MOX-Report No. 69/2026

Complexity-controlled QUBO formulation for topology optimization via anisotropic mesh adaptation

Zecchi, A. A.; Ferro, N.; Perotto, S.

MOX, Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

mox-dmat@polimi.it https://mox.polimi.it

Complexity-controlled QUBO formulation for topology optimization via anisotropic mesh adaptation

Alessandro Andrea Zecchi[1], Nicola Ferro[2] and Simona Perotto[1]

September 4, 2026

[1]MOX - Dipartimento di Matematica, Politecnico di Milano
Piazza L. da Vinci, 32, 20133 Milano, Italy
{alessandroandrea.zecchi, simona.perotto}@polimi.it

[2]Dipartimento di Scienze Molecolari e Nanosistemi, Università Ca' Foscari
Via Torino, 155, 30170 Venezia Mestre, Italy
nicola.ferro@unive.it

Abstract

QUBO formulation provides a natural route to address topology optimization of continuum structures with classical heuristic, quantum-inspired, and quantum annealing solvers. Their practical applicability, however, is still limited by the number of binary variables that can be handled, especially on near-term quantum hardware. In this work, we tackle this bottleneck by proposing QUBOSIMPATY algorithm which combines a QUBO formulation of the SIMP minimum compliance problem with anisotropic recovery-based mesh adaptation, enriched with metric-based complexity control. Specifically, we exploit anisotropic mesh adaptation to reduce the number of density degrees of freedom, while retaining directional resolution along the sharp material-void interfaces that characterize optimized layouts. A metric rescaling procedure is applied to prescribe the number of mesh vertices, and hence the number of binary variables, so as to comply with the severe size constraints imposed by quantum annealing hardware. Numerical tests in two and three dimensions are carried out with classical QUBO heuristics, including simulated annealing and tabu search. Moreover, a proof-of-concept experiment is performed on quantum annealing hardware. The results show that the proposed strategy improves the suppression of intermediate densities while keeping the optimization problem within the admissible size range.

1 Introduction

Topology optimization is a rigorous mathematical framework for the optimal distribution of material within a prescribed design domain, subject to governing physical laws, boundary conditions, loads, and design constraints [1]. In structural mechanics, one of the most representative formulations is the minimum compliance problem, where the goal is to maximize stiffness, or equivalently minimize compliance, under a constraint on the available material volume.

This paradigm has become a standard tool for the design of lightweight and high-performance structures, with applications ranging from mechanical and aerospace engineering to additive manufacturing and sustainable product design [2, 3].
In most computational settings, topology optimization is addressed through density-based methods. Among them, the solid isotropic material with penalization (SIMP) approach is arguably one of the most widely used strategies [4, 1, 5]. In SIMP, the original binary material-void formulation is relaxed by introducing a fictitious density variable, which interpolates between absence and presence of material and allows one to exploit continuous optimization techniques. The resulting optimization problem is typically solved by gradient-based methods, coupled with a finite element discretization of the governing elasticity equations. This combination has proved extremely effective, but the underlying problem remains, at its core, a discrete design problem. Indeed, after discretization of the design domain, each element or degree of freedom is associated with the binary choice of placing material or void. From this viewpoint, topology optimization can be interpreted as a nonlinear mixed 0-1 optimization problem. This observation has motivated several attempts to address topology optimization through genuinely discrete or combinatorial formulations, including alternative density interpolation schemes [6] and decomposition-based methods for the direct treatment of the binary problem [7].
In the context of combinatorial optimizations, the quadratic unconstrained binary optimization (QUBO) model has emerged as a particularly relevant formulation [8]. Many classical combinatorial problems, such as maximum cut and graph coloring, can be reformulated in QUBO form [9], which consists in minimizing a quadratic objective function over binary variables and is known to be NP-hard. A further reason for the current interest in QUBO formulations is their formal equivalence to Ising models from statistical physics. This equivalence provides a direct connection between binary optimization and spin systems, making QUBO problems naturally amenable to quantum annealing, quantum-inspired methods, and specialized classical heuristics.

The use of QUBO and quantum-compatible formulations in computational mechanics is an active research direction, with topology optimization representing a particularly appealing benchmark setting (see, e.g., [10, 11]). Indeed, the presence of many local minima, together with the binary nature of the ideal material distribution, makes topology optimization a natural candidate for combinatorial approaches. Recent works have investigated QUBO or quantum-annealing strategies for topology optimization along different directions [12, 13, 14, 15]. For instance, some methods introduce binary design updates within a SIMP-like optimization loop, thus replacing the continuous update step by the solution of a QUBO subproblem. Other approaches avoid the standard SIMP sensitivity-analysis framework and formulate the design variables directly in binary form. Splitting or decomposition strategies have also been considered, where the original topology optimization problem is recast as a sequence of smaller mixed-integer or binary optimization problems amenable to QUBO-based or quantum-annealing solvers.
Regardless of the specific strategy adopted, a common bottleneck is the growth of the number of binary variables with the discretization size, typically arising in application contexts. This issue is particularly critical for near-term quantum devices, where the number of available qubits and the hardware connectivity impose severe restrictions on the admissible problem size [16, 17]. In topology optimization, this limitation has a direct impact on the design process, since the binary variables are associated with the design degrees of freedom. Working with too few variables

may compromise the ability to resolve relevant structural features, such as thin members, holes, or sharp material-void interfaces, ultimately affecting both the quality and the reliability of the optimized layout. Thus, controlling the number of design degrees of freedom becomes a key requirement to make QUBO formulations compatible with current solvers, especially from a quantum-computing perspective, without excessively degrading the expressive capability of the design space.

This work addresses this issue by exploiting anisotropic mesh adaptation as a complexity-control tool for QUBO-based topology optimization. Anisotropic adaptation has long been used to improve finite element accuracy by modifying the size, orientation, and shape of mesh elements according to the directional features of the numerical solution [18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29]. This is particularly relevant in topology optimization, where optimized layouts typically exhibit sharp material-void interfaces. This feature is exploited in the SIMPATY algorithm [27, 30, 31], which combines the SIMP method with anisotropic mesh adaptation. Indeed, by aligning stretched elements with such interfaces, anisotropic meshes can preserve an accurate description of the design while reducing the number of density degrees of freedom. In the present setting, this classical property takes on a quantum-oriented role: since each density degree of freedom is associated with a binary variable in the QUBO formulation, reducing the design complexity makes the resulting optimization problem more compatible with the limited problem sizes currently accessible to quantum-computing hardware.

To match this goal, we set a two-phase process. First, we recast the SIMP minimum compliance problem into a QUBO framework, while retaining the information provided by the standard compliance sensitivities. This is motivated by the fact that sensitivity information is inexpensive to compute in density-based topology optimization and plays a key role in avoiding the slow convergence often observed in non-gradient approaches [32]. Second, we introduce anisotropic recovery-based mesh adaptation, combined with a metric rescaling strategy, as a complexity-control mechanism for QUBO-based topology optimization. The metric rescaling allows one to prescribe or bound the number of mesh vertices, and hence the number of binary variables, so that the generated QUBO instances remain compatible with the admissible problem size of the selected solver.

The resulting algorithm, named QUBOSIMPATY, combines SIMP-based topology optimization, additive binary density updates, QUBO solvers, anisotropic mesh adaptation, and metric-based complexity control. The QUBO subproblems can be solved by different classes of samplers, including classical heuristics such as simulated annealing and tabu search, as well as quantum annealing. Numerical experiments in two and three dimensions are used to assess the proposed formulation and to compare the resulting designs with those obtained by standard topology optimization procedures. A proof-of-concept experiment on D-Wave quantum annealing hardware is also included to validate the overall pipeline in a real quantum computing setting.

The paper is organized as follows. Section 2 recalls the basic QUBO formulation and the main classes of solvers considered in this work, including classical heuristics and quantum annealing. Section 3 introduces the SIMP minimum compliance problem, while in section 4 we derive the associate QUBO formulation. Section 5 presents the anisotropic mesh adaptation strategy, the metric rescaling procedure used to control the number of design degrees of freedom, and the overall QUBOSIMPATY algorithm. Section 6 reports numerical experiments on benchmark test

cases, while section 7 summarizes the main findings and discusses future research directions.

2 Background on Ising and QUBO models

The Ising model is a mathematical framework for magnetism in statistical mechanics [33, 34], which describes a lattice of N spin variables, interacting with an external magnetic field and with neighboring sites. The energy of the system is represented by the classical Ising Hamiltonian, H_p, defined as

$$H_p(s_1, \ldots, s_N) = -\sum_{i=1}^{N} h_i s_i - \sum_{i=1}^{N} \sum_{j=i+1}^{N} J_{ij} s_i s_j, \tag{1}$$

where $s_i \in \{-1, +1\}$ is the $i-$th spin variable, $h_i \in \mathbb{R}$ represents the magnetic field interaction exerted on s_i, and $J_{ij} \in \mathbb{R}$ is the interaction strength between s_i and s_j, for $i, j = 1, \ldots, N$. To rewrite H_p in its quantum counterpart on the Hilbert space $\mathcal{H} = (\mathbb{C}^2)^{\otimes N}$ of an N-qubit system, as is customary in the quantum optimization literature, each classical spin variable s_i is replaced by the Pauli-z operator $\sigma_{i,z}$. This commutes Hamiltonian H_p into

$$\widehat{H}_p = -\sum_{i=1}^{N} h_i \sigma_{i,z} - \sum_{i=1}^{N} \sum_{j=i+1}^{N} J_{ij} \sigma_{i,z} \sigma_{j,z}.$$

The ground state of $\widehat{H}_p$ corresponds to a spin assignment attaining the minimum classical energy in (1). Computing ground states of Ising Hamiltonians is relevant in several applications, including Max-Cut problems, where nodes are partitioned so as to maximize the weight of cut edges; satisfiability problems, where one seeks a Boolean assignment satisfying a set of logical constraints; and scheduling tasks, where resources or operations must be arranged under prescribed constraints [9, 8, 35, 36].

The minimization of (1) can be reformulated into a QUBO form through the variable transformation $x_i = (1 + s_i)/2$, where $x_i \in \{0, 1\}$ is a binary variable. More in detail, QUBO is an unconstrained combinatorial optimization problem defined as the minimization of a quadratic objective function f, given by

$$\min_{\mathbf{x} \in \{0,1\}^N} f(\mathbf{x}) = \min_{\mathbf{x} \in \{0,1\}^N} \mathbf{x}^T \mathrm{Q} \mathbf{x}, \tag{2}$$

where $\mathbf{x} = [x_1, \ldots, x_N]^T \in \{0, 1\}^N$ is a vector of binary variables and Q is a square upper triangular[1] real matrix of order N, with generic entry $[\mathrm{Q}]_{i,j} = q_{ij}$. Using the identity $x_i = x_i^2$ for binary variables, the objective function can be decomposed into linear and quadratic contributions as

$$\min_{\mathbf{x} \in \{0,1\}^N} f(\mathbf{x}) = \min_{\mathbf{x} \in \{0,1\}^N} \left(\sum_{i=1}^{N} q_{ii} x_i + \sum_{i=1}^{N} \sum_{j=i+1}^{N} q_{ij} x_i x_j \right), \tag{3}$$

[1] Other equivalent formulations rely on a symmetric matrix Q, which may provide computational advantages in the solution of the minimization problem [37].

this representing the well-known QUBO formulation.
The QUBO problem is widely used in scheduling, routing, finance, machine learning and numerical quantum computing [38, 39, 40, 41]. In its general form, QUBO is NP-hard and no polynomial-time algorithm is known for solving it efficiently [37]. Consequently, practical approaches typically rely on high-quality approximate solutions. These methods can be broadly divided into classical heuristics and quantum or quantum-inspired algorithms.
While quantum computing offers a promising route for tackling QUBO problems [42], current near-term devices still impose severe limitations in terms of available qubits and connectivity. For instance, recent advances in simulating spin-glass dynamics involve up to a few thousand qubits [43], but remain constrained by the fixed topologies of the underlying quantum processors.

2.1 Classical heuristics

Classical heuristic algorithms are often the default choice for large-scale practical QUBO problems, as they are well-established and widely available. The most employed algorithms are simulated annealing and tabu search [44], in addition to general purpose optimization software and other metaheuristics [37].

The simulated annealing algorithm consists in the repeated stochastic selection of a neighbor $\mathbf{x}_{new}$ of the current state $\mathbf{x}$. Then, by employing an acceptance probability P that depends on a fictitious temperature T and on the evaluation of the cost function f, the system either evolves into the new state or stays at $\mathbf{x}$. This procedure is inspired by metallurgical annealing, and it mimics the process of heating and controlled cooling of metal to reach a low-energy, ordered state, avoiding local minima by occasionally accepting worse solutions. Typically, the move step is repeated until the system reaches a state that is good enough for the application, or until a given computation budget has been exhausted. The algorithm ultimately lead the system to move to states of lower energy.

The tabu search procedure is a heuristic search method which employs a random neighbor selection. Differently from simulated annealing, this optimization method also uses a tabu list of recently explored states that it does not revisit in order to escape local minima.
In general, the tabu search algorithm performs better in higher dimensional scenario with respect to simulated annealing [45]. This justifies the adoption of tabu search as the primary tool for the solution of the QUBO problems in section 6.

2.2 Quantum and quantum-inspired algorithms

Quantum-based QUBO solvers are primarily represented by two paradigms: the quantum approximate optimization algorithm and quantum annealing.
Qantum approximate optimization algorithm operates within the gate-based model of quantum computation, but its current applicability is restricted by the limited scale and reliability of available quantum devices.
Quantum annealing offers a more direct physical route to QUBO optimization. The problem is encoded into an energy landscape, so that solving the QUBO formulation (2)-(3) corresponds to searching for the ground state of the Ising Hamiltonian in (1) [46]. This principle is implemented by quantum annealers, specialized devices designed to perform optimization through annealing

dynamics [47]. The approach is theoretically motivated by the adiabatic theorem, according to which a quantum system can remain close to its instantaneous ground state when its Hamiltonian evolves sufficiently slowly and a nonzero spectral gap is preserved [48].

Alongside quantum hardware-based methods, quantum-inspired approaches have also been developed. A notable example is digital annealing, which combines dedicated digital circuit design with physics-inspired optimization mechanisms to escape local minima by emulating selected features of quantum dynamics [49].

3 Topology optimization with the SIMP method

Topology optimization aims at determining the optimal distribution of material within a prescribed design domain by minimizing an objective function under suitable physical and design constraints.

We consider a design domain $\Omega \subset \mathbb{R}^d$, with $d \in \{2, 3\}$, in a small displacement regime, such that the linear elasticity equation

$$-\nabla \cdot \sigma(\mathbf{u}) = \mathbf{0} \qquad \text{in } \Omega \tag{4}$$

governs the structure statics, with $\mathbf{u}$ the structure displacement and where $\sigma(\mathbf{u}) = 2\mu\varepsilon(\mathbf{u}) + \lambda \mathrm{tr}(\varepsilon(\mathbf{u}))\mathrm{I}$ is the stress tensor, with $\varepsilon(\mathbf{u}) = (\nabla\mathbf{u} + \nabla\mathbf{u}^T)/2$ the strain, I the identity tensor, $\mathrm{tr}(\cdot)$ the trace operator,

$$\lambda = \begin{cases} \dfrac{E\nu}{1-\nu^2} & \text{if } d = 2 \\ \dfrac{E\nu}{(1+\nu)(1-2\nu)} & \text{if } d = 3 \end{cases} \qquad \mu = \frac{E}{2(1+\nu)} \tag{5}$$

the Lamé coefficients depending on the Young's modulus E and on the Poisson's ratio ν of the material contained in Ω.

A standard setting in topology optimization is the minimization of the structural compliance under a prescribed volume fraction constraint, when a traction $\mathbf{f}$ is applied on the Neumann portion Γ_N of the boundary $\partial\Omega$ of Ω. To this aim, we adopt the widely used density-based SIMP approach, where an auxiliary density field $\rho \in L^\infty(\Omega)$ describes the material distribution in Ω [1]. Ideally, ρ should take only the values 0, corresponding to void, and 1, corresponding to full material. In practice, however, intermediate values in $[0, 1]$ may occur and are penalized through a power law ρ^p, with $p > 1$.

The density field enters the mechanical model by modifying the stress tensor as

$$\sigma_\rho(\mathbf{u}) = \rho^p \left[2\mu\varepsilon(\mathbf{u}) + \lambda\, \mathrm{tr}(\varepsilon(\mathbf{u}))\mathrm{I}\right].$$

The SIMP formulation of the minimum compliance problem then reads

$$\min_{\rho \in L^\infty(\Omega)} c(\mathbf{u}(\rho)) : \begin{cases} a_\rho(\mathbf{u}(\rho), \mathbf{v}) = c(\mathbf{v}) \quad \forall \mathbf{v} \in U \\ \displaystyle\int_\Omega \rho \, d\Omega \le \alpha|\Omega| \\ 0 < \underline{\rho} \le \rho \le 1, \end{cases} \tag{6}$$

where

$$c(\mathbf{u}(\rho)) = \int_{\Gamma_N} \mathbf{f} \cdot \mathbf{u}(\rho)\, d\Gamma$$

is the objective function coinciding with the static compliance,

$$a_\rho(\mathbf{u}(\rho), \mathbf{v}) = \int_\Omega \sigma_\rho(\mathbf{u}(\rho)) : \varepsilon(\mathbf{v})\, d\Omega$$

is the SIMP counterpart of the bilinear form associated with (4), for $\mathbf{u} = \mathbf{u}(\rho)$ and $\mathbf{v}$ belonging to a suitable function space U, $\alpha \in (0, 1)$ is a prescribed maximum portion of the volume $|\Omega|$, and $\underline{\rho}$ is a positive lower bound on the density, introduced to avoid degeneracy of the stiffness tensor. Finally, we complete the physical model constraining (6) with the boundary conditions

$$\mathbf{u}(\rho) = \mathbf{0} \quad \text{on } \Gamma_D, \qquad \sigma_\rho(\mathbf{u}(\rho))\mathbf{n} = \begin{cases} \mathbf{f} & \text{on } \Gamma_N \\ \mathbf{0} & \text{on } \Gamma_F, \end{cases}$$

with $\mathbf{n}$ the unit outward normal vector to the boundary, and Γ_D, Γ_N and Γ_F portions of $\partial\Omega$ such that $\partial\Omega = \Gamma_D \cup \Gamma_N \cup \Gamma_F$ and $\mathring{\Gamma}_D \cap \mathring{\Gamma}_N \cap \mathring{\Gamma}_F = \emptyset$. Thus, the function space in (6) is chosen as $U = \{\mathbf{v} \in [H^1(\Omega)]^d : \mathbf{v} = \mathbf{0} \text{ on } \Gamma_D\}$.

