

Standard Quadratic Formulations of Many NP Problems: A Simplex-Based Compilation Framework for Combinatorial Optimization

Mohammad-Ali Miri*, Babak Emami, Nicholas Chancellor, PoJen Wang

Quantum Computing Inc (QCi), 5 Marine View Plaza, Hoboken, NJ 07030, USA

amiri@quantumcomputinginc.com

September 29, 2026

Abstract

The standard quadratic program (StQP) minimizes a quadratic form over nonnegative variables that sum to one. We compose classical graph reductions with regularized Motzkin–Straus clique formulations to express discrete optimization problems in this continuous domain. The graph matrix has diagonal entries τ , zeros on edges, and ones on nonedges. For $0 < \tau < 1$, its minimum is $\tau/\omega(G)$, where $\omega(G)$ is the clique number. Its strict local minimizers are precisely the uniform distributions on maximal cliques, and its global minimizers encode maximum cliques. At $\tau = 1/2$, integer scaling gives coefficients in $\{0, 1, 2\}$ and minimum $1/\omega(G)$, yielding an NP-complete StQP threshold problem with a restricted coefficient alphabet. We give explicit formulations for satisfiability, coloring, Hamiltonian cycles, independent set, vertex cover, set packing, three-dimensional matching, and graph isomorphism. A regularized weighted clique formulation combined with local-state compatibility graphs gives an exact compiler for finite-domain factor models specified by complete local tables, including QUBO, with at most four simplex coordinates per binary pair factor. The catalog covers Karp’s 21 problems: twelve use direct graph formulations, and nine use factor-state formulations, including six obtained through binary-linear feasibility. For each route we record dimensions, coefficient structure, and recovery rules. We analyze interaction count, coefficient range, objective separation, perturbation tolerance, support recovery, and decoding overhead. The separation bounds quantify the effects of clique size, factor weights, and offsets. In the complete factor-state construction, every assignment, including each suboptimal assignment, is a strict local minimum.

Keywords: standard quadratic programming; simplex; NP-completeness; maximum clique; constraint satisfaction; QUBO; Ising model; combinatorial optimization; analog optimization; simplex quadratic annealer.

Contents

1	Introduction	4
1.1	Scope and contributions	4
1.2	Relation to prior work	5
1.3	Computational scope	5
1.4	Organization	6
2	Preliminaries and formulation standards	6
2.1	The standard quadratic program	6
2.2	Representation metrics	7
2.3	Decision thresholds and certificates	7
2.4	Compilation pipeline	7
3	A three-level graph-to-StQP compiler	8
3.1	The graph matrix	8
3.2	An integer coefficient alphabet	9
3.3	Objective separation and perturbation tolerance	10
3.4	Landscape interpretation	11
3.5	Relative couplings	11
4	Compatibility graphs for constraint satisfaction and assignment	11
4.1	Finite-domain constraint satisfaction	11
4.2	CNF satisfiability	12
4.3	Graph coloring and clique cover	13
4.4	Hamiltonian cycles	13
4.5	Graph isomorphism and association graphs	13
4.6	Arbitrary-arity hard constraints	14
5	Weighted clique and a generic factor-state compiler	14
5.1	A weighted clique StQP theorem	14
5.2	Finite-domain factor models	17
5.3	Sparsity of the incompatibility matrix	19
5.4	QUBO as a special case	19
5.5	Worked example: weighted Max-Cut	20
5.6	Approximation distortion	20
5.7	Objective separation	21
6	Selection, packing, covering, and matching formulations	23
6.1	Clique, independent set, and vertex cover	23
6.2	Set packing	23
6.3	Three-dimensional matching	24
6.4	Exact cover through weighted clique	24
6.5	Hitting set and set cover	25
6.6	Dominant sets and weighted graph structure	25
7	A generic binary-linear feasibility compiler	25
7.1	From inequalities to a zero-penalty QUBO	25
7.2	Size accounting	26
7.3	Adding an objective or decision threshold	27
7.4	Precision of the binary-linear route	27
7.5	Representative constructions	28
8	A catalog for Karp's 21 NP-complete problems	29
8.1	Direct versus generic routes	34
8.2	Threshold normalization	34

9	Representation cost and comparison with QUBO/Ising	35
9.1	State semantics	35
9.2	Coefficient structure	35
9.3	Connectivity and baseline removal	35
9.4	Dimension	36
9.5	Precision and dynamic range	37
9.6	Landscape and local solvers	37
9.7	Decoding and verification	38
9.8	Summary of representation costs	38
10	A hardware-neutral simplex quadratic optimizer/analyzer	38
10.1	Input contract	39
10.2	State contract	39
10.3	Energy contract	39
10.4	Search contract	39
10.5	Output contract	40
10.6	Timing and energy accounting	40
10.7	A minimal graph-coded device mode	40
10.8	Evaluation metrics	41
11	Limitations and open research questions	41
11.1	Universality can hide poor formulations	41
11.2	Continuous precision is a computational resource	41
11.3	Approximation theory	41
11.4	Local landscape beyond cliques	42
11.5	Alternative simplex lifts	42
11.6	Global certification	43
11.7	Formulation search as an optimization problem	43
11.8	Empirical question	43
12	Conclusion	43
A	Compilation procedures	44
B	Detailed decoding and purification	44
B.1	Purifying an approximate graph solution	44
B.2	Factor-state decoding	45
C	Explicit formulations for selected catalog entries	45
C.1	Independent set and vertex cover	45
C.2	Coloring matrix blocks	45
C.3	Hamiltonian matrix blocks	45
C.4	Max-Cut local states	45
D	A note on copositive certification	46
E	Notation summary	46

1 Introduction

The Ising Hamiltonian and its equivalent quadratic unconstrained binary optimization (QUBO) form serve as common compilation targets for optimization algorithms and hardware. A large body of work maps logical, graph, packing, covering, routing, assignment, and scheduling problems to pairwise binary interactions, and those mappings have helped define the interface of quantum annealers, coherent Ising machines, simulated bifurcation systems, and related platforms [1–9]. Lucas’s catalog of Ising formulations for many NP problems, including Karp’s 21 classical NP-complete problems, is an especially influential expression of this viewpoint [3, 10].

That interface is binary by construction. For a device whose native state is a normalized, nonnegative vector, such as an intensity or population profile, a simplex formulation may provide a more direct interface. Converting to a binary representation can introduce discretization, one-hot constraints, or consistency penalties, depending on the encoding; these costs are not inherent to every Ising model. This paper develops a formulation catalog for the simplex and compares its costs with those of binary alternatives.

The continuous model considered here is

$$\nu(Q) := \min_{x \in \Delta_N} x^T Q x, \quad \Delta_N := \left\{ x \in \mathbb{R}^N : x \geq 0, \mathbf{1}^T x = 1 \right\}, \quad (1)$$

where $Q \in \mathbb{S}^N$. This is the standard quadratic program (StQP) [11]. The state is continuous but compact, nonnegative, and normalized. A coordinate may be interpreted as mass, probability, population, intensity, or allocation, and Q_{uv} measures a pairwise interaction between coordinates u and v .

At first sight, replacing binary variables by simplex variables seems poorly suited to exact combinatorial encoding. Normalization is what makes it work: it turns positive cross terms into incompatibility penalties. If two coordinates represent mutually inconsistent local choices, assigning mass to both incurs a quadratic cost. If compatible choices interact with coefficient zero, global minimization seeks a mutually compatible support. A diagonal term controls the concentration of mass on that support. For a carefully chosen three-level matrix, the optimal support is a maximum clique and the optimal vector is its barycenter.

In the formulations studied here, discrete choices are recovered from the support of a normalized continuous state. The matrix specifies incompatible pairs, while diagonal coefficients determine the optimum on each compatible support. We compare the dimension, connectivity, precision, and decoding costs of this representation with those of problem-specific binary formulations.

1.1 Scope and contributions

The regularized unweighted and weighted clique formulations provide the mathematical foundations [12–14]. We state and prove them in a common minimization normalization so that the matrix entries, objective values, and decoding rules can be used throughout the catalog. The paper develops four aspects of this formulation framework.

C1. Explicit formulations and a complete catalog. We compose classical compatibility-graph constructions with regularized clique StQP formulations. For Karp’s 21 problems, twelve routes use direct unweighted or weighted graphs and nine use factor-state formulations, six of them through binary-linear feasibility. Each entry specifies its route, dimension or encoding-bit count, decision or recovery rule, and principal representation cost.

C2. A common factor-state formulation and its landscape. We combine positive local-state weights, consistency graphs, and a weighted clique matrix with binary off-diagonal entries. For complete rational factor tables the construction is polynomial in their explicit encoding length; a

binary model with n_u unary factors and m pair factors uses at most $2n_u + 4m$ simplex coordinates. We give an exact recovery map and show that every complete assignment becomes a strict local minimum (Corollary 5.8).

- C3. Objective separation and decoding guarantees.** We record the NP-complete threshold problem associated with the coefficient alphabet $\{0, 1, 2\}$ and derive graph decision-value separations, perturbation tolerances, and factor-model separation bounds. We analyze how offsets and attainable energy spacings enter these bounds, and connect objective accuracy to a specified decoding procedure.
- C4. Representation-cost comparison.** We compare simplex dimension, interaction count, coefficient range, objective resolution, support recovery, and decoding overhead with strong problem-specific binary formulations. These quantities define the costs that a computational evaluation must measure.

1.2 Relation to prior work

The reduction of NP decision problems to StQP follows by composing the classical reduction to CLIQUE [10, 15] with the Motzkin–Straus formulation and its regularizations [12, 13, 16, 17]. StQP itself is a well-established class in global optimization [11, 18–20]. The direct SAT, coloring, Hamiltonian, and packing routes use classical reductions and compatibility constructions; composing them with the clique StQP supplies explicit matrices and thresholds rather than new representability results.

Compatibility and microstructure graphs are classical in constraint processing [21–23]. Association graphs underlie matching and subgraph isomorphism [24, 25], and dual microstructures yield the allowed-tuple construction for arbitrary-arity constraints [26]. Reductions of finite-domain maximum a posteriori estimation to maximum-weight independent set also predate this work [27]. The factor-state compiler combines this type of consistency construction with a weighted clique StQP. Weighted and regularized clique formulations have substantial precedents [14, 28–33]; the normalization comparison after Theorem 5.3 specifies the relation to Bomze et al. [14].

Two formulation catalogs provide useful points of comparison. Lucas [3] uses binary Ising interactions, whereas Cai et al. [34] develops first-order and multiphase Potts formulations and a three-state oscillator optimizer. The present catalog targets a quadratic objective over one globally normalized, nonnegative state. Its emphasis is the explicit matrix and decoding description, together with dimension, coefficient, separation, and landscape costs; it does not introduce an oscillator architecture or establish solver performance.

Relaxation labeling provides another relevant comparison. Its per-object label distributions have separate normalization constraints, with quadratic compatibility objectives and well-studied dynamics [35]. Such products of simplexes differ from the single simplex used here. Section 11.5 compares these feasible sets and explains why their local-minimum structures need not agree.

1.3 Computational scope

An exact formulation does not imply an efficient solution method. Unless $P = NP$, no polynomial-time algorithm can solve every encoded instance. Nor does a compact coefficient alphabet guarantee a large physical energy gap, robust analog behavior, or quantum advantage. The local-minimum correspondences characterize the optimization landscape; they do not determine the convergence rate or success probability of a solver.

We use *StQP* for the mathematical optimization model. When discussing a hypothetical hardware platform, we use *simplex quadratic optimizer/annealer* and abbreviate it as *StQA*. Nothing in the formulations assumes that annealing is the only or best physical solution mechanism.

1.4 Organization

Section 2 defines exactness and representation metrics. Section 3 develops the graph compiler and its complexity and precision consequences. Section 4 gives direct compatibility formulations. Section 5 develops the weighted factor-state and QUBO compilers. Section 6 treats graph selection, packing, and exact cover. Section 7 gives a generic binary-linear route. Section 8 catalogs Karp’s 21 problems. Section 9 compares representation costs with QUBO/Ising. Section 10 states a hardware-neutral StQA specification and evaluation metrics. Section 11 collects the mathematical, design, and empirical limitations, and Section 12 concludes. The appendices give compilation procedures (Section A), decoding and purification (Section B), explicit formulations for selected catalog entries (Section C), the copositive connection (Section D), and notation (Section E).

2 Preliminaries and formulation standards

2.1 The standard quadratic program

For $Q \in \mathbb{S}^N$, define

$$\text{StQP}(Q) : \quad \nu(Q) = \min_{x \in \Delta_N} x^\top Q x. \quad (2)$$

The simplex Δ_N has N vertices and affine dimension $N - 1$. We use the convention that the superscript or subscript records the number of coordinates, not the intrinsic dimension. We write $\mathbb{S}^N$ for the space of real symmetric $N \times N$ matrices, $\mathbf{1}$ for the all-ones vector, I_N for the identity matrix, and e_i for the i th coordinate vector. We use $[n] = \{1, \dots, n\}$. For $x \in \Delta_N$, its support is $\text{supp}(x) := \{i : x_i > 0\}$. A nonempty graph has at least one vertex; it need not have an edge. Empty instances, empty domains, and trivial or out-of-range decision thresholds are handled before compilation so that the simplex and the displayed reciprocal thresholds are well defined.

The model is invariant under several elementary transformations. Replacing Q by $(Q + Q^\top)/2$ does not alter the objective. Adding $\beta \mathbf{1}\mathbf{1}^\top$ shifts every feasible objective by β and leaves the minimizers unchanged. Multiplication by a positive scalar preserves minimizers and rescales all objective gaps. These operations are important when selecting a hardware-friendly coefficient alphabet.

Definition 2.1 (Encoding). Let $\mathcal{I}$ be an instance of a discrete decision or optimization problem. An *StQP encoding* consists of a polynomial-time map $\mathcal{I} \mapsto (Q_{\mathcal{I}}, \theta_{\mathcal{I}})$, together with a polynomial-time decoder, such that the desired answer can be recovered from $\nu(Q_{\mathcal{I}})$, from a global minimizer, or from both. Here $\theta_{\mathcal{I}}$ denotes any threshold or affine recovery data.

Definition 2.2 (Levels of exactness). We distinguish the following properties.

- (i) *decision exactness*: a yes/no answer is equivalent to an objective threshold;
- (ii) *value exactness*: the original optimum is recovered by an explicit function of $\nu(Q)$;
- (iii) *solution exactness*: a global minimizer can be decoded to an original optimum;
- (iv) *landscape exactness*: relevant local optima or barriers correspond under the mapping;
- (v) *approximation preservation*: a stated objective error in StQP implies a controlled error in the original problem.

Most reductions in complexity theory require only decision exactness. Hardware evaluation needs more. A formulation may be perfectly exact at the global optimum while creating exponentially many local minima, compressing decision values into a narrow interval, or requiring delicate support thresholding.




Figure 1: Compilation routes. Direct compatibility and selection graphs (green) preserve the problem’s structure and yield the simplest StQP matrices: three-level for clique, $\{0, 1\}$ off-diagonal for weighted clique. The generic routes (orange) reach StQP through the same weighted-clique step. The binary-linear route is polynomial in input bit length; the factor-model route is polynomial in the total explicit local-table encoding length.

2.2 Representation metrics

For an encoded instance, we report:

$$\begin{aligned}
 N_{\Delta} &:= \text{number of simplex coordinates,} \\
 M_Q &:= \# \{(u, v) : u \leq v, Q_{uv} \neq 0\}, \\
 d_Q &:= \max_u \# \{v \neq u : Q_{uv} \neq 0\}, \\
 \mathcal{A}_Q &:= \{Q_{uv} : 1 \leq u \leq v \leq N_{\Delta}\}, \\
 \kappa_{\text{coef}} &:= \frac{\max \{|q| : q \in \mathcal{A}_Q\}}{\min \{|q| : q \in \mathcal{A}_Q \setminus \{0\}\}}.
 \end{aligned}$$

Additional quantities are the target support size, objective separation between adjacent decision classes, entrywise perturbation tolerance, decoder complexity, and number or structure of local minima.

The raw nonzero count should be interpreted carefully. Since $\mathbf{1}^T x = 1$, adding a constant all-to-all matrix $\beta \mathbf{1}\mathbf{1}^T$ is physically and mathematically redundant up to an objective shift. A platform may implement such a baseline globally rather than through $O(N^2)$ individual couplers. We therefore distinguish the literal matrix Q from its *relative interaction matrix*, obtained after removing a convenient multiple of $\mathbf{1}\mathbf{1}^T$.

2.3 Decision thresholds and certificates

For graph-coded instances in [Section 3](#), an optimal support itself gives a polynomial certificate. For general rational StQP, complexity membership requires more care because minimizers are real. Quadratic programming with rational data admits polynomially verifiable certificates under standard formulations [\[36, 37\]](#), but our NP-completeness statement will avoid unnecessary subtleties by restricting to an explicit graph matrix class that is exactly equivalent to CLIQUE.

2.4 Compilation pipeline

[Figure 1](#) summarizes the routes developed below. Direct routes preserve combinatorial structure and usually give the cleanest coefficient matrices. Generic routes trade compactness for universality.

3 A three-level graph-to-StQP compiler

Let $G = (V, E)$ be a simple undirected graph with $V = [n]$, adjacency matrix A_G , and clique number $\omega(G)$. For a nonempty clique $C \subseteq V$, define its barycenter

$$x(C) := \frac{1}{|C|} \sum_{i \in C} e_i \in \Delta_n. \quad (3)$$

The classical Motzkin–Straus theorem expresses $\omega(G)$ through quadratic optimization over Δ_n [16]. Regularized formulations sharpen the relationship between local optima and maximal cliques [12, 13, 17]. We state the result in a minimization form designed for direct StQP compilation. [Algorithm 1](#) in [Section A](#) gives the compilation steps, and [Section C](#) works out several catalog entries explicitly.

3.1 The graph matrix

For $0 < \tau < 1$, define

$$Q_\tau(G) := \mathbf{1}\mathbf{1}^\top - A_G - (1 - \tau)I. \quad (4)$$

Its entries are

$$(Q_\tau)_{ij} = \begin{cases} \tau, & i = j, \\ 0, & i \neq j \text{ and } \{i, j\} \in E, \\ 1, & i \neq j \text{ and } \{i, j\} \notin E. \end{cases} \quad (5)$$

Thus graph edges represent compatible co-activation and nonedges represent repulsive pair interactions.

A clique is a set of pairwise adjacent vertices. It is maximal if no additional vertex can be added while preserving the clique property, and maximum if its cardinality is largest among all cliques of the graph.

