Treewidth and the complexity of box-constrained quadratic programs

Alberto Del Pia [*] Aida Khajavirad [†]

Abstract

We consider the problem of minimizing a sparse quadratic function over the unit hypercube. In binary quadratic programming, treewidth of the interaction graph is a central parameter for tractability: bounded treewidth yields polynomial-time solvability. Motivated by this fact, we investigate whether treewidth plays a similar role when the binary domain is replaced by the unit hypercube. We show that the situation is strikingly different. If the interaction graph is a forest, we give a strongly polynomial-time algorithm based on dynamic programming and a structural analysis of the resulting univariate value functions, which are shown to be concave and piecewise quadratic with linearly many pieces. However, the problem becomes strongly $\mathcal{NP}$-hard already when the interaction graph has treewidth two. We then identify substantially more general polynomial-time solvable classes whose interaction graphs may have unbounded treewidth. Our approach exploits the fact that there exists an optimal solution in which every variable with a nonpositive coefficient for its square term is binary-valued, thereby separating the problem into a combinatorial part and a genuinely continuous part. We show that tractability can be recovered by controlling the complexity of these two parts and their interaction, rather than the treewidth of the entire interaction graph.

Keywords: box-constrained quadratic programming, sparsity, treewidth, dynamic programming, polynomial-time algorithm.

1 Introduction

We consider a box-constrained quadratic program:

$$\begin{aligned} \min \quad & x^\top Q x + c^\top x \\ \text{s.t.} \quad & x \in [0,1]^n, \end{aligned} \tag{QP}$$

where $c \in \mathbb{Q}^n$ and $Q = (q_{ij}) \in \mathbb{Q}^{n \times n}$ is a symmetric matrix. It is well known that Problem (QP) is $\mathcal{NP}$-hard in general, as it contains the maximum cut problem as a special case [21]. In this paper, we are interested in obtaining sufficient conditions under which Problem (QP) can be solved in polynomial time. Throughout this paper, we consider the standard Turing bit model as our model of computation. If the number of variables is fixed, then the problem can be solved in strongly polynomial time using a simple face enumeration argument [40, 10]. If Q is positive semidefinite, then Problem (QP) is a convex optimization problem, and can be solved in polynomial time using the ellipsoid method [39]. If Q is positive semidefinite and tridiagonal, Pang and Han [31] give a

[*]Department of Industrial and Systems Engineering & Wisconsin Institute for Discovery, University of Wisconsin-Madison. E-mail: `delpia@wisc.edu`.

[†]Department of Industrial and Systems Engineering, Lehigh University. E-mail: `aida@lehigh.edu`.

strongly polynomial algorithm for Problem (QP). In [20], the authors prove that Problem (QP) can be solved in polynomial time if the rank of Q is fixed. If $q_{ii} \leq 0$ for some $i \in [n] := \{1, \ldots, n\}$, then it can be checked that there exists an optimal solution x^* of Problem (QP) at which $x_i^* \in \{0, 1\}$. It then follows that if $q_{ii} \leq 0$ for all $i \in [n]$, to find an optimal solution of Problem (QP), we can solve the *binary quadratic program*:

$$\begin{aligned} \min \quad & x^\top Q x + c^\top x \\ \text{s.t.} \quad & x \in \{0,1\}^n. \end{aligned} \tag{BQP}$$

Problem (BQP) has been extensively studied by the discrete optimization community [30, 8, 4]. If the objective function is submodular, i.e., $q_{ij} \leq 0$ for all $1 \leq i < j \leq n$, then Problem (BQP) can be solved in strongly polynomial time [36]. An important line of research exploits the sparsity of the objective function of Problem (BQP) to obtain sufficient conditions for its polynomial-time solvability. To this end, we define, for Problem (QP) or Problem (BQP), the *interaction graph* as the graph $G = (V, E)$ with vertex set $V := [n]$, containing one vertex i for each variable x_i, and with edge set E, where two distinct vertices i and j are adjacent if and only if $q_{ij} \neq 0$. The most general result in this literature states that if the graph G has bounded treewidth, then Problem (BQP) can be solved in strongly polynomial time [8]. In fact, the results of [5, 11] imply that bounded treewidth is essentially a necessary and sufficient condition for polynomial-time solvability of Problem (BQP). See [7, 28, 32, 9] for additional classes of polynomial-time solvable binary quadratic programs.

These results naturally raise the question of whether treewidth plays a similar role when the binary domain is replaced by the unit hypercube. Our goal is therefore to understand *the relationship between the treewidth of the interaction graph and the computational complexity of sparse box-constrained quadratic programs.* The work most closely related to ours is Khajavirad [22], who studies polynomial-time solvability of sparse polynomial optimization problems over the unit hypercube. Specializing to quadratic objectives, the paper observes that variables with nonpositive diagonal quadratic coefficients can be restricted to binary values at optimality, and exploits this structure to separate the problem into a discrete part and a collection of continuous components. The continuous components are eliminated individually under a condition controlling their degree of nonconvexity, yielding a structured binary quadratic optimization problem. Polynomial-time solvability is then obtained under suitable bounds on the treewidth of the resulting interaction graph. In this paper, we substantially extend this result. See [24, 37, 33, 12, 23] for sufficient conditions under which certain semidefinite programming relaxations of Problem (QP) are tight; namely, the optimal value of the relaxation coincides with that of the original problem.

Our main contributions can be summarized as follows.

(*i*) We prove that Problem (QP) can be solved in strongly polynomial time if its interaction graph has treewidth at most one, that is, if it is a forest. Our algorithm is based on dynamic programming over the forest and a structural analysis of the resulting univariate value functions, which are shown to be concave and piecewise quadratic with linearly many pieces. The resulting algorithm performs $O(n^2)$ arithmetic operations and comparisons (see Theorem 1).

(*ii*) We prove that the above tractability result is essentially sharp with respect to treewidth: Problem (QP) is strongly $\mathcal{NP}$-hard when the interaction graph has treewidth two. The hardness result holds even when all coefficients are integral and uniformly bounded, namely $\|Q\|_{\max} \leq 5$ and $\|c\|_\infty \leq 4$ (see Theorem 3).

(*iii*) We obtain more general sufficient conditions for polynomial-time solvability of Problem (QP) that apply to interaction graphs of potentially unbounded treewidth. We first exploit the fact that there exists an optimal solution in which every variable with a nonpositive diagonal quadratic coefficient is binary-valued, thereby separating the problem into a combinatorial part and a continuous part. We then derive tractability conditions based on the complexity of the connected components induced by these two parts and on the structure of their interaction (see Theorem 4 and Theorem 6).

Our results reveal a sharp transition in the role of treewidth for box-constrained quadratic programs: if the interaction graph is a forest, the problem is strongly polynomial-time solvable, while strong $\mathcal{NP}$-hardness already arises at treewidth two. This contrasts with binary quadratic programming, which is polynomial-time solvable on graphs of bounded treewidth. This difference can be attributed to the genuinely continuous part of the problem: variables with nonpositive square coefficients can be chosen binary at optimality, whereas variables with positive square coefficients cannot, in general, be reduced in this way. This simple observation separates the problem into a combinatorial part and a genuinely continuous part. Our results then show that tractability can be recovered by controlling the complexity of these two parts and their interaction, rather than the treewidth of the entire graph. Perhaps surprisingly, our results indicate that, at treewidth two, enlarging the feasible set from $\{0,1\}^n$ to $[0,1]^n$ while keeping the same quadratic objective function can turn a polynomial-time solvable problem into a strongly $\mathcal{NP}$-hard one.

It is important to note that all of the results in this paper extend directly to the mixed-binary setting; that is, the problem of minimizing a quadratic function over $x_i \in [0,1]$ for all $i \in I_1$ and $x_i \in \{0,1\}$ for all $i \in I_2$, where I_1 and I_2 partition $[n]$. Indeed, consider a binary variable x_i for some $i \in I_2$. Replace $q_{ii}x_i^2 + c_i x_i$ in the objective function by $(q_{ii} + c_i + 1)x_i - x_i^2$, and relax $x_i \in \{0,1\}$ to $x_i \in [0,1]$. Apply this transformation successively to all binary variables. It can be checked that the set of optimal solutions of the original mixed-binary problem coincides with that of the box-constrained quadratic program. The key to this transformation is that since we only modify diagonal and linear coefficients, the interaction graph remains unchanged.