3.1 The discrete SIMP formulation

In order to numerically compute the solution to (6), we introduce a conforming computational mesh, $\mathcal{T}_h$, consisting of a collection, $\{K\}$, of simplicial elements such that $\overline{\Omega} = \bigcup_{K \in \mathcal{T}_h} K$, with $\#\mathcal{T}_h$ the mesh cardinality.
We associate with tessellation $\mathcal{T}_h$ the finite element space V_h^r of the piecewise polynomials with degree $r \geq 0$. The displacement and density fields in (6) are approximated in discrete spaces that may differ. Indeed, the discrete SIMP minimum compliance topology optimization problem is

$$\min_{\rho_h \in V_h^r} c(\mathbf{u}_h(\rho_h)) : \begin{cases} a_\rho(\mathbf{u}_h(\rho_h), \mathbf{v}_h) = c(\mathbf{v}_h), \quad \forall \mathbf{v}_h \in [V_h^s]^d \cap U \\ \displaystyle\sum_{i=1}^{N_\rho} \frac{V_i(\rho_i)}{V_0} = 1 \\ 0 < \underline{\rho} \leq \rho_i \leq 1 \quad i = 1, \ldots, N_\rho, \end{cases} \tag{7}$$

with $N_\rho = \dim(V_h^r)$ and $\mathbf{u}_h(\rho_h) \in [V_h^s]^d \cap U$ the finite element displacement. The two design constraints require some care, since their discrete form is not a direct transcription of the corresponding continuous conditions in (6). To this aim, we write the discrete density in terms of the basis $\{\varphi_{h,i}\}_{i=1}^{N_\rho}$ of V_h^r as

$$\rho_h(\mathbf{r}) = \sum_{i=1}^{N_\rho} \rho_i\, \varphi_{h,i}(\mathbf{r}),$$

where $\mathbf{r} \in \Omega$ denotes the spatial variable, and the coefficients $\rho_i \in \mathbb{R}$ are the unknown density degrees of freedom (DOFs), collected in vector $\boldsymbol{\rho} = [\rho_1, \ldots, \rho_{N_\rho}]^T \in \mathbb{R}^{N_\rho}$. By exploiting this

representation, we rewrite the left-hand side of the volume fraction constraint in (6) as

$$\int_\Omega \rho_h \, d\Omega = \sum_{i=1}^{N_\rho} \rho_i \int_\Omega \varphi_{h,i} \, d\Omega = \sum_{i=1}^{N_\rho} \rho_i w_i = \sum_{i=1}^{N_\rho} V_i(\rho_i),$$

with $V_i(\rho_i)$ the material volume associated with the i-th DOF. Weights w_i are also instrumental to redefine the right-hand side of the volume fraction constraint, being

$$V_0 = \alpha|\Omega| = \alpha \sum_{i=1}^{N_\rho} \int_\Omega \varphi_{h,i} \, d\Omega = \alpha \sum_{i=1}^{N_\rho} w_i,$$

this leading to the inequality

$$\sum_{i=1}^{N_\rho} V_i(\rho_i) \leq V_0.$$

We observe that, in this minimum compliance setting, the volume constraint is equivalently written as the equality constraint

$$\sum_{i=1}^{N_\rho} V_i(\rho_i) = V_0,$$

as stated in (7), being the minimum of the structural compliance reached in correspondence with the maximum volume attainable.
Additionally, the discrete notation now adopted leads to commute the two-sided inequality in (6) into the N_ρ discrete controls in (7).

Topology optimization problems are commonly solved by gradient-based iterative algorithms, such as the optimality criteria (OC) method and the method of moving asymptotes (MMA) [50, 1, 51], which generate a sequence of approximate vectors $\{\boldsymbol{\rho}^{(s)}\}_{s\geq 0}$ describing the evolution of the design from the initial guess to the optimized structure. For the minimum compliance problem, these two approaches involve essentially the same class of computational operations, since both rely on a Lagrangian treatment of the constraints and require the update of the associated multipliers through an inner loop. Other optimization strategies include evolutionary approaches [52], and genetic algorithms [53]. Regardless of the chosen optimization strategy, the discrete spaces V_h^r and V_h^s have to be selected carefully, together with suitable filtering techniques, such as Helmholtz and Heaviside filters, to prevent numerical artifacts including checkerboard patterns, staircase effects, and greyness [54, 55].

4 QUBO for topology optimization

We now recast the minimum compliance topology optimization problem in (7) within a QUBO framework. This is achieved by expressing the density DOFs as a linear function of binary variables, so that the optimization step becomes a combinatorial search over admissible configurations.

4.1 Minimum compliance formulation

A QUBO formulation of (7) first requires expressing the compliance as a function of the binary variables in $\mathbf{x}$, with a dependence that is at most quadratic in form, according to (3). This requirement constrains the selection of the penalization exponent p in the definition of the compliance

$$c(\mathbf{u}_h(\rho_h)) = \int_\Omega \rho_h^p \, [2\mu\varepsilon(\mathbf{u}_h(\rho_h)) + \lambda \, \mathrm{tr}(\varepsilon(\mathbf{u}_h(\rho_h)))\mathrm{I}] : \varepsilon(\mathbf{u}_h(\rho_h)) \, d\Omega$$

to be set equal to 2. However, this value may lead to inaccurate design processes in different settings. To get rid of this limitation on the SIMP exponent p, let us focus on a generic SIMP optimization step updating the density vector $\boldsymbol{\rho}^{(k)}$ into $\boldsymbol{\rho}^{(k+1)}$. The compliance is then understood as a function of the density vector, i.e., $c(\mathbf{u}_h(\rho_h)) = c(\boldsymbol{\rho})$, and approximated by a first-order Taylor expansion around $\boldsymbol{\rho}^{(k)}$, namely

$$c(\boldsymbol{\rho}^{(k+1)}) \approx c_1(\boldsymbol{\rho}^{(k+1)}) = c(\boldsymbol{\rho}^{(k)}) + \sum_{i=1}^{N_\rho} \frac{\partial c}{\partial \rho_i}(\boldsymbol{\rho}^{(k)})(\rho_i^{(k+1)} - \rho_i^{(k)}), \tag{8}$$

with $\boldsymbol{\rho}^{(s)}$, for $s = k, \, k+1$, the vector containing the density DOFs associated with the s-th iteration of the optimization process, and $\rho_i^{(s)}$ the corresponding $i-$th component. Concerning the evaluation of the components, $\partial c/\partial \rho_i$, of the sensitivity $\nabla_\rho c$, we resort to a standard Lagrangian approach. If density filters are used, the sensitivities are computed with respect to the filtered densities and then mapped back to the original design variables by applying the chain rule [1].

To express the components of $\boldsymbol{\rho}^{(k+1)}$ in terms of $\boldsymbol{\rho}^{(k)}$ and of the binary design variable $\mathbf{x}$, we consider the generic update rule

$$\boldsymbol{\rho}^{(k+1)} = \boldsymbol{\rho}^{(k+1)}(\mathbf{x}; \boldsymbol{\rho}^{(k)}) \qquad \text{with} \quad \rho_i^{(k+1)} = \rho_i^{(k+1)}(x_i; \rho_i^{(k)}), \quad i = 1, \dots N_\rho. \tag{9}$$

Following [13], a possible choice for such a rule is the multiplicative update scheme

$$\rho_i^{(k+1)}(x_i; \rho_i^{(k)}) = (\theta_a x_i + \theta_b)\rho_i^{(k)} = \begin{cases} (\theta_a + \theta_b)\rho_i^{(k)} & \text{if } x_i = 1 \\ \theta_b \rho_i^{(k)} & \text{if } x_i = 0, \end{cases} \tag{10}$$

where $\theta_a, \theta_b > 0$ are real parameters suitably chosen to ensure a gradual update of the design variable (for instance, $\theta_a = 0.2$ and $\theta_b = 0.9$ increase or decrease $\rho_i^{(k)}$ by 10%, depending on the value of x_i).
However, the update scheme (10) suffers from a major drawback, as density DOFs $\rho_i^{(k)}$ that are zero, or close to zero, cannot be directly upgraded to solid material if needed.
To overcome this issue and improve the stability of the numerical scheme, we employ the additive law

$$\rho_i^{(k+1)}(x_i; \rho_i^{(k)}) = \rho_i^{(k)} - \delta_{a,i} + \delta_{b,i} \, x_i, \tag{11}$$

with

$$\delta_{a,i} = \begin{cases} \delta & \text{if } \rho_i^{(k)} = 1 \\ 0 & \text{if } \rho_i^{(k)} = \underline{\rho} \\ \delta & \text{otherwise,} \end{cases} \qquad \delta_{b,i} = \begin{cases} \delta & \text{if } \rho_i^{(k)} = 1 \\ \delta & \text{if } \rho_i^{(k)} = \underline{\rho} \\ 2\delta & \text{otherwise,} \end{cases} \tag{12}$$

and $\delta > 0$ a real parameter. The update rule in (11) also ensures that the box constraint in (7) is tightly enforced. Moreover, δ may vary across optimization iterations, analogously to trust region optimization methods [56].
Plugging the update scheme (11) into the first-order expansion (8) yields

$$c_1(\mathbf{x}) = c_1\big(\boldsymbol{\rho}^{(k+1)}(\mathbf{x}; \boldsymbol{\rho}^{(k)})\big) = c(\boldsymbol{\rho}^{(k)}) + \sum_{i=1}^{N_\rho} \frac{\partial c}{\partial \rho_i}(\boldsymbol{\rho}^{(k)}) \left(-\delta_{a,i} + \delta_{b,i}\, x_i\right). \tag{13}$$

Thus, at the k-th optimization step, the minimization of the linearized compliance is recast as a binary optimization problem in the variable $\mathbf{x} \in \{0,1\}^{N_\rho}$, rather than in terms of the density field $\rho_h \in V_h^r$.
As a next step to derive the QUBO formulation, we handle the volume constraint through a penalty term added to the objective function, so that, with a straightforward extension of the notation, we have

$$\min_{\mathbf{x}\in\{0,1\}^{N_\rho}} \left[\xi\, c_1(\mathbf{x}) + \zeta\, g(\mathbf{x})\right] \;:\; a_\rho(\mathbf{u}_h(\boldsymbol{\rho}^{(k)}), \mathbf{v}_h) = c(\mathbf{v}_h) \quad \forall \mathbf{v}_h \in [V_h^s]^d \cap U \tag{14}$$

with

$$g(\mathbf{x}) = \left[\sum_{i=1}^{N_\rho} \frac{w_i \rho_i^{(k+1)}(x_i; \rho_i^{(k)})}{V_0} - 1\right]^2 \tag{15}$$

the constraint function, and where ζ is the penalty parameter, while ξ is a scaling factor introduced to make ξc_1 and $\zeta\, g$ comparable in terms of order of magnitude. The values of ζ and ξ are problem-dependent and mutually related (we refer the reader to section 6 for the values adopted in the numerical tests).

To cast problem (14) into the form (2) (i.e., (3)), we need to define the matrix Q, starting from the expressions in (13)-(15). In particular, we observe that the constant terms in the objective function can be neglected, since they do not affect the minimization, so that function c_1 is replaced by

$$\tilde{c}_1(\mathbf{x}) = \sum_{i=1}^{N_\rho} \frac{\partial c}{\partial \rho_i}(\boldsymbol{\rho}^{(k)})\, \delta_{b,i}\, x_i,$$

while g reduces to

$$\tilde{g}(\mathbf{x}) = \frac{1}{V_0^2} \sum_{i=1}^{N_\rho} w_i^2 \delta_{b,i}^2 x_i + \frac{2}{V_0^2} \sum_{i=1}^{N_\rho} \sum_{j=i+1}^{N_\rho} w_i w_j \delta_{b,i} \delta_{b,j} x_i x_j + \frac{2\bar{V}}{V_0} \sum_{i=1}^{N_\rho} w_i \delta_{b,i} x_i,$$

with

$$\bar{V} = \frac{1}{V_0} \sum_{i=1}^{N_\rho} w_i (\rho_i^{(k)} - \delta_{a,i}) - 1.$$

Thus, the entries of matrix Q are given by

$$[\mathrm{Q}]_{i,j} = q_{ij} = \begin{cases} \xi \dfrac{\partial c}{\partial \rho_i}(\boldsymbol{\rho}^{(k)})\, \delta_{b,i} + \zeta \left(\dfrac{1}{V_0^2} w_i^2 \delta_{b,i}^2 + \dfrac{2\bar{V}}{V_0} w_i \delta_{b,i} \right) & \text{if } i = j \\ \zeta \dfrac{2}{V_0^2} w_i w_j \delta_{b,i} \delta_{b,j} & \text{if } j > i \\ 0 & \text{otherwise,} \end{cases} \tag{16}$$

for $i, j = 1, \ldots, N_\rho$.
The minimization problem (14) can thus be equivalently expressed in the QUBO formulation (2) (i.e., (3)) for $N = N_\rho$, additionally subject to the state equation $a_\rho(\mathbf{u}_h(\boldsymbol{\rho}^{(k)}), \mathbf{v}_h) = c(\mathbf{v}_h)$ for any $\mathbf{v}_h \in [V_h^s]^d \cap U$.

4.2 Implementation aspects

We now discuss some computational aspects featuring the QUBO formulation associated with the minimum compliance setting.
The matrix Q is stored in upper-triangular form, with most pairwise interaction coefficients q_{ij}, for $j > i$, being nonzero. This structure affects the computational cost of the QUBO solution process and becomes particularly critical for near-term quantum annealers, where it increases the complexity of the embedding on the quantum processing unit. Indeed, the all-to-all connectivity induced by the QUBO matrix must be mapped onto the sparse topology of the quantum processing unit, through minor embedding [57], which represents logical variables as chains of coupled physical qubits with consistent measurement outcomes. The use of such chains increases the number of physical qubits needed to represent the problem and therefore limits the maximum attainable problem size.
Another relevant issue concerns the numerical range and precision of the QUBO coefficients. QUBO instances with high-precision linear and quadratic coefficients are challenging for quantum annealing hardware, due to the limited coefficient precision that can be reliably implemented in the presence of control errors [58]. For this reason, values h_i and J_{ij} of the equivalent Ising formulation (1), have to lie within the admissible hardware range, typically $[-1, 1]$, while remaining sufficiently large in magnitude with respect to the device precision. Consequently, before being submitted to the hardware, matrix Q has to be rescaled by a suitable global factor. Thus, the choice of the scaling parameter ξ and of the penalty parameter ζ turns out to be relevant for balancing the objective and the constraint violation, as well as for obtaining a numerically meaningful hardware implementation.
We finally remark that the pattern of Q becomes less restrictive for other classes of QUBO solvers, such as digital annealers, simulated annealing, tabu search, or general-purpose classical optimization algorithms, where no minor embedding onto a sparse quantum architecture is required.

5 Anisotropic recovery-based mesh adaptation

Anisotropic mesh adaptation provides a framework to iteratively modify a computational grid by adjusting the size, orientation, and shape of its elements according to directional features of the numerical solution, typically guided by mathematically grounded error estimators. Successful applications have been reported in several areas, including ocean modelling, ice-flow simulations, biomedical and aerospace engineering, seismic wave propagation, material design, and sustainability-oriented applications [59, 60, 61, 62, 63, 64, 65, 66].
Unlike isotropic adaptation, which only controls the element size while preserving a fixed element shape, anisotropic adaptation offers greater flexibility and efficiency, either maximizing the solution accuracy for a fixed number of elements or reducing the number of elements needed to achieve a prescribed tolerance. This controlled reduction in the number of degrees of freedom is particularly relevant for the development of quantum-ready formulations of engineering optimization problems and, more specifically, structural design, since reducing the problem size is a key requirement for the effective use of current quantum computing technologies, whose applicability is still limited by the number of available information-processing units.
For this reason, in this section we enrich the QUBO formulation for the minimum compliance problem with the anisotropic mesh adaptation tool for structural topology optimization proposed with the SIMPATY (SIMP with AdaptiviTY) algorithm in [27]. Here anisotropic adapted meshes are employed to better describe the sharp material-void interface typically arising in optimized topologies.