Theorem 3.1 (Three-level graph compiler). *Let G be a nonempty graph and $0 < \tau < 1$. Then*

$$\nu(Q_\tau(G)) = \frac{\tau}{\omega(G)}. \quad (6)$$

Moreover:

- (i) $x \in \Delta_n$ is a global minimizer if and only if $x = x(C)$ for a maximum clique C ;
- (ii) $x \in \Delta_n$ is a local minimizer if and only if $x = x(C)$ for a maximal clique C ;
- (iii) every local minimizer is strict and isolated.

Remark 3.2 (Relation to regularized Motzkin–Straus). [Theorem 3.1](#) is a minimization-form restatement of the regularized Motzkin–Straus theorem rather than a new result. Since $\mathbf{1}^\top x = 1$ on Δ_n ,

$$x^\top Q_\tau(G) x = 1 - x^\top (A_G + (1 - \tau)I) x, \quad (7)$$

so minimizing $Q_\tau(G)$ over Δ_n is equivalent to maximizing the regularized Motzkin–Straus quadratic form [12], whose $\tau = 1/2$ case gives $\max_{\Delta_n} x^\top (A_G + \frac{1}{2}I) x = 1 - 1/(2\omega(G))$ together with the bijection between strict local maximizers and maximal cliques. The τ -parametrized family is the regularization studied in [13], and $\hat{Q}(G) = 2Q_{1/2}(G)$ is a positive rescaling of the same object. We give the proof because the minimization normalization, the nonnegative three-level alphabet, and the explicit gradient constants $2\tau/k$ and $2r_j(C)/k$ are used repeatedly in [Sections 4](#) to [6](#).

Proof. We divide the argument into support purification, minimization on a clique, and local optimality.

Suppose $x \in \Delta_n$ has positive coordinates $x_i, x_j > 0$ for a nonedge $\{i, j\} \notin E$. Along

$$x(t) = x + t(e_i - e_j), \quad -x_i \leq t \leq x_j,$$

the second derivative of $f(t) = x(t)^\top Q_\tau x(t)$ is

$$f''(t) = 2(Q_{ii} + Q_{jj} - 2Q_{ij}) = 4(\tau - 1) < 0.$$

Hence f is strictly concave on this segment. An interior point cannot minimize it locally, and at least one endpoint has strictly smaller objective. Repeating this elimination shows that every global minimizer has clique support, and any feasible point can be transformed without increasing the objective into a point supported on a clique.

If x is supported on a clique C , all off-diagonal terms within C vanish, so

$$x^\top Q_\tau x = \tau \sum_{i \in C} x_i^2 \geq \frac{\tau}{|C|},$$

with equality exactly when $x_i = 1/|C|$ for $i \in C$. Since $|C| \leq \omega(G)$, the global lower bound is $\tau/\omega(G)$, attained exactly at barycenters of maximum cliques.

It remains to characterize local minima. The strict-concavity argument shows that local minima must have clique support. On the face generated by a clique C , the unique face minimizer is $x(C)$ and the reduced Hessian is positive definite. Let $k = |C|$. At $x(C)$, the gradient components are

$$\nabla_i f = \frac{2\tau}{k} \quad (i \in C), \quad \nabla_j f = \frac{2r_j(C)}{k} \quad (j \notin C),$$

where $r_j(C)$ is the number of vertices in C not adjacent to j . If C is maximal, then $r_j(C) \geq 1$ for every $j \notin C$, and since $\tau < 1$,

$$\nabla_j f > \frac{2\tau}{k}.$$

Thus the Karush–Kuhn–Tucker (KKT) inequalities hold with strict complementarity, and $x(C)$ is a strict local minimizer. Conversely, if C is not maximal, some $j \notin C$ is adjacent to every vertex of C , so $r_j(C) = 0$ and moving a small amount of mass toward e_j decreases the objective. Hence $x(C)$ is not locally minimizing. Finiteness of the maximal-clique family gives isolation. $\square$

Remark 3.3 (Why regularization matters). At $\tau = 1$, the matrix is $\mathbf{1}\mathbf{1}^\top - A_G$, and the global value reduces to the classical Motzkin–Straus formula. The strict support-purification and the exact local-minimum correspondence can fail in this unregularized case. The range $0 < \tau < 1$ removes nonclique local minima and makes every maximal clique a strict local minimum.

3.2 An integer coefficient alphabet

Set $\tau = 1/2$ and scale by two:

$$\widehat{Q}(G) := 2Q_{1/2}(G) = 2\mathbf{1}\mathbf{1}^\top - 2A_G - I. \quad (8)$$

Then

$$\widehat{q}_{ij} = \begin{cases} 1, & i = j, \\ 0, & \{i, j\} \in E, \\ 2, & i \neq j \text{ and } \{i, j\} \notin E, \end{cases} \quad \nu(\widehat{Q}(G)) = \frac{1}{\omega(G)}. \quad (9)$$

This formulation uses no negative coefficients and only the alphabet $\{0, 1, 2\}$. The discrete solution is read from the support; the positive coordinates all equal $1/\omega(G)$.

Corollary 3.4 (Restricted StQP threshold is NP-complete). *Consider the following problem.*

THREE-LEVEL STQP THRESHOLD: *Given a symmetric matrix $Q \in \{0, 1, 2\}^{n \times n}$ satisfying $q_{ii} = 1$ and $q_{ij} \in \{0, 2\}$ for $i \neq j$, and an integer $k \geq 1$, decide whether*

$$\min_{x \in \Delta_n} x^\top Q x \leq \frac{1}{k}.$$

This problem is NP-complete.

Proof. Given a graph G , construct $\hat{Q}(G)$ in polynomial time. By [Theorem 3.1](#),

$$\nu(\hat{Q}(G)) \leq \frac{1}{k} \iff \omega(G) \geq k.$$

This proves NP-hardness by CLIQUE. Membership in NP follows directly from the matrix pattern: a certificate is a set C of k indices with $q_{ij} = 0$ for all distinct $i, j \in C$. Such a set is a k -clique in the graph defined by the zero off-diagonal entries and can be checked in polynomial time. $\square$

Corollary 3.5 (Universality for NP). *Every decision language in NP has a polynomial-time many-one reduction to THREE-LEVEL STQP THRESHOLD.*

Proof. By the Cook–Levin theorem and Karp reductions, every language in NP reduces in polynomial time to CLIQUE [\[10, 15\]](#). Apply [Corollary 3.4](#). $\square$

Caution 3.6. [Corollary 3.5](#) is a representability statement, not an algorithmic speedup. The encoded matrices form a restricted NP-complete family, so a general exact polynomial-time StQP solver would imply $P = NP$.

3.3 Objective separation and perturbation tolerance

Compact coefficients do not imply a constant decision gap. If the relevant distinction is between clique number k and $k - 1$, with $k \geq 2$, then for Q_τ the adjacent optimal values differ by

$$g_k(\tau) := \frac{\tau}{k-1} - \frac{\tau}{k} = \frac{\tau}{k(k-1)}. \quad (10)$$

For the integer-scaled matrix $\hat{Q}$, the separation is $1/[k(k-1)]$.

Proposition 3.7 (Entrywise perturbation bound). *Let $Q, E \in \mathbb{S}^N$ and $\|E\|_{\max} := \max_{u,v} |E_{uv}| \leq \varepsilon$. Then*

$$|\nu(Q + E) - \nu(Q)| \leq \varepsilon. \quad (11)$$

Proof. For every $x \in \Delta_N$,

$$|x^\top E x| \leq \sum_{u,v} |E_{uv}| x_u x_v \leq \varepsilon \left(\sum_u x_u \right)^2 = \varepsilon.$$

Apply this inequality at minimizers of Q and $Q + E$. $\square$

Thus a midpoint threshold between the k and $k - 1$ value levels remains valid under arbitrary entrywise perturbations smaller than $g_k(\tau)/2$. The required objective discrimination decays as $O(k^{-2})$. Support recovery has a related scale: the ideal nonzero coordinate of a k -clique barycenter is $1/k$, so coordinate noise substantially below $1/(2k)$ permits simple threshold decoding.

3.4 Landscape interpretation

The graph compiler converts combinatorial maximality into continuous local optimality. If $\mathcal{M}(G)$ denotes the family of maximal cliques, then

$$\#\{\text{strict local minima of } \text{StQP}(Q_\tau(G))\} = |\mathcal{M}(G)|. \quad (12)$$

Graphs can have exponentially many maximal cliques. Hence a simplex-native local relaxer can face an exponentially multimodal landscape even though the coefficient alphabet is fixed and nonnegative.

This observation cuts both ways. The landscape is undesirable for exact optimization but useful as a transparent diagnostic. A returned local minimum is not arbitrary: its support is a maximal clique, and its objective immediately reports that clique's size. Global exploration mechanisms can therefore be studied in a controlled setting where all local solutions have clear combinatorial meaning. Maximum-clique algorithms based on continuous dynamics exploit precisely this relationship [12, 14, 38].

3.5 Relative couplings

The literal expression $Q_\tau = \mathbf{1}\mathbf{1}^\top - A_G - (1 - \tau)I$ contains a dense all-ones term. On the simplex, however, $x^\top \mathbf{1}\mathbf{1}^\top x = 1$ is constant. Therefore the same minimizers are obtained from

$$\tilde{Q}_\tau(G) := -A_G - (1 - \tau)I, \quad (13)$$

with all objective values shifted by -1 . A physical platform that supplies a global normalization constraint need only realize the relative edge and diagonal interactions. Conversely, the nonnegative form (5) may be preferable when repulsive incompatibility couplings are natural. The two forms are mathematically equivalent but can have very different implementation costs.

4 Compatibility graphs for constraint satisfaction and assignment

Many discrete problems ask for one local choice from each of several groups, subject to pairwise consistency. Such problems have a natural compatibility graph: vertices represent local choices, and edges join choices that can coexist. The graph-to-StQP theorem then converts consistency into a simplex quadratic.

Choices belonging to the same group are represented by mutually nonadjacent vertices. A clique can therefore contain at most one choice from each group. When there are n groups, a clique of size n selects exactly one choice from every group, while adjacency between groups enforces consistency. Smaller cliques represent partial selections.

4.1 Finite-domain constraint satisfaction

A binary constraint satisfaction problem (CSP) consists of variables $z_1, \dots, z_n$ with finite domains $D_1, \dots, D_n$ and binary constraints $R_{ij} \subseteq D_i \times D_j$. Here binary refers to constraints involving two variables; the domains need not be Boolean. Unary restrictions may be absorbed into the domains. The microstructure graph is classical in constraint processing [21–23].

Construction 4.1 (Variable–value compatibility graph). Create one vertex (i, a) for every $i \in [n]$ and $a \in D_i$. Two distinct vertices (i, a) and (j, b) are adjacent if and only if $i \neq j$ and either no constraint links i and j , or $(a, b) \in R_{ij}$. Vertices belonging to the same variable are nonadjacent.

Proposition 4.2 (Binary CSP as StQP). *Let H be the graph from [Construction 4.1](#). Then the CSP is satisfiable if and only if H contains a clique of size n . Equivalently,*

$$\min_{x \in \Delta_N} x^\top \widehat{Q}(H)x = \frac{1}{n}, \quad N = \sum_{i=1}^n |D_i|. \quad (14)$$

When satisfiable, the support of a global minimizer contains one variable–value vertex for each i and decodes directly to a satisfying assignment.

Proof. A clique contains at most one vertex from each variable group. A clique of size n therefore selects exactly one value per variable, and adjacency enforces all binary constraints. Conversely, a satisfying assignment selects n pairwise adjacent vertices. Since no clique can exceed n , [Theorem 3.1](#) gives (14). $\square$

The selected variable–value pairs are identified by the support of the StQP solution. The coordinates themselves are continuous masses, with one global normalization; the size- n clique condition supplies the one-choice-per-variable requirement.

4.2 CNF satisfiability

Continuous formulations of satisfiability have a long history, from global-optimization models of SAT [39] to recent Fourier and polynomial relaxations [40, 41]; the construction below is exact and keeps the three-level alphabet. Let a conjunctive normal form (CNF) formula $\Phi = C_1 \wedge \cdots \wedge C_m$ contain N_{lit} literal occurrences. A particularly compact construction uses one vertex per occurrence, rather than two vertices per Boolean variable.

Construction 4.3 (Clause–literal graph). Create a vertex (r, ℓ) for every literal occurrence $\ell \in C_r$. Join (r, ℓ) and (s, ℓ') when $r \neq s$ and the two literals are not complementary. Vertices from the same clause are nonadjacent.

Proposition 4.4 (CNF-SAT). *The formula Φ is satisfiable if and only if the graph in [Construction 4.3](#) has a clique of size m . Consequently,*

$$\Phi \text{ satisfiable} \iff \nu(\widehat{Q}(H_\Phi)) = \frac{1}{m}. \quad (15)$$

The formulation has $N_\Delta = N_{\text{lit}}$ coordinates; for 3-SAT, $N_{\text{lit}} \leq 3m$.

Proof. A satisfying assignment makes at least one literal true in every clause. Selecting one true occurrence per clause gives pairwise noncomplementary literals and hence an m -clique. Conversely, an m -clique contains one occurrence from every clause and contains no complementary pair. Assigning all selected positive literals true and all selected negative literals false is consistent and can be extended arbitrarily to unselected variables, satisfying every clause. $\square$

A local minimum of the associated StQP corresponds to a maximal set of mutually consistent clause literals. Such a support may omit clauses and therefore need not certify satisfiability. This gives a concrete example of why local minimization alone is insufficient.

Example 4.5 (A 3-SAT decoding rule). Suppose a global minimizer of $\widehat{Q}(H_\Phi)$ has support

$$\{(1, x_2), (2, \neg x_5), (3, x_2), \dots, (m, x_7)\}.$$

All nonzero coordinates equal $1/m$. The support cannot contain both x_i and $\neg x_i$. Set each selected positive literal to true and each selected negative literal to false; assign all remaining variables arbitrarily. The result satisfies Φ .

4.3 Graph coloring and clique cover

For a graph $G = (V, E)$ and integer k , create vertices (v, c) with $v \in V$ and $c \in [k]$. Two assignments (v, c) and (u, d) are compatible when $v \neq u$ and either $\{u, v\} \notin E$ or $c \neq d$.

Proposition 4.6 (k -colorability). *Let $H_{G,k}$ be the compatibility graph above. Then*

$$\chi(G) \leq k \iff \omega(H_{G,k}) = |V| \iff \nu(\hat{Q}(H_{G,k})) = \frac{1}{|V|}. \quad (16)$$

The StQP uses $N_\Delta = k|V|$ coordinates, and a global support selects one color for every vertex.

Covering G by at most k cliques is equivalent to coloring the complement graph $\overline{G}$ with at most k colors. Hence the same construction yields a clique-cover formulation with $k|V|$ coordinates.

4.4 Hamiltonian cycles

Let $G = (V, E)$ have $n \geq 2$ vertices in the directed case and $n \geq 3$ in the simple undirected case; smaller instances are handled separately. Use vertices (p, v) meaning that graph vertex v occupies tour position $p \in \{0, \dots, n-1\}$. Two such assignments are compatible if they use different positions and different graph vertices, and if assignments to consecutive positions follow an edge of G .

For a directed graph, more explicitly, (p, u) and (q, v) are adjacent when:

- (i) $p \neq q$ and $u \neq v$;
- (ii) if $q \equiv p+1 \pmod{n}$, then $(u, v) \in E$;
- (iii) if $p \equiv q+1 \pmod{n}$, then $(v, u) \in E$.

For nonconsecutive positions, no graph-edge condition is imposed.

Proposition 4.7 (Hamiltonian cycle). *The directed graph G contains a Hamiltonian cycle if and only if the compatibility graph contains a clique of size n . Therefore*

$$G \text{ Hamiltonian} \iff \nu(\hat{Q}(H_G^{\text{Ham}})) = \frac{1}{n}. \quad (17)$$

The formulation has $N_\Delta = n^2$ coordinates. The undirected case uses the same construction with undirected adjacency.

Proof. An n -clique contains at most one vertex for each position and at most one occurrence of each graph vertex, so it defines a permutation of V . Compatibility at consecutive positions makes that permutation a Hamiltonian cycle. The converse follows by selecting the n position–vertex assignments from any Hamiltonian cycle. $\square$

The construction is closely related to one-hot QUBO encodings of tours, but its constraint logic is expressed solely through the compatibility graph. There is no separate penalty-scale choice: all conflicts receive the same coefficient in $\hat{Q}$.

4.5 Graph isomorphism and association graphs

Graph isomorphism is not known to be NP-complete, but it is a useful illustration of the association-graph principle and has a well-established continuous clique interpretation [42]; association graphs underlie continuous methods for matching, common subgraphs, and shape correspondence [43, 44], and clique and constraint models are established approaches to maximum common subgraph [45]. Let $G = (V_G, E_G)$ and $H = (V_H, E_H)$ have n vertices. Create one vertex (u, v) for each possible mapping $u \in V_G$ to $v \in V_H$. Two mapping vertices (u, v) and (u', v') are adjacent when $u \neq u'$, $v \neq v'$, and

$$\{u, u'\} \in E_G \iff \{v, v'\} \in E_H. \quad (18)$$

Proposition 4.8 (Graph isomorphism). *The graphs G and H are isomorphic if and only if the association graph has a clique of size n , equivalently if the associated three-level StQP has optimum $1/n$. The formulation uses n^2 coordinates.*

For subgraph isomorphism from G into H , allow different graph orders, replace the equivalence in (18) by the implication $\{u, u'\} \in E_G \Rightarrow \{v, v'\} \in E_H$, and use $|V_G|$ as the target clique size. Keeping the equivalence while allowing partial injective supports instead gives the maximum common *induced* subgraph formulation. Weighted vertex scores can be included through the weighted clique construction below.

4.6 Arbitrary-arity hard constraints

For constraints involving more than two variables, use one vertex for every allowed tuple of every constraint. Vertices from different constraints are adjacent when the tuples agree on all shared variables; vertices from the same constraint are nonadjacent. A clique containing one tuple from every constraint is equivalent to a globally consistent satisfying assignment. If there are F constraints and A_α allowed tuples for constraint α , then

$$N_\Delta = \sum_{\alpha=1}^F |A_\alpha|. \quad (19)$$

This construction is polynomial in the explicit allowed-tuple list, but expanding succinct high-arity constraints can require exponentially many tuples. The tuple graph is the feasibility counterpart of the weighted factor-state compiler developed next. The two are complementary rather than nested: pruning to allowed tuples is exact here, but falls outside the hypotheses of the weighted compiler (Remark 5.7).