Organization. The remainder of the paper is structured as follows. In Section 2 we present a strongly polynomial-time algorithm for box-constrained quadratic programs whose interaction graph is a forest. In Section 3, we show that if the treewidth of the interaction graph is two, then Problem (QP) is strongly $\mathcal{NP}$-hard, in general. In Section 4 we obtain sufficient conditions for polynomial-time solvability of box-constrained quadratic programs whose interaction graphs may have unbounded treewidth.

Notation. Throughout the paper, we make use of the following notation. For two subsets $S, S' \subseteq [n]$, we denote by $Q_{S,S'}$ the submatrix of Q with rows indexed by S and columns indexed by S', and we abbreviate $Q_S := Q_{S,S}$ for the principal submatrix of Q indexed by S. In particular, for $i \in [n]$, the vector $Q_{S,i}$ is the column of Q indexed by i, restricted to the rows in S. For $S \subseteq [n]$, we denote by c_S and x_S the restriction of c and x, respectively, to the coordinates in S. Given a graph $G = (V, E)$ and $S \subseteq V$, we denote by G_S the subgraph of G induced by S. We denote by $\langle \cdot \rangle$ the *lengths* (also known as encoding lengths or sizes) of rational numbers, vectors and matrices, as defined in [19, 35]; the *input length* of Problem (QP) is $\mathcal{L} := \langle Q \rangle + \langle c \rangle$. Throughout the paper, log denotes the base-2 logarithm.

2 Treewidth one: A strongly polynomial-time algorithm

In this section, we focus on the case where the interaction graph G of Problem (QP) has treewidth at most one, that is, G is a forest, and we present a strongly polynomial-time algorithm for it. Recall that an algorithm is *strongly polynomial* when the number of arithmetic operations and comparisons that it performs is bounded by a polynomial in the number of input numbers, independently of their values, and every intermediate number has length polynomial in the input length.

Theorem 1. *Assume that the interaction graph G is a forest. Then an optimal solution of Problem* (QP) *can be found by a strongly polynomial algorithm that performs $O(n^2)$ arithmetic operations and comparisons.*

Theorem 1 is proved by a leaf-to-root dynamic programming algorithm whose univariate value functions are concave and piecewise quadratic, with a number of arcs that adds, rather than multiplies, at vertices with several children. The algorithm performs only arithmetic operations $(+, -, \times, \div)$ and comparisons on rational numbers; throughout this section, an *operation* is one of these. The value functions are described by rational numbers and by algebraic numbers of degree at most two over $\mathbb{Q}$, which are in general irrational and are represented exactly as $\tau_0 + \tau_1\sqrt{D}$ with rational τ_0, τ_1, D; auxiliary numbers of degree at most four, represented by nested square roots, arise while a single value function is computed and are discarded afterwards. No root is ever extracted: comparing any two of these numbers takes $O(1)$ operations.

We denote by $f(x)$ the objective function of Problem (QP). Problem (QP) decomposes on the connected components $V_1, \ldots, V_k$ of G, and since G is a forest, each resulting subproblem is an instance of Problem (QP) whose interaction graph is a tree. As these components can be identified with $O(n^2)$ operations and $\sum_{i=1}^{k} |V_i|^2 \le n^2$, it suffices to prove Theorem 1 when G is a tree. Henceforth, we assume that G is a tree, and we write $T = (V, E) := G$.

The rest of this section is organized as follows. Section 2.1 sets up the dynamic program; Section 2.2 is a self-contained study of convex envelopes of piecewise quadratic functions with concave kinks; Section 2.3 proves that the value functions are piecewise quadratic with a linear number of arcs, whose coefficients are rationals of polynomial length; Section 2.4 shows how to compute them exactly and proves Theorem 1. Section 2.5 relates our approach to the literature, and in particular compares it with that of [1]. Finally, Section 2.6 shows that a similar tractability result does not hold for higher degree polynomials; namely, minimizing a quartic over the unit hypercube is strongly $\mathcal{NP}$-hard even when the interaction graph is a path.

2.1 Value functions and the dynamic program

Root T at an arbitrary vertex r; for $v \neq r$, write $p(v)$ for the *parent* of v, write $\mathrm{ch}(v) := \{u \in V : p(u) = v\}$ for the set of *children* of v, and write T_v for the vertex set of the subtree rooted at v. The *depth* of u is the number of edges on the path from u to r, and the *height* of v is the largest number of edges on a path from v to a vertex of T_v. Separate f along the tree:

$$f(x) = \sum_{v \in V} \left(q_{vv} x_v^2 + c_v x_v\right) + \sum_{v \neq r} 2q_{p(v)v}\, x_{p(v)} x_v.$$

We abbreviate $s_v := 2q_{p(v)v}$ for $v \neq r$, which is nonzero since $\{p(v), v\}$ is an edge of G, and set $s_r := 0$. For $v \in V$ let $e_v \in \mathbb{R}^{T_v}$ be the v-th unit vector, and define

$$F_v(y;t) := y^\top Q_{T_v}\, y + \left(c_{T_v} + s_v t\, e_v\right)^\top y, \qquad y \in \mathbb{R}^{T_v},\ t \in \mathbb{R}. \tag{1}$$

Then $F_v(\cdot\,; t)$ is the objective of the subproblem on the subtree rooted at v, with the parent variable set to t, and $F_r(y; 0) = f(y)$. Since the sets T_j, $j \in \mathrm{ch}(v)$, partition $T_v \setminus \{v\}$, and the edges of T between v and $T_v \setminus \{v\}$ are the edges $\{v, j\}$, $j \in \mathrm{ch}(v)$, we have

$$F_v(y;t) = \left(q_{vv} y_v^2 + c_v y_v\right) + \sum_{j \in \mathrm{ch}(v)} F_j\left(y_{T_j}; y_v\right) + s_v\, t\, y_v. \tag{2}$$

We define two families of univariate functions, by induction on the height of v: the *value function* μ_v, one for each $v \neq r$, and an auxiliary function φ_v, one for each $v \in V$, given by

$$\varphi_v(x) \;:=\; q_{vv}x^2 + c_v x + \sum_{j\in \mathrm{ch}(v)} \mu_j(x), \qquad \mu_v(t) \;:=\; \min_{x\in[0,1]} \Big[\varphi_v(x) + s_v\, t\, x\Big]. \tag{3}$$

The functions φ_v are defined on all of $\mathbb{R}$, but only their restriction to $[0,1]$ enters (3).

We will need some terminology for univariate functions. We call a continuous function $g : I \to \mathbb{R}$ on an interval $I \subseteq \mathbb{R}$ (possibly unbounded) *piecewise quadratic* (pw-quadratic) if I splits into finitely many nondegenerate intervals (*arcs*), each closed in I and possibly unbounded, on each of which g is a polynomial of degree ≤ 2; the shared arc endpoints interior to I are *breakpoints*. We always take the arcs to be maximal, so that no two consecutive ones carry the same polynomial; the breakpoints are then exactly the points of I at which g fails to be a polynomial of degree ≤ 2 on a neighborhood, and the decomposition into arcs is unique. A breakpoint x is a *kink* if $g'(x^-) \neq g'(x^+)$; a kink is *concave* if $g'(x^-) > g'(x^+)$. An arc is *strictly convex*, *affine*, or *strictly concave* according to the sign of its quadratic coefficient.

We first record a standard fact about concave functions; see Theorems 10.1 and 24.1 in [34], applied to $-g$.

Lemma 1. *A concave function $g : \mathbb{R} \to \mathbb{R}$ is continuous, has one-sided derivatives everywhere, and satisfies $g'(x^-) \geq g'(x^+)$ for every x.*

Lemma 2. *For every $v \in V$, the function φ_v of (3) is continuous on $\mathbb{R}$. For every $v \neq r$ and every $t \in \mathbb{R}$, the minimum in (3) is attained, and $\mu_v : \mathbb{R} \to \mathbb{R}$ is concave.*

Proof. Induction on the height of v. Let $v \in V$ and assume the lemma for all $j \in \mathrm{ch}(v)$, a vacuous assumption if v is a leaf. Each μ_j is then continuous by Lemma 1, hence so is φ_v, and the minimum in (3) is attained by compactness of $[0,1]$. Moreover, μ_v is the pointwise minimum of the affine functions $t \mapsto \varphi_v(x) + s_v t x$, $x \in [0,1]$, and it is finite, hence concave. □

The next lemma justifies the recursion (3).

Lemma 3. *For every $v \neq r$ and every $t \in \mathbb{R}$,*

$$\mu_v(t) \;=\; \min\big\{F_v(y;t) \;:\; y \in [0,1]^{T_v}\big\},$$

and $\min_{x\in[0,1]^n} f = \min_{x\in[0,1]} \varphi_r(x)$.