5.1 From the PDE problem to the estimator

As a preliminary step, we recall the anisotropic framework used to formalize the adaptive mesh procedure. Following [67], each element K of the partition $\mathcal{T}_h$ is characterized through the standard affine map $T_K : \widehat{K} \to K$ from the reference element $\widehat{K}$ to the generic element K, namely

$$\mathbf{r} = T_K(\widehat{\mathbf{r}}) = M_K \widehat{\mathbf{r}} + \mathbf{t}_K,$$

with $\mathbf{r} \in K$, $\widehat{\mathbf{r}} \in \widehat{K}$, $M_K \in \mathbb{R}^{d\times d}$, and $\mathbf{t}_K \in \mathbb{R}^d$.
By factorizing the matrix M_K as $B_K Z_K$ through the polar decomposition, the geometric features of the element K can be described in terms of the spectral properties of B_K, being $B_K = R_K^T \Lambda_K R_K$, where $\Lambda_K = \mathrm{diag}(\lambda_{1,K}, \ldots, \lambda_{d,K})$ collects the eigenvalues of B_K, and $R_K = [\mathbf{r}_{1,K}, \ldots, \mathbf{r}_{d,K}]$ contains the corresponding eigenvectors. In particular, the eigenvalues measure the lengths of the semi-axes of the ellipsoid circumscribed to K, whereas the eigenvectors identify the corresponding directions. This provides a complete characterization of the three distinguishing features of an anisotropic element, namely the size, related to the values $\lambda_{i,K}$, the orientation, determined by the vectors $\mathbf{r}_{i,K}$, and the shape, described by the associated aspect ratios

$$s_{i,K} = \lambda_{i,K}^{\frac{2(d-1)}{d}} \left[\prod_{j=1, j\neq i}^{d} \lambda_{j,K} \right]^{-\frac{2}{d}} \qquad i = 1, \ldots, d. \tag{17}$$

The isotropic case is included in this definition, corresponding to the limiting situation in which the ellipsoid reduces to a sphere with radius $\lambda_{1,K} = \ldots = \lambda_{d,K}$ and the aspect ratios satisfy the

relation $s_{i,K} = 1$, for $i = 1, \dots, d$.

The anisotropic setting above provides the framework used to enrich the QUBO formulation of the minimum compliance problem with a mesh adaptation module. To this end, we employ an a posteriori anisotropic error estimator for the $H^1(\Omega)$-seminorm of the discretization error $e_h = \rho - \rho_h$ associated with the density field. This choice is motivated by the need to accurately track regions with large density gradients, typically located at the solid-void interface, so that the adaptive procedure promotes a sharp mesh resolution of the optimized design.
The error estimator selected to drive the mesh adaptation procedure is

$$\eta^2 = \sum_{K \in \mathcal{T}_h} \eta_K^2 \quad \text{with} \quad \eta_K^2 = \left[\prod_{j=1}^{d} \lambda_{j,K} \right]^{-\frac{2}{d}} \sum_{i=1}^{d} \lambda_{i,K}^2 \left(\mathbf{r}_{i,K}^T G_{\Delta_K}(E_\nabla(\rho_h)) \mathbf{r}_{i,K} \right), \tag{18}$$

proposed in [60, 68], which generalizes the well-established isotropic recovery-based error estimators originally proposed by O.C. Zienkiewicz and J.Z. Zhu [69, 70, 71] to an anisotropic environment.
In (18), $E_\nabla(\rho_h) = P(\nabla\rho_h) - \nabla\rho_h$ is the recovered error, which surrogates the error on the gradient, $\nabla\rho - \nabla\rho_h$, in the evaluation of the seminorm $|e_h|_{H^1(\Omega)}$, with $P(\nabla\rho_h)$ the recovered gradient. Quantity $G_{\Delta_K}(\cdot)$ denotes the symmetric semidefinite positive matrix with entries

$$[G_{\Delta_K}(\mathbf{q})]_{l,m} = \sum_{T \in \Delta_K} \int_T q_l q_m \, dT \quad l, m = 1, \dots, d, \tag{19}$$

for any vector-valued function $\mathbf{q} = [q_1, \dots, q_d]^T \in [L^2(\Omega)]^d$, and $\Delta_K = \{T \in \mathcal{T}_h : T \cap K \neq \emptyset\}$ the patch of elements associated with K.
Concerning the specific choice for the recovered gradient, in the numerical validation below we adopt the formula in [68], so that

$$P(\nabla\rho_h)(\mathbf{r}) = \frac{1}{|\Delta_K|} \sum_{T \in \Delta_K} |T| \, \nabla\rho_h\big|_T \quad \text{for } \mathbf{r} \in K \tag{20}$$

coincides with a volume-weighted average of the discrete gradient $\nabla\rho_h$.
Of course, depending on the specific application, other scalar or vector physical quantities may be used in place of ρ_h in (20) (see, for instance, [19] and section 6 therein).
Finally, we observe that the sum appearing in the definition of η_K^2 in (18) provides the anisotropic counterpart of the standard $H^1(\Omega)$-seminorm. In line with the original Zienkiewicz-Zhu (ZZ) estimator, the scaling factor multiplying this sum ensures the consistency of the anisotropic quantity with the corresponding isotropic version.

5.2 From the estimator to the adapted mesh

As a next step, we translate the error estimator η into an operational criterion for adapting the current computational mesh. We pursue three simultaneous goals, namely enforcing a prescribed accuracy τ for the discrete solution, reducing the mesh cardinality, $\#\mathcal{T}_h$, and achieving an approximately uniform distribution of the estimated error across all elements of the mesh. To

this aim, we adopt a metric-based approach, where the error estimator η is used to define the spacing of the new adapted mesh through a metric tensor $\mathcal{M}$, i.e., a symmetric positive definite tensor that modifies the notion of length in the computational domain [72].
Following [68, 60], it can be proved that, at each iteration of the adaptation procedure, we solve the elementwise constrained minimization problem

$$\min_{s_{i,K},\mathbf{r}_{i,K}} \sum_{i=1}^{d} s_{i,K}\left(\mathbf{r}_{i,K}^T \widehat{G}_{\Delta_K}(E_\nabla(\rho_h))\mathbf{r}_{i,K}\right) : \quad \begin{cases} \mathbf{r}_{i,K}\cdot\mathbf{r}_{j,K} = \delta_{ij} \\ s_{1,K} \geq \cdots \geq s_{d,K} \\ \prod_{j=1}^{k} s_{j,K} = 1, \end{cases} \tag{21}$$

with δ_{ij} the Kronecher symbol and $\widehat{G}_{\Delta_K}(E_\nabla(\rho_h)) = |\Delta_K|^{-1} G_{\Delta_K}(E_\nabla(\rho_h))$ the matrix in (19) scaled with respect to the measure of the patch Δ_K.
The solution to problem (21) is provided in [60] through an explicit formula. In particular, the optimal values for $s_{i,K}$ and $\mathbf{r}_{i,K}$ are

$$s_{i,K}^* = \left[\prod_{i=1}^{d} g_{i,K}\right]^{\frac{1}{d}} g_{d+1-i}^{-1}, \quad \mathbf{r}_{i,K}^* = \mathbf{g}_{d+1-i} \quad \text{with } i = 1,\ldots,d, \tag{22}$$

where $\{g_i, \mathbf{g}_i\}_{i=1,\ldots,d}$ are the eigenpairs associated with matrix $\widehat{G}_{\Delta_K}(E_\nabla(\rho_h))$, with $g_1 \geq \ldots \geq g_d > 0$ and $\{\mathbf{g}_i\}_{i=1,\ldots,d}$ orthonormal vectors. In addition, the error equidistribution criterion which imposes

$$\eta_K^2 = \frac{\tau^2}{\#\mathcal{T}_h},$$

allows us to identify the optimal anisotropic lengths

$$\lambda_{i,K}^* = \left(\frac{\tau^2}{d\,\#\mathcal{T}_h|\widehat{\Delta}_K|}\right)^{\frac{1}{d}} \left[\prod_{i=1}^{d} g_{i,K}\right]^{\frac{d-2}{2d^2}} g_{d+1-i}^{\frac{1}{2}}, \tag{23}$$

where $\widehat{\Delta}_K = T_K^{-1}(\Delta_K)$ denotes the patch Δ_K pulled-back through the map T_K.
The optimal quantities $\lambda_{i,K}^*$ and $\mathbf{r}_{i,K}^*$ are instrumental to define the metric tensor $\mathcal{M}$, approximated on the mesh $\mathcal{T}_h$ by the piecewise constant matrices $\mathcal{M}_K = \mathcal{M}|_K = (R_K^*)^T(\Lambda_K^*)^{-2}R_K^*$ with $R_K^* = [\mathbf{r}_{1,K}^*, \ldots, \mathbf{r}_{d,K}^*]$ and $\Lambda_K^* = \mathrm{diag}(\lambda_{1,K}^*, \ldots, \lambda_{d,K}^*)$.
After being suitably projected onto the vertices of the mesh $\mathcal{T}_h$, the new metric tensor $\mathcal{M}_V$ is provided as input to a metric-based mesh generator, to generate the anistropic adapted mesh $\mathcal{T}_h^*$.

5.3 Embedding anisotropic mesh adaptation into the QUBO formulation

The adaptive procedure described in the previous section may generate a computational mesh $\mathcal{T}_h$ whose number of vertices, depending on the prescribed tolerance τ, exceeds the admissible problem size for the subsequent QUBO solution step, both on classical and quantum solvers.
To control the complexity of the resulting binary optimization problem, we adopt the strategy

in [73, 74], which iteratively rescales the metric tensor until the generated mesh attains the target number of vertices N_v^* within an assigned tolerance. The rationale behind this procedure relies on the notion of the complexity of the metric tensor $\mathcal{M}$, defined as

$$C(\mathcal{M}) = \int_\Omega \sqrt{\det(\mathcal{M})}\, d\Omega. \tag{24}$$

In particular, the number, N_v, of vertices and the cardinality, $\#\mathcal{T}_h$, of a mesh $\mathcal{T}_h$ associated with the metric $\mathcal{M}$ are directly proportional to this complexity, namely

$$N_v = \gamma_v\, C(\mathcal{M}), \qquad \#\mathcal{T}_h = \gamma_t\, C(\mathcal{M}) \tag{25}$$

with γ_v and γ_t positive constants.
We now analyze the effect of the uniform rescaling

$$\mathcal{M}_S = S^{\frac{2}{d}}\mathcal{M} \tag{26}$$

of the prescribed metric $\mathcal{M}$ on the number of vertices and on the mesh cardinality, with $S > 0$. It follows that the element orientation and shape are preserved, while the metric complexity modifies as

$$C(\mathcal{M}_S) = \int_\Omega \sqrt{\det(S^{\frac{2}{d}}\mathcal{M})}\, d\Omega = S\, C(\mathcal{M}). \tag{27}$$

Consequently, by (25), the number, $N_{v,S}$, of vertices and the cardinality, $\#\mathcal{T}_{h,S}$, of the mesh associated with the scaled metric $\mathcal{M}_S$ modify into

$$N_{v,S} \approx S\, N_v, \qquad \#\mathcal{T}_{h,S} \approx S\, \#\mathcal{T}_h. \tag{28}$$

These relations are instrumental to select the scaling factor S which keeps the number of QUBO unknowns within the admissible problem size, without altering the anisotropic features of the mesh. Once the admissible problem size has been identified, either with a target number of vertices N_v^* or with a target mesh cardinality $\#\mathcal{T}_h^*$, the scaling factor S is determined through the iterative mesh adaptation process. In more detail, let $\mathcal{M}_S^{(0)} = \mathcal{M}^{(0)}$ be the metric associated with the initial mesh $\mathcal{T}_h^{(0)}$. Inspired by (26) and by the vertex-based relation in $(28)_1$, we update the metric according to the rules

$$\mathcal{M}_S^{(j+1)} = \left(S^{(j)}\right)^{\frac{2}{d}} \mathcal{M}_S^{(j)}, \qquad S^{(j+1)} = S^{(j)}\, \frac{N_v^*}{N_{v,S}^{(j)}}, \tag{29}$$

where $\mathcal{M}_S^{(k)}$ and $S^{(k)}$ denote the metric and the scaling factor at the k-th adaptive iteration, while $N_{v,S}^{(k)}$ is the corresponding number of vertices, for $k = j, j+1$. Alternatively, when the admissible problem size is prescribed in terms of the target mesh cardinality $\#\mathcal{T}_h^*$, the element-based relation in $(28)_2$ yields

$$\mathcal{M}_S^{(j+1)} = \left(S^{(j)}\right)^{\frac{2}{d}} \mathcal{M}_S^{(j)}, \qquad S^{(j+1)} = S^{(j)}\, \frac{\#\mathcal{T}_h^*}{\#\mathcal{T}_{h,S}^{(j)}}, \tag{30}$$

with $\#\mathcal{T}_{h,S}^{(j)}$ the cardinality of mesh $\mathcal{T}_{h,S}^{(j)}$. In both cases, the iterations are stopped when $N_{v,S}^{(j)}$, or $\#\mathcal{T}_{h,S}^{(j)}$, coincides with the target value N_v^*, or $\#\mathcal{T}_h^*$, up to the tolerance θ, namely

$$N_{v,S}^{(j)} \in [N_v^*(1-\theta),\, N_v^*(1+\theta)] \qquad \text{or} \qquad \#\mathcal{T}_{h,S}^{(j)} \in [\#\mathcal{T}_h^*(1-\theta),\, \#\mathcal{T}_h^*(1+\theta)].$$

Algorithm 1 summarizes the QUBO version of the SIMPATY procedure, enriched with the metric rescaling step described above.
The algorithm takes as input parameters: the initial discretization, $\mathcal{T}_h^{(0)}$, of the domain Ω and the associated density vector, $\boldsymbol{\rho}^{(0)}$; volume/density-related quantities, α and $\underline{\rho}$; the QUBO coefficients δ, ξ and ζ; the sampler settings `sampler`, n_S, and $\overline{T}$; the filtering data η, β, $\beta_{\max}$, R; the parameters for the mesh adaptation, τ, $h_{\min}$ and $h_{\max}$ and for the metric rescaling, N_v^* and θ; the tolerances $\mathrm{tol}_{\mathrm{TO}}$ and $\mathrm{tol}_{\mathrm{A}}$ for the stopping criteria; the maximum iteration counts, $k_{\max}$ and $i_{\max}$.
QUBOSIMPATY returns the adapted mesh $\mathcal{T}_h$, the standard, $\boldsymbol{\rho}$, and the filtered $\widetilde{\boldsymbol{\rho}}$, density vectors.
QUBOSIMPATY procedure is organized into two nested loops. The inner loop performs topology optimization on the current adapted mesh $\mathcal{T}_h^{(i)}$. At each topology optimization iteration k, the density field is filtered by combining the standard Helmholtz filter and a smoothed Heaviside projection, controlled by the parameters R and by η, β, respectively, through the function Filter. The resulting filtered density is then used to compute the compliance sensitivities by means of the function Sensitivities. Finally, the QUBO problem (14) is solved to update the design variables according to (11). This step is carried out by the function QUBOoptimize, which constructs the QUBO matrix in (16), initializes the binary variables from the current density distribution, and calls a prescribed sampler to approximate the minimum-energy configuration. The sampler can be chosen among classical, quantum-inspired, or quantum-based QUBO solvers, such as simulated annealing, tabu search, or quantum annealing, with n_S and $\overline{T}$ denoting the number of samples and the maximum computational time allowed for the sampling procedure, respectively. The inner loop is stopped when the relative error in the $L^2(\Omega)$-norm on the density field, computed by function

$$\textsc{ErrorDensity}(\rho_h^{(k)}, \rho_h^{(k+1)}) = \frac{\|\rho_h^{(k+1)} - \rho_h^{(k)}\|_{L^2(\Omega)}}{\|\rho_h^{(k)}\|_{L^2(\Omega)}}, \tag{31}$$

falls below the tolerance $\mathrm{tol}_{\mathrm{TO}}$, or when the maximum number of topology optimization iterations, $k_{\max}$, is achieved.
Once the topology optimization loop has converged on the current mesh $\mathcal{T}_h^{(i)}$, the resulting density field is used to drive anisotropic mesh adaptation. Before adaptation, an additional filtering step is performed with a progressively increased Heaviside projection parameter β_A. This produces a sharper density transition and therefore promotes anisotropic elements aligned with the solid-void interface. We emphasize that this continuation on β_A is used here to improve the mesh adaptation step, while it is not required by the QUBO optimization procedure to obtain a binary final design.
The adapted mesh is generated by the function AdaptMesh, which computes the anisotropic

Algorithm 1 QUBOSIMPATY

Input: $\mathcal{T}_h^{(0)}$, $\boldsymbol{\rho}^{(0)}$, α, $\underline{\rho}$, δ, ξ, ζ, `sampler`, n_S, $\overline{T}$, η, β, $\beta_{\max}$, R, τ, $h_{\min}$, $h_{\max}$, N_v^*, θ, $\mathrm{tol}_{\mathrm{TO}}$, $\mathrm{tol}_{\mathrm{A}}$, $k_{\max}$, $i_{\max}$