5 Weighted clique and a generic factor-state compiler

The unweighted graph compiler captures feasibility and cardinality. Optimization objectives require weights. We first prove a weighted clique StQP formula, then use it to compile arbitrary finite-domain factor objectives.

5.1 A weighted clique StQP theorem

Let $G = (V, E)$ have positive vertex weights $w_i > 0$. For a clique C , write

$$w(C) := \sum_{i \in C} w_i, \quad \omega_w(G) := \max \{w(C) : C \text{ is a clique}\}.$$

Let $w_{\min} := \min_i w_i$, $w_{\max} := \max_i w_i$, and choose $0 < \tau < 1$. Define $Q_{\tau, w}(G)$ entrywise by

$$q_{ij} = \begin{cases} \tau \frac{w_{\min}}{w_i}, & i = j, \\ 0, & i \neq j \text{ and } \{i, j\} \in E, \\ 1, & i \neq j \text{ and } \{i, j\} \notin E. \end{cases} \quad (20)$$

This is a regularized version of weighted Motzkin–Straus formulations [14, 28].

Theorem 5.1 (Weighted graph compiler). *For the matrix in (20),*

$$\nu(Q_{\tau, w}(G)) = \frac{\tau w_{\min}}{\omega_w(G)}. \quad (21)$$

There exists a global minimizer supported on a maximum-weight clique C^* , with

$$x_i^* = \begin{cases} \frac{w_i}{w(C^*)}, & i \in C^*, \\ 0, & i \notin C^*. \end{cases} \quad (22)$$

Every local minimizer, and therefore every global minimizer, has clique support. An arbitrary feasible point can be purified to clique support without increasing the objective.

Proof. If positive coordinates x_i, x_j correspond to a nonedge, then along $x + t(e_i - e_j)$ the second derivative is

$$2 \left(\tau \frac{w_{\min}}{w_i} + \tau \frac{w_{\min}}{w_j} - 2 \right) \leq 4(\tau - 1) < 0.$$

Thus a local minimizer cannot contain a nonedge in its support, and an arbitrary feasible point can be purified by endpoint moves until its support is a clique.

For a fixed clique C , off-diagonal terms vanish and

$$x^\top Q_{\tau,w} x = \tau w_{\min} \sum_{i \in C} \frac{x_i^2}{w_i}.$$

Cauchy–Schwarz gives

$$1 = \left(\sum_{i \in C} x_i \right)^2 \leq \left(\sum_{i \in C} \frac{x_i^2}{w_i} \right) \left(\sum_{i \in C} w_i \right),$$

so the face minimum is $\tau w_{\min}/w(C)$, attained exactly at (22). Minimizing over cliques selects a maximum-weight clique and proves (21). $\square$

Remark 5.2. The diagonal values encode the vertex weights through reciprocals, while incompatibility remains binary. Scaling all w_i by a common positive constant changes neither the matrix nor the solution. The coefficient dynamic range on the diagonal is $w_{\max}/w_{\min}$.

Theorem 5.1 locates the global minimum. The next theorem describes every local minimizer, extending **Theorem 3.1** to positive weights.

Theorem 5.3 (Weighted landscape). *Let G be a nonempty graph with positive vertex weights w , let $0 < \tau < 1$, and let $Q_{\tau,w}(G)$ be as in (20). For a clique C write $x(C, w) := \sum_{i \in C} (w_i/w(C)) e_i \in \Delta_n$ for its weighted barycenter; (22) is $x(C^*, w)$. Then:*

- (i) $x \in \Delta_n$ is a local minimizer if and only if $x = x(C, w)$ for a maximal clique C ;
- (ii) every local minimizer is strict and isolated;
- (iii) x is a global minimizer if and only if $x = x(C, w)$ for a maximum-weight clique.

In particular the number of strict local minimizers equals $|\mathcal{M}(G)|$, the number of maximal cliques, independently of the weights.

Proof. By the purification argument in the proof of **Theorem 5.1**, every local minimizer has clique support. If x is a local minimizer with support exactly the clique C , then x lies in the relative interior of the face generated by C and minimizes $f = x^\top Q_{\tau,w} x$ locally there; on that face $f = \tau w_{\min} \sum_{i \in C} x_i^2/w_i$ is strictly convex, so $x = x(C, w)$ by **Theorem 5.1**. Local minimizers are therefore among the weighted barycenters of cliques.

Fix a clique C and put $x^* = x(C, w)$. For $j \notin C$ let

$$r_j^w(C) := \sum_{\substack{i \in C \\ \{i,j\} \notin E}} w_i \quad (23)$$

be the total weight of vertices of C not adjacent to j . Since off-diagonal entries vanish inside C ,

$$(Q_{\tau,w}x^*)_i = \tau \frac{w_{\min}}{w_i} \cdot \frac{w_i}{w(C)} = \frac{\tau w_{\min}}{w(C)} \quad (i \in C), \quad (24)$$

which is constant on C , while for $j \notin C$ only nonedges to C contribute, each with coefficient one, so

$$(Q_{\tau,w}x^*)_j = \frac{r_j^w(C)}{w(C)}. \quad (25)$$

Suppose C is maximal. Every $j \notin C$ then has a non-neighbour in C , so $r_j^w(C) \geq w_{\min} > \tau w_{\min}$ because $\tau < 1$; comparing (24) and (25), the KKT conditions hold with strict complementarity. For strictness, let $x \in \Delta_n$ with $d := x - x^* \neq 0$. Then $\mathbf{1}^\top d = 0$, and $d_j = x_j \geq 0$ for $j \notin C$ because $x_j^* = 0$ there; put $s := \sum_{j \notin C} d_j \geq 0$. Using $\sum_{i \in C} d_i = -s$,

$$d^\top Q_{\tau,w}x^* = \frac{1}{w(C)} \sum_{j \notin C} d_j (r_j^w(C) - \tau w_{\min}) \geq \frac{(1 - \tau)w_{\min}}{w(C)} s \geq 0. \quad (26)$$

Split $d = u + v$ with u supported on C and $v \geq 0$ supported off C , so $\|v\|_1 = s$ and $\|u\|_1 \leq \|d\|_1$. The matrix $Q_{\tau,w}$ is entrywise nonnegative with entries at most one and is diagonal on C , so $u^\top Q_{\tau,w}u \geq 0$, $v^\top Q_{\tau,w}v \geq 0$ and $|u^\top Q_{\tau,w}v| \leq \|u\|_1 s$. Hence

$$f(x) - f(x^*) = 2d^\top Q_{\tau,w}x^* + d^\top Q_{\tau,w}d \geq 2s \left(\frac{(1 - \tau)w_{\min}}{w(C)} - \|d\|_1 \right).$$

If $s > 0$ and $\|d\|_1 < (1 - \tau)w_{\min}/w(C)$, the right-hand side is positive. If $s = 0$, then $d = u \neq 0$, the linear term vanishes, and $f(x) - f(x^*) = \tau w_{\min} \sum_{i \in C} u_i^2/w_i > 0$. Thus $f(x) > f(x^*)$ for every $x \in \Delta_n$ with $0 < \|x - x^*\|_1 < (1 - \tau)w_{\min}/w(C)$: the minimizer is strict.

If C is not maximal, some $j \notin C$ is adjacent to every vertex of C , so $r_j^w(C) = 0$ and $(e_j - x^*)^\top Q_{\tau,w}x^* = -\tau w_{\min}/w(C) < 0$; moving mass toward e_j decreases the objective and $x(C, w)$ is not locally minimizing. This proves (i) and strictness. Since the graph has finitely many maximal cliques, (i) also gives isolation, completing (ii). For (iii), a global minimizer exists by compactness and is a local minimizer, so by (i) it equals $x(C, w)$ for a maximal clique C , with value $\tau w_{\min}/w(C)$. This value is minimal exactly when $w(C) = \omega_w(G)$, and every maximum-weight clique is maximal because the weights are positive. $\square$

Relation to the weighted regularization of Bomze–Pelillo–Stix. The local/global clique correspondence and strictness are already established for a class of weighted regularizations in Bomze et al. [14]. In their normalization, a matrix C has $C_{ii} = 1/(2w_i)$, zero edge entries, and $C_{ij} \geq C_{ii} + C_{jj}$ on nonedges. The rescaling $C = Q_{\tau,w}/(2\tau w_{\min})$ satisfies this condition when

$$\tau w_{\min} \left(\frac{1}{w_i} + \frac{1}{w_j} \right) \leq 1 \quad \text{for every nonedge } \{i, j\},$$

in particular for $0 < \tau \leq 1/2$. The proof above verifies the binary off-diagonal normalization for the full range $0 < \tau < 1$, without requiring that stronger matrix-class condition. We use its explicit constants in the factor-state and precision analyses; the weighted clique correspondence itself is not a new principle.

Remark 5.4 (The landscape is weight-independent). Weights move the *values* of the local minima but not their *number* or their combinatorial meaning: the strict local minimizers of $\text{StQP}(Q_{\tau,w}(G))$ stand in bijection with $\mathcal{M}(G)$ for every positive weight vector, exactly as in the unweighted case (12). The consequences for complete factor-state models are given in Corollary 5.8.

5.2 Finite-domain factor models

Let $z = (z_1, \dots, z_n)$, where z_i belongs to a finite domain D_i . Consider

$$E(z) = \sum_{\alpha \in \mathcal{F}} \phi_\alpha(z_{S_\alpha}), \quad (27)$$

where $S_\alpha \subseteq [n]$ is the scope of factor α , and z_{S_α} denotes the local assignment on that scope. No sign assumption is imposed on ϕ_α . The identities below allow real table values. For polynomial-time compilation, the factors are supplied as explicit complete rational tables, with rational offsets of polynomial encoding length; a fixed rational τ may be used. Complexity is measured in the total table encoding length, not in a succinct description of arbitrary-arity factors. A single factor on r binary variables has 2^r local states.

Construction 5.5 (Factor-state compatibility graph). For each factor α and each local state

$$a \in \mathcal{A}_\alpha := \prod_{i \in S_\alpha} D_i,$$

create a vertex $v = (\alpha, a)$. Vertices belonging to the same factor are nonadjacent. Vertices (α, a) and (β, b) with $\alpha \neq \beta$ are adjacent when a and b agree on every variable in $S_\alpha \cap S_\beta$.

Choose constants

$$B_\alpha > \max_{a \in \mathcal{A}_\alpha} \phi_\alpha(a) \quad (28)$$

and give vertex (α, a) the positive weight

$$w_{\alpha,a} := B_\alpha - \phi_\alpha(a). \quad (29)$$

Write $B := \sum_\alpha B_\alpha$.

Theorem 5.6 (Exact factor-state compiler). *Assume every variable that matters to the objective appears in at least one factor, and that H_E carries a vertex for every local state $a \in \mathcal{A}_\alpha = \prod_{i \in S_\alpha} D_i$ of every factor, with no states pruned. Let H_E be the weighted compatibility graph in [Construction 5.5](#). Then*

$$\omega_w(H_E) = B - E^*, \quad E^* := \min_z E(z). \quad (30)$$

Consequently, for the weighted StQP matrix $Q_E = Q_{\tau,w}(H_E)$,

$$\nu(Q_E) = \frac{\tau w_{\min}}{B - E^*}, \quad E^* = B - \frac{\tau w_{\min}}{\nu(Q_E)}. \quad (31)$$

Every global minimizer decodes to a globally optimal assignment. The simplex dimension is

$$N_\Delta = \sum_{\alpha \in \mathcal{F}} |\mathcal{A}_\alpha|. \quad (32)$$

Proof. Any global assignment z chooses one local state (α, z_{S_α}) from each factor. These vertices agree on overlaps and form a clique C_z with weight

$$w(C_z) = \sum_\alpha (B_\alpha - \phi_\alpha(z_{S_\alpha})) = B - E(z).$$

Thus $\omega_w(H_E) \geq B - E^*$.

Conversely, vertices in a clique are pairwise consistent, so their local states combine into one partial global assignment. If a clique omits factor β , choose a local state for β that agrees with this partial

assignment on already assigned variables and assign arbitrary values to the remaining variables in S_β . Because all local states are present, this vertex exists and is adjacent to the clique. Its weight is positive, so no maximum-weight clique can omit a factor. Hence every maximum-weight clique contains exactly one state from every factor and corresponds to a global assignment z , with weight $B - E(z)$. This proves (30). Equation (31) follows from Theorem 5.1. The support states combine to the optimal assignment. $\square$

Remark 5.7 (Hard constraints must not prune states). The completeness hypothesis on $\mathcal{A}_\alpha$ is essential to the stated theorem. The proof extends a clique that omits factor β by exhibiting a state of β consistent with the partial assignment already fixed; that state is guaranteed to exist precisely because $\mathcal{A}_\beta$ is the full product of domains. If disallowed states are deleted instead of penalized, the extension can fail: a maximum-weight clique may then legitimately omit a factor, and (30) is false. A minimal example has one variable $z_1 \in \{0, 1\}$ and two factors with scope $\{z_1\}$: $\phi_\alpha(0) = 10$, $\phi_\alpha(1) = 0$, $B_\alpha = 11$, and β enforces $z_1 = 0$ by keeping only the state 0, with $\phi_\beta(0) = 0$, $B_\beta = 1$. The states $(\alpha, 1)$ and $(\beta, 0)$ disagree, so the cliques are $\{(\alpha, 0), (\beta, 0)\}$ of weight 2 and $\{(\alpha, 1)\}$ of weight 11. Thus $\omega_w = 11$, whereas $B - E^* = 2$, and the decoded assignment $z_1 = 1$ violates the constraint.

The proof uses less than completeness: it suffices that, for every factor β , every assignment to the variables that S_β shares with other factors extends to a retained state of β . States that differ only on variables private to S_β may therefore be pruned, and (30) then holds with E^* the minimum over retained assignments.

To retain the complete-state guarantee without a separate extendibility argument, encode hard constraints through the values ϕ_α rather than deleting states. Write $\phi_\alpha = \psi_\alpha + P \mathbf{1}[\text{constraint } \alpha \text{ violated}]$ with soft parts ψ_α . If some assignment satisfies every constraint and

$$P > S_{\text{soft}} := \sum_{\gamma} \left(\max_a \psi_{\gamma}(a) - \min_a \psi_{\gamma}(a) \right),$$

then every minimizer of E is feasible, since a feasible z has $E(z) \leq \sum_{\gamma} \max \psi_{\gamma}$ and a violating one has $E(z) \geq P + \sum_{\gamma} \min \psi_{\gamma}$. Remark 5.13 discusses how larger penalties can affect coefficient ratios and conservative separation bounds. In particular the unweighted tuple graph of Section 4.6, which carries one vertex per *allowed* tuple, is not a special case of Construction 5.5; it is the complementary construction, exact for pure feasibility but outside the scope of Theorem 5.6.

Corollary 5.8 (Every assignment is a strict local minimizer). *Under the hypotheses of Theorem 5.6, suppose every variable lies in the scope of some factor. Then the maximal cliques of H_E are exactly the sets*

$$C_z := \{(\alpha, z_{S_\alpha}) : \alpha \in \mathcal{F}\}, \quad z \in \prod_i D_i.$$

Consequently $\text{StQP}(Q_E)$ has exactly $\prod_i |D_i|$ strict local minimizers, one for each assignment z , with value $\tau w_{\min}/(B - E(z))$. In particular the encoding in Corollary 5.11 of an n -variable QUBO in which every variable appears in some term has 2^n strict local minimizers.

Proof. C_z is a clique by the proof of Theorem 5.6, and it is maximal because any other state (α, a) with $a \neq z_{S_\alpha}$ is nonadjacent to (α, z_{S_α}) . Conversely, the extension step in the same proof shows that a clique omitting a factor is not maximal, and a clique with one state per factor is C_z for a unique z because every variable is covered. Apply Theorem 5.3(i)–(ii). $\square$

Remark 5.9 (What the landscape does not filter). The continuous landscape of a compiled factor model eliminates no configuration: every assignment, feasible or not, optimal or not, is a strict local minimum, and only the ordering of the local values $\tau w_{\min}/(B - E(z))$ reflects the objective. This is a weaker filter than a QUBO's, whose local minima under single-bit flips are only the 1-flip-stable

assignments; near $x(C_z, w)$ there is no descent direction that changes even one variable. On the binary-linear route of [Section 7](#), every infeasible configuration is likewise a strict local minimizer. Any advantage of a continuous method on these encodings must therefore come from its trajectory and basins of attraction from the interior of the simplex, not from an absence of spurious local minima.

Corollary 5.10 (Decision threshold). *For any $T < B$,*

$$E^* \leq T \iff \nu(Q_E) \leq \frac{\tau w_{\min}}{B - T}. \quad (33)$$

5.3 Sparsity of the incompatibility matrix

Although the compatibility graph can be dense, the off-diagonal nonzeros of Q_E correspond to *incompatibilities*. Factors with disjoint scopes are fully compatible and produce zero cross blocks. Nonzero interactions occur only:

- (i) within each factor block, to prevent selecting two states of the same factor;
- (ii) between factor blocks that share a variable, and only for state pairs that disagree on that variable.

A crude upper bound is

$$M_Q = O \left(\sum_{\alpha} |\mathcal{A}_{\alpha}|^2 + \sum_{i=1}^n \left(\sum_{\alpha: i \in S_{\alpha}} |\mathcal{A}_{\alpha}| \right)^2 \right). \quad (34)$$

For uniformly bounded domain sizes, factor arity, and variable incidence, $M_Q = O(N_{\Delta})$ and $d_Q = O(1)$. Without a bound on local-state counts, even a unary factor can produce a dense incompatibility block.

5.4 QUBO as a special case

Consider a binary quadratic objective

$$E(z) = c + \sum_{i=1}^n h_i z_i + \sum_{\{i,j\} \in \mathcal{E}} J_{ij} z_i z_j, \quad z_i \in \{0, 1\}. \quad (35)$$

Use one unary factor for every retained $h_i z_i$ and one pair factor for every nonzero $J_{ij} z_i z_j$. Constants are handled outside the StQP recovery map.

Corollary 5.11 (QUBO-to-StQP). *A QUBO with n_u unary factors and m nonzero pair factors admits an exact weighted StQP encoding with*

$$N_{\Delta} \leq 2n_u + 4m. \quad (36)$$

For a dense n -variable QUBO this is $O(n^2)$; for a sparse interaction graph it is $O(n + m)$. The off-diagonal coefficient alphabet remains $\{0, 1\}$, while objective values appear in the positive diagonal entries.