Proof. Induction on the height of v. By (2), minimizing $F_v(y;t)$ over $y \in [0,1]^{T_v}$ amounts to fixing $y_v = x \in [0,1]$ and minimizing each $F_j(\cdot\,;x)$ over $[0,1]^{T_j}$ separately, the blocks T_j, $j \in \mathrm{ch}(v)$, being disjoint; by the inductive hypothesis the j-th of these minima is $\mu_j(x)$, so that the whole minimum equals $\min_{x\in[0,1]}[\varphi_v(x) + s_v t x] = \mu_v(t)$ by (3). The root identity is this computation for $v = r$, where $s_r = 0$ and $F_r(y;0) = f(y)$. □

2.2 The convex envelope of a piecewise quadratic function with concave kinks

This subsection is self-contained univariate convex analysis, with no reference to the tree: the results below concern the class $\mathcal{Q}$ of pw-quadratic functions $\varphi : [0,1] \to \mathbb{R}$ all of whose kinks are concave, and the symbols φ and μ no longer refer to the specific φ_v and μ_v of (3).

We recall that the *convex envelope* $\mathrm{env}\,\varphi$ of a continuous function $\varphi : [0,1] \to \mathbb{R}$ is the greatest convex function $\leq \varphi$ on $[0,1]$, and we say that $x_0 \in [0,1]$ is a *contact point* of φ if $(\mathrm{env}\,\varphi)(x_0) =$

$\varphi(x_0)$; the set of all contact points is the *contact set* of φ. The *conjugate* of a continuous function $g : [0,1] \to \mathbb{R}$ is $g^*(\sigma) := \max_{x\in[0,1]} \big(\sigma x - g(x)\big)$, $\sigma \in \mathbb{R}$, which is the Fenchel conjugate of g extended by $+\infty$ outside $[0,1]$. These notions, and all the results of this subsection, hold verbatim on an arbitrary compact interval $[\lambda,\rho]$ with $\lambda < \rho$: the affine bijection $x = \lambda + (\rho-\lambda)\hat{x}$ onto $[0,1]$ preserves convexity, arcs, kinks and their concavity.

We first collect the basic properties of the convex envelope and describe its contact set with φ; we then use this description for the function obtained by minimizing $\varphi(x) + stx$ over $[0,1]$, which depends on φ only through $\operatorname{env}\varphi$ (Lemma 9). We conclude with two facts used in Section 2.4: where the minimum of a function of $\mathcal{Q}$ is attained (Lemma 10), and the replacement of strictly concave arcs by their chords (Lemma 11).

Lemma 4. *Let $\varphi : [0,1] \to \mathbb{R}$ be continuous and $H := \operatorname{env}\varphi$. Then H is continuous on $[0,1]$, $H(0) = \varphi(0)$, $H(1) = \varphi(1)$, and*

$$H(x) = \min\big\{\omega\varphi(x_1) + (1-\omega)\varphi(x_2) \;:\; \omega \in [0,1],\ x_1, x_2 \in [0,1],\ \omega x_1 + (1-\omega)x_2 = x\big\}. \qquad (4)$$

On every maximal open interval on which $H < \varphi$, H is affine, and $H = \varphi$ at the endpoints of that interval.

Proof. Extend φ by $+\infty$ outside $[0,1]$, so that H is the convex hull of φ in the sense of [34]. By Corollary 17.1.5 in [34], applied with $n = 1$, $H(x)$ is the infimum of $\omega\varphi(x_1) + (1-\omega)\varphi(x_2)$ over all $\omega \in [0,1]$ and $x_1, x_2 \in [0,1]$ with $\omega x_1 + (1-\omega)x_2 = x$, and the infimum is attained since φ is continuous on the compact interval $[0,1]$; this is (4). Since φ is real-valued and continuous on the compact set $[0,1]$, Corollary 17.2.1 in [34] shows that H is a closed proper convex function, and $[0,1]$ is a segment, hence a locally simplicial subset of the effective domain of H; by Theorem 10.2 in [34], H is continuous relative to $[0,1]$. In (4) with $x = 0$, every point carrying a positive weight equals 0, since 0 is an extreme point of $[0,1]$; hence $H(0) = \varphi(0)$, and similarly $H(1) = \varphi(1)$.

Now let (α,β) be a maximal open interval with $H < \varphi$, and $x_0 \in (\alpha,\beta)$. Let (ω, x_1, x_2) attain the minimum in (4) at $x = x_0$, labelled so that $x_1 \le x_2$; then $\omega \in (0,1)$ and $x_1 < x_0 < x_2$, since otherwise $H(x_0) = \varphi(x_0)$. Let ℓ be the affine function with $\ell(x_i) = \varphi(x_i)$, $i = 1,2$; then $H(x_0) = \ell(x_0)$. By (4), for every $x \in [x_1, x_2]$, $H(x) \le \ell(x)$ (take the combination of x_1, x_2 realizing x), while $\ell(x_0) = H(x_0) \le \omega H(x_1) + (1-\omega)H(x_2) \le \omega\varphi(x_1) + (1-\omega)\varphi(x_2) = \ell(x_0)$ forces $H(x_1) = \varphi(x_1)$, $H(x_2) = \varphi(x_2)$, and, by the equality case of convexity at the interior point x_0, $H = \ell$ on $[x_1, x_2]$. In particular $x_1 \le \alpha$ and $x_2 \ge \beta$, as x_1, x_2 are contact points, so H is affine on $[\alpha,\beta]$; and $H = \varphi$ at α and β by the maximality of (α,β) and the continuity of H and φ. ☐

By Lemma 4, for continuous φ the contact set is closed, being the zero set of the continuous function $\varphi - \operatorname{env}\varphi$, and its complement in $[0,1]$ is an at most countable union of open intervals on each of which $\operatorname{env}\varphi$ is affine; we call the closures of these intervals *bridges*. If moreover φ is pw-quadratic, we call *contact piece* any nonempty intersection of a maximal interval of the contact set of φ with an arc of φ; a contact piece is *degenerate* if it is a single point.

Lemma 5. *Let $\lambda < \xi < \rho$, let $g : [\lambda,\rho] \to \mathbb{R}$ be continuous, and set $H_1 := \operatorname{env}\big(g|_{[\lambda,\xi]}\big)$, $H_2 := \operatorname{env}\big(g|_{[\xi,\rho]}\big)$ and $H := \operatorname{env}(g)$. Then there exist $\eta \in [\lambda,\xi]$ and $\zeta \in [\xi,\rho]$ such that $H = H_1$ on $[\lambda,\eta]$, H is affine on $[\eta,\zeta]$, and $H = H_2$ on $[\zeta,\rho]$. Moreover, let $w \in (\lambda,\xi]$ be such that $H_1'(w^-)$ is finite, and let ℓ_0 be the line through $(w, H_1(w))$ with slope $H_1'(w^-)$. Then $H = H_1$ on $[\lambda,w]$ if and only if $\ell_0 \le g$ on $[\xi,\rho]$.*

Proof. By Lemma 4 we have $H_1(\xi) = g(\xi) = H_2(\xi)$, so the function $\tilde{H}$ equal to H_1 on $[\lambda,\xi]$ and to H_2 on $[\xi,\rho]$ is well defined and continuous. We claim $\operatorname{env}\tilde{H} = H$. Indeed $\tilde{H} \le g$, so $\operatorname{env}\tilde{H} \le H$;

conversely H is convex with $H \le g$, so its restriction to $[\lambda, \xi]$ is a convex function $\le g|_{[\lambda,\xi]}$, and therefore $H \le H_1$ there, and similarly $H \le H_2$ on $[\xi, \rho]$, that is $H \le \tilde{H}$, so $H \le \operatorname{env} \tilde{H}$.

Apply now Lemma 4 to $\tilde{H}$: on every maximal open interval on which $H < \tilde{H}$, the function H is affine and agrees with $\tilde{H}$ at the endpoints. Suppose such a component (α_1, β_1) were contained in $[\lambda, \xi]$. Then H is the chord of H_1 over $[\alpha_1, \beta_1]$, and since H_1 is convex that chord is at least $H_1 = \tilde{H}$ on $[\alpha_1, \beta_1]$, contradicting $H < \tilde{H}$ there. The same argument excludes components contained in $[\xi, \rho]$. Hence every such component contains ξ, so there is at most one; taking $[\eta, \zeta]$ to be its closure, or $\eta = \zeta = \xi$ if there is none, proves the first part.

