcations in nonlinear Hamiltonian systems. *Sci. Adv.* **2019**, *5*, eaav2372. [CrossRef]
22. Goto, H. High-performance combinatorial optimization based on classical mechanics. *Sci. Adv.* **2021**, *7*, eabe7953. [CrossRef]
23. Pawłowski, J.; Tarasiuk, P.; Tuziemski, J.; Pawela, Ł.; Gardas, B. Simulated Bifurcation Quantum Annealing. *arXiv* **2026**, arXiv:2604.01050. [CrossRef]
24. McGeoch, C.C.; Farré, P. *The D-Wave Advantage System: An Overview*; Technical Report; D-Wave Systems Inc.: Boca Raton, FL, USA, 2020.
25. Boothby, K.; Bunyk, P.; Raymond, J.; Roy, A. *Next-Generation Topology of D-Wave Quantum Processors*; Technical Report; D-Wave Systems Inc.: Boca Raton, FL, USA, 2020.
26. Lebret, H.; Boyd, S. Antenna array pattern synthesis via convex optimization. *IEEE Trans. Signal Process.* **1997**, *45*, 526–532. [CrossRef] [PubMed]
27. Isernia, T.; Bucci, O.M.; Fiorentino, N. Shaped beam antenna synthesis problems: Feasibility criteria and new strategies. *J. Electromagn. Waves Appl.* **1998**, *12*, 103–138. [CrossRef]
28. Basar, E.; Di Renzo, M.; de Rosny, J.; Debbah, M.; Alouini, M.-S.; Zhang, R. Wireless communications through reconfigurable intelligent surfaces. *IEEE Access* **2019**, *7*, 116753–116773. [CrossRef]
29. Mailloux, R.J. *Phased Array Antenna Handbook*, 2nd ed.; Artech House: Norwood, MA, USA, 2005.
30. Chew, W.C. *Waves and Fields in Inhomogeneous Media*; IEEE Press: New York, NY, USA, 2001.
31. Bucci, O.M.; Isernia, T. Electromagnetic inverse scattering: Retrievable information and measurement strategies. *Radio Sci.* **1997**, *32*, 2123–2137. [CrossRef]
32. Donelli, M.; Massa, A. An integrated multiscaling strategy based on a particle swarm algorithm for inverse scattering problems. *IEEE Trans. Antennas Propag.* **2005**, *53*, 861–869. [CrossRef]