The compiler expands a pairwise binary objective into local configurations. Consistency among repeated copies of each variable determines the zero/one off-diagonal pattern, while factor costs determine diagonal reciprocals. The resulting dimension can exceed that of the original QUBO substantially.

5.5 Worked example: weighted Max-Cut

Let $G = (V, E)$ have nonnegative edge weights c_e . A cut is represented by binary labels $z_i \in \{0, 1\}$ and has weight

$$\text{cut}_c(z) = \sum_{e=\{i,j\} \in E} c_e \mathbf{1}[z_i \neq z_j].$$

For each edge $e = \{i, j\}$, create four local states (e, a, b) with $(a, b) \in \{0, 1\}^2$. Two edge states are compatible if they agree on every shared endpoint. Give state (e, a, b) weight

$$w_{e,a,b} = L + c_e \mathbf{1}[a \neq b], \quad L > 0. \quad (37)$$

A full consistent clique contains one state per edge and has weight

$$L|E| + \text{cut}_c(z).$$

Therefore

$$\omega_w(H) = L|E| + \text{MaxCut}_c(G). \quad (38)$$

With $w_{\min} = L$, [Theorem 5.1](#) gives

$$\nu(Q) = \frac{\tau L}{L|E| + \text{MaxCut}_c(G)}. \quad (39)$$

For an unweighted graph and $L = 1$,

$$\text{MaxCut}(G) = \frac{\tau}{\nu(Q)} - |E|. \quad (40)$$

The formulation uses exactly $4|E|$ simplex coordinates before optional preprocessing.

5.6 Approximation distortion

Let Q_E be the factor-state matrix of [Theorem 5.6](#), let $\nu = \nu(Q_E) > 0$, and retain the recovery constants B , τ , and $w_{\min}$ from (31). For a positive scalar estimate $\hat{\nu} = \nu + \delta$, define the energy proxy

$$\hat{E} := B - \frac{\tau w_{\min}}{\hat{\nu}}.$$

The reciprocal recovery formula gives the exact identity

$$|\hat{E} - E^*| = \frac{\tau w_{\min} |\delta|}{\nu \hat{\nu}}. \quad (41)$$

For a scalar estimate alone, $\hat{E}$ need not be the energy of any discrete assignment, and δ can have either sign. A discrete-solution guarantee requires both a feasible simplex point and a specified decoder.

Suppose instead that $x \in \Delta_{N_\Delta}$ is returned and that $\hat{\nu} = x^\top Q_E x$ is evaluated using the intended matrix. Then $\delta \geq 0$. Repeatedly eliminate a nonedge from the positive support by moving all its mass to whichever endpoint gives the smaller objective, as in [Appendix B](#). The weighted purification argument in [Theorem 5.1](#) makes each such step objective-nonincreasing. Minimize on the resulting clique face, extend the clique to a maximal clique, and minimize on the enlarged face. Positive vertex weights ensure that extension cannot increase the face-minimum value.

Under the complete-state hypothesis of [Theorem 5.6](#), every maximal clique contains one consistent state from every factor and therefore determines an assignment z . Its weighted barycenter y satisfies

$$y^\top Q_E y = \frac{\tau w_{\min}}{B - E(z)} \leq x^\top Q_E x = \hat{\nu}.$$

Since $t \mapsto B - \tau w_{\min}/t$ is increasing for $t > 0$, this decoder gives

$$E^* \leq E(z) \leq \hat{E}, \quad 0 \leq E(z) - E^* \leq \frac{\tau w_{\min} \delta}{\nu(\nu + \delta)}.$$

The bound does not automatically apply to a decoder based on voting, coordinate ranking, or another repair rule that can increase the StQP objective.

Writing $W = B - E^* = \tau w_{\min}/\nu$, a multiplicative guarantee $x^\top Q_E x \leq \rho\nu$, with $\rho \geq 1$, similarly yields

$$E(z) - E^* \leq \left(1 - \frac{1}{\rho}\right) W.$$

Thus the factor-state route preserves a ratio for the positive transformed objective $B - E$, but generally not the same ratio for E itself, which may vanish or be negative. Both additive and multiplicative guarantees for the original energy depend on the offsets through W . If a constant c was omitted before compiling, the original objective is $c + E(z)$; restoring c changes neither the assignment nor its additive suboptimality. These bounds concern the intended objective evaluated on a feasible simplex point; coefficient and evaluation errors must be included separately.

5.7 Objective separation

For integer-valued factor tables, the reciprocal recovery map gives a separation between distinct assignment values. This is a gap between local-minimum values, not a positive separation between the global minimum and every other feasible simplex point.

Proposition 5.12 (Factor-model objective separation). *Assume the hypotheses of Theorem 5.6, with integer-valued ϕ_α and integer offsets B_α satisfying (28). Let $W := \omega_w(H_E) = B - E^*$ and assume $W \geq 2$. Every nonglobal local minimum of $\text{StQP}(Q_E)$ has value at least $\nu(Q_E) + g_W$, where*

$$g_W := \frac{\tau w_{\min}}{W-1} - \frac{\tau w_{\min}}{W} = \frac{\tau w_{\min}}{W(W-1)}. \quad (42)$$

The corresponding relative separation is at least $g_W/\nu(Q_E) = 1/(W-1)$. Equality is attained when a maximal clique has weight $W-1$; otherwise the separation between distinct local-minimum values is larger. If all local minima are global, the statement about nonglobal minima is vacuous.

Proof. Each vertex weight $w_{\alpha,a} = B_\alpha - \phi_\alpha(a)$ is a positive integer. The weighted landscape theorem gives a local-minimum value $\tau w_{\min}/w(C)$ for each maximal clique C . A nonglobal clique has integer weight $w(C) \leq W-1$, so its value is at least $\tau w_{\min}/(W-1)$. Subtracting $\nu(Q_E) = \tau w_{\min}/W$ proves the claim. $\square$

Under the complete-state hypothesis, $W = 1$ is degenerate: positive integer weights imply that there is only one factor and all its states have weight one, so every assignment is optimal. There is then no nonglobal assignment value to separate.

More generally, suppose every nonglobal assignment has energy at least $E^* + \Delta_E$, where $0 < \Delta_E < W$. The same argument gives

$$g_{W,\Delta_E} := \frac{\tau w_{\min} \Delta_E}{W(W - \Delta_E)}. \quad (43)$$

This is exact when the level $E^* + \Delta_E$ is attained. For integer factor tables, a positive common divisor of all within-table energy differences provides a computable spacing bound. Unlike the unit-step estimate, (43) is invariant when all weights and the known energy spacing are multiplied by a common positive constant, as required by Remark 5.2.

Interpretation of separation bounds. The quantities g_W and g_{W,Δ_E} are guaranteed lower bounds on the difference between the optimum and a nonglobal local-minimum value. They are exact only when the relevant next assignment level is attained. They do not separate the optimum from every other feasible simplex point. Accuracy tolerances derived from these bounds are sufficient guarantees for the specified decision or decoding procedure, rather than necessary resolution requirements for every instance. A smaller bound alone therefore establishes neither a smaller actual gap nor a change in stationary points or basins. Throughout the offset, penalty, and hardware comparisons, we use “separation bound” in this sense.

For a decision threshold $E^* \leq T$, with integer T and $K := B - T \geq 2$, the possible optimum values satisfy

$$E^* \leq T \implies \nu(Q_E) \leq \frac{\tau w_{\min}}{K}, \quad E^* > T \implies \nu(Q_E) \geq \frac{\tau w_{\min}}{K-1}.$$

Thus $g_K = \tau w_{\min}/[K(K-1)]$ is a lower bound on the separation of these decision classes. Proposition 3.7 implies that a midpoint test on the optimal value remains valid when the sum of the entrywise coefficient-error bound and the two-sided optimal-value estimation error is less than $g_K/2$. A feasible objective value without a suboptimality bound does not by itself provide such a test.

The same uniform perturbation bound gives a useful decoding condition. Suppose $\tilde{Q}_E = Q_E + \Delta Q$, $\|\Delta Q\|_{\max} \leq \epsilon$, and a feasible point satisfies

$$x^\top \tilde{Q}_E x \leq \nu(\tilde{Q}_E) + r, \quad r \geq 0.$$

Then $x^\top Q_E x \leq \nu(Q_E) + r + 2\epsilon$. Consequently, $r + 2\epsilon < g_W$ is sufficient for the objective-nonincreasing decoder of Section 5.6 to recover an optimal assignment. A valid g_{W,Δ_E} can replace g_W in this condition.

Remark 5.13 (Choice of offsets and penalties). Let $F := |\mathcal{F}| \geq 1$, and choose a common positive integer slack η in the offsets:

$$B_\alpha = \max_{a \in \mathcal{A}_\alpha} \phi_\alpha(a) + \eta, \quad \Phi := \sum_{\alpha \in \mathcal{F}} \max_{a \in \mathcal{A}_\alpha} \phi_\alpha(a), \quad D := \Phi - E^* \geq 0.$$

Then $w_{\min} = \eta$ and $W = D + F\eta$. For $W \geq 2$, the unit-step bound in (42) is exactly

$$g(\eta) = \frac{\tau\eta}{(D + F\eta)(D + F\eta - 1)}. \quad (44)$$

If $D = 0$, every assignment attains the same energy and there is no nonoptimal assignment level to separate; $\eta = 1$ is admissible. For $D \geq 1$, differentiation with respect to a real $\eta > 0$ gives

$$g'(\eta) = \frac{\tau[D(D-1) - F^2\eta^2]}{(D + F\eta)^2(D + F\eta - 1)^2}.$$

The integer maximizing this unit-step bound can be selected by evaluating it at the two neighboring integers of the constrained continuous maximizer:

$$\eta_c := \max \left\{ 1, \frac{\sqrt{D(D-1)}}{F} \right\}, \quad \mathcal{J} := \{\lfloor \eta_c \rfloor, \lceil \eta_c \rceil\}, \quad (45)$$

$$\eta^* \in \arg \max_{j \in \mathcal{J}} \frac{\tau j}{(D + Fj)(D + Fj - 1)}.$$

Nearest-integer rounding of D/F is not an exact selection rule. For example, with $D = 19$ and $F = 13$, $g(1) = \tau/992$ but $g(2) = \tau/990$.

For fixed $D \geq 1$ and fixed τ , sufficiently large F gives $\eta^* = 1$ and $g(1) = \Theta(F^{-2})$. When $D = \Theta(F)$, the optimized bound likewise scales as $\Theta(F^{-2})$. In the regime $D/F \rightarrow \infty$, the interior approximation gives $\eta^* \sim D/F$ and $g(\eta^*) \sim \tau/(4DF)$. For fixed D, F and excessive slack, $g(\eta) \sim \tau/(F^2\eta)$ as $\eta \rightarrow \infty$.

The exact D is not generally available before solving. A certified lower bound $E_{\text{lb}} \leq E^*$ gives the integer upper bound $\overline{D} := \lceil \Phi - E_{\text{lb}} \rceil \geq D$. Since (44) decreases as D increases, substituting $\overline{D}$ for D yields a conservative gap bound, and (45) can be used to maximize that bound. The factorwise bound $E_{\text{lb}} = \sum_{\alpha} \min_a \phi_{\alpha}(a)$ is always available; a relaxation may improve it. A feasible incumbent supplies an upper bound on E^* , not the lower bound required here.

Let $R_{\text{max}} := \max_{\alpha} (\max_a \phi_{\alpha}(a) - \min_a \phi_{\alpha}(a))$. The diagonal coefficient range is $w_{\text{max}}/w_{\text{min}} = 1 + R_{\text{max}}/\eta$. When the matrix contains a nonedge entry equal to one, its full nonzero coefficient range is $\kappa_{\text{coef}} = (\eta + R_{\text{max}})/(\tau\eta)$; otherwise this expression is an upper bound. Increasing η therefore improves coefficient range while eventually reducing the gap bound.

For hard constraints encoded by a common penalty P , retain the sufficient condition from the complete-state discussion: $P > S_{\text{soft}}$, provided a feasible assignment exists. For integer soft-factor data, $P = S_{\text{soft}} + 1$ satisfies this condition. Once feasibility is ensured, increasing P leaves the optimum over feasible assignments unchanged but can increase Φ , D , and the coefficient range. Its effect on the separation bound follows from (44), with the interpretation given above; Section 7.4 illustrates the dependence on energy spacings.

6 Selection, packing, covering, and matching formulations

The graph compiler is particularly direct when a feasible solution is a family of pairwise compatible objects. This section collects the most economical formulations of that kind.

6.1 Clique, independent set, and vertex cover

Maximum clique is represented directly by Theorem 3.1. The maximum independent-set number $\alpha(G)$ is the clique number of the complement graph $\overline{G}$, so

$$\alpha(G) = \frac{1}{\nu(\widehat{Q}(\overline{G}))}. \quad (46)$$

A vertex cover C has complement $V \setminus C$ independent. Thus G has a vertex cover of size at most k if and only if

$$\nu(\widehat{Q}(\overline{G})) \leq \frac{1}{n - k}. \quad (47)$$

These formulations use n coordinates and no auxiliary variables.

6.2 Set packing

Given a family $\mathcal{S} = \{S_1, \dots, S_m\}$, create one graph vertex per set and join i and j when $S_i \cap S_j = \emptyset$. Then cliques are pairwise disjoint subfamilies.

Proposition 6.1 (Set packing). *The maximum packing cardinality is*

$$p^* = \frac{1}{\nu(\widehat{Q}(H_{\mathcal{S}}))}. \quad (48)$$

In particular, a packing of at least k sets exists if and only if the StQP optimum is at most $1/k$. The formulation uses m simplex coordinates.

A weighted set-packing problem is obtained by assigning a positive weight to each candidate set and applying Theorem 5.1.

6.3 Three-dimensional matching

Let $T \subseteq X \times Y \times Z$ be a collection of triples. Create one graph vertex per triple and connect two vertices when the triples are disjoint in all three coordinates. A clique is a three-dimensional matching.

Proposition 6.2 (3D matching). *The maximum matching cardinality equals $1/\nu(\hat{Q}(H_T))$. If $|X| = |Y| = |Z| = q$, a perfect three-dimensional matching exists if and only if*

$$\nu(\hat{Q}(H_T)) = \frac{1}{q}.$$

The formulation uses $|T|$ coordinates.

6.4 Exact cover through weighted clique

Let U be a universe and $\mathcal{S} = \{S_1, \dots, S_m\}$ a family of nonempty subsets. Form the disjointness graph used for set packing, but give vertex i weight $w_i = |S_i|$. A clique is a disjoint subfamily, and its total weight equals the cardinality of its union. Therefore every clique has weight at most $|U|$.

Proposition 6.3 (Exact cover). *The family $\mathcal{S}$ contains an exact cover of U if and only if*

$$\omega_w(H_{\mathcal{S}}) = |U|. \quad (49)$$

Equivalently, with $w_{\min} = \min_i |S_i|$,

$$\nu(Q_{\tau, w}(H_{\mathcal{S}})) = \frac{\tau w_{\min}}{|U|}. \quad (50)$$

A clique-supported minimizer decodes to the exact-cover subfamily.

The natural binary comparison is maximum-weight independent set on the conflict graph of intersecting sets. Let

$$\mathcal{E}_c := \{\{i, j\} : i < j, S_i \cap S_j \neq \emptyset\}, \quad r_i := \frac{w_i}{w_{\max}} \in (0, 1].$$

The standard independent-set penalty construction [3, Section 4.2] gives

$$H_{\text{EC}}(z) = -\sum_{i=1}^m r_i z_i + 2 \sum_{\{i, j\} \in \mathcal{E}_c} z_i z_j, \quad z \in \{0, 1\}^m. \quad (51)$$

If a selected set i has $d_i \geq 1$ selected conflicts, deleting it changes H_{EC} by $r_i - 2d_i < 0$. Every minimizer is therefore conflict-free. On a disjoint subfamily the objective equals minus its covered cardinality divided by $w_{\max}$, so an exact cover exists precisely when

$$\min_z H_{\text{EC}}(z) = -\frac{|U|}{w_{\max}}.$$

Both formulations use m variables or coordinates and one pair interaction per intersecting set pair. The QUBO has a fixed conflict coefficient after normalizing its unary weights; the StQP stores reciprocal weights on its diagonal and uses one global normalization. Thus neither a dimension saving nor a general penalty-scale advantage follows from this comparison. In the displayed QUBO polynomial convention, if a conflict exists, the nonzero coefficient ratio is $2w_{\max}/w_{\min}$, equal to the full StQP coefficient ratio at $\tau = 1/2$. Converting the polynomial to a symmetric $z^T Q z$ convention halves its off-diagonal coefficients, so coefficient comparisons must specify the convention.

6.5 Hitting set and set cover

Karp’s HITTING SET asks for a set $W \subseteq S$ that meets every member of a family $U_1, \dots, U_m \subseteq S$ in *exactly* one element [10]. It is exact cover on the transposed incidence structure, and therefore weighted-graph-native. Assume every U_i is nonempty, and discard elements that lie in no U_i , since they never affect feasibility. Give each remaining element s the weight $\deg(s) := |\{i : s \in U_i\}|$, and join $s \neq s'$ when no U_i contains both.

Proposition 6.4 (Karp’s hitting set). *Let H_U be this weighted graph. A set W with $|W \cap U_i| = 1$ for all i exists if and only if $\omega_w(H_U) = m$, equivalently $\nu(Q_{\tau,w}(H_U)) = \tau w_{\min}/m$ with $w_{\min} = \min_s \deg(s)$.*

Proof. A clique W meets every U_i at most once, so $w(W) = \sum_{s \in W} \deg(s) = \sum_i |W \cap U_i| \leq m$, with equality exactly when every U_i is met once. The value statement is [Theorem 5.1](#). $\square$

The Garey–Johnson variant [46], which asks for at most k elements meeting every U_i *at least* once, is a covering problem and is treated next.

The Garey–Johnson hitting set and set cover do not generally reduce to pairwise compatibility as directly as packing. A feasible hitting set is not a family of pairwise compatible elements; its defining condition is collective coverage. Two routes are available:

- (i) reduce set cover to hitting set or vice versa and then to a binary linear system;
- (ii) build a local-state factor graph in which each coverage constraint is a factor over the incident selection variables, encoding infeasible coverage patterns as high-cost states rather than omitting them ([Remark 5.7](#)).