For the second part, suppose first that $\ell_0 \le g$ on $[\xi, \rho]$. By convexity, $\ell_0 \le H_1$ on $[\lambda, \xi]$, and $\ell_0 \le H_2$ on $[\xi, \rho]$ since ℓ_0 is affine and $\ell_0 \le g$ there; so $\ell_0 \le \tilde{H}$ and hence $\ell_0 \le H$; thus $H(w) = H_1(w)$. If $w > \eta$, then H is affine on $[\eta, w]$ and agrees with the convex function $H_1 \ge H$ at η and w, so $H = H_1$ on $[\eta, w]$; in either case $H = H_1$ on $[\lambda, w]$. Conversely, suppose that $H = H_1$ on $[\lambda, w]$, and let ℓ be a supporting line of H at w, which exists as $w < \rho$. Its slope is at least $H'(w^-) = H_1'(w^-)$, and $\ell(w) = \ell_0(w)$, so $\ell_0 \le \ell \le H \le g$ on $[w, \rho] \supseteq [\xi, \rho]$. □

Lemma 6. *Let $\varphi \in \mathcal{Q}$, let $H := \operatorname{env} \varphi$, and let $x_0 \in (0,1)$ be a contact point of φ. Then*

$$\varphi'(x_0^-) \ \le \ m \ \le \ \varphi'(x_0^+) \qquad \text{for every } m \in \partial H(x_0). \tag{5}$$

In particular, φ and H are differentiable at x_0, with $H'(x_0) = \varphi'(x_0)$.

Proof. The function H is convex and finite on $[0, 1]$ and $x_0 \in (0, 1)$, so $\partial H(x_0) \ne \emptyset$ by Theorem 23.4 in [34]. Let $m \in \partial H(x_0)$ and let $\ell(x) := H(x_0) + m(x - x_0)$ be the corresponding supporting line, so that $\ell \le H \le \varphi$, and $\ell(x_0) = H(x_0) = \varphi(x_0)$ because x_0 is a contact point. Hence $\varphi(x) - \varphi(x_0) \ge m(x - x_0)$ for every $x \in [0, 1]$; dividing by $x - x_0$ and letting $x \to x_0$ on either side yields (5), the one-sided derivatives existing because φ is pw-quadratic.

Since $\partial H(x_0) \ne \emptyset$, (5) forces $\varphi'(x_0^-) \le \varphi'(x_0^+)$. A kink of a function of $\mathcal{Q}$ is concave, that is, it satisfies $\varphi'(x_0^-) > \varphi'(x_0^+)$; so x_0 is not a kink, and φ is differentiable at x_0. The two bounds in (5) then coincide, so $\partial H(x_0)$ is the singleton $\{\varphi'(x_0)\}$; by Theorem 25.1 in [34], H is differentiable at x_0 with $H'(x_0) = \varphi'(x_0)$. □

Lemma 7. *Let $\varphi \in \mathcal{Q}$. Then no contact point of φ lies in the interior of a strictly concave arc.*

Proof. Write $H := \operatorname{env} \varphi$ and suppose, for a contradiction, that some contact point x_0 lies in the interior of an arc on which $\varphi(x) = ax^2 + bx + d$ with $a < 0$; let $N \subseteq (0, 1)$ be a neighborhood of x_0 contained in that arc. By Lemma 6, $m := \varphi'(x_0)$ is a subgradient of H at x_0, so $H(x) \ge \varphi(x_0) + m(x - x_0)$ for all $x \in [0, 1]$, as $H(x_0) = \varphi(x_0)$. For $x \in N \setminus \{x_0\}$, since φ is quadratic on N, $\varphi(x) - H(x) \le \varphi(x) - \varphi(x_0) - m(x - x_0) = a(x - x_0)^2 < 0$, contradicting $H \le \varphi$. □

Lemma 8. *Let $\varphi \in \mathcal{Q}$, $H := \operatorname{env} \varphi$, and let A be an arc of φ with nonnegative quadratic coefficient. If $x_1 < x_2$ are contact points of φ in A, then $H = \varphi$ on $[x_1, x_2]$. Consequently, the contact set of φ meets every arc with nonnegative quadratic coefficient in a single (possibly empty or degenerate) closed interval.*

Proof. Suppose, for a contradiction, that $H(x_0) < \varphi(x_0)$ for some $x_0 \in (x_1, x_2)$, and let (α, β) be the maximal open interval containing x_0 on which $H < \varphi$. Since x_1 and x_2 are contact points, neither belongs to (α, β), so that

$$x_1 \le \alpha < x_0 < \beta \le x_2,$$

and in particular $[\alpha, \beta] \subseteq [x_1, x_2] \subseteq A$. By Lemma 4, H is affine on $[\alpha, \beta]$ and $H = \varphi$ at α and at β; that is, H coincides on $[\alpha, \beta]$ with the chord of φ over that interval. But φ is convex on A, so its chord over $[\alpha,\beta]$ is at least φ there. Hence $H \geq \varphi$ on (α,β), contradicting $H < \varphi$ there. The final sentence follows, since the contact set is closed and, by the first assertion, its intersection with an arc of nonnegative quadratic coefficient is convex. $\square$

Lemma 9. *Let $\varphi \in \mathcal{Q}$ have K arcs, let $H := \operatorname{env}\varphi$, let $s \in \mathbb{R}\setminus\{0\}$, and define*

$$\mu(t) := \min_{x\in[0,1]} \bigl[\varphi(x) + stx\bigr], \qquad t \in \mathbb{R}.$$

Then μ is a concave pw-quadratic function on $\mathbb{R}$ with at most $K+2$ arcs. Specifically, $\mu(t) = \varphi(0)$ for $-st \leq H'(0^+)$, $\mu(t) = \varphi(1) + st$ for $-st \geq H'(1^-)$, and, for each nondegenerate contact piece $[w_1, w_2]$ inside a strictly convex arc on which $\varphi(x) = ax^2 + bx + d$,

$$\mu(t) = -\frac{(st+b)^2}{4a} + d \qquad \text{for all } t \text{ with } -st \in [\,2aw_1 + b,\ 2aw_2 + b\,]. \tag{6}$$

These ranges of t cover $\mathbb{R}$ except for finitely many points, at most one for each maximal nondegenerate interval on which H is affine.

Proof. We first observe that

$$\mu(t) = -\max_{x\in[0,1]} \bigl[(-st)x - \varphi(x)\bigr] = -\varphi^*(-st) = -H^*(-st).$$

The last equality holds because an affine function lies below φ if and only if it lies below H: extending φ by $+\infty$ outside $[0,1]$, its closed convex hull is H, and Corollary 12.1.1 in [34] applies. Hence $\varphi^* = H^*$, and it suffices to describe H^*.

By Lemma 4 and Lemmas 6–8, H is convex and pw-quadratic: it coincides with φ on the contact set (which meets each arc with nonnegative quadratic coefficient in at most one interval and avoids the interiors of strictly concave arcs and all kinks) and is affine on the bridges in between. By Lemma 6, every contact point in $(0,1)$ is a point at which φ is differentiable and H is differentiable with $H' = \varphi'$; on bridges H is affine; hence H is differentiable on $(0,1)$ and H' is continuous and nondecreasing, strictly increasing on contact pieces inside strictly convex arcs (where $H'(x) = 2ax + b$) and constant on intervals on which H is affine.

Fix $\sigma \in \bigl(H'(0^+), H'(1^-)\bigr)$. Since H' is continuous and nondecreasing on $(0,1)$, $\sigma = H'(x)$ for some $x \in (0,1)$, and x lies in a nondegenerate contact piece inside a strictly convex arc or in a maximal nondegenerate interval on which H is affine, as H is affine on bridges and on contact pieces inside affine arcs, and an isolated contact point in $(0,1)$ lies between two bridges, which have the same slope by Lemma 6. The maximum of the concave function $x \mapsto \sigma x - H(x)$ is attained exactly on $\{x : \sigma \in \partial H(x)\} = \{x : H'(x) = \sigma\}$. If σ lies in the interior of the slope range $[2aw_1 + b,\, 2aw_2 + b]$ of a contact piece inside a strictly convex arc, the maximizer $x = (\sigma - b)/2a$ is unique, and since $H = \varphi$ there, $H^*(\sigma) = \sigma x - (ax^2 + bx + d) = \frac{(\sigma-b)^2}{4a} - d$: one quadratic arc of H^* per such contact piece. If σ is the slope of a maximal nondegenerate interval on which H is affine, the maximizer set is that interval, which equals $\partial H^*(\sigma)$ by Corollary 23.5.1 in [34]; thus σ contributes no arc of H^*, and is a kink of H^*. Distinct such intervals and distinct contact pieces occupy slope ranges with disjoint interiors, consecutively ordered, since H' is nondecreasing; so the arcs of H^* on $\bigl(H'(0^+), H'(1^-)\bigr)$ are exactly one per nondegenerate strictly convex contact piece. For $\sigma \leq H'(0^+)$, $x = 0$ is a maximizer and $H^*(\sigma) = -H(0) = -\varphi(0)$; for $\sigma \geq H'(1^-)$, $x = 1$ is a maximizer and $H^*(\sigma) = \sigma - H(1) = \sigma - \varphi(1)$ (using Lemma 4 for the endpoint values): the two unbounded affine arcs.