$i \leftarrow 0$
$\boldsymbol{\rho} \leftarrow \boldsymbol{\rho}^{(0)}$
$\beta_A \leftarrow \beta$
$\mathrm{err}_{\mathrm{A}} \leftarrow \mathrm{tol}_{\mathrm{A}} + 1$
while $\mathrm{err}_{\mathrm{A}} > \mathrm{tol}_{\mathrm{A}}$ **and** $i < i_{\max}$ **do**
  $k \leftarrow 0$
  $\boldsymbol{\rho}^{(k)} \leftarrow \boldsymbol{\rho}$
  $\mathrm{err}_{\mathrm{TO}} \leftarrow \mathrm{tol}_{\mathrm{TO}} + 1$
  while $\mathrm{err}_{\mathrm{TO}} > \mathrm{tol}_{\mathrm{TO}}$ **and** $k < k_{\max}$ **do**
    $\widetilde{\boldsymbol{\rho}}^{(k)} \leftarrow \textsc{Filter}(\boldsymbol{\rho}^{(k)}; \eta, \beta, R)$
    $\nabla_\rho c \leftarrow \textsc{Sensitivities}(\widetilde{\boldsymbol{\rho}}^{(k)})$
    $\boldsymbol{\rho}^{(k+1)} \leftarrow \mathrm{QUBO}\textsc{optimize}(\boldsymbol{\rho}^{(k)}, \nabla_\rho c, \delta, \xi, \zeta, \underline{\rho}, \alpha, \texttt{sampler}, n_S, \overline{T})$
    $\mathrm{err}_{\mathrm{TO}} \leftarrow \textsc{ErrorDensity}(\rho_h^{(k)}, \rho_h^{(k+1)})$
    $k \leftarrow k + 1$
  end while
  $\boldsymbol{\rho} \leftarrow \boldsymbol{\rho}^{(k)}$
  $\beta_A \leftarrow \min\{\beta_{\max}, 2\beta_A\}$
  $\widetilde{\boldsymbol{\rho}}_A \leftarrow \textsc{Filter}(\boldsymbol{\rho}; \eta, \beta_A, R)$
  $\mathcal{T}_h^{(i+1)} \leftarrow \textsc{AdaptMesh}(\mathcal{T}_h^{(i)}, \widetilde{\boldsymbol{\rho}}_A, \tau, h_{\min}, h_{\max}, N_v^*, \theta)$
  $\widetilde{\boldsymbol{\rho}} \leftarrow \textsc{Filter}(\boldsymbol{\rho}; \eta, \beta, R)$
  $\boldsymbol{\rho} \leftarrow \textsc{Project}(\boldsymbol{\rho}, \mathcal{T}_h^{(i)}, \mathcal{T}_h^{(i+1)})$
  $\widetilde{\boldsymbol{\rho}} \leftarrow \textsc{Project}(\widetilde{\boldsymbol{\rho}}, \mathcal{T}_h^{(i)}, \mathcal{T}_h^{(i+1)})$
  $\mathrm{err}_{\mathrm{A}} \leftarrow \textsc{ErrorMesh}(\mathcal{T}_h^{(i)}, \mathcal{T}_h^{(i+1)})$
  $i \leftarrow i + 1$
end while
$\mathcal{T}_h \leftarrow \mathcal{T}_h^{(i)}$

Output: $\mathcal{T}_h$, $\boldsymbol{\rho}$, $\widetilde{\boldsymbol{\rho}}$

metric (22)-(23) from the filtered density field[2] and the prescribed tolerance τ, constrained by the minimum and maximum allowed mesh size parameters, $h_{\min}$ and $h_{\max}$. In addition, the metric rescaling strategy in (29) is applied so that the adapted mesh $\mathcal{T}_h^{(i+1)}$ has the prescribed number of vertices, N_v^*, up to the tolerance θ. After remeshing, the standard and filtered density fields are projected from $\mathcal{T}_h^{(i)}$ to $\mathcal{T}_h^{(i+1)}$.
The outer loop is then repeated until the mesh cardinality variation,

$$\textsc{ErrorMesh}(\mathcal{T}_h^{(i)}, \mathcal{T}_h^{(i+1)}) = \frac{|\#\mathcal{T}_h^{(i+1)} - \#\mathcal{T}_h^{(i)}|}{\#\mathcal{T}_h^{(i)}},$$

remains below the tolerance tol_A, or the maximum number of mesh adaptation cycles, $i_{\max}$, is reached.

6 Numerical results

In this section, we numerically assess the performance of the proposed QUBOSIMPATY algorithm on standard two (2D)- and three (3D)-dimensional benchmark test cases [75]. We first report the input parameters of the algorithm that are common to all the considered configurations. The test case-dependent quantities are instead specified in the dedicated sections.
We refer to a generic isotropic material, with Young modulus $E = 1$ and Poisson ratio $\nu = 0.3$ in (5). In (6), the SIMP penalization exponent is set to $p = 3$ for the 2D tests and to $p = 4$ for the 3D simulations, consistently with the literature [76]. Moreover, unless otherwise stated, the initial guess for the design variable is set to the constant vector $\boldsymbol{\rho}^{(0)} = [\alpha, \dots, \alpha]^T$, while the lower bound $\underline{\rho}$ is chosen as 10^{-6} for the 2D assessments, and as 10^{-3} in the 3D cases.
The filtering phase is governed by parameters $\eta = 0.5$, $\beta = 1$, $\beta_{\max} = 8$, while the radius for the Helmholtz filter is specified for each test.
For the QUBO solver, we consider the two values $\delta = 0.05$ and $\delta = 0.025$ for the density step size defined in (12), with the specific choice depending on the numerical test under consideration. Concerning the penalty parameter ζ and the scaling coefficient ξ in (14), we set $\zeta = 1000$ for the 2D cases, $\zeta = 500$ for the 3D tests, while we pick

$$\xi = \left\|\nabla_\rho c\right\|_{\ell^1}^{-1} = \left(\sum_{i=1}^{N} \left|\frac{\partial c}{\partial \rho_i}\right|\right)^{-1}, \tag{32}$$

independently of the domain dimension. The number of samples is set to $n_S = 1$ for each selected sampler [77]. When tabu search (TS) is used, we additionally prescribe the maximum computational time $\overline{T} = 1$ s, whereas no value of $\overline{T}$ is required for simulated annealing (SA).
Concerning the control of the adaptive mesh complexity, the target number of vertices N_v^* is enforced up to the tolerance $\theta = 5 \cdot 10^{-2}$.
Finally, the breaking conditions are tuned by tolerances $\text{tol}_\text{A} = 10^{-2}$, $\text{tol}_\text{TO} = 10^{-4}$ and by the

[2] We observe that the recovered gradient in (20) is computed from the filtered density. Indeed, regardless of the polynomial degree r used for the density space V_h^r, the filtered density is always taken as a piecewise linear finite element function, so that its gradient is well defined.

maximum number of allowed iterations $i_{\max} = 3$ and $k_{\max} = 100$ for the 2D simulations, while we pick $\text{tol}_\text{A} = 10^{-2}$, $\text{tol}_\text{TO} = 10^{-5}$, $i_{\max} = 3$ and $k_{\max} = 200$ in the 3D test cases.

The quality of the optimized designs is assessed through the value of the objective function (i.e., the compliance), the fulfillment of the prescribed volume constraint, and two additional indicators computed on the unfiltered density field ρ_h. The first indicator is the grayness index [78]

$$\text{G}(\rho_h) = \frac{4}{|\Omega|} \int_\Omega \rho_h (1 - \rho_h)\, d\Omega, \tag{33}$$

which measures the degree of binary separation of the layout, where $\text{G}(\rho_h) = 0$ corresponds to the optimal sharp 0-1 distribution. The second indicator is the total variation

$$\text{TV}(\rho_h) = \int_\Omega \|\nabla \rho_h\|_{\ell^1}\, d\Omega, \tag{34}$$

which is related to the perimeter of the final design in 2D and to the measure of the interfacial area in 3D [79]. Low values of $\text{TV}(\rho_h)$ indicate less oscillatory and geometrically simple designs.

The finite element solution of the linear elasticity problem in (6) is performed by the `dolfinx` Python library [80], following the approach adopted in the `FEniTop` package [81]. The generation and manipulation of anisotropic adapted meshes are delegated to `FreeFEM` for the computation of the metric tensor [82] and to `mmg`, as an open-source software package specifically designed for metric-based mesh adaptation [83]. The QUBO problem is constructed and solved through the D-Wave software libraries, which allow different choices for the `sampler` variable to be employed, including quantum annealing (QA) on remote D-Wave hardware [84], SA, and TS on standard hardware[3]. Finally, when the standard SIMP or SIMPATY algorithms are used for comparison purposes, MMA is adopted as the optimizer.

6.1 2D test case 1: the cantilever beam

The minimum compliance problem (6) is here solved in the design domain $\Omega = (0, 2) \times (0, 1)$, where a traction $\mathbf{f} = [0, -1]^T$ is applied on the boundary portion $\Gamma_N = \{(x, y) \in \partial\Omega : x = 2, |y - 0.5| \leq 0.1\}$ and homogenous Dirichlet data are assigned to the displacement on $\Gamma_D = \{(x, y) \in \partial\Omega : x = 0\}$ (see Figure 1, left, for a sketch). The design process is constrained by allowing a maximum volume fraction $\alpha = 0.6$. The domain Ω is discretized using an isotropic structured triangular mesh, $\mathcal{T}_h^{(0)}$, composed of 5000 elements.
In the filtering stage, the Helmholtz radius is set to $R = 0.048$. The QUBOptimize function is employed by selecting TS as sampler, with a density step size $\delta = 0.025$.
The anisotropic mesh adaptation procedure is driven by the tolerance $\tau = 6$, with minimum and maximum admissible element sizes equal to $h_{\min} = 10^{-2}$ and $h_{\max} = 10^{-1}$, respectively. In this test, the metric-scaling procedure is initially switched off.

The optimization procedure is first validated by comparing different polynomial degrees for the discrete space V_h^r in (7), namely by picking $r = 0$ or $r = 1$. Figure 2 presents the density field

[3]The numerical simulations have been performed on the Kami cluster at the Department of Mathematics, Politecnico di Milano. The cluster features 40 CPU nodes, each one powered by two 24-core CPUs with 512 GB RAM.

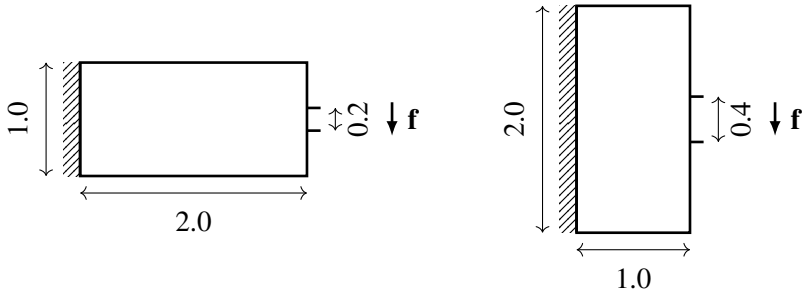

Figure 1: 2D test cases: sketch of the design setting for the cantilever beam (left) and the coat hanger (right) configurations.

ρ_h at different global iterations. We observe that the designs associated with the initial mesh (i.e., for $i = 0$) differ in terms of intermediate densities, with gray regions being more pronounced for $r = 1$. This discrepancy is damped in the converged layouts obtained for $i = 2$. Based on this comparison, we select $r = 0$ in all the checks performed in this section. This choice has to be interpreted in terms of the number of density DOFs. For $r = 0$, these DOFs are associated with the element barycenters, whereas, for $r = 1$, they coincide with the mesh vertices. In the 2D meshes considered here, the numbers of elements and vertices are relatively close, so that using $r = 0$ does not significantly vary the QUBO size with respect to the $r = 1$ choice. In 3D, instead, the number of vertices is much smaller than the number of tetrahedra, and the choice $r = 1$ becomes preferable to limit the number of binary variables (see section 6.3).

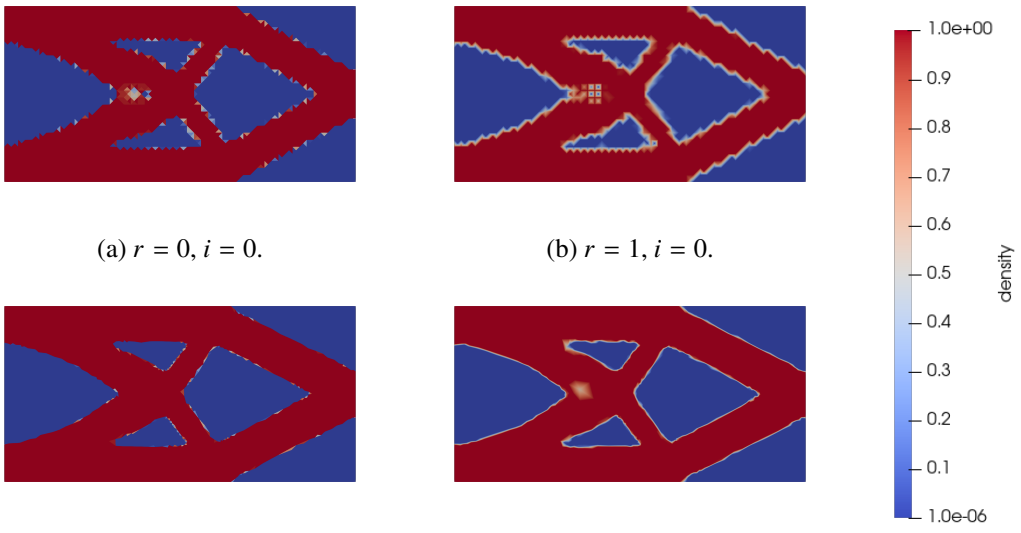

(a) $r = 0, i = 0$.

(b) $r = 1, i = 0$.

(c) $r = 0, i = 2$.

(d) $r = 1, i = 2$.

Figure 2: The 2D cantilever beam: sensitivity of the optimized design to the polynomial degree r (by column) at different global iterations i (by row).

Table 1: The 2D cantilever beam: sensitivity of the computed design to the selected sampler in terms of structure performance and quality, and required computational time.

i	`sampler`	VF	c	G	TV	T_{conv} [s]
5000 elements						
0	TS	0.602	1.3225	0.0179	9.8477	250
	SA	0.602	1.3222	0.0208	9.8449	2727
1	TS	0.602	1.3175	0.0103	8.8650	251
	SA	0.602	1.3170	0.0126	8.7500	2750
2	TS	0.602	1.3165	0.0033	8.6702	257
	SA	0.602	1.3159	0.0094	9.0231	2748
3200 elements						
0	TS	0.603	0.7897	0.0198	10.0524	160
	SA	0.603	0.7894	0.0219	9.9523	1124
1	TS	0.603	0.7872	0.0140	9.3001	163
	SA	0.603	0.7869	0.0156	9.1790	1143
2	TS	0.603	0.7864	0.0084	9.1948	163
	SA	0.603	0.7861	0.0112	9.1056	1144

As a second check, we investigate the effect of the selected sampler on the quality of the optimized design, measured through the indicators in (33) and (34), and on the achievement of the optimization goal under the constraint in (6). To this aim, we compare the TS and SA samplers on a fixed mesh, in order to avoid any bias introduced by the mesh adaptation procedure. Table 1 gathers the results of this analysis on two structured grids, consisting of 5000 (top panel) and 3200 (bottom panel) triangles. The quantities monitored over the 3 global iterations to convergence are the achieved volume fraction (VF), compliance (c), grayness (G), and total variation (TV). For each fixed mesh, the two samplers lead to very similar structural performances, while small discrepancies are observed in terms of grayness and total variation, generally in favor of TS at convergence. On the contrary, the time to convergence, T_{conv}, reported in the last column of the table, exhibits a pronounced difference. This quantity is reduced by about one order of magnitude when using TS, leading to a considerably faster computation of the optimized design. This observation is consistent with the well-known behavior of SA, which may require a large number of iterations in high-dimensional spaces and, when stopped after a prescribed maximum time, may remain trapped in suboptimal local minima [45], as already observed in section 2.1.

We now assess the benefits of the additive density update rule proposed in (11) with respect to the multiplicative formula in (10). To this end, we consider two choices of the initial density field, $\rho_h^{(0)}$, namely a constant and a random initialization. The layouts obtained at convergence with the additive rule (11), shown in Figure 3 (top), indicate that the final topology is essentially independent of the selected initial guess. In contrast, when the multiplicative rule (10) is adopted, the QUBOSIMPATY algorithm does not reach convergence within the prescribed maximum

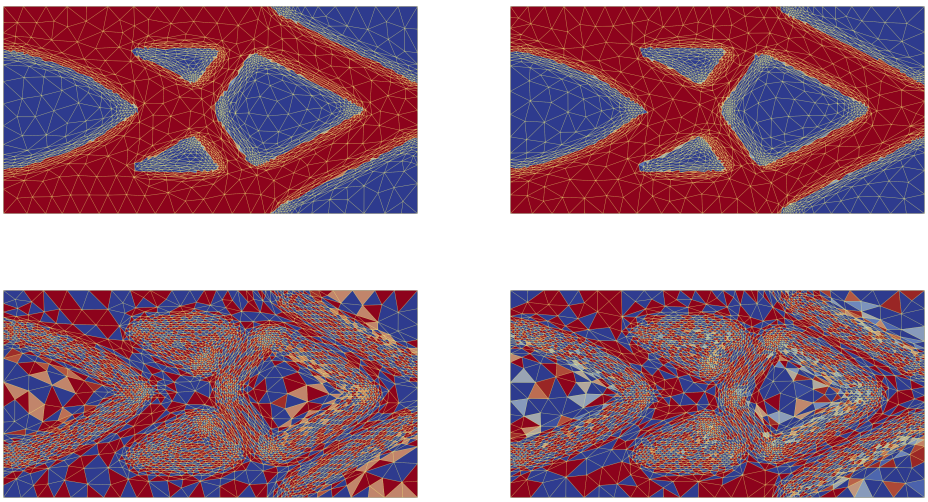

Figure 3: The 2D cantilever beam: comparison between the additive (11) (top) and the multiplicative (10) (bottom) density update rules, for a constant (left) and a random (right) initial density $\rho_h^{(0)}$. Final density distributions are overlapped to the corresponding anisotropic adapted meshes.

number of iterations independently of the selected initial density (see Figure 3, bottom).