The second route can be compact when each set or element has bounded incidence, but a factor of degree d has 2^d local states in the factor-state compiler, none of which may be deleted ([Remark 5.7](#)). The binary-linear compiler of [Section 7](#) remains polynomial in the ordinary bit model and is preferable for unbounded incidence.

6.6 Dominant sets and weighted graph structure

Continuous simplex quadratics are also used natively for weighted clustering. Maximizing $x^T A x$ over the simplex can identify coherent weighted vertex groups, with local solutions related to dominant sets [47], with edge-weighted and hypergraph extensions [48, 49]. Here the simplex quadratic is the model itself rather than a reduction from a discrete NP problem, and it belongs in this survey because the same StQP hardware primitive can support both exact combinatorial encodings and native continuous graph models.

7 A generic binary-linear feasibility compiler

Direct compatibility formulations are preferable when available. For completeness, we now give a generic polynomial route from binary linear feasibility to StQP. It covers every problem with a polynomial-size 0–1 integer programming formulation and supplies a uniform route for the remaining members of Karp’s list.

7.1 From inequalities to a zero-penalty QUBO

Consider

$$Az \leq b, \quad z \in \{0, 1\}^n, \tag{52}$$

with integer data. Rational data can be scaled to integers with polynomial bit-length growth. For row r , let

$$L_r := \sum_{i: a_{ri} < 0} a_{ri}, \quad U_r := b_r - L_r. \quad (53)$$

If $U_r < 0$, the row is infeasible for every binary z and the instance can be rejected immediately. Otherwise, if the row is feasible, its slack $s_r = b_r - a_r^\top z$ is an integer in $[0, U_r]$. Introduce

$$\ell_r := \lceil \log_2(U_r + 1) \rceil \quad (54)$$

slack bits $y_{r0}, \dots, y_{r, \ell_r - 1}$ and encode

$$s_r(y) = \sum_{t=0}^{\ell_r - 1} 2^t y_{rt}.$$

Define

$$P(z, y) := \sum_{r=1}^m \left(a_r^\top z + s_r(y) - b_r \right)^2. \quad (55)$$

Because binary variables satisfy $u^2 = u$, expansion of (55) is a QUBO [50]; higher-degree residuals can be quadratized with few ancillas [51], and penalty weights admit alternatives to the dominant-coefficient rule [52].

Theorem 7.1 (Binary-linear feasibility compiler). *The system (52) is feasible if and only if*

$$\min_{z, y \in \{0, 1\}} P(z, y) = 0. \quad (56)$$

Combining the expanded QUBO with Theorem 5.6 gives a polynomial-size StQP threshold encoding of binary-linear feasibility.

Proof. If z is feasible, every integer slack $b_r - a_r^\top z \in [0, U_r]$ has a binary representation using ℓ_r bits, and all squared residuals vanish. Conversely, if $P = 0$, every residual vanishes and $a_r^\top z = b_r - s_r(y) \leq b_r$. $\square$

Equalities require no slack. A constraint $a^\top z \geq b$ is multiplied by -1 . Bounded integer variables are encoded in binary together with range constraints if the binary box is larger than the intended domain.

7.2 Size accounting

Let

$$p = n + \sum_{r=1}^m \ell_r \quad (57)$$

be the number of binary variables after adding slack bits. Let t_r be the number of binary variables appearing in residual r . Expansion of P contains at most

$$M_2 \leq \sum_{r=1}^m \binom{t_r}{2} \quad (58)$$

nonzero pair terms before duplicate aggregation. The basic factor-state compiler therefore uses at most

$$N_\Delta \leq 2p + 4M_2. \quad (59)$$

This is polynomial in the binary input length because each ℓ_r is logarithmic in an integer whose bit length is part of the input. The bound is intentionally simple rather than tight. Factor aggregation, quadratization reuse, and direct CSP constructions can reduce it substantially.

7.3 Adding an objective or decision threshold

A 0–1 optimization problem

$$\min \left\{ c^\top z : Az \leq b, z \in \{0, 1\}^n \right\}$$

can be converted to a sequence of feasibility questions $c^\top z \leq T$, or to one QUBO using a sufficiently dominant feasibility penalty. In the second route, minimize $MP(z, y) + c^\top z$ with

$$M > \sum_{i=1}^n |c_i|,$$

which suffices because $P \geq 1$ at every infeasible point while the objective varies by at most $\sum_i |c_i|$. For integer objective data, the feasibility-sequence route avoids that coefficient and preserves exact decision thresholds, using $O(\log(2 + \sum_i |c_i|))$ StQP solves after checking feasibility.

7.4 Precision of the binary-linear route

The residual squares in (55) carry the input numbers into the factor values. The resulting matrix still has polynomial encoding length, but its coefficient ratios and value separations depend on the factorization, offsets, and attainable energy levels. We use the separation bounds and accuracy conditions of Section 5.7.

For Partition with $n \geq 2$ positive items, let $A = \sum_i a_i$. Expanding $(2 \sum_i a_i z_i - A)^2$ gives the constant A^2 , unary coefficients $4a_i(a_i - A)$, and pair coefficients $8a_i a_j$. Retaining these $F = n(n+1)/2$ factors gives, on a yes-instance,

$$D = 5A^2 - 4 \sum_i a_i^2.$$

When $D/F \rightarrow \infty$, the offset rule (45) yields the unit-step estimate $g(\eta^*) \sim \tau/(4DF)$, whereas unit slack gives exactly $g(1) = \tau/[(D+F)(D+F-1)]$. These bounds can be inverse-polynomial in integer magnitudes, and hence exponentially small in their bit lengths. A tolerance of $2^{-O(b)}$ requires $O(b)$ binary precision bits; physical resolution and bit complexity are distinct costs.

A common scaling illustrates the distinction. Take $a_1 = a_2 = L$ and slack $\eta = 4L^2$, where $L \geq 1$ is an integer. Each unary block has weights $(4, 8)L^2$, and the pair block has weights $(12, 12, 12, 4)L^2$. Thus $w_{\min} = 4L^2$ and $W = 24L^2$. The two distinct local-minimum values are $\tau/6$ and $\tau/5$, so the actual gap is $\tau/30$, independent of L . The matrix is also independent of L , with full nonzero coefficient range $3/\tau$. Nevertheless, the unit-step bound is

$$g_W = \frac{\tau}{6(24L^2 - 1)} \rightarrow 0.$$

Using the attainable spacing $\Delta_E = 4L^2$ in (43) recovers the exact gap. Integer magnitudes alone therefore do not determine the physical precision burden.

For a dominant feasibility penalty M , if F and τ are fixed and $D(M) = \Theta(M)$, optimizing the common slack makes the unit-step bound scale as $\Theta(M^{-1})$; fixed slack instead gives $\Theta(M^{-2})$. These scalings apply to the bound defined in (44).

An alternative is to encode arithmetic in Boolean constraints and compose a polynomial reduction to 3-SAT with Construction 4.3. For $m \geq 2$ clauses this gives the fixed alphabet $\{0, 1, 2\}$ and a target decision-class separation of at least $1/[m(m-1)]$, with m polynomial in input bit length. This offers an inverse-polynomial sufficient tolerance, but there is no universal dimension ordering between this route and the residual-square route. The comparison must report matrix size, coefficient ratios, and a valid separation bound for the chosen representation.

7.5 Representative constructions

Subset sum and Karp’s knapsack variant. Given positive integers $a_1, \dots, a_n$ and target t , define

$$P(z) = \left(\sum_{i=1}^n a_i z_i - t \right)^2. \quad (60)$$

Then $P^* = 0$ exactly when a target-sum subset exists. Expanding P gives a QUBO, dense in the worst case, to which [Corollary 5.11](#) applies.

Partition. Let $A = \sum_i a_i$. A partition into equal sums exists exactly when

$$\min_{z \in \{0,1\}^n} \left(2 \sum_i a_i z_i - A \right)^2 = 0. \quad (61)$$

Set cover. Use one bit z_j per candidate set. The decision whether at most k sets cover U is

$$\sum_{j: u \in S_j} z_j \geq 1 \quad (u \in U), \quad \sum_j z_j \leq k. \quad (62)$$

Hitting set (Garey–Johnson variant). Use one bit z_u per universe element and impose

$$\sum_{u \in S_j} z_u \geq 1 \quad (j = 1, \dots, m), \quad \sum_{u \in U} z_u \leq k. \quad (63)$$

Job sequencing (Karp). Karp’s JOB SEQUENCING has one machine and jobs $j \in [n]$ with integer execution times T_j , deadlines D_j and penalties P_j ; it asks for an order in which the total penalty of late jobs is at most k [10]. Let $y_j = 1$ mean that job j is on time. A set of on-time jobs can be scheduled before their deadlines exactly when earliest-deadline-first order does so, which gives

$$\sum_{i: D_i \leq D_j} T_i y_i \leq D_j \quad (j \in [n]), \quad \sum_{j=1}^n P_j (1 - y_j) \leq k. \quad (64)$$

The late jobs are appended in any order. Using (54), this uses $p = n + \sum_j \lceil \log_2(D_j + 1) \rceil + \lceil \log_2(k + 1) \rceil$ bits.

Parallel-machine job sequencing (Lucas variant). For the integer-length makespan version used in Ising formulation work [3], let job i have length L_i , let $a \in [M]$ index machines, and let T be the makespan threshold. Binary assignment variables z_{ia} satisfy

$$\sum_{a=1}^M z_{ia} = 1 \quad (i \in [n]), \quad \sum_{i=1}^n L_i z_{ia} \leq T \quad (a \in [M]). \quad (65)$$

This uses nM assignment bits plus $\lceil \log_2(T + 1) \rceil$ slack bits for each machine-load constraint.

Feedback arc set. For a directed graph, let $y_{uv} = 1$ mean arc (u, v) is deleted. Assign each vertex an integer order $\pi_v \in \{0, \dots, n - 1\}$, binary encoded. Impose

$$\pi_u + 1 \leq \pi_v + n y_{uv} \quad ((u, v) \in E), \quad \sum_{(u,v) \in E} y_{uv} \leq k. \quad (66)$$

If at most k arcs can be removed to make the graph acyclic, the remaining arcs admit such a topological order; the converse is immediate. Each arc row has slack range $O(n)$ and needs $O(\log n)$ slack bits, so $p = \Theta((n + |E|) \log n)$.

Feedback vertex set. Let $y_v = 1$ mean vertex v is removed. One valid order formulation is

$$\pi_u + 1 \leq \pi_v + n(y_u + y_v) \quad ((u, v) \in E), \quad \sum_v y_v \leq k. \quad (67)$$

Every remaining arc then respects a topological order. As for arc deletion, $p = \Theta((n + |E|) \log n)$.

Steiner tree. Let $R \subseteq V$ be terminals, let edge costs satisfy $c_e \geq 0$, and choose a root $r \in R$. Use binary edge variables y_e and directed flow variables f_{uv} , binary encoded, with

$$0 \leq f_{uv} \leq (|R| - 1)y_{\{u,v\}}, \quad (68)$$

$$\sum_v f_{rv} - \sum_v f_{vr} = |R| - 1, \quad (69)$$

$$\sum_v f_{vu} - \sum_v f_{uv} = 1 \quad (u \in R \setminus \{r\}), \quad (70)$$

$$\sum_v f_{vu} - \sum_v f_{uv} = 0 \quad (u \notin R), \quad (71)$$

$$\sum_e c_e y_e \leq K. \quad (72)$$

A feasible solution contains a connected subgraph spanning all terminals; cycles can be removed without increasing nonnegative cost. With $\lceil \log_2 |R| \rceil$ bits per directed flow, $p = |E| + O(|E| \log |R|) + \lceil \log_2(K + 1) \rceil$.

These formulations are standard integer-programming patterns rather than claims of novel polyhedral strength. Their role here is to close the compilation chain from Karp’s problems to StQP.

8 A catalog for Karp’s 21 NP-complete problems

Karp’s 1972 paper established the polynomial equivalence of a broad class of satisfiability, covering, packing, routing, assignment, sequencing, and numerical problems [10]. Lucas later provided explicit Ising constructions for all 21 problems and analyzed spin counts and coupling requirements [3]. Table 1 summarizes the corresponding StQP routes developed in this paper. Rather than reproduce Karp’s original reduction tree, we choose modern formulations when they are more direct. Linear-growth interreductions among subsets of the 21 problems may provide alternative compilers with lower overhead for particular applications [53]; standard statements of the problems follow Garey and Johnson [46]. The Hamiltonian rows use the position–vertex compatibility construction of Section 4. Controlled-Markov-chain formulations provide a different continuous approach to Hamiltonicity [54].

Among the 21 original numbered rows, ten use unweighted compatibility graphs and two use weighted graphs. Six use binary-linear feasibility followed by residual-square QUBO and factor-state compilation; the knapsack target-sum and Partition rows start directly from residual-square QUBOs, and Max-Cut uses a factor-state formulation directly. Rows 15’ and 19’ give additional variants and are not included in these counts. The catalog therefore establishes coverage through several routes, with different dimension and precision costs.

For the table, “direct graph” refers to (9), “weighted graph” to (20), and “BILP” to Theorem 7.1 followed by Theorem 5.6. For every residual-square route write $P = c + E$, where c is the omitted constant and E is the compiled factor energy. The original penalty is recovered as

$$P^* = c + B - \frac{\tau w_{\min}}{\nu(Q_E)}; \quad P^* = 0 \iff E^* = -c. \quad (73)$$

Trivial instances and out-of-range thresholds are handled before compilation.

The size column (“Basic N_Δ ”) reports the number of simplex coordinates before problem-specific simplification. For generic binary-linear routes, p is the number of binary variables after slack and bounded-integer encoding, and M_2 is the number of nonzero pair terms in the resulting zero-penalty QUBO; the bound $2p + 4M_2$ follows from (59). The column gives p for rows 7, 8, 16, 19 and 19’, where the encoding fixes it; for rows 2, 6 and 15’ it depends on the instance’s coefficient ranges through (53). A row over q variables contributes up to $\binom{q}{2}$ pair terms, so a single cardinality budget $\sum y \leq k$ already makes $M_2 = \Theta(q^2)$. These generic bounds are polynomial but are not necessarily optimal.

Table 1: StQP compilation routes for Karp’s 21 problems. Rows 15 and 19 use Karp’s original definitions; primed rows give additional variants. Route labels, size conventions, and constant recovery are defined in the preceding text.

#	Problem	Preferred StQP route	Basic N_Δ	Exact decision or recovery rule	Main representation issue
1	Satisfiability	Clause–literal compatibility graph, Construction 4.3	N_{lit} , the total number of literal occurrences	SAT iff $\nu(\widehat{Q}) = 1/m$, where $m \geq 1$ is the number of clauses.	At the target value, the support selects one consistent literal per clause; local supports may omit clauses.
2	0–1 integer programming	BILP residual squares, (55), then factor-state StQP	$\leq 2p + 4M_2$	Feasible iff $P^* = 0$, equivalently $E^* = -c$; restore the omitted constant using (73).	Slack bit length and pair terms generated by squaring residuals.
3	Clique	Direct graph compiler on G	$n = V $	$\omega(G) = 1/\nu(\widehat{Q}(G))$. A clique of size at least k exists iff $\nu \leq 1/k$.	Every maximal clique is a strict local minimum.
4	Set packing	Disjointness graph on candidate sets	m candidate sets	Maximum packing cardinality is $1/\nu(\widehat{Q})$.	Pairwise intersection tests; positive packing weights require the weighted formulation.
5	Vertex cover	Independent set in G , equivalently clique in $\overline{G}$	n	A cover of size at most $k < n$ exists iff $\nu(\widehat{Q}(\overline{G})) \leq 1/(n - k)$.	Complementation changes the interaction pattern; the decoder returns the complement of an independent set.
6	Set covering	BILP constraints (62)	$\leq 2p + 4M_2$	A cover using at most k candidate sets exists iff $P^* = 0$.	High-incidence coverage constraints and the budget row generate pair interactions.
7	Feedback node set	Vertex-deletion bits and a binary-encoded topological order, (67)	$\leq 2p + 4M_2$; $p = O((n + E) \log(n + 1))$	At most k vertices can be deleted to leave a directed acyclic graph iff $P^* = 0$.	Order and slack bits; order coefficients scale with n . The deletion budget can generate dense pairs.
8	Feedback arc set	Arc-deletion bits and a binary-encoded topological order, (66)	$\leq 2p + 4M_2$; $p = O((n + E) \log(n + 1))$	At most k arcs can be deleted to leave a directed acyclic graph iff $P^* = 0$.	Order encoding and slack bits; the arc budget can couple every pair of deletion variables.
9	Directed Hamiltonian circuit	Position–vertex compatibility graph	n^2	A Hamiltonian circuit exists iff $\nu(\widehat{Q}) = 1/n$.	Conflicts enforce distinct positions, distinct vertices, and adjacency in the required direction.

Continued on next page

Table 1 – continued

#	Problem	Preferred StQP route	Basic N_Δ	Exact decision or recovery rule	Main representation issue
10	Undirected Hamiltonian circuit	Position–vertex compatibility graph	n^2	A Hamiltonian circuit exists iff $\nu(\widehat{Q}) = 1/n$.	The same position lift, with undirected adjacency constraints.
11	3-SAT	Clause–literal compatibility graph	$N_{\text{lit}} \leq 3m$	Satisfiable iff $\nu(\widehat{Q}) = 1/m$.	For $m \geq 2$, the adjacent target levels are separated by $1/[m(m-1)]$.
12	Chromatic number	Variable–color graph for the decision $\chi(G) \leq k$	nk ; nontrivial $k < n$	k -colorable iff $\nu(\widehat{Q}) = 1/n$.	Color-label symmetry; finding $\chi(G)$ requires testing candidate values of k .
13	Clique cover	k -colorability of $\overline{G}$	nk ; nontrivial $k < n$	A cover by at most k cliques exists iff $\nu(\widehat{Q}) = 1/n$.	Complement density and permutation symmetry among cover labels.
14	Exact cover	Weighted disjointness graph with $w_j = S_j $	m nonempty candidate sets	Exact cover iff $\omega_w = U $, equivalently $\nu(Q_{\tau,w}) = \tau w_{\min}/ U $.	Diagonal coefficients contain reciprocals of set sizes.
15	Hitting set: $ W \cap U_i = 1$ (Karp)	Weighted element graph; $w_s = \{i : s \in U_i\} $	$\leq S $ retained elements	Such a set W exists iff $\omega_w = m$, equivalently $\nu(Q_{\tau,w}) = \tau w_{\min}/m$.	Elements sharing a set conflict. Discard elements of zero incidence; weights are incidence degrees.
15'	Hitting set: at most k elements (Garey–Johnson)	Coverage/budget BILP, (63)	$\leq 2p + 4M_2$	A set of at most k elements meeting every member at least once exists iff $P^* = 0$.	Large incidence rows. A full-table factor lift can be large when incidence is unbounded.
16	Steiner tree	Flow BILP, (68)–(72)	$\leq 2p + 4M_2$; $p = O(E (1 + \log R) + \log(K + 1))$	For nonnegative edge costs, a tree spanning R with cost at most K exists iff $P^* = 0$.	Flow and capacity-slack bits, conservation equations, and a cost-budget row.
17	3-dimensional matching	Disjointness graph on triples	$ \mathcal{T} $	For $ X = Y = Z = q$, a perfect matching exists iff $\nu(\widehat{Q}) = 1/q$.	Overlapping triples generate conflicts; the support identifies the selected triples.
18	Knapsack: target-sum form (Karp)	Square residual (60), then QUBO compiler	$\leq 2n^2$	Target sum t is achievable iff $P^* = 0$. Here $c = t^2$, so $E^* = -t^2$.	Coefficient magnitudes depend on item values and target; pair interactions can be dense.