Thus H^* is convex pw-quadratic, with at most two arcs more than the number of strictly convex arcs of φ, each of which contains at most one contact piece by Lemma 8, hence with at most $K+2$

arcs. Substituting $\sigma = -st$ and negating yields the statement, including (6); concavity of μ is immediate since $t \mapsto H^*(-st)$ is convex. $\square$

Lemma 10. *Let $\varphi \in \mathcal{Q}$ have arcs $[\xi_{j-1}, \xi_j]$, $j \in [K]$, where $0 = \xi_0 < \xi_1 < \cdots < \xi_K = 1$, with $\varphi(x) = a_j x^2 + b_j x + d_j$ on the j-th arc. Then φ attains its minimum over $[0,1]$ at 0, at 1, or at $-b_j/(2a_j)$ for some j with $a_j > 0$ and $-b_j/(2a_j) \in [\xi_{j-1}, \xi_j]$.*

Proof. Let $\mathcal{M}$ be the set of minimizers of φ over $[0,1]$, which is nonempty and compact by continuity, and let $x^\star := \max \mathcal{M}$. If $x^\star \in \{0,1\}$ we are done, so assume $x^\star \in (0,1)$. Then $x^\star$ is not a kink of φ: all kinks are concave, that is $\varphi'(x^{\star -}) > \varphi'(x^{\star +})$, whereas minimality in the interior gives $\varphi'(x^{\star -}) \le 0 \le \varphi'(x^{\star +})$. Hence φ is differentiable at $x^\star$ and $\varphi'(x^\star) = 0$. Let $[\xi_{j-1}, \xi_j]$ be the arc immediately to the right of $x^\star$, that is, the arc with $\xi_{j-1} = x^\star$ if $x^\star$ is a breakpoint and the arc containing $x^\star$ in its interior otherwise; thus $2a_j x^\star + b_j = 0$. If $a_j < 0$ then φ is strictly decreasing immediately to the right of $x^\star$, contradicting minimality; if $a_j = 0$ then $b_j = 0$ and φ is constant on that arc, so every point of $[x^\star, \xi_j]$ is a minimizer, contradicting the maximality of $x^\star$. Hence $a_j > 0$ and $x^\star = -b_j/(2a_j) \in [\xi_{j-1}, \xi_j]$. $\square$

The last result of this subsection allows us to dispense with strictly concave arcs when computing convex envelopes.

Lemma 11. *Let $\varphi \in \mathcal{Q}$ have arcs $[\xi_{j-1}, \xi_j]$, $j \in [K]$, where $0 = \xi_0 < \xi_1 < \cdots < \xi_K = 1$, and let $\hat\varphi$ be obtained from φ by replacing every strictly concave arc by its chord. Then $\hat\varphi \in \mathcal{Q}$, $\hat\varphi$ is strictly convex or affine on each $[\xi_{j-1}, \xi_j]$, $\hat\varphi(\xi_j) = \varphi(\xi_j)$ for every j, and $\operatorname{env}(\hat\varphi|_{[0,\xi_j]}) = \operatorname{env}(\varphi|_{[0,\xi_j]})$ for every $j \in [K]$.*

Proof. Clearly $\hat\varphi$ is continuous, agrees with φ at every ξ_j, and is strictly convex or affine on each $[\xi_{j-1}, \xi_j]$; its breakpoints are among the ξ_j. On a strictly concave arc $[\alpha, \beta]$, the slope γ of the chord satisfies $\varphi'(\alpha^+) > \gamma > \varphi'(\beta^-)$. Hence at every breakpoint ξ of φ the left derivative of $\hat\varphi$ is at least $\varphi'(\xi^-)$ and the right derivative is at most $\varphi'(\xi^+)$; as $\varphi'(\xi^-) \ge \varphi'(\xi^+)$, every kink of $\hat\varphi$ is concave, and $\hat\varphi \in \mathcal{Q}$. Finally fix j. Since $\hat\varphi \le \varphi$, we have $\operatorname{env}(\hat\varphi|_{[0,\xi_j]}) \le \operatorname{env}(\varphi|_{[0,\xi_j]})$. Conversely, $\operatorname{env}(\varphi|_{[0,\xi_j]})$ is convex and at most $\varphi = \hat\varphi$ at the endpoints of every arc, hence at most every chord, and so at most $\hat\varphi$ on $[0, \xi_j]$; therefore $\operatorname{env}(\varphi|_{[0,\xi_j]}) \le \operatorname{env}(\hat\varphi|_{[0,\xi_j]})$. $\square$

2.3 Structure of the value functions

We now apply the results of Section 2.2 to the value functions of (3). We write K_v for the number of arcs of $\varphi_v|_{[0,1]}$.

Proposition 1. *For every $v \in V$, the restriction of φ_v to $[0,1]$ is pw-quadratic with $K_v \le 2|T_v| - 1$ arcs and all its kinks are concave; that is, $\varphi_v|_{[0,1]} \in \mathcal{Q}$. Moreover, for every $v \ne r$, the value function μ_v is a concave pw-quadratic function on $\mathbb{R}$ with at most $2|T_v| + 1$ arcs.*

Proof. Induction on the height of v. Let $v \in V$ and assume the proposition for all $j \in \operatorname{ch}(v)$, a vacuous assumption if v is a leaf. Each μ_j is then pw-quadratic on $\mathbb{R}$ with at most $2|T_j|$ breakpoints, so $\varphi_v|_{[0,1]}$ is pw-quadratic and its breakpoints are among the breakpoints of the μ_j lying in $(0,1)$, of which there are at most $\sum_{j \in \operatorname{ch}(v)} 2|T_j| = 2(|T_v| - 1)$; hence $K_v \le 2|T_v| - 1$. At every breakpoint, the jump of φ_v' is the sum of the jumps of the μ_j', each nonpositive by Lemmas 2 and 1; so every kink of φ_v is concave, and $\varphi_v|_{[0,1]} \in \mathcal{Q}$.

Let now $v \ne r$. By (3), μ_v is the function μ of Lemma 9 for $(\varphi, s) = (\varphi_v|_{[0,1]}, s_v)$, where $s_v \ne 0$. Hence μ_v is a concave pw-quadratic function with at most $K_v + 2 \le 2|T_v| + 1$ arcs, which completes the induction. $\square$

Remark 1. The bounds of Proposition 1 are attained, even when $Q \succ 0$. For $n \geq 2$ and $w_i := 4^{i-1}$, let the objective of Problem (QP) be $\sum_{i=1}^{n-1} w_i(x_i - 3x_{i+1} + 1)^2 + w_n(x_n - \frac{1}{2})^2$ minus its constant term, so that the interaction graph is the path $1, 2, \ldots, n$, which we root at n. It can be checked that every φ_i is strongly convex and every breakpoint of μ_i lies in $(0, 1)$, and hence that $K_i = 2i-1$ for every $i \in [n]$ and μ_i has $2i + 1$ arcs for every $i < n$. In particular, $\sum_{v \in V} K_v = n^2$.

We next show that all arcs produced by the recursion (3) have rational coefficients.