Figure 4 shows the convergence histories of the compliance, c, and of the volume fraction, VF, for the standard SIMPATY algorithm and for QUBOSIMPATY with the TS sampler, using the two values $\delta = 0.025$ and $\delta = 0.05$. SIMPATY enforces the volume constraint exactly, whereas the QUBO formulation allows for a marginal deviation from the prescribed value. On the other hand, QUBOSIMPATY exhibits a faster decrease of the compliance for $i = 0$, reaching lower values in fewer iterations for both choices of δ, with the best performance obtained for $\delta = 0.05$.

This comparison is complemented by the grayness and total variation values reported in Table 2 for the three optimization runs throughout the global iterations. To ensure a fair comparison, both indicators are computed after projecting the different density fields onto a high-resolution fixed grid with 80000 triangles. It turns out that QUBOSIMPATY with $\delta = 0.05$ yields the lowest value of G, whereas TV does not show a clear trend and remains comparable across the different runs.

The improvement brought by anisotropic mesh adaptation within the QUBO formulation is assessed in Table 3. The table compares QUBOSIMPATY with the standard QUBO algorithm applied on a fixed mesh with 5000 triangles, in terms of structural performance and design quality over the first three global iterations. As previously done, to avoid mesh-dependent effects in the evaluation of the reported indicators, all quantities are computed after transferring the corresponding density fields onto the a common high-resolution fixed mesh with 80000 elements. QUBO and QUBOSIMPATY methods yield comparable values for the monitored

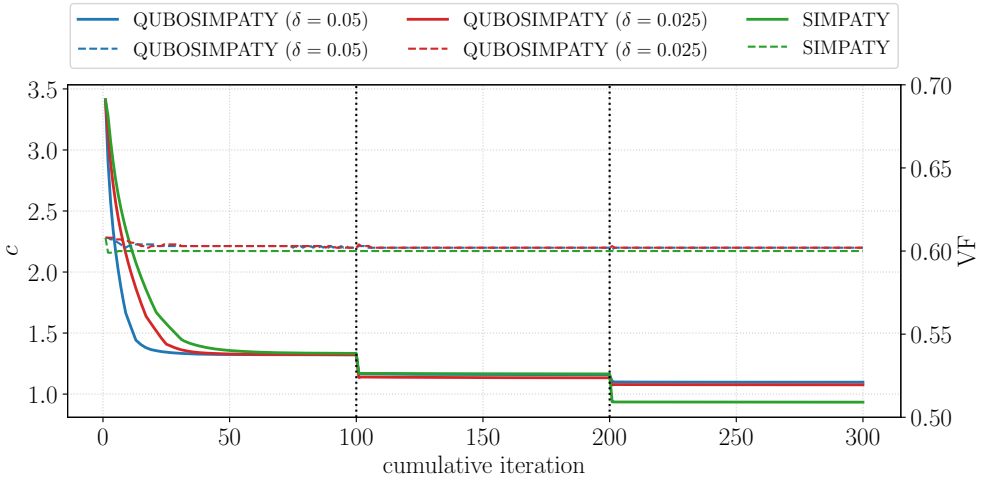

Figure 4: The 2D cantilever beam: convergence history of the compliance (solid lines) and of the volume fraction (dotted lines) for the standard SIMPATY and the QUBOSIMPATY algorithm, with two values of parameter δ. Mesh adaptation takes place in correspondence with the vertical lines.

metrics, but at different computational costs, here measured by the mesh cardinality reported in the last column of the table. In particular, QUBOSIMPATY improves the compliance by 0.21% and decreases the total variation by 3.88% with respect to the fixed mesh QUBO algorithm, while using a smaller number of elements, i.e., of density DOFs (approximately 3700 versus 5000 triangles).

Finally, we assess the effectiveness of the QUBOSIMPATY algorithm when activating the metric rescaling module, to obtain a mesh with a number of vertices close to the prescribed target N_v^*. To this end, we vary N_v^* in the set $\{1000, 2000, 3000, 4000\}$. For the selected value of the tolerance θ, the rescaling algorithm converges within 5 iterations in all cases.
Figure 5 provides the meshes generated at convergence by QUBOSIMPATY, together with the actual number of vertices, N_v, and the maximum and the average value of the aspect ratio, denoted by $s_{K,\max}$, and $s_{K,\text{avg}}$, respectively. As expected, the four meshes exhibit similar anisotropic patterns, with stretched elements aligned along the material-void interface. The average aspect ratio values remain essentially unchanged across the different target sizes, consistently with the uniform rescaling of the metric tensor in (29). This behavior confirms that metric rescaling is instrumental to a quantum-oriented setting, since reducing the number of QUBO variables to a hardware-compatible size, while retaining the anisotropic mesh features needed to represent the optimized layout.

6.2 2D test case 2: the coat hanger

We now consider a simplified version of the optimization setting discussed in the previous case, in order to assess the potential computational benefits of a quantum-based approach using the QA sampler on remote D-Wave hardware. To this aim, the reference optimization context coincides

Table 2: The 2D cantilever beam: comparison in terms of structure quality between the standard SIMPATY and the QUBOSIMPATY algorithms, for two values of parameter δ.

i	algorithm	δ	G	TV
0	QUBOSIMPATY	0.025	0.0164	9.9076
	QUBOSIMPATY	0.050	0.0104	9.8244
	SIMPATY	-	0.0359	9.2903
1	QUBOSIMPATY	0.025	0.0089	9.3674
	QUBOSIMPATY	0.050	0.0037	9.3163
	SIMPATY	-	0.0230	9.5640
2	QUBOSIMPATY	0.025	0.0051	9.1649
	QUBOSIMPATY	0.050	0.0007	8.8939
	SIMPATY	-	0.0163	9.4308

Table 3: The 2D cantilever beam: comparison between the standard QUBO and the QUBOSIMPATY algorithms, in terms of structure performance and quality, and mesh cardinality.

i	algorithm	VF	c	G	TV	$\#\mathcal{T}_h^{(i)}$
0	QUBOSIMPATY	0.603	1.4730	0.0179	9.8477	5000
	QUBO	0.603	1.4730	0.0179	9.8477	5000
1	QUBOSIMPATY	0.602	1.4591	0.0097	9.5697	3830
	QUBO	0.602	1.4628	0.0103	8.8650	5000
2	QUBOSIMPATY	0.602	1.4561	0.0051	9.0069	3726
	QUBO	0.602	1.4592	0.0033	8.6702	5000

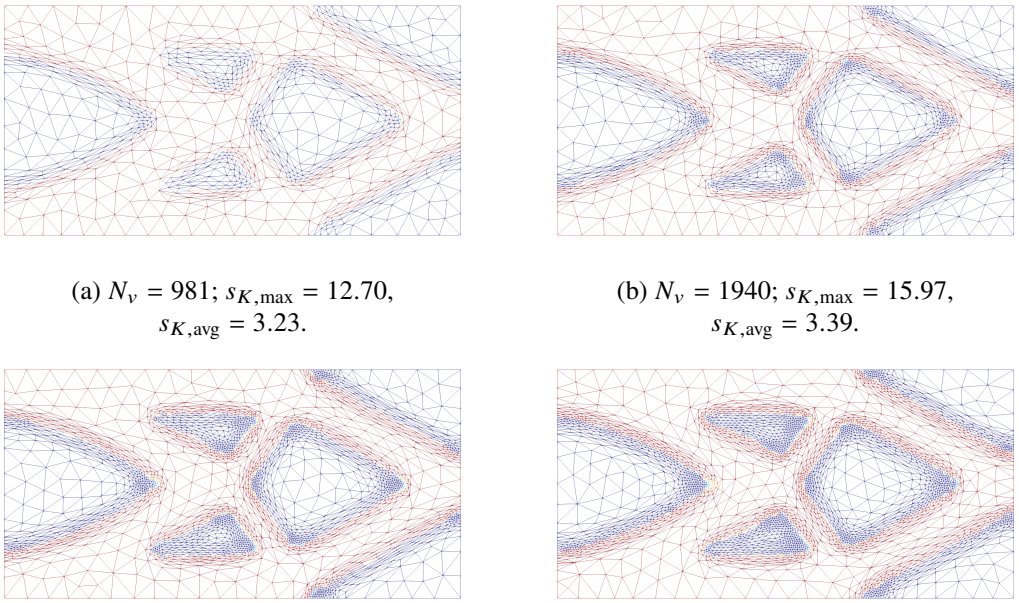

(a) $N_v = 981$; $s_{K,\max} = 12.70$, $s_{K,\text{avg}} = 3.23$.

(b) $N_v = 1940$; $s_{K,\max} = 15.97$, $s_{K,\text{avg}} = 3.39$.

(c) $N_v = 2870$; $s_{K,\max} = 19.12$, $s_{K,\text{avg}} = 2.83$.

(d) $N_v = 3873$; $s_{K,\max} = 17.24$, $s_{K,\text{avg}} = 2.38$.

Figure 5: The 2D cantilever beam: anisotropic adapted meshes delivered by QUBOSIMPATY algorithm for a metric rescaling associated with $N_v^* = 1000, 2000, 3000, 4000$ (from (a) to (d)).

with the coat hanger configuration (see Figure 1, right, for a sketch).
A load $\mathbf{f} = [0, -1]^T$ is applied on the boundary portion $\Gamma_N = \{(x, y) \in \partial\Omega : x = 1, |y-1| \leq 0.2\}$, while the structure is clamped along the left edge $\Gamma_D = \{(x, y) \in \partial\Omega : x = 0\}$. The prescribed volume fraction is fixed to $\alpha = 0.6$. The domain Ω is discretized using an isotropic triangular mesh, $\mathcal{T}_h^{(0)}$, with cardinality equal to 3200.
We filter the density variable by picking a radius $R = 0.06$, while the TS sampler is adopted in the function QUBOoptimize, for $\delta = 0.05$. The parameters $\tau = 3$, $h_{\min} = 10^{-2}$ and $h_{\max} = 10^{-1}$ drive the anisotropic mesh adaptation step.

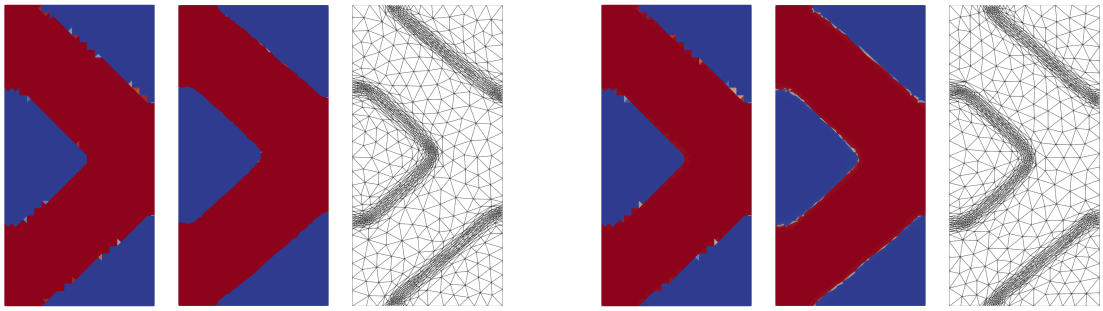

(a) (b) (c) (d) (e) (f)

Figure 6: The coat hanger: final design provided by the standard QUBO algorithm on $\mathcal{T}_h^{(0)}$ (a) and by the QUBOSIMPATY algorithm (b) on the final anisotropic adapted mesh (c); final design provided by the standard SIMP algorithm on $\mathcal{T}_h^{(0)}$ (d) and by the SIMPATY algorithm (e) on the final anisotropic adapted mesh (f).

Figure 6 presents a comparative visualization of the final density field ρ_h obtained using the QUBO and the SIMP formulations, potentially enriched with anisotropic mesh adaptation (i.e., QUBOSIMPATY and SIMPATY algorithms, respectively). The results show that the binary nature of the QUBO formulation does not, by itself, prevent the occurrence of intermediate densities. This is due to the fact that the binary variables control the design update, but do not coincide directly with the density field. As a consequence, a suitably adapted mesh is still needed to obtain a sharp material-void interface, as shown by the comparison between panels (a) and (b). A similar behavior is observed for the SIMP formulation, when comparing the standard SIMP and SIMPATY results in panels (d) and (e). In both the QUBO and SIMP settings, the activation of anisotropic mesh adaptation removes jagged boundaries and leads to a sharper representation of the optimized layout.
Table 4 compares the QUBO method, for the two values of δ considered above, with the SIMP formulation in terms of objective value, constraint fulfillment, and structural quality. The QUBO reformulation of problem (6) yields results that are fully comparable with those obtained by the SIMP method combined with MMA and, for some choices of δ, slightly better. An exception concerns the volume fraction constraint, which is enforced exactly by the SIMP method only, consistently with the results of the previous section.
Similar considerations apply to the comparison between QUBOSIMPATY and SIMPATY at

comparable computational complexity, measured in terms of the number of density DOFs involved in the optimization problem. This conclusion is supported by the values in Table 5 associated with the convergence of the optimization procedure, after projecting the density fields onto a common high-resolution mesh with 80000 elements.

Table 4: The coat hanger: comparison between the standard QUBO and SIMP algorithms, in terms of structure performance and quality.

i	algorithm	δ	VF	c	G	TV
0	QUBO	0.025	0.601	0.2483	0.0019	4.0739
	QUBO	0.050	0.601	0.2483	0.0026	4.1724
	SIMP	-	0.600	0.2491	0.0152	4.0679
1	QUBO	0.025	0.601	0.2483	0.0019	4.0860
	QUBO	0.050	0.601	0.2483	0.0026	4.1724
	SIMP	-	0.600	0.2490	0.0081	4.0931
2	QUBO	0.025	0.601	0.2483	0.0018	4.0860
	QUBO	0.050	0.601	0.2483	0.0026	4.1724
	SIMP	-	0.600	0.2490	0.0069	4.0865

Table 5: The coat hanger: comparison at convergence in terms of structure performance and quality, and mesh cardinality between the standard SIMPATY and the QUBOSIMPATY algorithms, for two values of parameter δ.

i	algorithm	δ	VF	c	G	TV	$\#\mathcal{T}_h$
2	QUBOSIMPATY	0.025	0.601	0.3386	0.0006	4.3204	2053
	QUBOSIMPATY	0.050	0.601	0.3385	0.0003	4.3718	2092
	SIMPATY	-	0.600	0.3402	0.0123	4.5725	1978

We now consider a downscaled version of the optimization framework described above, in order to validate QUBOSIMPATY on D-Wave Advantage quantum processing units (QPUs) [85], based on the Pegasus topology [86]. To this end, the admissible volume fraction is set to $\alpha = 0.5$ and the QUBO problem is solved using the QA sampler, with $n_S = 100$, $\overline{T} = 10^{-1}$ s, and $\delta = 0.025$.
The parameters ζ and ξ in the optimization problem (14) are set to $\zeta = 5000$ and

$$\xi = \left\|\nabla_\rho c\right\|_{\ell^\infty}^{-1} = \left(\max_{i=1,\dots,N}\left|\frac{\partial c}{\partial \rho_i}\right|\right)^{-1}. \tag{35}$$

In addition, the filter radius is increased to $R = 0.1$, while the anisotropic mesh adaptation is constrained by the parameters $\tau = 4$, $h_{\min} = 10^{-1}$, and $h_{\max} = 4 \cdot 10^{-1}$. Then, to balance

computational cost and convergence, we allow a larger number of global adaptation iterations, namely $i_{\max} = 4$, and reduce the maximum number of topology optimization iterations to $k_{\max} = 50$.
The most challenging aspect remains the maximum problem size that can be handled by the QPU. In this respect, building on the observations on the finite element degree r in section 6.1, and exploiting the metric rescaling strategy included in QUBOSIMPATY, we set $r = 1$ and discretize the domain with an initial coarse triangular mesh of cardinality $\#\mathcal{T}_h^{(0)} = 200$ (i.e., $N_v^{(0)} = 116$ vertices). In addition, we choose the target number of vertices for metric rescaling as $N_v^* = 160$. The configurations returned by the QPU are finally processed by the QA sampler through a simple greedy selection step, which identifies the best candidates among the sampled solutions. This post-processing follows the standard workflow adopted when using D-Wave QPUs [84].

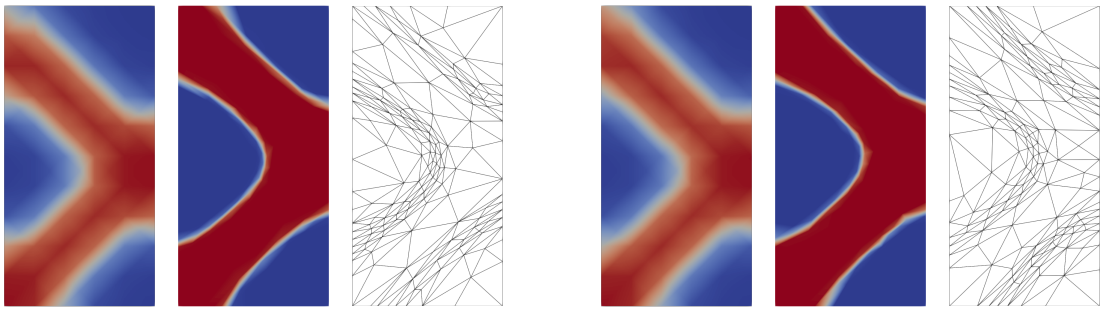

(a) $i = 0$. (b) $i = 3$. (c) $i = 3$. (d) $i = 0$. (e) $i = 3$. (f) $i = 3$.