Continued on next page

Table 1 – continued

#	Problem	Preferred StQP route	Basic N_Δ	Exact decision or recovery rule	Main representation issue
19	Job sequencing: one machine, deadlines and penalties	Earliest-deadline BILP in Section 7.5	$\leq 2p + 4M_2;$ $p = n + \sum_j \ell_j + \ell_k,$ $\ell_j = \lceil \log_2(D_j + 1) \rceil,$ $\ell_k = \lceil \log_2(k + 1) \rceil$	An order with total late-job penalty at most k exists iff $P^* = 0$.	Processing times, deadlines, and late penalties affect bit counts and coefficients.
19'	Multiprocessor makespan (Lucas)	Assignment/load BILP in Section 7.5	$\leq 2p + 4M_2;$ $p = nM + M\ell_T,$ $\ell_T = \lceil \log_2(T + 1) \rceil$	A schedule on M machines with makespan at most T exists iff $P^* = 0$.	Machine-label symmetry, numerical job lengths, and load-row pair expansion.
20	Partition	Square residual (61), then QUBO compiler	$\leq 2n^2$	With $A = \sum_i a_i$, equal partition exists iff $P^* = 0$. Here $c = A^2$, so $E^* = -A^2$.	Offset choices, weight ratios, and attainable energy spacings determine the numerical burden.
21	Max cut	Four local states per edge, (37)	$4 E $	$\text{MaxCut}_c(G) = \tau L / \nu(Q) - L E $, with $L > 0$.	Reciprocal recovery, offset choice, and repeated local copies of vertex labels.

8.1 Direct versus generic routes

The table shows four distinct regimes.

Graph-native problems. Clique, independent set, vertex cover, set packing, three-dimensional matching, coloring, clique cover, Hamiltonian cycle, and satisfiability admit direct unweighted compatibility graphs. Their matrices use coefficients in $\{0, 1, 2\}$, and an optimal support encodes a maximum compatible selection. For feasibility problems, a satisfying assignment is obtained when that support reaches the specified target size.

Weighted graph-native problems. Exact cover, Karp’s hitting set, and weighted packing are naturally maximum-weight clique problems. Their off-diagonal entries remain binary, while weights enter through reciprocal diagonal terms. Max-Cut reaches a weighted clique after lifting to local edge configurations; this is the factor-state construction, so [Corollary 5.8](#) applies to row 21.

Globally constrained problems. Set cover, the Garey–Johnson hitting set, feedback sets, Steiner tree, and general 0–1 feasibility involve collective constraints not captured by pairwise compatibility among the original objects. The binary-linear route is polynomial in input bit length. A local-state route is polynomial in explicit table size but can expand succinct high-arity constraints exponentially. Either representation may be too large or numerically demanding for a particular device.

Numerical problems. Knapsack, Partition, and the job-sequencing variants carry numerical data into the residual-square coefficients and factor weights. Their precision costs depend on attainable energy spacings, common scaling, and offset choices, as illustrated in [Section 7.4](#).

Exactness of the routes. The direct graph routes recover maximum compatible selections and, when the target threshold is attained, solutions of the original feasibility problem. Their local minima correspond exactly to maximal cliques of the compatibility graph; this need not preserve a different local-move neighborhood on the original variables. The factor-state routes are decision-, value-, and solution-exact for the compiled objective: [Theorem 5.6](#) provides both reciprocal value recovery and an optimal-assignment decoder. On a residual-square feasibility route, restore the omitted constant and require $P^* = 0$ to obtain a feasible witness. These routes do not generally preserve the original discrete local-search landscape, because every assignment becomes a strict local minimum. Approximation guarantees also require a specified objective normalization and decoder; see [Section 11.3](#).

8.2 Threshold normalization

For direct graph formulations, a decision problem typically takes the form $\omega(H) \geq K$. With Q_τ ,

$$\omega(H) \geq K \iff \nu(Q_\tau(H)) \leq \frac{\tau}{K}. \quad (74)$$

If the construction guarantees $\omega(H) \leq K$, as in SAT, coloring, or Hamiltonian cycle, feasibility is equivalent to the equality $\nu = \tau/K$. Integer scaling at $\tau = 1/2$ gives threshold $1/K$.

For weighted formulations, a target $\omega_w(H) \geq W$ is equivalent to

$$\nu(Q_{\tau,w}(H)) \leq \frac{\tau w_{\min}}{W}. \quad (75)$$

These formulas provide a direct recovery rule and illustrate the reciprocal compression of large combinatorial values.

9 Representation cost and comparison with QUBO/Ising

Both StQP and QUBO are universal under polynomial reductions, so the relevant comparison is how a particular representation uses the resources available to a solver or physical platform.

9.1 State semantics

A QUBO state is a vertex of a binary hypercube. An StQP state is a point of a simplex. In the direct graph compiler, the discrete solution is encoded by the support of x , while the normalized magnitudes are determined analytically. In the factor-state compiler, a selected local state receives mass proportional to its positive vertex weight. These representations produce different failure modes:

- a QUBO device can violate a logical constraint when penalty weights are insufficient;
- a graph-coded StQP state is always normalized, but an approximate solution may assign small mass to incompatible coordinates;
- QUBO readout is discrete but may require chain repair after embedding [55];
- StQP readout is continuous and requires feasibility and support tolerances.

9.2 Coefficient structure

The graph compiler has an unusually small coefficient alphabet:

$$\hat{q}_{ij} \in \{0, 1, 2\}.$$

There is no instance-dependent penalty coefficient. The magnitude two always marks incompatibility, zero marks compatibility, and one is the diagonal concentration term. This simplicity should not be confused with numerical robustness: the decision-value separation still decreases with the target support size.

The weighted compiler retains off-diagonal entries in $\{0, 1\}$ but uses diagonal values $\tau w_{\min}/w_i$. Thus objective data are localized to onsite coefficients. The factor-state construction can therefore be viewed as a separation between:

$$\begin{aligned} \text{factor-state consistency} &\longleftrightarrow \text{off-diagonal pattern,} \\ \text{objective and constraint penalties} &\longleftrightarrow \text{diagonal values.} \end{aligned} \tag{76}$$

QUBO generally distributes both constraint and objective information across linear and pairwise coefficients. That can be more compact, especially when the original problem is already pairwise binary. Conversely, the direct graph compilers of [Sections 3](#) and [4](#) have no penalty scale to balance. Factor-compiled hard constraints do reintroduce a penalty scale ([Remark 5.7](#)); [Remark 5.13](#) gives a sufficient choice. Fixed normalized conflict coefficients are also available for the exact-cover QUBO in [\(51\)](#) and the clique QUBO below.

9.3 Connectivity and baseline removal

Literal matrix density is not always the correct physical cost. As in [Section 3.5](#), the dense baseline in $Q_\tau(G)$ may be removed, leaving the attractive gauge $\tilde{Q}_\tau(G)$ of [\(13\)](#), which is sparse when G is sparse; the repulsive gauge $Q_\tau(G)$ instead uses nonnegative interactions and highlights conflicts. A hardware report should state which gauge is implemented and whether the normalization circuitry implements a global all-to-all baseline at negligible marginal cost.

In factor-state graphs, nonzero off-diagonal blocks occur within each factor and between overlapping factors. This preserves bounded interaction degree when domain sizes, arity, and incidence are uniformly

Table 2: Representative size behavior. Counts are formulation-level and exclude hardware embedding auxiliaries.

Problem class	Direct StQP coordinates	Typical binary variables before hardware embedding	Comment
Clique / independent set	n	n	Same state count; different geometry; the fixed QUBO below has the same clique supports
3-SAT with m clauses	at most $3m$	original Boolean variables plus quadratization auxiliaries, when needed	Which is smaller depends on clause-to-variable ratio
k -coloring	nk	nk in one-hot QUBO	Similar lifted state count
Hamiltonian cycle	n^2	n^2 in standard one-hot QUBO	Similar assignment lift; constraint encoding differs
Sparse QUBO with n variables, m pairs	at most $2n + 4m$	n	Generic StQP compiler expands local configurations
Set packing with m candidate sets	m	m	Direct compatibility is equally compact
Exact cover with m candidate sets	m	m	Same candidate count and conflict graph; compare normalized unary weights with reciprocal diagonal weights; see (51)
Max-Cut	$4 E $ by factor-state route	$ V $	QUBO/Ising is normally much more compact

bounded. By contrast, a one-hot QUBO often introduces dense penalty blocks inside each assignment group and additional objective couplings. The one-hot baseline is itself a choice: domain-wall encoding uses $n(k - 1)$ variables for n variables with k values, with a different penalty connectivity [56–58], so the counts in Table 2 should be read against the QUBO encoding actually used. Neither representation dominates uniformly.

9.4 Dimension

Table 2 summarizes typical state counts.

The remaining metrics of Section 2.2 depend on both the route and the chosen gauge. For the literal integer graph matrix, $\mathcal{A}_Q \subseteq \{0, 1, 2\}$ and $\kappa_{\text{coef}} \leq 2$, with equality in the latter when a nonedge exists. The nonzero count is $M_Q = N_\Delta + \#\{\text{nonedges}\}$; after baseline removal it is $N_\Delta + |E|$. Factor-state matrices satisfy $\mathcal{A}_Q \subseteq \{0, 1\} \cup \{\tau w_{\min}/w_i\}$. Their diagonal range is $w_{\max}/w_{\min}$, and their full nonzero range is $w_{\max}/(\tau w_{\min})$ when a unit nonedge entry exists; otherwise it is $w_{\max}/w_{\min}$. Equation (34) bounds M_Q , while d_Q depends on local-state counts and factor overlaps.

StQP is therefore not a general compression of QUBO. The basic factor-state compiler expands local configurations, whereas clique, set packing, and exact cover admit direct simplex formulations with one coordinate per original candidate.

9.5 Precision and dynamic range

The integer graph compiler has a fixed coefficient alphabet, but distinguishing clique sizes k and $k - 1$ by their optimal values requires resolving a difference of order k^{-2} . For weighted matrices, the diagonal range is $w_{\max}/w_{\min}$; the full coefficient range also depends on the presence of unit nonedge entries. Reciprocal recovery compresses assignment values. The local-value and decision-class bounds of [Section 5.7](#) quantify this compression, with refinements when energy spacings are known. A hardware comparison should report:

- (i) programmed coefficient bit depth and realized coefficient error;
- (ii) objective-estimation error or sample variance;
- (iii) feasibility error $|\mathbf{1}^\top x - 1| + \|x_-\|_1$;
- (iv) support-recovery threshold and stability;
- (v) the relevant decision or local-value separation, identifying whether it is measured exactly or bounded conservatively.

An accurate support can remain useful even when objective readout is insufficient for a threshold decision. Conversely, an accurate scalar objective does not guarantee clean support recovery. Both outputs should be retained.

9.6 Landscape and local solvers

The regularized graph compiler makes landscape complexity explicit: maximal cliques are strict local minima, and by [Theorem 5.3](#) the same holds with weights. In QUBO, penalty choices can create additional infeasible local minima whose interpretation is less direct [\[59\]](#). The comparison is not symmetric, however. On any factor-compiled instance every assignment is a strict local minimizer ([Corollary 5.8](#)), whereas the local minima of the corresponding QUBO under single-bit flips are only the 1-flip-stable assignments, a subset. On these routes the StQP landscape therefore filters less, not more. Together the two facts suggest a classification of traps:

Feasibility traps

local solutions that violate the intended discrete constraints;

Combinatorial traps

valid but suboptimal discrete solutions, such as maximal nonmaximum cliques;

Continuous traps

local minima whose support does not represent a compatible discrete selection;

Embedding traps

artifacts introduced by hardware connectivity or auxiliary chains.

For $0 < \tau < 1$, the graph construction excludes nonclique *local minima*, but not nonclique stationary points. For the two-vertex edgeless graph, $Q = \begin{pmatrix} \tau & 1 \\ 1 & \tau \end{pmatrix}$ and $x = (1/2, 1/2)$ satisfy the simplex KKT conditions, yet x is a strict maximum. Thus stationarity alone does not imply a clique-supported solution. On complete factor-state instances every local minimum represents an assignment, but infeasible and suboptimal assignments remain local minima by [Corollary 5.8](#).

A fixed-coefficient QUBO baseline. For unweighted clique, the standard conflict-penalty construction, equivalently independent set in the complement graph [\[3, Section 4.2\]](#), permits

$$H(z) = - \sum_i z_i + 2 \sum_{\substack{i < j \\ \{i,j\} \notin E}} z_i z_j, \quad z \in \{0, 1\}^n. \quad (77)$$

Deleting a selected vertex with $d \geq 1$ selected nonneighbors changes H by $1 - 2d < 0$. A nonmaximal clique admits an addition of cost -1 , whereas every single-bit move from a maximal clique increases H . Hence the global minima encode maximum cliques, and the strict single-bit local minima encode exactly maximal cliques. The polynomial coefficients are fixed: -1 on unary terms and 2 on nonedge terms. Fixed coefficients and meaningful clique local minima therefore do not distinguish StQP from every QUBO formulation. The relevant differences include continuous geometry, coefficient signs, readout precision, and the implementation of normalization and interactions.

9.7 Decoding and verification

A direct graph-coded StQP candidate can be processed as follows:

- (1) project or renormalize to the simplex if needed;
- (2) rank coordinates by mass;
- (3) extract a compatible support using the graph;
- (4) greedily extend it to a maximal clique;
- (5) evaluate the resulting discrete objective;
- (6) compare against a certified upper or lower bound when available.

This ranking-based procedure is heuristic. The objective-nonincreasing purification in [Section B](#), followed by face minimization and extension, is the decoder required for the guarantees in [Sections 5.6](#) and [11.3](#); first-order and active-set methods for StQP provide the continuous half [\[60, 61\]](#). The continuous vector therefore also serves as a heuristic score over local choices. In weighted factor-state models, consistency repair combines local states that agree on shared variables and then locally reoptimizes conflicting factors.

9.8 Summary of representation costs

The basic factor-state compiler expands an existing pairwise binary objective into local configurations; it can therefore be substantially larger than a direct QUBO, as in the Max-Cut example. For clique and set packing, direct StQP and binary formulations use the same number of coordinates or variables. Both can have fixed coefficients and maximal-clique local minima, as [\(77\)](#) shows. For exact cover, the weighted independent-set QUBO in [\(51\)](#) also avoids elementwise equality penalties and uses the same candidate count and conflict graph. The case for StQP is its alignment with a normalized, nonnegative native state and its explicit support-based decoding. Whether that alignment offsets continuous precision, readout, and search costs depends on the device and instance family. Comparisons should use strong problem-specific binary formulations and report dimension, connectivity, coefficient range, decision separation, decoding overhead, and end-to-end time-to-target. The formulations establish these comparison points; they do not select a universally superior model.

10 A hardware-neutral simplex quadratic optimizer/analyzer

Analog and oscillator-based platforms for continuous optimization are advancing quickly [\[34, 62\]](#). The formulations motivate an abstract hardware interface. We call a device satisfying this interface an StQA ([Section 1](#)).

10.1 Input contract

The logical input is a symmetric matrix $Q \in \mathbb{S}^N$, represented by:

- diagonal onsite terms Q_{ii} ;
- programmable relative pair interactions Q_{ij} ;
- a declared global scaling and additive simplex gauge;
- coefficient precision, range, and sign restrictions.

For graph-coded NP instances, a minimal device mode needs only diagonal values and binary compatibility or incompatibility couplings. Weighted factor instances require programmable positive diagonal ratios.

10.2 State contract

The physical state must map to

$$x_i \geq 0, \quad \sum_{i=1}^N x_i = 1. \quad (78)$$

Normalization may be enforced exactly by a conservation law or approximately by feedback. If only approximate, the device should report the unnormalized state so that feasibility correction is accounted for as postprocessing rather than hidden inside the solver.

A scaled simplex $\sum_i x_i = R$ is equivalent after $x = Ry$, but it rescales the quadratic energy by R^2 . Hardware comparisons must distinguish logical normalization from physical amplitude.

10.3 Energy contract

The effective Lyapunov or energy function should be proportional, up to a simplex-constant shift, to

$$E_Q(x) = x^\top Q x. \quad (79)$$

Pairwise physical interactions may realize Q directly, realize a relative gauge $Q - \beta \mathbf{1}\mathbf{1}^\top$, or implement a factorized approximation. The realized matrix Q^{real} should be measured or calibrated, not merely inferred from requested settings.

10.4 Search contract

The local minima characterized in [Section 9.6](#) motivate search mechanisms beyond monotone local descent. Candidate mechanisms include:

- controlled noise or temperature schedules;
- nonlocal support moves;
- multistart and restart policies;
- parameter continuation in τ or in coupling strength, the continuous analogue of annealed replication [\[63, 64\]](#);
- parallel replicas and exchange;
- hybrid branch-and-bound or cutting-plane control;
- a mechanism that produces certified lower bounds.

None of these mechanisms is analyzed here as a global solution algorithm; its success rate and certification capability require separate evaluation.