Lemma 12. *Every arc of every μ_v, $v \neq r$, and of every $\varphi_v|_{[0,1]}$ has rational coefficients. Every breakpoint τ of μ_v is a root of $\pi_1 - \pi_2$, where $\pi_1 \neq \pi_2$ are the polynomials of the two arcs of μ_v meeting at τ; this is a nonzero polynomial of degree at most two with rational coefficients, which we call the* defining quadratic *of τ. Every breakpoint of $\varphi_v|_{[0,1]}$ is a breakpoint of some μ_j, $j \in \mathrm{ch}(v)$, and we take its defining quadratic to be that of this breakpoint of μ_j.*

Proof. By Proposition 1 these functions are pw-quadratic. We argue by induction on the height of v. Each arc of $\varphi_v|_{[0,1]}$ coincides on a nondegenerate subinterval, and hence as a polynomial, with the sum of $q_{vv}x^2 + c_v x$ and of one arc of each μ_j, $j \in \mathrm{ch}(v)$, so it has rational coefficients by induction; and the breakpoints of $\varphi_v|_{[0,1]}$ are breakpoints of the μ_j. By Lemma 9, each arc of μ_v is $\varphi_v(0)$, or $\varphi_v(1) + s_v t$, or $-\frac{(s_v t+b)^2}{4a} + d$ with (a, b, d) the coefficients of an arc of φ_v; since $\varphi_v(0)$ and $\varphi_v(1)$ are values of arcs of φ_v at 0 and 1, and $s_v \in \mathbb{Q}$, it has rational coefficients. Let now τ be a breakpoint of μ_v and let π_1, π_2 be the polynomials of the two arcs of μ_v meeting at τ. They have rational coefficients by what we just proved, and $\pi_1 \neq \pi_2$ because arcs are maximal. By continuity of μ_v, $\pi_1(\tau) = \pi_2(\tau)$, so τ is a root of the nonzero polynomial $\pi_1 - \pi_2$, which has rational coefficients and degree at most two. $\square$

We represent a real number τ of degree at most two over $\mathbb{Q}$ by rationals τ_0, τ_1, D with $D \geq 0$ and $\tau = \tau_0 + \tau_1\sqrt{D}$, and call (τ_0, τ_1, D) a *representation* of τ. If τ is a root of a nonzero polynomial π of degree at most two with rational coefficients, a representation of τ is obtained from the coefficients of π by the quadratic formula with $O(1)$ operations, once it is known which root of π is τ; this fixes the sign of τ_1.

We next record an observation that holds for every box-constrained quadratic program, not only on trees; in it, the empty matrix is regarded as positive definite. It is implicit in the proof of NP membership of quadratic programming by Vavasis [40].

Lemma 13. *Let $Q \in \mathbb{R}^{n \times n}$ be symmetric and $c \in \mathbb{R}^n$. Then $\min\{x^\top Qx + c^\top x : x \in [0, 1]^n\}$ has an optimal solution $x^\star$ such that $Q_S \succ 0$, where $S := \{i \in [n] : 0 < x_i^\star < 1\}$.*

Proof. Write $f(x) := x^\top Qx + c^\top x$. Among the optimal solutions, which exist by compactness, let $x^\star$ have the largest number of coordinates in $\{0, 1\}$, and let S be as in the statement. For every $d \in \mathbb{R}^n$ with $d_{[n]\setminus S} = 0$, the point $x^\star + td$ is feasible for all t in a neighborhood of 0, and $f(x^\star + td) = f(x^\star) + t\,\nabla f(x^\star)^\top d + t^2\, d_S^\top Q_S d_S$; optimality of $x^\star$ forces $\nabla_S f(x^\star) := (\nabla f(x^\star))_S = 0$ and $d_S^\top Q_S d_S \geq 0$, so $Q_S \succeq 0$. If Q_S were singular, pick $0 \neq d_S \in \ker Q_S$, extended by zero outside S. Then $f(x^\star + td) = f(x^\star)$ for every t. Let $\bar{t}$ be the largest $t \geq 0$ with $x^\star + td \in [0, 1]^n$, which is finite as $d \neq 0$; then $x^\star + \bar{t}d$ is an optimal solution with more coordinates in $\{0, 1\}$ than $x^\star$, a contradiction. Hence $Q_S \succ 0$. $\square$

We now give a closed-form description of the arcs of the value functions, which bypasses the recursion, along which lengths could a priori double at every level. In it, S collects the coordinates that are left free, while the coordinates outside S are fixed at the bounds prescribed by β.

Lemma 14. *For every $v \neq r$, every arc of μ_v coincides, as a polynomial in t, with the function*

$$t \longmapsto \min\left\{F_v(y;t) \;:\; y \in \mathbb{R}^{T_v},\ y_{T_v\setminus S} = \beta\right\} \tag{7}$$

for some $S \subseteq T_v$ with $Q_S \succ 0$ and some $\beta \in \{0,1\}^{T_v\setminus S}$.

Proof. For $S \subseteq T_v$ with $Q_S \succ 0$ and $\beta \in \{0,1\}^{T_v\setminus S}$, denote by $\bar{F}_{S,\beta}$ the function (7). The restriction of $F_v(\cdot\,;t)$ to $\{y : y_{T_v\setminus S} = \beta\}$ is a strictly convex quadratic in y_S whose coefficients are affine in t, so $\bar{F}_{S,\beta}$ is a polynomial in t of degree at most two.

Now fix $t \in \mathbb{R}$. By Lemma 13, applied to the matrix Q_{T_v} and the vector $c_{T_v} + s_v t e_v$, the problem $\min\{F_v(y;t) : y \in [0,1]^{T_v}\}$ has a minimizer $y^\star$ with $Q_S \succ 0$ for $S := \{u \in T_v : 0 < y^\star_u < 1\}$; put $\beta := y^\star_{T_v\setminus S} \in \{0,1\}^{T_v\setminus S}$. Since $y^\star$ is a minimizer and $y^\star_S$ lies in the interior of $[0,1]^S$, we have $\nabla_S F_v(y^\star;t) = 0$; as $Q_S \succ 0$, the point $y^\star$ minimizes $F_v(\cdot\,;t)$ over $\{y \in \mathbb{R}^{T_v} : y_{T_v\setminus S} = \beta\}$, so $\mu_v(t) = F_v(y^\star;t) = \bar{F}_{S,\beta}(t)$ by Lemma 3.

Finally, let π be an arc of μ_v, a polynomial of degree at most two that coincides with μ_v on a nondegenerate interval I. As there are finitely many pairs (S,β) and infinitely many $t \in I$, some pair satisfies $\pi(t) = \bar{F}_{S,\beta}(t)$ for infinitely many t, and then $\pi = \bar{F}_{S,\beta}$ as polynomials. □

We can now bound the lengths of all the numbers that describe the value functions.

Lemma 15. *All arc coefficients of all μ_v and $\varphi_v|_{[0,1]}$, and the coefficients of the defining quadratics and the representations of their breakpoints, are rationals of length polynomial in $\mathcal{L}$.*

Proof. All the quantities in question are rational by Lemma 12; it remains to bound their lengths. We use that sums and products of polynomially many rationals of length polynomial in $\mathcal{L}$ have length polynomial in $\mathcal{L}$.

Consider first an arc of a value function μ_v with $v \neq r$, and let S and β be as furnished by Lemma 14. Fixing $y_{T_v\setminus S} = \beta$ in (1) and separating the terms according to whether they involve the free block or not, we obtain

$$F_v(y;t) = y_S^\top Q_S\, y_S + \big(\theta + t\psi\big)^\top y_S + \big(\gamma_0 + \gamma_1 t\big) \qquad \text{for every } y \in \mathbb{R}^{T_v} \text{ with } y_{T_v\setminus S} = \beta,$$

where, for $u \in S$,

$$\theta_u = c_u + \sum_{\substack{w \in T_v\setminus S\\ \{u,w\}\in E}} 2q_{uw}\,\beta_w, \qquad \psi = \begin{cases} s_v\,(e_v)_S & v \in S,\\ 0 & v \notin S,\end{cases}$$

and γ_0, γ_1 are obtained from the data and from β in the same manner. Each β_w lies in $\{0,1\}$, so each entry of θ, ψ, γ_0 and γ_1 is, up to a factor two, a sum of at most $2n^2$ entries of Q and c, and therefore has length polynomial in $\mathcal{L}$.

Since $Q_S \succ 0$, the minimum in (7) is attained at $y_S = -\frac{1}{2}Q_S^{-1}(\theta + t\psi)$, and the arc formula is

$$\min\left\{F_v(y;t) \;:\; y \in \mathbb{R}^{T_v},\ y_{T_v\setminus S} = \beta\right\} = \gamma_0 + \gamma_1 t - \tfrac{1}{4}\big(\theta + t\psi\big)^\top Q_S^{-1}\big(\theta + t\psi\big). \tag{8}$$

By Corollary 3.2a in [35], the entries of Q_S^{-1} have length polynomial in the length of Q_S, a submatrix of Q, and hence polynomial in $\mathcal{L}$. Expanding (8), each coefficient of the arc is a sum of at most $n^2 + 1$ products of at most three such rationals, and therefore has length polynomial in $\mathcal{L}$. The same bound holds for the arcs of φ_v, since by (3) each of them is a sum of the quadratic $q_{vv}x^2 + c_v x$ and of at most n arc formulas of value functions. Finally, by Lemma 12 the defining quadratic of a breakpoint is the difference of two arc polynomials, so its coefficients have length polynomial in $\mathcal{L}$ as well, and so do the rationals of a representation, which are obtained from them by $O(1)$ operations. □

2.4 Exact computation of the value functions

Computing the value functions requires comparing their breakpoints, and also auxiliary numbers, obtained from them by one further square root, that arise while a single value function is computed (Lemma 17). The next lemma covers both; it is folklore, and we include a proof for completeness.