Figure 7: The coat hanger: design provided by the QUBOSIMPATY algorithm at the first (a) and last (b) global iteration and final anisotropic adapted mesh (c); design provided by the SIMPATY algorithm at the first (d) and last (e) global iteration and final anisotropic adapted mesh (f).

Figure 7 shows the optimized layouts obtained with QUBOSIMPATY for $i = 0$ (panel (a)) and $i = 3$ (panel (b)), associated with the initial uniform mesh and the final anisotropic adapted grid (panel (c)). The use of a very small number of vertices makes the effect of mesh adaptation particularly evident in the geometry of the final design. In this test, the number of vertices in the adapted mesh exactly matches the prescribed target value $N_v^* = 160$, further supporting the effectiveness of the metric rescaling procedure. For comparison, the corresponding SIMPATY results run on standard hardware (panels (d)–(f)) are also reported, confirming the consistency of the QUBO-based SIMP formulation, when endowed with a mesh adaptation procedure.

Comparing the performance of the QPU with classical optimization procedures is not straightforward in general, since the execution on D-Wave quantum computers involves several distinct time contributions. We denote by T_c the service time required to access the remote hardware, including communication latency. This contribution is typically of the order of 10^{-1} s and is often excluded when the goal is to isolate the solver performance on the QPU.
A more representative measure of the computational cost associated with the quantum device is the QPU access time $T = T_p + T_s + T_\Delta$, where T_p is the QPU programming time, T_s is the sampling time, and T_Δ collects additional overheads. The sampling time can be further decomposed into

the quantum annealing time, T_a, and the readout time, T_r, both multiplied by the number of samples n_S, namely $T_s \approx n_S\,(T_a + T_r)$. Since the annealing stage represents the core quantum part of the computation, we focus on the contribution $n_S T_a$. With the standard annealing time $T_a = 20\,\mu$s and $n_S = 100$ samples, this contribution amounts to 2 ms. For the same test case, the classical MMA optimization step employed by SIMPATY typically requires about 25 ms per iteration. Thus, when only the annealing contribution is considered, the QPU sampling phase is faster than the corresponding classical optimization step by approximately one order of magnitude.

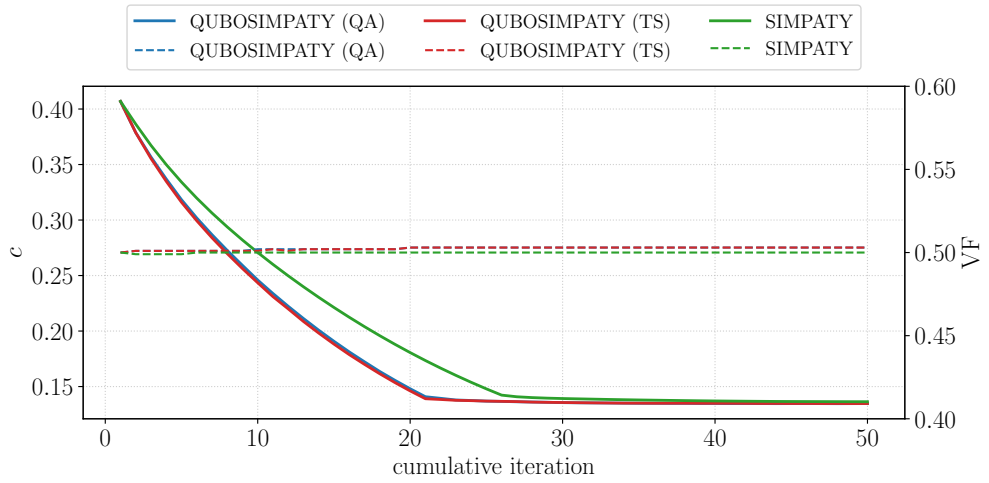

Figure 8: The coat hanger: convergence history of the compliance (solid lines) and the volume fraction (dotted lines) for the standard SIMPATY and the QUBOSIMPATY algorithm with QA and TS samplers.

On the other hand, the total wall-clock time of a QA call, approximately 0.8 s in the present test, includes the QUBO matrix construction, the QPU access time, and remote-service overheads. This value is larger than the time required by one classical SIMPATY optimization step on the local machine, taking approximately 0.025 s. Therefore, the present experiment should not be interpreted as an end-to-end speedup of QUBOSIMPATY over SIMPATY. Rather, it separates two different aspects of the computation: the intrinsic annealing contribution has a favorable time scale, whereas the overall execution time is still dominated by overheads associated with the current hardware-access workflow. From this perspective, anisotropic mesh adaptation is meant to control the number of QUBO variables and to make the resulting instance compatible with the size restrictions of the quantum annealer, while retaining the mesh features needed to describe the optimized layout.

Figure 8 compares the convergence curves, restricted to the first optimization cycle, associated with the QUBO problems solved by QA and TS, with the SIMPATY algorithm. The QA and TS curves are almost indistinguishable and exhibit a faster decrease of the compliance during the first global iteration, reaching a nearly stationary value earlier than SIMPATY. The volume fraction remains close to the prescribed target for all methods, with a slight overestimation in the

QUBO-based runs.

To sum up, at comparable problem size, the QUBO minimization step is consistent with the classical optimization strategy and can lead to a faster reduction of the objective. At the same time, the overall computational balance depends on both the cost of a single optimization step and the number of iterations required to reach a satisfactory design. In the present experiment, this balance is still affected by QUBO construction, remote QPU access, and hardware-execution overheads. Nevertheless, the observed convergence behavior suggests a possible regime in which faster QUBO minimization, combined with mesh complexity control, may become advantageous as hardware access and problem embedding improve. Along this line, we expect the potential benefit of quantum annealing to become more evident for larger-scale problems, whose treatment on currently available QPUs remains challenging.

6.3 3D test cases

The performance of QUBOSIMPATY is now investigated in 3D design settings close to realistic applications. In the context of quantum computing, however, 3D topology optimization does not represent a standard benchmark, as it is in direct contrast with the limited computational resources currently available on quantum hardware. To mitigate this limitation, we consider configurations in which geometric and loading symmetries can be exploited. Indeed, when the design domain, boundary conditions, and loads are symmetric, the optimal topology is expected to inherit the same symmetry. Accordingly, the results reported below are obtained by simulating only a symmetric portion of the full geometry, with the condition $\mathbf{u} \cdot \mathbf{n} = 0$ imposed on symmetry planes to prevent spurious normal displacements. Consistently with this reduced computational setting, the values in the tables and convergence plots refer to the simulated symmetric portion of the domain.

We consider three optimization settings and use linear finite elements in order to limit the number of density DOFs, coherently with the discussion in sections 6.1-6.2. The QUBOSIMPATY algorithm is run with $\delta = 0.025$.

The structures shown in the following paragraphs are obtained through a post-processing that retains the elements for which the filtered density is greater than 0.5.

3D test case 1: the cantilever beam We select the design domain $\Omega = (0, 0.5)\times(0, 0.5)\times(0, 2)$ and apply the load $\mathbf{f} = [0, -1, 0]^T$ to the boundary portion $\Gamma_N = \{(x, y, z) \in \partial\Omega : z = 2,\ y \leq 0.1\}$, while we clamp the structure on $\Gamma_D = \{(x, y, z) \in \partial\Omega : z = 0\}$ (see Figure 9, left, for a sketch). A maximum volume fraction $\alpha = 0.3$ is allowed during the design process. Owing to symmetry, the simulation is carried out on half of the domain, which is discretized by a uniform tetrahedral mesh, $\mathcal{T}_h^{(0)}$, with 12000 elements.
The filter radius is set to $R = 0.05$. The mesh adaptation procedure is driven by the tolerance $\tau = 6$ and by the admissible element size range $[h_{\min}, h_{\max}] = [5 \cdot 10^{-3}, 2.5 \cdot 10^{-1}]$. Finally, the metric rescaling procedure is enabled with target number of vertices $N_v^* = 4000$.

Table 6 and Figure 10 provide qualitative and quantitative information on the performance of the QUBOSIMPATY algorithm, which converges in $i = 2$ iterations. The optimized structure is consistent with the literature and with the imposed loads and boundary conditions.

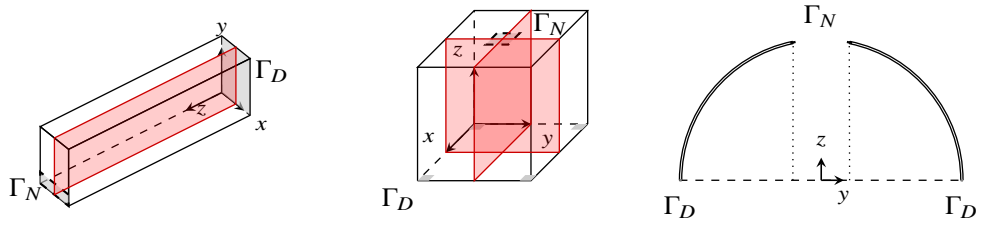

Figure 9: 3D test cases: sketch of the design domain for the cantilever (left) and the stool (center), and a cross-sectional view of the dome (right). The symmetry planes for the cantilever and stool test cases are highlighted in red.

The effect of anisotropic mesh adaptation is clearly visible. Despite the limited number of vertices, namely $N_v = 3965$ at convergence, enforced by the metric rescaling procedure, the adapted mesh concentrates tetrahedra along the material-void interface and yields a predominantly smooth structure.
The values in Table 6 confirm the better performance of the TS sampler compared with SA in terms of time to convergence. In particular, the computational time is reduced by a factor of about 3 for $i = 0$ and of about 7 for $i = 1, 2$. Differences in grayness, total variation, compliance minimization and volume fraction are instead negligible or absent. Finally, Figure 11 compares the convergence histories of QUBOSIMPATY, for the two samplers, with that of the SIMPATY algorithm over the whole optimization process. In contrast to what was observed in the 2D case, no significant difference in convergence speed is detected between QUBOSIMPATY and SIMPATY, as indicated by the nearly overlapping solid lines. Moreover, the volume fraction constraint is exactly satisfied at convergence by all three schemes. The jumps occurring at the iterations where mesh adaptation is performed are consistent with the behavior typically observed when combining mesh adaptation procedures with optimization cycles [66].

Table 6: The 3D cantilever beam: sensitivity of the QUBOSIMPATY design to the selected sampler in terms of structure performance and quality.

i	sampler	VF	c	G	TV	T_{conv}
0	TS	0.296	2.7469	0.1594	1.5225	412
	SA	0.296	2.7150	0.1548	1.4971	1478
1	TS	0.296	1.8978	0.0887	1.5913	579
	SA	0.296	1.8777	0.0893	1.6077	3612
2	TS	0.299	0.9063	0.0744	1.5746	547
	SA	0.299	0.9060	0.0724	1.5634	3401

3D test case 2: the stool We consider the cube $\Omega = (0, 1)^3$ as the design domain that we load with the surface traction $\mathbf{f} = [0, 0, 1]^T$ on $\Gamma_N = \{(x, y, z) \in \partial\Omega : z = 1, |x - 0.5| \leq$

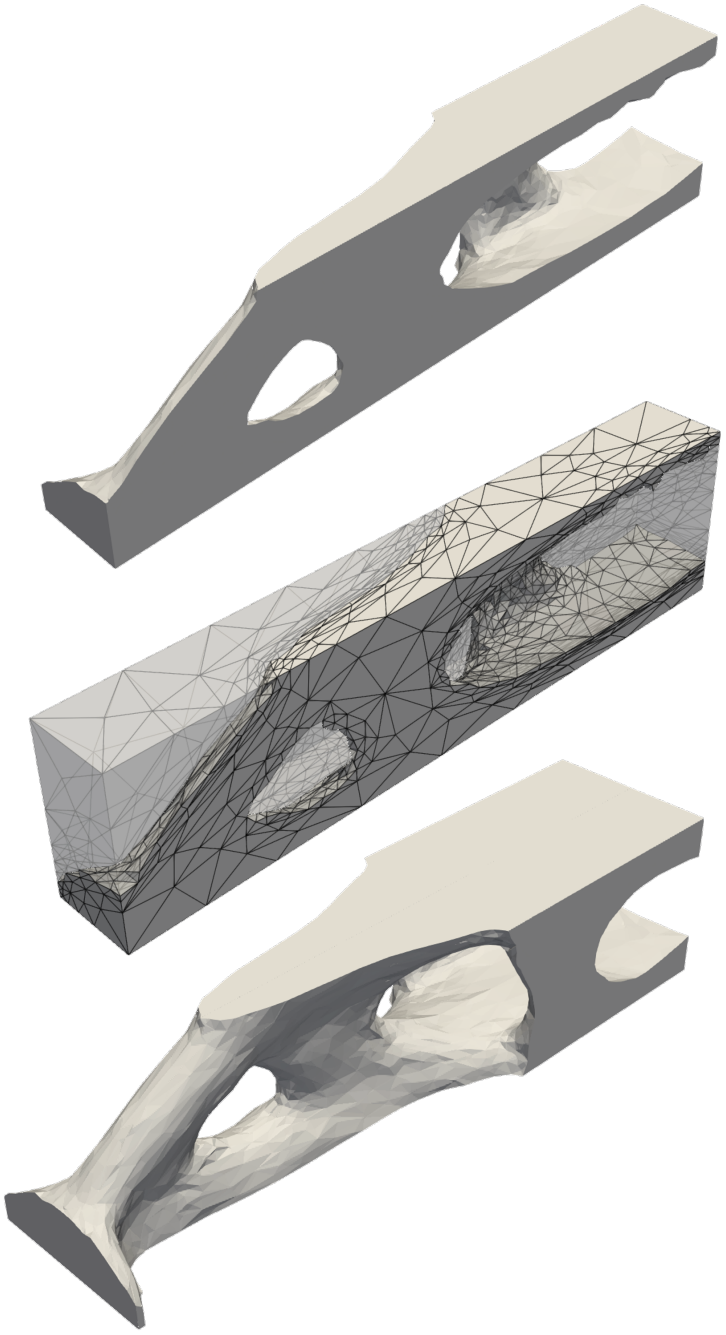

Figure 10: The 3D cantilever beam: half design (top), corresponding anisotropic adapted mesh (center), and final structure (bottom).

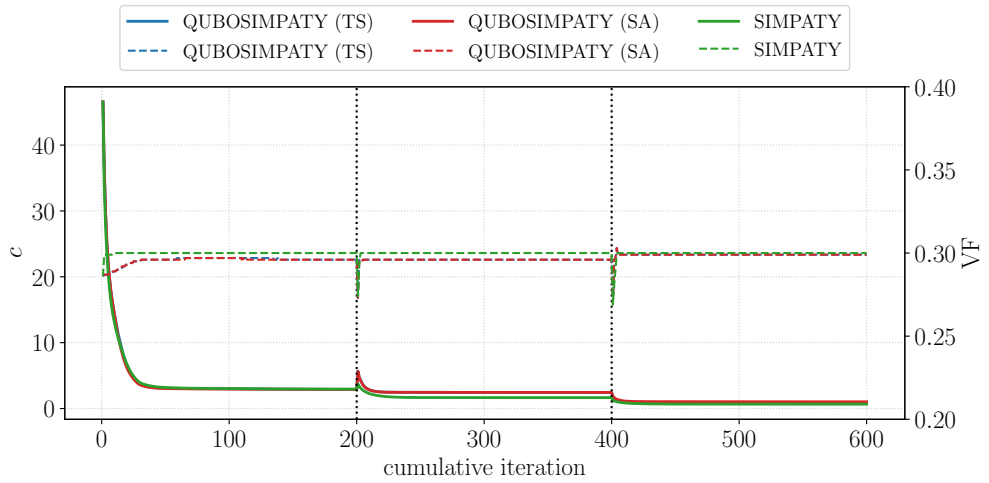

Figure 11: The 3D cantilever beam: convergence history of the compliance (solid lines) and of the volume fraction (dotted lines) for the standard SIMPATY and the QUBOSIMPATY algorithm, with two samplers. Mesh adaptation takes place in correspondence with the vertical lines.

$0.1, |y - 0.5| \leq 0.1\}$ (see Figure 9, center, for a sketch). The structure is clamped at the four corners of the bottom face, being $\Gamma_D = \{(x, y, z) \in \partial\Omega : z = 0, |x - 0.5| \geq 0.4, |y - 0.5| \geq 0.4\}$. The prescribed volume fraction is $\alpha = 0.2$.
Exploiting symmetry, the optimization is carried out on one quarter of Ω, discretized with a uniform tetrahedral mesh characterized by 12000 elements.
We set the filter radius R to 0.03, while we control the mesh adaptation procedure through parameters $\tau = 6$, $h_{\min} = 5 \cdot 10^{-3}$ and $h_{\max} = 2.5 \cdot 10^{-1}$. The resulting metric is successively rescaled in order to guarantee a final number of vertices equal to $N_v^* = 6000$. Finally, the TS sampler is employed.

We use this configuration to further compare QUBOSIMPATY and SIMPATY, with particular attention to structural performance and quality. The two algorithms produce fully comparable and remarkably smooth final layouts, as shown in Figure 12. This behavior is enabled by the use of anisotropic adapted meshes, which have approximately the same number of vertices and comparable maximum aspect ratios (see Table 7).
The quantitative results in Table 7 also show that the volume fraction constraint is exactly satisfied by both methods, although QUBOSIMPATY reaches this trend only at convergence, as for the 3D cantilever test. Finally, the compliance and total variation values do not exhibit significant differences between the two approaches, whereas, at convergence, the grayness obtained with QUBOSIMPATY is reduced by a factor of about 3 compared with SIMPATY.