10.5 Output contract

At minimum, a run should return:

$$x^{\text{raw}}, \quad t_{\text{device}}, \quad \text{programmed } Q, \quad \text{estimated realized } Q^{\text{real}}, \quad (80)$$

plus all random seeds or control traces needed for reproducibility. Derived quantities should include

$$r_{\text{feas}} = |\mathbf{1}^\top x - 1| + \|x_-\|_1, \quad (81)$$

$$r_{\text{PG}}(x; \eta) = \frac{\|x - \Pi_{\Delta_N}(x - \eta \nabla E_Q(x))\|}{\eta}, \quad (82)$$

$$E_{\text{logical}} = x^\top Q x, \quad E_{\text{realized}} = x^\top Q^{\text{real}} x. \quad (83)$$

Here $(x_-)_i = \max\{-x_i, 0\}$, Π_{Δ_N} is Euclidean projection onto the simplex, and the residual uses a positive step size η and the Euclidean norm. For feasible x , $r_{\text{PG}} = 0$ certifies first-order stationarity only, not local or global minimality (Section 9.6). For direct graph instances, report decoded clique size and whether the support is maximal. For factor instances, report assignment consistency and original objective value.

10.6 Timing and energy accounting

Results should separate:

- (i) matrix compilation;
- (ii) device programming and calibration;
- (iii) physical evolution time;
- (iv) repeated sampling or restarts;
- (v) readout;
- (vi) feasibility and support repair;
- (vii) classical verification or certification.

Both native-device and end-to-end times are useful; neither should be reported in place of the other without clarification. Similar care is already necessary for Ising and QUBO platforms, where embedding and chain repair can dominate the logical solve [9, 65, 66].

10.7 A minimal graph-coded device mode

The integer-scaled graph family suggests a sharply defined first hardware target:

$$\min_{x \in \Delta_N} \left(\sum_i x_i^2 + 4 \sum_{\{i,j\} \notin E, i < j} x_i x_j \right). \quad (84)$$

The expression follows from $x^\top \hat{Q} x$ because each off-diagonal matrix entry contributes twice. A device needs one common diagonal strength and one common repulsive strength on programmable conflict pairs. Its global optimum is $1/\omega(G)$. This family is restrictive enough for clean physical experiments and expressive enough to remain NP-complete.

Caution 10.1 (No automatic advantage). A fixed-alphabet NP-complete input family is an attractive hardware target because it simplifies programming. It does not remove NP-hardness or guarantee efficient global search. Claims of advantage require end-to-end comparison on declared instance distributions, including precision and success probability.

10.8 Evaluation metrics

For stochastic or analog solvers, report time-to-target at several target levels rather than only best-of-run objective. For direct graph instances, natural targets are:

$$\text{find a clique of size at least } k, \quad \text{find a maximum clique}, \quad \text{certify maximum size.} \quad (85)$$

These are increasingly demanding and should not be conflated. A solver may rapidly find large maximal cliques but rarely find or certify the maximum.

For independent runs of equal duration with success probability $0 < p_s < 1$, a time-to-solution estimate at confidence $1 - \delta$, where $0 < \delta < 1$, is [8, 9]:

$$T_{1-\delta} = t_{\text{run}} \max \left\{ 1, \left\lceil \frac{\log \delta}{\log(1 - p_s)} \right\rceil \right\}, \quad (86)$$

For $p_s = 1$ one run suffices; for $p_s = 0$ no finite number of runs attains positive success confidence. Programming, readout, and postprocessing costs should be reported separately.

11 Limitations and open research questions

11.1 Universality can hide poor formulations

Composing a reduction to CLIQUE with the graph compiler proves universality but may destroy application structure. Even a polynomial expansion can be impractical. A useful formulation theory must therefore seek direct encodings, low-rank or sparse matrix representations, and problem-specific decoders rather than treating universal reducibility as the endpoint.

11.2 Continuous precision is a computational resource

A real-valued state should not be treated as carrying infinite information at zero cost. The exact graph compiler uses simple input coefficients, but its objective levels approach one another as solution size grows. Weighted compilers may introduce large diagonal dynamic ranges and reciprocal value recovery. Future work should characterize bit complexity, noise sensitivity, and finite-temperature behavior for each formulation family.

11.3 Approximation theory

The direct clique formulations preserve multiplicative approximation under objective-nonincreasing decoding. Indeed, let $f(x) = x^\top Q_\tau(G)x$ and suppose a feasible point satisfies $f(x) \leq \rho \nu(Q_\tau(G))$ for some $\rho \geq 1$. Purification followed by minimization on the resulting clique face produces a clique C such that

$$\frac{\tau}{|C|} \leq f(x) \leq \rho \frac{\tau}{\omega(G)}, \quad \text{hence} \quad |C| \geq \frac{\omega(G)}{\rho}.$$

Conversely, the barycenter of a clique of size at least $\omega(G)/\rho$ is a ρ -approximate StQP solution. The weighted formulation gives the same statement with $|C|$ and $\omega(G)$ replaced by $w(C)$ and $\omega_w(G)$. Positive scaling, including the integer-scaled matrix $\hat{Q}$, preserves these ratios; adding a simplex-constant objective shift does not generally preserve them.

Consequently, for every fixed $0 < \varepsilon < 1$, finding a feasible $N^{1-\varepsilon}$ -approximate solution of the restricted N -coordinate StQP family with diagonal entries one and off-diagonal entries in $\{0, 2\}$ is NP-hard, by the corresponding maximum-clique inapproximability result [67]. A polynomial-time

algorithm with this guarantee would imply $P = NP$. This statement concerns multiplicative error relative to the positive minimum value in the specified normalization.

A different guarantee is available for general StQP. Bomze and de Klerk [19] establish a polynomial-time approximation scheme in objective-range-normalized additive error. For rational input data and every fixed $\varepsilon > 0$, a feasible point can be found in polynomial time with

$$f(x) - f_{\min} \leq \varepsilon(f_{\max} - f_{\min}), \quad f_{\min} := \min_{y \in \Delta_N} f(y), \quad f_{\max} := \max_{y \in \Delta_N} f(y).$$

Here the dimension is variable, not fixed. The guarantee is invariant under positive affine changes of the objective, but it does not bound $f(x)/f_{\min}$ when the minimum is small compared with the objective range. There is therefore no conflict with the clique hardness result.

For the direct graph formulation, an additive guarantee $f(x) \leq \tau/\omega(G) + \delta$, with $\delta \geq 0$, gives

$$|C| \geq \frac{\omega(G)}{1 + \delta \omega(G)/\tau}$$

after the same decoding. Fixed additive error thus gives a progressively weaker relative guarantee as the clique number grows. For factor-state models, [Section 5.6](#) instead bounds the decoded energy error through $B - E^*$; no uniform multiplicative guarantee for the original energy follows without additional assumptions. Likewise, the binary-linear feasibility route preserves exact decisions, but an approximate residual minimum need not satisfy the original constraints. Approximation claims must therefore specify the original objective, its normalization, and the decoder; exact representability alone does not determine them.

11.4 Local landscape beyond cliques

The local-minimum correspondences in [Theorem 5.3](#) and [Corollary 5.8](#) specify the admissible endpoints of local descent but do not determine their basins of attraction. The unresolved quantities include basin volumes under a declared initialization distribution, convergence times, the quality of decoded assignments, and their dependence on offsets and search dynamics. The existence of a local minimum for every assignment does not imply that all assignments are reached with equal probability.

11.5 Alternative simplex lifts

Relaxation labeling assigns a probability vector to each object and optimizes compatibility between labels [35]. With variables indexed by i and labels $a \in D_i$, its state space has the form

$$\mathcal{P} = \left\{ (p_{ia}) : p_{ia} \geq 0, \quad \sum_{a \in D_i} p_{ia} = 1 \text{ for every } i \right\} = \prod_i \Delta_{|D_i|}.$$

Quadratic optimization over products of simplexes is a related established class [68]. Each block has a prescribed mass, whereas the single-simplex constructions here allow mass to move between blocks. Rescaling $x_{ia} = p_{ia}/n$ embeds $\mathcal{P}$ into a single simplex only with the additional constraints $\sum_a x_{ia} = 1/n$. Dropping those constraints changes the feasible set. Consequently the maximal-clique and every-assignment local-minimum results do not automatically transfer to product-simplex formulations.

Other representations use marginal polytopes, low-rank factorizations, or entropy regularization. A useful comparison must account for both the number of local states and the cost of enforcing multiple normalization constraints.

11.6 Global certification

StQP has exact copositive and completely positive reformulations [19, 69, 70], summarized in [Section D](#). This suggests hybrid architectures in which an unconventional device supplies primal candidates while classical conic, semidefinite, linear, or branch-and-bound methods supply lower bounds. For graph-coded instances, combinatorial bounds on clique number offer another certification path [38, 71, 72], and exact StQP solvers based on mixed-integer and conic reformulations give a second [32, 73–75]. The division of labor between candidate generation and certification may be more realistic than expecting a physical relaxer to certify globality internally.

11.7 Formulation search as an optimization problem

An instance may admit many equivalent matrices related by scaling, gauge shifts, variable duplication, factor grouping, and alternative reduction routes. Selecting a formulation for a given device can itself be posed as a multiobjective optimization problem over:

$$N_{\Delta}, \quad M_Q, \quad d_Q, \quad \kappa_{\text{coef}}, \quad \text{decision gap}, \quad \text{decoder robustness}.$$

Automated formulation search is a promising counterpart to compiler optimization in conventional computing.

11.8 Empirical question

This paper provides formulations, proofs, and representation-cost analysis. It does not report solver-performance experiments, basin measurements, or hardware benchmarks. A numerical study could compare projected-gradient or replicator dynamics with random restarts on direct graph and factor-state instances, reporting decoded objective values, success probabilities, and sensitivity to offsets. Comparisons with problem-specific binary solvers would then test whether the continuous representation offers a practical benefit for the chosen instance family.

12 Conclusion

The catalog expresses discrete problems as quadratic optimization over one normalized, nonnegative state, using established clique formulations and compatibility graphs. It records explicit matrices, dimensions, thresholds, and decoding rules for direct graph routes and for generic factor-state routes. Twelve of Karp’s 21 problems use direct graph formulations; the other nine use factor-state compilation, including six through binary-linear feasibility.

The representation analysis identifies the cost of these choices. Direct graph routes have simple coefficients and support-based decoding. Generic factor-state compilation can expand dimension and creates a strict local minimum for every complete assignment. Reciprocal objective recovery introduces a dependence on weights, offsets, and attainable energy spacings. The resulting separation bounds and decoding guarantees specify accuracy conditions under the stated assumptions.

Strong binary formulations can match the candidate count and use fixed normalized conflict coefficients for clique, packing, and exact cover. The practical comparison therefore depends on the implementation of normalization and interactions, continuous precision, search dynamics, and readout. Establishing performance for any device requires measurements beyond the exact formulations developed here.

A Compilation procedures

Algorithm 1 Three-level graph-to-StQP compiler

Input: A nonempty simple undirected graph $G = (V, E)$ and $0 < \tau < 1$.

Output: A symmetric matrix $Q \in \mathbb{S}^{|V|}$.

1. Set $Q_{ii} = \tau$ for every $i \in V$.
 2. For each unordered pair $\{i, j\}$, set $Q_{ij} = Q_{ji} = 0$ if it is an edge, and set $Q_{ij} = Q_{ji} = 1$ otherwise.
 3. Return Q .
-

Algorithm 2 Clause–literal compatibility graph

Input: Nonempty CNF clauses $C_1, \dots, C_m$, with $m \geq 1$.

Output: A clause–literal compatibility graph H .

1. Create one vertex $v_{r,\ell}$ for every literal occurrence $\ell \in C_r$.
 2. Join $v_{r,\ell}$ to $v_{s,\ell'}$ exactly when $r < s$ and ℓ' is not the complement of ℓ .
 3. Return H .
-

Algorithm 3 Factor-state weighted StQP compiler

Input: Nonempty finite domains D_i , a nonempty collection of complete factor tables ϕ_α on scopes S_α , offsets $B_\alpha > \max_a \phi_\alpha(a)$, and $0 < \tau < 1$.

Output: A matrix Q and recovery constants $(B, w_{\min}, \tau)$.

1. For every factor α and assignment $a \in \prod_{i \in S_\alpha} D_i$, create a vertex $v = (\alpha, a)$ with weight $w_v = B_\alpha - \phi_\alpha(a)$.
 2. Mark distinct vertices (α, a) and (β, b) compatible exactly when $\alpha \neq \beta$ and a, b agree on $S_\alpha \cap S_\beta$.
 3. Compute $w_{\min} = \min_v w_v$ and set $Q_{vv} = \tau w_{\min} / w_v$.
 4. For $u \neq v$, set $Q_{uv} = 0$ on compatible pairs and $Q_{uv} = 1$ otherwise.
 5. Return Q , $B = \sum_\alpha B_\alpha$, $w_{\min}$, and τ .
-

B Detailed decoding and purification

B.1 Purifying an approximate graph solution

Given $x \in \Delta_n$ and graph G , a simple decoder orders vertices by decreasing x_i and greedily builds a clique. A more objective-aware purification repeatedly identifies a nonedge $\{i, j\}$ in the positive support and moves all mass from one endpoint to the other, choosing the endpoint move with smaller StQP objective. By the proof of [Theorem 3.1](#), this never increases the objective and terminates at clique support in at most $n - 1$ eliminations.

Once a clique C is obtained, replace its masses by the exact face minimizer:

$$x_i \leftarrow \frac{1}{|C|} \quad (i \in C) \tag{87}$$

for the unweighted compiler, or

$$x_i \leftarrow \frac{w_i}{\sum_{j \in C} w_j} \quad (i \in C) \tag{88}$$

for the weighted compiler. A greedy extension can then make C maximal. None of these operations certifies maximum weight.

B.2 Factor-state decoding

At any clique-supported local minimizer, each factor contributes exactly one local state ([Corollary 5.8](#)), so the selected states agree on overlaps and merge into a complete assignment; only its value distinguishes a global minimizer from the rest. For an approximate nonclique support, one may:

- (i) select the largest-mass state of each factor;
- (ii) resolve variable disagreements by weighted voting;
- (iii) perform local coordinate or factor updates in the original objective.

This repair is heuristic and should be timed and reported separately.

C Explicit formulations for selected catalog entries

C.1 Independent set and vertex cover

This is [Section 6.1](#) in matrix form. For a graph G , let $A_{\overline{G}}$ be the adjacency matrix of its complement. Define

$$\widehat{Q}_{\text{IS}} = 2\mathbf{1}\mathbf{1}^\top - 2A_{\overline{G}} - I.$$

Then

$$\min_{x \in \Delta_n} x^\top \widehat{Q}_{\text{IS}} x = \frac{1}{\alpha(G)}.$$

If x^* has support S , then S is a maximum independent set and $V \setminus S$ is a minimum vertex cover.

C.2 Coloring matrix blocks

Index StQP coordinates by (v, c) . The integer-scaled matrix has diagonal one. For distinct coordinates, the coefficient is zero exactly when they are compatible:

$$\widehat{q}_{(v,c),(u,d)} = 0 \iff v \neq u \text{ and } (\{u, v\} \notin E \text{ or } c \neq d). \quad (89)$$

All other off-diagonal coefficients equal two. The same-vertex $k \times k$ blocks are conflict cliques in the matrix, and an edge $\{u, v\}$ contributes conflicts only between equal-color pairs.

C.3 Hamiltonian matrix blocks

Coordinates are (p, v) . Same-position and same-vertex pairs receive coefficient two. For consecutive positions, a pair receives zero only if the corresponding graph edge exists. For nonconsecutive positions with different vertices, the coefficient is zero. This block structure permits procedural matrix evaluation without materializing all n^4 potential assignment pairs.

C.4 Max-Cut local states

For edge $e = \{i, j\}$, the four local states are 00, 01, 10, 11. Their weights are $L, L + c_e, L + c_e, L$. States on disjoint edges are always compatible. States on incident edges are incompatible exactly when they assign different labels to the shared endpoint. Thus each pair of incident edge factors contributes a sparse 4×4 conflict block.

D A note on copositive certification

For any $Q \in \mathbb{S}^N$,

$$\nu(Q) = \max \left\{ \lambda : Q - \lambda \mathbf{1}\mathbf{1}^\top \in \mathcal{COP}_N \right\}, \quad (90)$$

where $\mathcal{COP}_N$ is the copositive cone [19, 69]. Indeed, $Q - \lambda \mathbf{1}\mathbf{1}^\top$ is copositive if and only if

$$y^\top Q y \geq \lambda (\mathbf{1}^\top y)^2 \quad \forall y \geq 0.$$

Normalizing any nonzero $y \geq 0$ to $x = y/(\mathbf{1}^\top y) \in \Delta_N$ gives $x^\top Q x \geq \lambda$. This identity provides an exact global lower-bound certificate, although optimizing over the copositive cone is itself computationally hard. Tractable inner approximations, semidefinite hierarchies, and combinatorial bounds can nevertheless certify many instances.

For a graph-coded matrix $\hat{Q}(G)$, proving $\omega(G) \leq k$ is equivalent to proving

$$\hat{Q}(G) - \frac{1}{k} \mathbf{1}\mathbf{1}^\top \in \mathcal{COP}_n. \quad (91)$$

Thus the formulation connects the physical energy minimum directly to a classical conic certificate.

E Notation summary

symbol	meaning
Δ_N	probability simplex $\{x \geq 0 : \mathbf{1}^\top x = 1\}$
$\nu(Q)$	optimum of $\min_{x \in \Delta_N} x^\top Q x$
A_G	adjacency matrix of graph G
$\omega(G)$	maximum clique cardinality
$\omega_w(G)$	maximum total vertex weight of a clique
$Q_\tau(G)$	three-level graph StQP matrix with diagonal τ
$\hat{Q}(G)$	integer-scaled graph matrix $2Q_{1/2}(G)$
$x(C)$	uniform barycenter of clique C
N_Δ	number of simplex coordinates in an encoding
M_Q	number of nonzero upper-triangular matrix entries
$E(z)$	original finite-domain factor objective
B_α	factor offset satisfying (28)
$w_{\alpha,a}$	factor-state vertex weight $B_\alpha - \phi_\alpha(a)$
$P(z, y)$	zero-penalty QUBO for binary-linear feasibility

Competing interests

The authors are affiliated with Quantum Computing Inc. (QCi).

References

- [1] Francisco Barahona. On the computational complexity of Ising spin glass models. *Journal of Physics A: Mathematical and General*, 15(10):3241–3253, 1982. doi: 10.1088/0305-4470/15/10/028.
- [2] Gary Kochenberger, Jin-Kao Hao, Fred Glover, Mark Lewis, Zhipeng Lü, Haibo Wang, and Yang Wang. The unconstrained binary quadratic programming problem: A survey. *Journal of Combinatorial Optimization*, 28(1):58–81, 2014. doi: 10.1007/s10878-014-9734-0.