Lemma 16. *Let k and d be fixed positive integers. Let $E_1 \in \mathbb{Q}$ and, for $i = 2, \ldots, k$, let E_i be a polynomial of total degree at most d with rational coefficients in $\sqrt{E_1}, \ldots, \sqrt{E_{i-1}}$, where every E_i is a nonnegative real number. Then the sign of a polynomial of total degree at most d with rational coefficients in $\sqrt{E_1}, \ldots, \sqrt{E_k}$ can be determined with $O(1)$ arithmetic operations and comparisons on rational numbers, starting from E_1 and the coefficients of the given polynomials.*

Proof. Using $\sqrt{E_i}^{\,2} = E_i$, every such polynomial α can be written as $\alpha = P + R\sqrt{E_k}$, where P and R are polynomials with rational coefficients in $\sqrt{E_1}, \ldots, \sqrt{E_{k-1}}$; recursively, α is represented by 2^k rational numbers, and sums and products of such numbers cost $O(1)$ operations. As all the given polynomials have total degree at most d, these representations are obtained with $O(1)$ operations. We argue by induction on k, the case $k = 0$ being a comparison of a rational number with zero. Let $k \geq 1$ and determine the signs of P and R by induction. If $R = 0$, or if P and R are both positive or both negative, the sign of α is that of P. If $P = 0 \neq R$, the sign of α is the product of the signs of R and E_k, and the sign of E_k is determined by induction. Otherwise P and R have opposite signs, and

$$\operatorname{sign}\left(P + R\sqrt{E_k}\right) = \operatorname{sign}(P) \cdot \operatorname{sign}\left(P^2 - R^2 E_k\right),$$

since for $P > 0 > R$ we have $P + R\sqrt{E_k} > 0$ if and only if $P > -R\sqrt{E_k}$, that is $P^2 > R^2 E_k$, and symmetrically for $P < 0 < R$. As E_k is a polynomial in $\sqrt{E_1}, \ldots, \sqrt{E_{k-1}}$, so is $P^2 - R^2 E_k$, and its sign is determined by induction. Altogether the number of operations is bounded by a constant depending only on k and d. □

We compute a value function from the corresponding function φ in two steps: first the convex envelope of φ, then the value function itself. For $\varphi \in \mathcal{Q}$ with arcs $[\xi_{j-1}, \xi_j]$, let $\hat{\varphi}$ be as in Lemma 11, so that $\operatorname{env} \varphi = \operatorname{env} \hat{\varphi}$. By Lemmas 4 and 6–8, applied to $\hat{\varphi}$, this envelope is determined by the contact pieces of $\hat{\varphi}$, consecutive ones being either adjacent or joined by a bridge. We call *envelope list* of φ the list of these contact pieces in increasing order, each given by the index j of the arc containing it and by its two endpoints; the *carrier* of such a contact piece is the polynomial of degree at most two that coincides with $\hat{\varphi}$ on the arc j, that is, the polynomial of the arc j of φ, or its chord if that arc is strictly concave. The sweep in the proof below follows the scheme of the linear-time convex envelope algorithm of Gardiner and Lucet [17] for univariate piecewise linear-quadratic functions, which is designed for real arithmetic. What the next lemma adds is that, for functions of $\mathcal{Q}$ with rational data, the sweep can be organized so that every endpoint it computes is obtained from the data with at most two nested square roots; this is what allows it to be carried out exactly, with $O(1)$ operations on rational numbers per step.

Lemma 17. *Let $\varphi \in \mathcal{Q}$ have K arcs with rational coefficients and breakpoints given by representations. Then the envelope list of φ can be computed with $O(K)$ arithmetic operations and comparisons on rational numbers. Each endpoint in it is given as a polynomial with rational coefficients in at most two square roots, nested to depth at most two.*

Proof. Throughout, every comparison performed is the determination of the sign of a polynomial of bounded degree with rational coefficients in a bounded number of square roots, nested to depth at most two, of the kind described in the next paragraph; each of them costs $O(1)$ operations and comparisons on rational numbers by Lemma 16, and we do not repeat this below. Some of these tests involve quotients, such as slopes of lines through two stored points; their signs are obtained from those of numerators and denominators.

Write $0 = \xi_0 < \cdots < \xi_K = 1$ for the endpoints of the arcs of φ, with representations $\xi_i = \xi_{i,0} + \xi_{i,1}\sqrt{D_i}$, and $(a_j, b_j, d_j) \in \mathbb{Q}^3$ for the coefficients of its j-th arc. We work with $\hat{\varphi}$ throughout, which is strictly convex or affine on each arc, and represent the convex envelopes $H_j := \operatorname{env}(\varphi|_{[0,\xi_j]}) = \operatorname{env}(\hat{\varphi}|_{[0,\xi_j]})$ computed below by lists of contact pieces of $\hat{\varphi}|_{[0,\xi_j]}$, as in the envelope list. Every carrier is determined by the index of its arc, so no new coefficients are stored; the chord of a strictly concave arc j is used only through the points $(\xi_{j-1}, \varphi(\xi_{j-1}))$ and $(\xi_j, \varphi(\xi_j))$. The endpoints are of three kinds: breakpoints ξ_i of φ; abscissae of the tangency points of a common tangent of two parabolas with rational coefficients, which are again of the form $\tau_0 + \tau_1\sqrt{D}$ with $\tau_0, \tau_1, D \in \mathbb{Q}$; and abscissae η of the tangency point of the tangent from a point $(\xi_i, \varphi(\xi_i))$ to a parabola $y = ax^2 + bx + d$ with rational coefficients, which satisfy $a\eta^2 - 2a\xi_i\eta + \varphi(\xi_i) - b\xi_i - d = 0$ and are therefore of the form $\eta = \xi_i \pm \sqrt{\Delta}$ with Δ a polynomial with rational coefficients in ξ_i, hence in $\sqrt{D_i}$. We store every endpoint as such an expression. It is then a polynomial with rational coefficients in at most two square roots nested to depth at most two, and every test below is a polynomial in the square roots of $O(1)$ endpoints, which can be ordered so that each radicand is a polynomial in the preceding ones, as Lemma 16 requires. Every endpoint is computed from the arc coefficients and the breakpoints of φ alone, never from previously computed endpoints, so the depth of nesting never exceeds two. The endpoints computed during the sweep may have degree four over $\mathbb{Q}$.

We compute $H_1, \ldots, H_K$ in turn by a sweep over the arcs from left to right; $H_K = \operatorname{env}\varphi$. The envelope H_1 is $\hat{\varphi}$ on the arc 1. For the update, apply Lemma 5 to $g := \hat{\varphi}|_{[0,\xi_{j+1}]}$ with $\xi := \xi_j$; the three envelopes of that lemma are then H_j, the restriction of $\hat{\varphi}$ to the arc $j+1$, which is convex, and H_{j+1}. Hence there are $\eta \le \xi_j \le \zeta$ with

$$H_{j+1} = H_j \text{ on } [0, \eta], \qquad H_{j+1} \text{ affine on } [\eta, \zeta], \qquad H_{j+1} = \hat{\varphi} \text{ on } [\zeta, \xi_{j+1}].$$

In particular the update deletes the contact pieces and bridges of H_j lying to the right of η and leaves the rest untouched. The arc $j+1$, which we call the incoming arc, is the graph of a strictly convex parabola or a segment. We scan the contact pieces of H_j from right to left. Let P be the rightmost contact piece, lying in the arc i, and let w be its left endpoint. If $w = 0$, then P survives. Otherwise let ℓ_0 be the line through $(w, H_j(w))$ with slope $H_j'(w^-)$, which is the line of the bridge ending at w if there is one, and the tangent to the preceding contact piece at w otherwise. By Lemma 5, $H_{j+1} = H_j$ on $[0, w]$ if and only if ℓ_0 lies below the incoming arc, a test of $O(1)$ sign determinations: a line lies below the graph of a strictly convex parabola or of an affine function over a compact interval if and only if it does so at the two endpoints and, for a parabola, at the point of the interval where the two slopes agree, if there is one. If it does, P survives. Otherwise H_{j+1} differs from H_j somewhere on $[0, w]$, so $\eta < w$, and P is popped, together with the bridge ending at w if there is one; the scan then continues with the preceding contact piece.