3D test case 3: the dome We consider a dome-like geometry, defined on a non-Cartesian design domain consisting of a hemispherical shell with outer radius 1.25 and thickness 0.02. A top hole is obtained by removing a cylinder of radius 0.25, aligned with the z-axis and passing

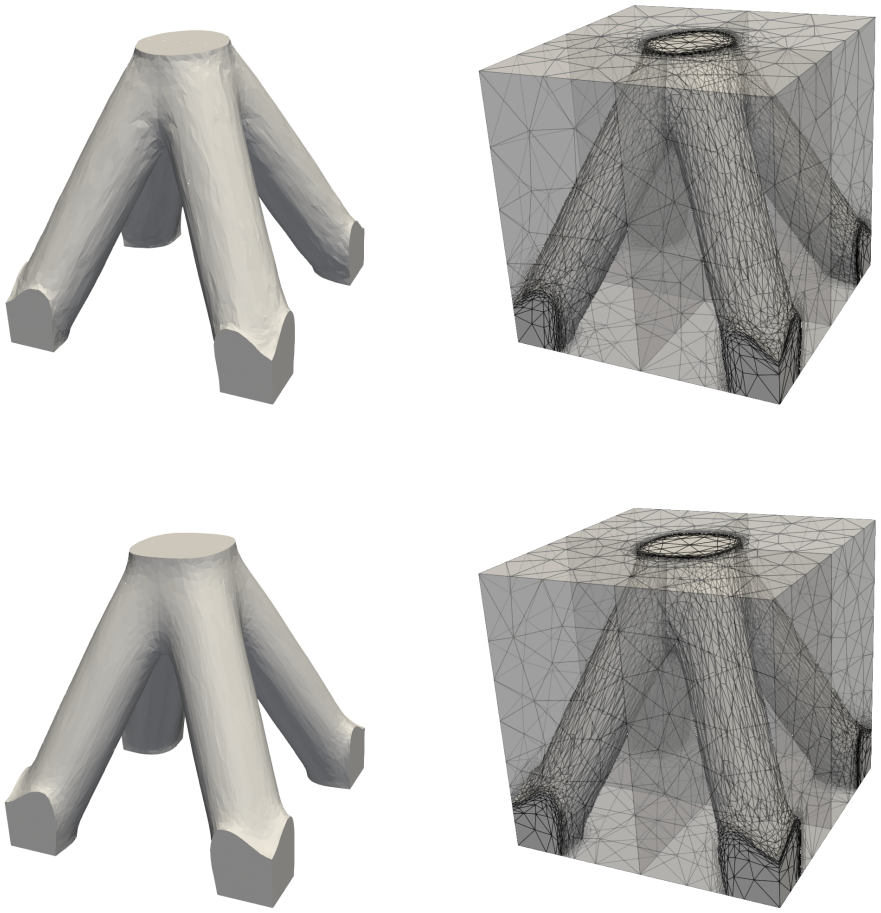

Figure 12: The stool: final design (left) provided by the QUBOSIMPATY (top) and SIMPATY (bottom) algorithms and corresponding anisotropic adapted meshes (right).

Table 7: The stool: comparison between the QUBOSIMPATY and SIMPATY algorithms in terms of structure performance and quality, and mesh cardinality and maximum aspect ratio.

i	algorithm	VF	c	G	TV	N_v	$s_{K,\max}$
0	QUBOSIMPATY	0.199	0.0915	0.0877	1.0490	2541	2.71
	SIMPATY	0.200	0.0906	0.0879	1.0434	2541	2.71
1	QUBOSIMPATY	0.199	0.0713	0.0392	1.0810	6079	18.69
	SIMPATY	0.200	0.0709	0.0465	1.0623	6086	16.47
2	QUBOSIMPATY	0.200	0.0577	0.0286	1.0594	5811	45.78
	SIMPATY	0.200	0.0590	0.0763	1.0118	5828	60.47

through the origin. The resulting domain is

$$\Omega = \left\{(x, y, z) \in \mathbb{R}^3 : 1.23^2 < x^2 + y^2 + z^2 < 1.25^2,\ z > 0,\ x^2 + y^2 > 0.25^2\right\}.$$

The structure is clamped at the bottom $\Gamma_D = \{(x, y, z) \in \partial\Omega : z = 0\}$, while the load $\mathbf{f} = [0, 0, -1]^T$ is applied on $\Gamma_N = \{(x, y, z) \in \partial\Omega : \sqrt{x^2 + y^2} = 0.25\}$ (see Figure 9, right, for a sketch).
The optimization process is performed on a quarter of the whole design domain by symmetry, discretized through an isotropic computational mesh consisting of 11918 tetrahedra. The filtering process is performed with a filter radius of $R = 0.05$, while we use parameters $\tau = 4$, $h_{\min} = 5 \cdot 10^{-3}$ and $h_{\max} = 2.5 \cdot 10^{-1}$ to adapt the mesh. The metric rescaling procedure is switched on with $N_v^* = 3000$ and we select TS as sampler.

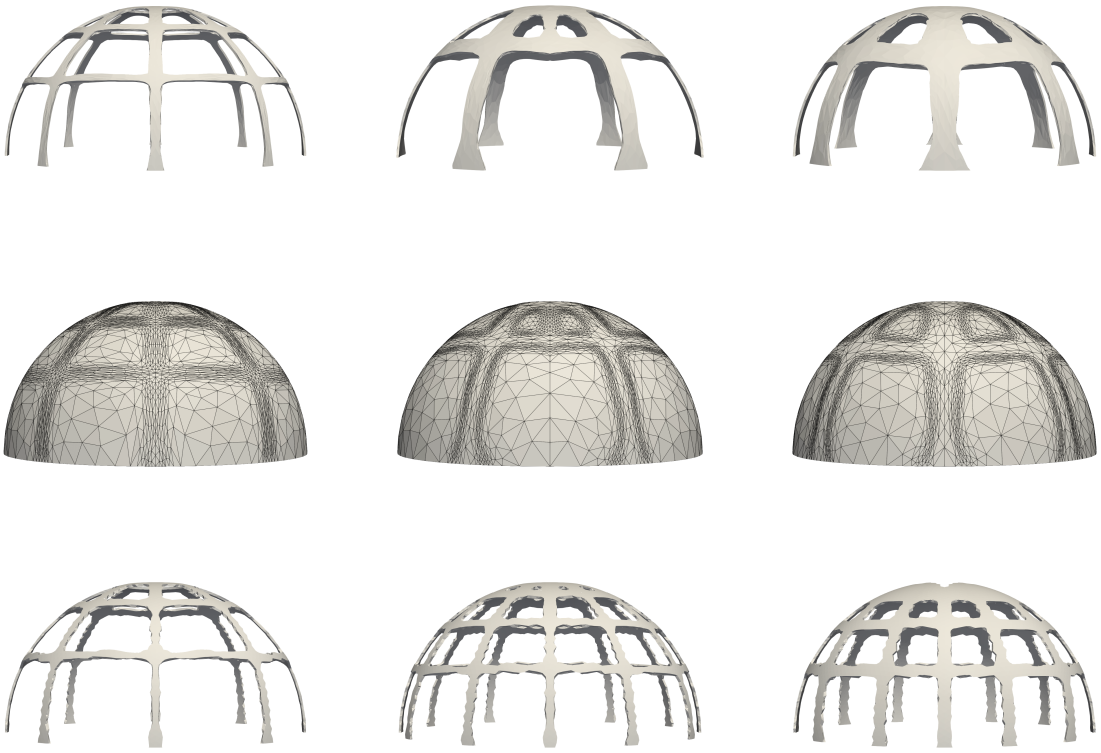

Figure 13: The dome: final design for $\alpha = 0.2$ (left), $\alpha = 0.3$ (center), and $\alpha = 0.4$ (right) provided by QUBOSIMPATY algorithm (first row) on the corresponding final anisotropic adapted meshes (second row) and by the standard QUBO algorithm (third row).

To assess the robustness of QUBOSIMPATY for shell-like optimized layouts, we consider three values of the maximum volume fraction, namely $\alpha \in \{0.2, 0.3, 0.4\}$. Figure 13 shows the optimized structures in the first row and the corresponding anisotropic adapted meshes in the

second row. The target number of vertices is reached within the prescribed tolerance θ, with $N_v = 3097$ for $\alpha = 0.2$, $N_v = 3119$ for $\alpha = 0.3$, and $N_v = 3060$ for $\alpha = 0.4$. It follows that, even for the smallest amount of available material, the algorithm successfully produces a meaningful optimized design. The adapted meshes exhibit stretched tetrahedra aligned with the struts and boundaries of the optimized dome. The maximum aspect ratio values are $s_{K,\max} = 88.47$, 57.58, and 63.27, while the corresponding average values are $s_{K,\text{avg}} = 10.84$, 9.11, and 9.24, for $\alpha = 0.2$, 0.3, and 0.4, respectively.

The regularizing effect of anisotropic mesh adaptation is evident by comparing the structures in the first row with those in the third row of Figure 13, obtained with the baseline QUBO algorithm on a fixed mesh with 3964 vertices. Despite the use of about 1000 additional vertices, the fixed mesh results do not provide the same quality of the final designs, which exhibit less regular material-void interfaces and a more fragmented structural layout. This effect is particularly visible for $\alpha = 0.3$ and $\alpha = 0.4$, where the QUBO algorithm produces more complex topologies whose accurate representation would likely require a substantially larger number of elements on a fixed mesh.

7 Discussion and conclusions

In this work, we have introduced QUBOSIMPATY, a QUBO-based topology optimization algorithm that combines a SIMP formulation of the minimum compliance problem with anisotropic mesh adaptation and a metric rescaling-based complexity control. The proposed strategy is motivated by the need to reduce the number of binary variables in QUBO formulations, a key requirement for their practical use with classical heuristics, quantum-inspired solvers, and near-term quantum annealing hardware.

In more detail, the QUBO reformulation is based on an additive binary update of the density variables and exploits standard compliance sensitivities, thereby retaining a close connection with classical density-based topology optimization. The anisotropic mesh adaptation component is used to concentrate the discretization along material-void interfaces, while the metric rescaling procedure allows the number of mesh vertices, and hence the number of QUBO variables, to be prescribed within a given tolerance. In this sense, mesh adaptation acts as a complexity-control mechanism rather than only as a discretization improvement tool.

The numerical results provide an extensive assessment of the proposed strategy on two- and three-dimensional benchmark configurations. Overall, QUBOSIMPATY produces optimized layouts that are consistent with those obtained by classical SIMP-based approaches, while keeping the number of density degrees of freedom under control. The observed trends are broadly coherent when moving from two to three dimensions, although the computational balance becomes more delicate in three dimensions because of the larger design spaces and the stronger impact of the QUBO size.

These observations support the use of mesh-complexity control as a structural component of QUBO-based topology optimization, particularly when the admissible problem size is dictated by quantum-oriented hardware restrictions. This role is highlighted by the downscaled experiment performed on D-Wave quantum annealing hardware. The results should not be interpreted as an end-to-end speedup over classical topology optimization, since the overall wall-clock time is

still affected by QUBO construction, hardware access, and communication overheads. Instead, the experiment shows that QUBOSIMPATY can produce hardware-compatible QUBO instances while preserving the mesh resolution needed to describe the optimized layout. Within this setting, the intrinsic annealing contribution exhibits a favorable time scale, and the convergence behavior suggests that QUBO-based minimization may become advantageous in regimes where faster sampling can compensate for the current overheads.

Several directions remain open for future work. A first issue concerns the dense structure of the present QUBO matrices, which affects both the classical solution cost and the embedding complexity on quantum annealing hardware. Alternative formulations leading to sparser QUBO structures, possibly through the introduction of auxiliary variables [87], represent a natural continuation of this study. Further developments concern more systematic strategies for selecting the density increment and the penalty parameters, as well as the integration of adaptive encoding techniques for real-valued design variables [88]. Finally, the combination of QUBOSIMPATY with other quantum, quantum-inspired, or hybrid optimization strategies [89, 90] will be important to assess the potential of quantum-ready topology optimization beyond the problem sizes currently accessible to available hardware.

Acknowledgments

All the authors are members of the Gruppo Nazionale Calcolo Scientifico Istituto Nazionale di Alta Matematica (GNCS–INdAM).
A.A.Z. and S.P. acknowledge the INdAM-GNCS 2026 project “Metodi Numerici Avanzati per Problemi di Fluidodinamica Computazionale”.
N.F. acknowledges the INdAM-GNCS 2026 project “Metodi Computazionali Avanzati per Sistemi Complessi: Approcci Ibridi Numerici e Machine Learning”.
The simulations contained in this paper have been (partially) carried out on the HPC cluster of the Mathematics Department of Politecnico di Milano. We acknowledge the CINECA award under the ISCRA C initiative, for the availability of quantum computing resources and support.

References

[1] Martin P. Bendsøe and Ole Sigmund. *Topology Optimization: Theory, Methods, and Applications*. Springer, Berlin, Heidelberg, 2004.

[2] Li Peng, Shen Xiuli, Dong Shaojing, and Fang Yufan. Topology optimization methods and its applications in aerospace: a review. *Struct. Multidiscip. Optim.*, 68(5):105, 2025.

[3] Jikai Liu, Andrew T. Gaynor, Shikui Chen, Zhan Kang, Krishnan Suresh, Akihiro Takezawa, Lei Li, Junji Kato, Jinyuan Tang, Charlie C. L. Wang, Lin Cheng, Xuan Liang, and Albert C. To. Current and future trends in topology optimization for additive manufacturing. *Struct. Multidiscip. Optim.*, 57(6):2457–2483, 2018.

[4] Martin P Bendsøe. *Optimization of Structural Topology, Shape, and Material*, volume 414. Springer, Berlin, 1995.

[5] George I. N. Rozvany. A critical review of established methods of structural topology optimization. *Struct. Multidiscip. Optim.*, 37(3):217–237, 2009.

[6] Mathias Stolpe and Krister Svanberg. An alternative interpolation scheme for minimum compliance topology optimization. *Struct. Multidiscip. Optim.*, 22(2):116–124, 2001.

[7] Eduardo Muñoz and Mathias Stolpe. Generalized benders' decomposition for topology optimization problems. *J. Global Optim.*, 51(1):149–183, 2011.

[8] Fred Glover, Gary Kochenberger, Rick Hennig, and Yu Du. Quantum bridge analytics I: a tutorial on formulating and using QUBO models. *Ann. Oper. Res.*, 314(1):141–183, July 2022.

[9] Andrew Lucas. Ising formulations of many NP problems. *Front. Physics*, 2:1–15, 2014.

[10] Yuki Sato, Ruho Kondo, Satoshi Koide, and Seiji Kajita. Quantum topology optimization of ground structures for near-term devices. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, page 168–176. IEEE, 2023.

[11] Giorgio Tosti Balducci, Boyang Chen, Matthias Möller, Marc Gerritsma, and Roeland De Breuker. Review and perspectives in quantum computing for partial differential equations in structural mechanics. *Front. Mech. Eng.*, 8:1–26, 2022.

[12] Zisheng Ye, Xiaoping Qian, and Wenxiao Pan. Quantum topology optimization via quantum annealing. *IEEE Trans. Quantum Eng.*, 4:1–15, 2023.

[13] Naruethep Sukulthanasorn, Junsen Xiao, Koya Wagatsuma, Reika Nomura, Shuji Moriguchi, and Kenjiro Terada. A novel design update framework for topology optimization with quantum annealing: Application to truss and continuum structures. *Comput. Methods Appl. Mech. Engrg.*, 437:117746, 2025.

[14] Yudai Suzuki, Shiori Aoki, Fabian Key, Katsuhiro Endo, Yoshiki Matsuda, Shu Tanaka, Marek Behr, and Mayu Muramatsu. An Ising machine formulation for design updates in topology optimization of flow channels. *Eng. Comput.*, 42(1):15, 2026.

[15] Zisheng Ye and Wenxiao Pan. Towards quantum accelerated large-scale topology optimization. *Comput. Methods Appl. Mech. Engrg.*, 453:118819, 2026.

[16] John Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018.

[17] Sheir Yarkoni, Elena Raponi, Thomas Bäck, and Sebastian Schmitt. Quantum annealing for industry applications: introduction and review. *Rep. Prog. Phys.*, 85(10):104001, sep 2022.

[18] Julien Dompierre, Marie-Gabrielle Vallet, Yves Bourgault, Michel Fortin, and Wagdi G Habashi. Anisotropic mesh adaptation: towards user-independent, mesh-independent and solver-independent CFD. Part III. Unstructured meshes. *Int. J. Numer. Methods Fluids*, 39(8):675–702, 2002.

[19] Simona Perotto, Nicolò Ferro, Giacomo Speroni, and Erika Temellini. Anisotropic recovery-based error estimators and mesh adaptation for real-life engineering innovation. In Franz Chouly, Stéphane P. A. Bordas, Roland Becker, and Pascal Omnes, editors, *Error Control, Adaptive Discretizations, and Applications, Part 4*, volume 61 of *Advances in Applied Mechanics*, pages 129–173. Elsevier, 2025.

[20] E. Stein and S. Ohnimus. Anisotropic discretization- and model-error estimation in solid mechanics by local Neumann problems. *Comput. Methods Appl. Mech. Engrg.*, 176(1-4):363–385, 1999.