- [3] Andrew Lucas. Ising formulations of many NP problems. *Frontiers in Physics*, 2:5, 2014. doi: 10.3389/fphy.2014.00005.
- [4] Fred Glover, Gary Kochenberger, and Yu Du. Quantum bridge analytics I: A tutorial on formulating and using QUBO models. *4OR*, 17(4):335–371, 2019. doi: 10.1007/s10288-019-00424-y.
- [5] Alireza Marandi, Zhe Wang, Kenta Takata, Robert L. Byer, and Yoshihisa Yamamoto. Network of time-multiplexed optical parametric oscillators as a coherent Ising machine. *Nature Photonics*, 8(12):937–942, 2014. doi: 10.1038/nphoton.2014.249.
- [6] Peter L. McMahon, Alireza Marandi, Yoshitaka Haribara, Ryan Hamerly, Carsten Langrock, Shuhei Tamate, Takahiro Inagaki, Hiroki Takesue, Shoko Utsunomiya, Kazuyuki Aihara, Robert L. Byer, Martin M. Fejer, Hideo Mabuchi, and Yoshihisa Yamamoto. A fully programmable 100-spin coherent Ising machine with all-to-all connections. *Science*, 354(6312):614–617, 2016. doi: 10.1126/science.aah5178.
- [7] Hayato Goto, Kosuke Tatsumura, and Alexander R. Dixon. Combinatorial optimization by simulating adiabatic bifurcations in nonlinear Hamiltonian systems. *Science Advances*, 5(4):eaav2372, 2019. doi: 10.1126/sciadv.aav2372.
- [8] Naeimeh Mohseni, Peter L. McMahon, and Tim Byrnes. Ising machines as hardware solvers of combinatorial optimization problems. *Nature Reviews Physics*, 4(6):363–379, 2022. doi: 10.1038/s42254-022-00440-8.
- [9] Byron Tasseff, Tameem Albash, Zachary Morrell, Marc Vuffray, Andrey Y. Lokhov, Sidhant Misra, and Carleton Coffrin. On the emerging potential of quantum annealing hardware for combinatorial optimization. *Journal of Heuristics*, 30(5–6):325–358, 2024. doi: 10.1007/s10732-024-09530-5.
- [10] Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller, James W. Thatcher, and Jean D. Bohlinger, editors, *Complexity of Computer Computations*, The IBM Research Symposia Series, pages 85–103. Plenum Press, New York, 1972. doi: 10.1007/978-1-4684-2001-2_9.
- [11] Immanuel M. Bomze. On standard quadratic optimization problems. *Journal of Global Optimization*, 13(4):369–387, 1998. doi: 10.1023/A:1008369322970.
- [12] Immanuel M. Bomze. Evolution towards the maximum clique. *Journal of Global Optimization*, 10(2):143–164, 1997. doi: 10.1023/A:1008230200610.
- [13] James T. Hungerford and Francesco Rinaldi. A general regularized continuous formulation for the maximum clique problem. *Mathematics of Operations Research*, 44(4):1161–1173, 2019. doi: 10.1287/moor.2018.0954.
- [14] Immanuel M. Bomze, Marcello Pelillo, and Volker Stix. Approximating the maximum weight clique using replicator dynamics. *IEEE Transactions on Neural Networks*, 11(6):1228–1241, 2000. doi: 10.1109/72.883403.
- [15] Stephen A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the Third Annual ACM Symposium on Theory of Computing (STOC 1971)*, pages 151–158, New York, 1971. ACM. doi: 10.1145/800157.805047.
- [16] Theodore S. Motzkin and Ernst G. Straus. Maxima for graphs and a new proof of a theorem of Turán. *Canadian Journal of Mathematics*, 17:533–540, 1965. doi: 10.4153/CJM-1965-053-6.
- [17] Luana E. Gibbons, Donald W. Hearn, Panos M. Pardalos, and Motakuri V. Ramana. Continuous characterizations of the maximum clique problem. *Mathematics of Operations Research*, 22(3):754–768, 1997. doi: 10.1287/moor.22.3.754.
- [18] Katta G. Murty and Santosh N. Kabadi. Some NP-complete problems in quadratic and nonlinear programming. *Mathematical Programming*, 39(2):117–129, 1987. doi: 10.1007/BF02592948.
- [19] Immanuel M. Bomze and Etienne de Klerk. Solving standard quadratic optimization problems via linear, semidefinite and copositive programming. *Journal of Global Optimization*, 24(2):163–185, 2002. doi: 10.1023/A:1020209017701.
- [20] Immanuel M. Bomze, Werner Schachinger, and Reinhard Ullrich. The complexity of simple models—a study of worst and typical hard cases for the standard quadratic optimization problem. *Mathematics of Operations Research*, 43(2):651–674, 2018. doi: 10.1287/moor.2017.0877.
- [21] Philippe Jégou. Decomposition of domains based on the micro-structure of finite constraint-satisfaction problems. In *Proceedings of the Eleventh National Conference on Artificial Intelligence (AAAI 1993)*, pages 731–736. AAAI Press, 1993.

- [22] Rina Dechter. *Constraint Processing*. Morgan Kaufmann, San Francisco, 2003. doi: 10.1016/B978-1-55860-890-0.X5000-2.
- [23] Francesca Rossi, Peter van Beek, and Toby Walsh, editors. *Handbook of Constraint Programming*, volume 2 of *Foundations of Artificial Intelligence*. Elsevier, Amsterdam, 2006.
- [24] Harry G. Barrow and Rod M. Burstall. Subgraph isomorphism, matching relational structures and maximal cliques. *Information Processing Letters*, 4(4):83–84, 1976. doi: 10.1016/0020-0190(76)90049-1.
- [25] Giorgio Levi. A note on the derivation of maximal common subgraphs of two directed or undirected graphs. *Calcolo*, 9(4):341–352, 1973. doi: 10.1007/BF02575586.
- [26] Achref El Mouelhi, Philippe Jégou, and Cyril Terrioux. Microstructures for CSPs with constraints of arbitrary arity. In *Proceedings of the Tenth Symposium on Abstraction, Reformulation, and Approximation (SARA 2013)*. AAAI Press, 2013.
- [27] Sujay Sanghavi, Devavrat Shah, and Alan S. Willsky. Message passing for max-weight independent set. In *Advances in Neural Information Processing Systems*, volume 20, 2007.
- [28] Stanislav Busygin. A new trust region technique for the maximum weight clique problem. *Discrete Applied Mathematics*, 154(15):2080–2096, 2006. doi: 10.1016/j.dam.2005.04.010.
- [29] Panos M. Pardalos, Oleg A. Prokopyev, and Stanislav Busygin. Continuous approaches for solving discrete optimization problems. In *Handbook on Modelling for Discrete Optimization*, volume 88 of *International Series in Operations Research & Management Science*, pages 39–60. Springer, 2006. doi: 10.1007/0-387-32942-0_2.
- [30] William W. Hager and James T. Hungerford. Continuous quadratic programming formulations of optimization problems on graphs. *European Journal of Operational Research*, 240(2):328–337, 2015. doi: 10.1016/j.ejor.2014.05.042.
- [31] Mykyta Makovenko, Sergiy Butenko, and Miltiades Pardalos. Regularized standard polynomial programming formulations for the maximum clique problem. *Computational Optimization and Applications*, 92(3):923–949, 2025. doi: 10.1007/s10589-025-00674-z.
- [32] Andrea Scozzari and Fabio Tardella. A clique algorithm for standard quadratic programming. *Discrete Applied Mathematics*, 156(13):2439–2448, 2008. doi: 10.1016/j.dam.2007.09.020.
- [33] Fabio Tardella. Connections between continuous and combinatorial optimization problems through an extension of the fundamental theorem of linear programming. *Electronic Notes in Discrete Mathematics*, 17:257–262, 2004. doi: 10.1016/j.endm.2004.03.054.
- [34] Wenxiao Cai, Zongru Li, Yu-Neng Wang, Sara Achour, and Thomas H. Lee. Oscillator formulations of many NP problems. arXiv:2509.13560, 2025.
- [35] Marcello Pelillo. The dynamics of nonlinear relaxation labeling processes. *Journal of Mathematical Imaging and Vision*, 7(4):309–323, 1997. doi: 10.1023/A:1008255111261.
- [36] Stephen A. Vavasis. Quadratic programming is in NP. *Information Processing Letters*, 36(2):73–77, 1990. doi: 10.1016/0020-0190(90)90100-C.
- [37] Stephen A. Vavasis. *Nonlinear Optimization: Complexity Issues*. Oxford University Press, New York, 1991.
- [38] Qinghua Wu and Jin-Kao Hao. A review on algorithms for maximum clique problems. *European Journal of Operational Research*, 242(3):693–709, 2015. doi: 10.1016/j.ejor.2014.09.064.
- [39] Jun Gu. Global optimization for satisfiability (SAT) problem. *IEEE Transactions on Knowledge and Data Engineering*, 6(3):361–381, 1994. doi: 10.1109/69.334864.
- [40] Anastasios Kyrillidis, Anshumali Shrivastava, Moshe Y. Vardi, and Zhiwei Zhang. FourierSAT: A Fourier expansion-based algebraic framework for solving hybrid Boolean constraints. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020)*, volume 34, pages 1552–1560. AAAI Press, 2020. doi: 10.1609/aaai.v34i02.5515.
- [41] Yunuo Cen, Zixuan Wang, Jintao Zhang, Zhiwei Zhang, and Xuanyao Fong. FourierCSP: Differentiable constraint satisfaction problem solving by Walsh–Fourier expansion. arXiv:2510.04480, 2025.
- [42] Marcello Pelillo. Replicator equations, maximal cliques, and graph isomorphism. *Neural Computation*, 11(8):1933–1955, 1999. doi: 10.1162/089976699300016034.

- [43] Marcello Pelillo, Kaleem Siddiqi, and Steven W. Zucker. Matching hierarchical structures using association graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(11):1105–1120, 1999. doi: 10.1109/34.809105.
- [44] Marius Leordeanu and Martial Hebert. A spectral technique for correspondence problems using pairwise constraints. In *Proceedings of the Tenth IEEE International Conference on Computer Vision (ICCV 2005)*, volume 2, pages 1482–1489. IEEE, 2005. doi: 10.1109/ICCV.2005.20.
- [45] Ciaran McCreesh, Samba Ndojh Ndiaye, Patrick Prosser, and Christine Solnon. Clique and constraint models for maximum common (connected) subgraph problems. In *Principles and Practice of Constraint Programming (CP 2016)*, volume 9892 of *Lecture Notes in Computer Science*, pages 350–368. Springer, 2016. doi: 10.1007/978-3-319-44953-1_23.
- [46] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, San Francisco, 1979.
- [47] Massimiliano Pavan and Marcello Pelillo. Dominant sets and pairwise clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(1):167–172, 2007. doi: 10.1109/TPAMI.2007.250608.
- [48] Massimiliano Pavan and Marcello Pelillo. Generalizing the Motzkin–Straus theorem to edge-weighted graphs, with applications to image segmentation. In *Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR 2003)*, volume 2683 of *Lecture Notes in Computer Science*, pages 485–500. Springer, 2003. doi: 10.1007/978-3-540-45063-4_31.
- [49] Samuel Rota Bulò and Marcello Pelillo. A generalization of the Motzkin–Straus theorem to hypergraphs. *Optimization Letters*, 3(2):287–295, 2009. doi: 10.1007/s11590-008-0108-3.
- [50] Endre Boros and Peter L. Hammer. Pseudo-Boolean optimization. *Discrete Applied Mathematics*, 123(1–3): 155–225, 2002. doi: 10.1016/S0166-218X(01)00341-9.
- [51] Endre Boros, Yves Crama, and Elisabeth Rodríguez-Heck. Compact quadratizations for pseudo-Boolean functions. *Journal of Combinatorial Optimization*, 39(3):687–707, 2020. doi: 10.1007/s10878-019-00511-0.
- [52] J. A. Montañez-Barrera, Dennis Willsch, A. Maldonado-Romo, and Kristel Michielsen. Unbalanced penalization: A new approach to encode inequality constraints of combinatorial problems for quantum optimization algorithms. *Quantum Science and Technology*, 9(2):025022, 2024. doi: 10.1088/2058-9565/ad35e4.
- [53] Jerzy A. Filar, Michael Haythorpe, and Richard Taylor. Linearly-growing reductions of Karp’s 21 NP-complete problems. *Numerical Algebra, Control and Optimization*, 8(1):1–16, 2018. doi: 10.3934/naco.2018001.
- [54] Vivek S. Borkar, Vladimir Ejov, and Jerzy A. Filar. Directed graphs, Hamiltonicity and doubly stochastic matrices. *Random Structures & Algorithms*, 25(4):376–395, 2004. doi: 10.1002/rsa.20034.
- [55] Vicky Choi. Minor-embedding in adiabatic quantum computation: II. minor-universal graph design. *Quantum Information Processing*, 10(3):343–353, 2011. doi: 10.1007/s11128-010-0200-3.
- [56] Jesse Berwald, Nicholas Chancellor, and Raouf Dridi. Understanding domain-wall encoding theoretically and experimentally. *Philosophical Transactions of the Royal Society A*, 381(2241):20210410, 2023. doi: 10.1098/rsta.2021.0410.
- [57] Federico Domínguez, Josua Unger, Matthias Traube, Barry Mant, Christian Ertler, and Wolfgang Lechner. Encoding-independent optimization problem formulation for quantum computing. *Frontiers in Quantum Science and Technology*, 2:1229471, 2023. doi: 10.3389/frqst.2023.1229471.
- [58] Philippe Codognet. Comparing QUBO models for quantum annealing: Integer encodings for permutation problems. *International Transactions in Operational Research*, 32(1):18–37, 2025. doi: 10.1111/itor.13471.
- [59] Amit Verma and Mark Lewis. Penalty and partitioning techniques to improve performance of QUBO solvers. *Discrete Optimization*, 44:100594, 2022. doi: 10.1016/j.disopt.2020.100594.
- [60] Immanuel M. Bomze, Francesco Rinaldi, and Damiano Zeffiro. Active set complexity of the away-step Frank–Wolfe algorithm. *SIAM Journal on Optimization*, 30(3):2470–2500, 2020. doi: 10.1137/19M1309419.
- [61] Guglielmo Beretta, Alessandro Torcinovich, and Marcello Pelillo. On generalized KKT points for the motzkin–strauss program. *Journal of Global Optimization*, 91(3):535–557, 2025. doi: 10.1007/s10898-024-01457-2. Preprint arXiv:2305.08519.

- [62] Kirill P. Kalinin, Jannes Gladrow, Jiaqi Chu, James H. Clegg, Daniel Cletheroe, Douglas J. Kelly, Babak Rahmani, Grace Brennan, Burcu Canakci, Fabian Falck, Michael Hansen, Jim Kleewein, Heiner Kremer, Greg O'Shea, Lucinda Pickup, Saravan Rajmohan, Ant Rowstron, Victor Ruhle, Lee Braine, Shrirang Khedekar, Natalia G. Berloff, Christos Gkantsidis, Francesca Parmigiani, and Hitesh Ballani. Analog optical computer for AI inference and combinatorial optimization. *Nature*, 645(8080):354–361, 2025. doi: 10.1038/s41586-025-09430-z.
- [63] Immanuel M. Bomze, Marco Budinich, Marcello Pelillo, and Claudio Rossi. Annealed replication: A new heuristic for the maximum clique problem. *Discrete Applied Mathematics*, 121(1–3):27–49, 2002. doi: 10.1016/S0166-218X(01)00233-5.
- [64] Leonard E. Baum and John A. Eagon. An inequality with applications to statistical estimation for probabilistic functions of Markov processes and to a model for ecology. *Bulletin of the American Mathematical Society*, 73(3):360–363, 1967. doi: 10.1090/S0002-9904-1967-11751-8.
- [65] Mario S. Könz, Wolfgang Lechner, Helmut G. Katzgraber, and Matthias Troyer. Embedding overhead scaling of optimization problems in quantum annealing. *PRX Quantum*, 2(4):040322, 2021. doi: 10.1103/PRXQuantum.2.040322.
- [66] Elijah Pelofske. Comparing three generations of D-Wave quantum annealers for minor embedded combinatorial optimization problems. *Quantum Science and Technology*, 10(2):025025, 2025. doi: 10.1088/2058-9565/adb029. Preprint arXiv:2301.03009.
- [67] David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing*, 3(1):103–128, 2007. doi: 10.4086/toc.2007.v003a006. Preliminary version in STOC 2006, pp. 681–690, doi:10.1145/1132516.1132612.
- [68] Immanuel M. Bomze and Werner Schachinger. Multi-standard quadratic optimization: Interior point methods and cone programming reformulation. *Computational Optimization and Applications*, 45(2): 237–256, 2010. doi: 10.1007/s10589-009-9243-8.
- [69] Immanuel M. Bomze, Mirjam Dür, Etienne de Klerk, Cornelis Roos, Arie J. Quist, and Tamás Terlaky. On copositive programming and standard quadratic optimization problems. *Journal of Global Optimization*, 18(4):301–320, 2000. doi: 10.1023/A:1026583532263.
- [70] Samuel Burer. On the copositive representation of binary and continuous nonconvex quadratic programs. *Mathematical Programming*, 120(2):479–495, 2009. doi: 10.1007/s10107-008-0223-z.
- [71] Panos M. Pardalos and Jue Xue. The maximum clique problem. *Journal of Global Optimization*, 4(3): 301–328, 1994. doi: 10.1007/BF01098364.
- [72] Etienne de Klerk and Dmitrii V. Pasechnik. Approximation of the stability number of a graph via copositive programming. *SIAM Journal on Optimization*, 12(4):875–892, 2002. doi: 10.1137/S1052623401383248.
- [73] Jacek Gondzio and E. Alper Yıldırım. Global solutions of nonconvex standard quadratic programs via mixed integer linear programming reformulations. *Journal of Global Optimization*, 81(2):293–321, 2021. doi: 10.1007/s10898-021-01017-y.
- [74] Immanuel M. Bomze, Marco Locatelli, and Fabio Tardella. New and old bounds for standard quadratic optimization: Dominance, equivalence and incomparability. *Mathematical Programming*, 115(1):31–64, 2008. doi: 10.1007/s10107-007-0138-0.
- [75] Mirjam Dür and Franz Rendl. Conic optimization: A survey with special focus on copositive optimization and binary quadratic problems. *EURO Journal on Computational Optimization*, 9:100021, 2021. doi: 10.1016/j.ejco.2021.100021.