[21] L. Formaggia, S. Micheletti, and S. Perotto. Anisotropic mesh adaption with application to CFD problems. In *Proceedings of the Fifth World Congress on Computational Mechanics (WCCM V)*, pages 1481–1493, 2002.

[22] R. El Khaoulani and P. O. Bouchard. An anisotropic mesh adaptation strategy for damage and failure in ductile materials. *Finite Elem. Anal. Des.*, 59:1–10, 2012.

[23] Y. Belhamadia, A. Fortin, and Y. Bourgault. On the performance of anisotropic mesh adaptation for scroll wave turbulence dynamics in reaction–diffusion systems. *J. Comput. Appl. Math.*, 271:233–246, 2014.

[24] J. Carpio, J. L. Prieto, and M. Vera. A local anisotropic adaptive algorithm for the solution of low-Mach transient combustion problems. *J. Comput. Phys.*, 306:19–42, 2016.

[25] N. Ferro, S. Micheletti, and S. Perotto. Anisotropic mesh adaptation for crack propagation induced by a thermal shock in 2D. *Comput. Methods Appl. Mech. Engrg.*, 331:138–158, 2018.

[26] B. Lèye, J. Koko, S. Kane, and M. Sy. Numerical simulation of saltwater intrusion in coastal aquifers with anisotropic mesh adaptation. *Math. Comput. Simul.*, 154:1–18, 2018.

[27] Stefano Micheletti, Simona Perotto, and Luca Soli. Topology optimization driven by anisotropic mesh adaptation: Towards a free-form design. *Comput. Struct.*, 214:60–72, 2019.

[28] Federico Gatti, Carlo de Falco, Simona Perotto, Luca Formaggia, and Manuel Pastor. A scalable well-balanced numerical scheme for the modeling of two-phase shallow granular landslide consolidation. *J. Comput. Phys.*, 501:112798, 2024.

[29] Francesco Vaccaro, Aurelio G. Mauri, Simona Perotto, Stefano Brivio, and Sabina Spiga. Modeling and simulation of electrochemical and surface diffusion effects in filamentary cation-based resistive memory devices. *Appl. Math. Model.*, 134:591–610, 2024.

[30] Nicola Ferro, Stefano Micheletti, and Simona Perotto. An optimization algorithm for automatic structural design. *Comput. Methods Appl. Mech. Engrg.*, 372:113335, 2020.

[31] Nicola Ferro, Stefano Micheletti, and Simona Perotto. Compliance–stress constrained mass minimization for topology optimization on anisotropic meshes. *SN Applied Sciences*, 2:1196, 2020.

[32] O. Sigmund. On the usefulness of non-gradient approaches in topology optimization. *Struct. Multidiscip. Optim.*, 43(5):589–596, May 2011.

[33] T. D. Lee and C. N. Yang. Statistical theory of equations of state and phase transitions. II. lattice gas and Ising model. *Phys. Rev.*, 87:410–419, Aug 1952.

[34] Barry A. Cipra. An introduction to the Ising model. *Amer. Math. Monthly*, 94(10):937–959, 1987.

[35] Rafael Cacao, Lucas R. C. T. Cortez, Jackson Forner, Hamidreza Validi, Ismael R. de Farias, and Illya V. Hicks. The set partitioning problem in a quantum context. *Optim. Lett.*, 18(1):1–17, 2024.

[36] Salwa Shaglel, Markus Kirsch, Marten Winkler, Christian Muench, Stefan Walter, Fritz Schinkel, and Martin Kliesch. A comprehensive benchmark of an Ising machine on the Max-Cut problem. *New J. Phys.*, pages 1–48, 2026.

[37] Abraham P Punnen. *The Quadratic Unconstrained Binary Optimization Problem*. Springer Cham, 2022.

[38] Prasanna Date, Davis Arthur, and Lauren Pusey-Nazzaro. QUBO formulations for training machine learning models. *Sci. Rep.*, 11(1):10029, 2021.

[39] Zeynab Kaseb, Matthias Möller, Pedro P Vergara, and Peter Palensky. Power flow analysis using quantum and digital annealers: a discrete combinatorial optimization approach. *Sci. Rep.*, 14(1):23216, 2024.

[40] Hyunju Lee and Kyungtaek Jun. Range dependent Hamiltonian algorithms for numerical QUBO formulation. *Sci. Rep.*, 15(1):8819, 2025.

[41] Abhishek Awasthi, Nico Kraus, Florian Krellner, and David Zambrano. Real world application of quantum-classical optimization for production scheduling. In *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*, volume 02, pages 239–244, 2024.

[42] Seongmin Kim, Sang-Woo Ahn, In-Saeng Suh, Alexander W Dowling, Eungkyu Lee, and Tengfei Luo. Quantum annealing for combinatorial optimization: a benchmarking study. *NPJ Quantum Inf.*, 11(1):77, 2025.

[43] Andrew D. King, Jack Raymond, Trevor Lanting, Richard Harris, Alex Zucca, Fabio Altomare, Andrew J. Berkley, Kelly Boothby, Sara Ejtemaee, Colin Enderud, Emile Hoskinson, Shuiyuan Huang, Eric Ladizinsky, Allison J. R. MacDonald, Gaelen Marsden, Reza Molavi, Travis Oh, Gabriel Poulin-Lamarre, Mauricio Reis, Chris Rich, Yuki Sato, Nicholas Tsai, Mark Volkmann, Jed D. Whittaker, Jason Yao, Anders W. Sandvik, and Mohammad H. Amin. Quantum critical dynamics in a 5,000-qubit programmable spin glass. *Nature*, 617(7959):61–66, 2023.

[44] Endre Boros, Peter L Hammer, and Gabriel Tavares. Local search heuristics for quadratic unconstrained binary optimization (QUBO). *J. Heuristics*, 13(2):99–132, 2007.

[45] Alzira Mota, Paulo Ávila, João Bastos, Luís A. C. Roque, and António Pires. Comparative analysis of simulated annealing and tabu search for parallel machine scheduling. *Procedia Comput. Sci.*, 256:573–582, 2025.

[46] Tameem Albash and Daniel A. Lidar. Adiabatic quantum computation. *Rev. Mod. Phys.*, 90:015002, Jan 2018.

[47] Vrinda Mehta, Hans De Raedt, Kristel Michielsen, and Fengping Jin. Understanding the physics of D-Wave annealers: From Schrödinger to Lindblad to Markovian dynamics. *Phys. Rev. A*, 112:032616, Sep 2025.

[48] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. Quantum computation by adiabatic evolution, 2000. arXiv:quant-ph/0001106.

[49] Bruno Hideki Fukushima-Kimura, Noda Eitaro, Kawamoto Noe, and Sakai Akira. Mathematical aspects of the digital annealer's simulated annealing algorithm. *J. Stat. Phys.*, 190(12):190, 11 2023.

[50] O. Sigmund. A 99 line topology optimization code written in Matlab. *Struct. Multidiscip. Optim.*, 21(2):120–127, April 2001.

[51] Krister Svanberg. The method of moving asymptotes-a new method for structural optimization. *Int. J. Numer. Methods Eng.*, 24(2):359–373, 1987.

[52] Dustin R. Jantos, Philipp Junker, and Klaus Hackl. An evolutionary topology optimization approach with variationally controlled growth. *Comput. Methods Appl. Mech. Engrg.*, 310:780–801, 2016.

[53] S. Y. Wang and K. Tai. Structural topology design optimization using genetic algorithms with a bit-array representation. *Comput. Methods Appl. Mech. Engrg.*, 194(36):3749–3770, 2005.

[54] O. Sigmund and J. Petersson. Numerical instabilities in topology optimization: A survey on procedures dealing with checkerboards, mesh-dependencies and local minima. *Struct. Multidiscip. Optim.*, 16(1):68–75, August 1998.

[55] B. S. Lazarov and O. Sigmund. Filters in topology optimization based on Helmholtz-type differential equations. *Int. J. Numer. Methods Eng.*, 86(6):765–781, May 2011.

[56] Jorge Nocedal and Stephen J Wright. *Numerical Optimization*. Springer New York, NY, 2006.

[57] Vicky Choi. Minor-embedding in adiabatic quantum computation: I. The parameter setting problem. *Quantum Inf. Process.*, 7(5):193–209, 2008.

[58] Andrew D. King and Catherine C. McGeoch. Algorithm engineering for a quantum annealing platform, 2014. arXiv:1410.2628.

[59] Matthew D. Piggott, Patrick E. Farrell, Christopher R. Wilson, Gerard J. Gorman, and Christopher C. Pain. Anisotropic mesh adaptivity for multi-scale ocean modelling. *Philos. Trans. R. Soc. A*, 367(1907):4591–4611, 2009.

[60] Patrick E Farrell, Stefano Micheletti, and Simona Perotto. An anisotropic Zienkiewicz–Zhu-type error estimator for 3D applications. *Int. J. Numer. Methods Eng.*, 85(6):671–692, 2011.

[61] Kedar C. Chitale, Michel Rasquin, Onkar Sahni, Mark S. Shephard, and Kenneth E. Jansen. Anisotropic boundary layer adaptivity of multi-element wings. *Comput. Fluids*, 103:61–71, 2014.

[62] Martin van Driel, Christian Boehm, Lion Krischer, and Michael Afanasiev. Accelerating numerical wave propagation using wavefield adapted meshes. Part I: Forward and adjoint modelling. *Geophys. J. Int.*, 221(3):1580–1590, 2020.

[63] Davor Dundovic, Joseph G. Wallwork, Stephan C. Kramer, Fabien Gillet-Chaulet, Regine Hock, and Matthew D. Piggott. Anisotropic metric-based mesh adaptation for ice flow modelling in Firedrake. *Geosci. Model Dev.*, 18:4023–4044, 2025.

[64] Nicola Ferro, Francesco Mezzadri, Dario Carbonaro, Emanuele Galligani, Diego Gallo, Umberto Morbiducci, Claudio Chiastra, and Simona Perotto. Designing novel vascular stents with enhanced mechanical behavior through topology optimization of existing devices. *Finite Elem. Anal. Des.*, 244:104304, 2025.

[65] N. Ferro, S. Perotto, D. Bianchi, R. Ferrante, and M. Mannisi. Design of cellular materials for multiscale topology optimization: application to patient-specific orthopedic devices. *Struct. Multidiscip. Optim.*, 65(3):79, 2022.

[66] Giacomo Speroni, Nicolò Mondini, Nicola Ferro, and Simona Perotto. A topology optimization framework for scaffold design in soilless cultivation. *To appear in Eng. Comput.*, 2026.

[67] Luca Formaggia and Simona Perotto. New anisotropic a priori error estimates. *Numer. Math.*, 89(4):641–667, 2001.

[68] S. Micheletti and S. Perotto. Anisotropic adaptation via a Zienkiewicz–Zhu error estimator for 2D elliptic problems. In *Numerical Mathematics and Advanced Applications*, pages 645–653. Springer Berlin, Heidelberg, 2010.

[69] O. C. Zienkiewicz and J. Z. Zhu. The superconvergent patch recovery and a posteriori error estimates. part 1: The recovery technique. *Int. J. Numer. Methods Eng.*, 33(7):1331–1364, 1992.

[70] O. C. Zienkiewicz and J. Z. Zhu. The superconvergent patch recovery and a posteriori error estimates. part 2: Error estimates and adaptivity. *Int. J. Numer. Methods Eng.*, 33(7):1365–1382, 1992.

[71] O. C. Zienkiewicz and J. Z. Zhu. A simple error estimator and adaptive procedure for practical engineerng analysis. *Int. J. Numer. Methods Eng.*, 24(2):337–357, 1987.

[72] P. J. Frey and P. L. George. *Mesh Generation: Application to Finite Elements*. ISTE, John Wiley & Sons, London, 2008.

[73] Julia Docampo-Sánchez, Guillaume Houzeaux, and Abel Gargallo-Peiró. Iterative metric rescaling for mesh adaptation with prescribed node count. In *Proceedings of the 2026 SIAM International Meshing Roundtable*, pages 74–85, 2026.

[74] Francesco Clerici, Nicola Ferro, Stefania Marconi, Stefano Micheletti, Erika Negrello, and Simona Perotto. Anisotropic adapted meshes for image segmentation: Application to three-dimensional medical data. *SIAM J. Imaging Sci.*, 13(4):2189–2212, 2020.

[75] O. Sigmund. On benchmarking and good scientific practise in topology optimization. *Struct. Multidiscip. Optim.*, 65(11):315–324, 2022.

[76] Martin P Bendsøe and Ole Sigmund. Material interpolation schemes in topology optimization. *Arch. Appl. Mech.*, 69(9):635–654, 1999.

[77] Yuta Atobe, Masashi Tawada, and Nozomu Togawa. Hybrid annealing method based on subQUBO model extraction with multiple solution instances . *IEEE Trans. Comput.*, 71(10):2606–2619, 2022.

[78] Ole Sigmund. Morphology-based black and white filters for topology optimization. *Struct. Multidiscip. Optim.*, 33(4):401–424, 2007.

[79] Christian Clason, Florian Kruse, and Karl Kunisch. Total variation regularization of multi-material topology optimization. *ESAIM: Math. Model. Numer. Anal.*, 52(1):275–303, 2018.

[80] Igor A. Baratta, Joseph P. Dean, Jørgen S. Dokken, Michal Habera, Jack S. Hale, Chris N. Richardson, Marie E. Rognes, Matthew W. Scroggs, Nathan Sime, and Garth N. Wells. DOLFINx: The next generation FEniCS problem solving environment, 2025. doi:10.5281/zenodo.10447665.

[81] Yingqi Jia, Chao Wang, and Xiaojia Shelly Zhang. FEniTop: A simple FEniCSx implementation for 2D and 3D topology optimization supporting parallel computing. *Struct. Multidiscip. Optim.*, 67(140):140, August 2024.

[82] F. Hecht. New development in FreeFem++. *J. Numer. Math.*, 20(3-4):251–265, 2012.

[83] Charles Dapogny, Cécile Dobrzynski, and Pascal J. Frey. Three-dimensional adaptive domain remeshing, implicit domain meshing, and applications to free and moving boundary problems. *J. Comput. Phys.*, 262:358–378, 2014.

[84] D-Wave Systems Inc. Ocean software documentation, 2026. URL: https://docs.ocean.dwavesys.com/.

[85] Catherine McGeoch and Pau Farré. The D-Wave advantage system: An overview. *D-Wave Systems Inc., Burnaby, BC, Canada, Tech. Rep*, 2020. URL: https://www.dwavequantum.com/resources/white-paper/the-d-wave-advantage-system-an-overview/.

[86] Kelly Boothby, Paul Bunyk, Jack Raymond, and Aidan Roy. Next-generation topology of D-Wave quantum processors, 2020. arXiv:2003.00133.

[87] Kohei Suda, Soshun Naito, and Yoshihiko Hasegawa. Sparse QUBO formulation for efficient embedding via network-based decomposition of equality and inequality constraints, 2026. arXiv:2601.18108.

[88] Fabian Key, Lukas Freinberger, Mayu Muramatsu, and Norbert Hosters. Adaptive encoding strategy for quantum annealing in mixed-variable engineering optimization, 2026. arXiv:2603.17506.

[89] Osama Muhammad Raisuddin and Suvranu De. FEqa: Finite element computations on quantum annealers. *Comput. Methods Appl. Mech. Engrg.*, 395:115014, 2022.

[90] Naruethep Sukulthanasorn and Kenjiro Terada. Quantum neural network-assisted topology optimization: Concept and implementation with parameterized quantum circuits. *Comput. Methods Appl. Mech. Engrg.*, 448:118411, 2026.

MOX Technical Reports, last issues

Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

68/2026 Patanè, G.; Greven, S.; Menafoglio, A.
Random mixtures in Bayes Hilbert spaces

67/2026 Ciardulli, S.; Fontana, N.; Vantini, S.; Ieva F.
Generalized propensity score weighting for functional causal inference framework

66/2026 Antonietti, P.F.;Bonizzoni, F.;Corti, M.; De March, N.;Di Noto, S.;Regazzoni, F.
A stability-preserving polytopal discontinuous Galerkin method for the Fisher-Kolmogorov model with applications to neurodegenerative diseases

65/2026 De Sanctis, M.F.; Arnone, E.; Ieva, F.; Sangalli, L.M.
Modeling group heterogeneity in spatio-temporal data via physics-informed regression

64/2026 Radisic, I.; Tirotta, R.; Zingaro, A.; Pagani, S.; Dede', L.
Physics-Informed Neural Networks for the High-Resolution Reconstruction of Flow Measurement Indicators in Fluid Dynamics

63/2026 Ciaramella, G.; Gong, W.; Kwok, F., Tan, Z.
Uniform Convergence of the Schwarz Alternating Method for Optimal Control Problems

60/2026 Ciaramella, G.; Gander, M.J.; Van Criekingen, S.; Vanzan, T.
Algebraic and Two-Level Parallel Substructured Schwarz Methods

62/2026 Ciaramella, G.; Gong, W.; Kwok, F.; Tan, Z.
On the Uniform Convergence Analysis of the Schwarz Alternating Method for Optimal Control Problems

58/2026 Mapelli, A.; Massi, M.C.; Cuccuru, G.; Di Angelantonio, E.; Ieva, F.
Prior-informed conditional Gaussian graphical models: an application to protein interaction network reconstruction

57/2026 Fontana, N.; Secchi, P.; Di Angelantonio, E.; Ieva, F.
Modeling time-varying genetic effects on binary disease risk via functional Mendelian Randomization