If P survives, the new bridge is the common supporting line ℓ from below of the carrier of P over $[w, \xi_i]$ and of the incoming arc, both of which are the graph, over a compact interval, of a strictly convex parabola with rational coefficients or of an affine function. Indeed, the line of H_{j+1} on $[\eta, \zeta]$ lies below $H_j \le \hat{\varphi}$ and below the incoming arc, and touches both; and a line with these properties is unique, since between its two contact points it coincides with the convex envelope of the union of the two graphs, unless both contacts are at a common endpoint $\xi_i = \xi_j$, in which case $\eta = \zeta$ and no bridge is needed. The line ℓ touches each of them at an interior point, where it is tangent, or at an endpoint, and it is found among $O(1)$ candidates by $O(1)$ sign determinations: the common tangent of the two parabolas; the tangent from an endpoint $(\xi_j, \varphi(\xi_j))$ or $(\xi_{j+1}, \varphi(\xi_{j+1}))$ of the incoming arc to the parabola carrying P; the tangent from $(\xi_i, \varphi(\xi_i))$, or from $(0, \varphi(0))$ if $w = 0$, to the incoming parabola; and the line through one of these points on each side. Here a contact with the carrier of P at a point of $(0, \xi_i)$ is a tangency, even at w, because such a point is a contact point of H_j, where H_j is differentiable by Lemma 6. If ℓ touches the carrier of P or the incoming arc on a nondegenerate interval, which happens only if that carrier or arc is affine, we take η as large and ζ as small as possible, that is, $\eta = \xi_i$ or $\zeta = \xi_j$ in the respective case, so that the contact pieces stored are maximal. In every case ℓ touches the carrier of P at a point $\eta \in [w, \xi_i]$ and the incoming arc at a point ζ, both endpoints of the three kinds listed above; moreover $\eta \in P$, since ℓ lies below H_j and $\ell(\eta) = \hat{\varphi}(\eta)$, so that η is a contact point of H_j, and, by Lemma 8, the contact set of $\hat{\varphi}|_{[0,\xi_j]}$ meets the arc i only in P. Then P is replaced by $[w, \eta]$, the incoming arc contributes the contact piece $[\zeta, \xi_{j+1}]$, and ℓ is the bridge between them, unless $\eta = \zeta$, in which case no bridge is stored; this completes the update.

It remains to count the operations. Each contact piece is created once and popped at most once during the whole sweep, and each survival test, pop, and bridge computation costs $O(1)$ arithmetic operations and $O(1)$ comparisons on rational numbers, so the sweep costs $O(K)$ of each. □

The second step reads off the value function from the envelope list.

Lemma 18. *Let $\varphi \in \mathcal{Q}$ have K arcs with rational coefficients and breakpoints given by representations, and let $s \in \mathbb{Q} \setminus \{0\}$. Then the arcs of $\mu(t) = \min_{x \in [0,1]}[\varphi(x) + stx]$ can be computed exactly with $O(K)$ arithmetic operations and comparisons on rational numbers. The arcs produced have rational coefficients, and each breakpoint produced is given by its defining quadratic and a representation. Moreover the arcs of μ are produced in the order of their breakpoints, increasing if $s < 0$ and decreasing if $s > 0$.*

Proof. By Lemma 17, the envelope list of φ, which describes $H := \operatorname{env} \varphi$, can be computed with $O(K)$ operations; as there, every comparison below costs $O(1)$ operations by Lemma 16. We read off μ from H by Lemma 9, applied to the function $\hat{\varphi}$ of Lemma 11, which has the same convex envelope and the same values at 0 and 1 as φ, hence the same μ: each nondegenerate contact piece contained in a strictly convex arc j, with coefficients (a_j, b_j, d_j), nondegeneracy being one comparison, contributes the arc $\mu(t) = -\frac{(st+b_j)^2}{4a_j} + d_j$; each bridge, and each contact piece in the list that is degenerate or whose carrier is a segment, contributes at most a breakpoint; and the two unbounded arcs are $\mu(t) = \varphi(0)$ and $\mu(t) = \varphi(1) + st$. This yields the arcs of μ, in order and with rational coefficients, with $O(K)$ further operations. Consecutive arcs produced carry distinct polynomials, so no two need to be merged: two consecutive quadratic arcs come from contact pieces either in consecutive arcs of φ, which carry distinct polynomials, or separated by an interval on which H is affine, whose line is tangent to both parabolas at two distinct points, so that the parabolas differ; a quadratic arc has t^2-coefficient $-s^2/(4a_j) \neq 0$, while the two unbounded arcs are affine and distinct as $s \neq 0$. A breakpoint τ of μ separates two consecutive arcs with polynomials $\pi_1 \neq \pi_2$, and by continuity of μ it is a root of $\pi_1 - \pi_2$, which we call its defining quadratic, as in Lemma 12. Moreover $\tau = -\sigma/s$, where σ is the slope of H at the corresponding endpoint in the envelope list, or on the corresponding bridge, so deciding which root of $\pi_1 - \pi_2$ equals τ is one comparison, after which a representation of τ is obtained with $O(1)$ operations; the nested radicals are then discarded. Lastly, the contact pieces in the envelope list are ordered by increasing x, hence in increasing order of the slope $\sigma = H'$, and the arc of μ that a contact piece contributes is indexed by $t = -\sigma/s$; the arcs of μ are therefore produced in the order of their breakpoints, increasing if $s < 0$ and decreasing if $s > 0$. □

It remains to recover an optimal solution of Problem (QP).

Lemma 19. *Problem* (QP) *admits an optimal solution* $x^\star \in \mathbb{Q}^n$*, with* $\langle x^\star \rangle$ *polynomial in* $\mathcal{L}$*, computable by backtracking from the functions* φ_v *of* (3) *with* $O\big(\sum_v K_v\big)$ *additional operations.*

Proof. Consider first the root problem $\min_{x\in[0,1]} \varphi_r(x)$. By Lemma 10, applied to $\varphi_r|_{[0,1]} \in \mathcal{Q}$, the minimum is attained at one of the at most $K_r + 2$ points listed there, which we call candidates. Each candidate is 0, 1, or is obtained from the coefficients of one arc by $O(1)$ operations, and whether it belongs to its arc is decided by two comparisons against the endpoints of the arc, which are given by representations. Evaluating φ_r at the admissible candidates and comparing the values exactly selects $x_r^\star$.

Now descend. Given the value already chosen for the parent, $t^\star := x^\star_{p(v)}$, we choose $x_v^\star$ as a minimizer of $\varphi_v(x) + s_v t^\star x$ over $[0,1]$. Since adding an affine function changes no one-sided derivative jump, $x \mapsto \varphi_v(x) + s_v t^\star x$ lies in $\mathcal{Q}$, with the same arcs as $\varphi_v|_{[0,1]}$ and with b replaced by $b + s_v t^\star$ on each arc, so Lemma 10 applies to it, and $x_v^\star$ is selected as at the root. The total operation count is $O\big(\sum_v K_v\big)$.

The resulting $x^\star$ is optimal. Indeed, by induction on the height of $v \neq r$, the point $x^\star_{T_v}$ minimizes $F_v(\cdot\,; x^\star_{p(v)})$ over $[0,1]^{T_v}$: by (2), the inductive hypothesis and Lemma 3, $F_v(x^\star_{T_v}; x^\star_{p(v)}) = \varphi_v(x_v^\star) + s_v x^\star_{p(v)} x_v^\star$, which equals $\mu_v(x^\star_{p(v)})$ by the choice of $x_v^\star$ and (3). At the root, the same computation with $s_r = 0$ gives $f(x^\star) = F_r(x^\star; 0) = \varphi_r(x_r^\star) = \min_{x\in[0,1]} \varphi_r(x)$, which is the optimal value by Lemma 3.

All candidates are rational. To bound their lengths, write $h(\alpha) := \max\{\lceil \log_2(|\nu|+1)\rceil, \lceil \log_2(\delta+1)\rceil\}$ for a rational $\alpha = \nu/\delta$ in lowest terms, so that $\langle \alpha \rangle \le 2h(\alpha) + 2$. By Lemma 15, the arc coefficients have h bounded by some κ polynomial in $\mathcal{L}$, and we may take κ so large that $h(s_v) \le \kappa$ for all v. At the root, every candidate therefore has $h = O(\kappa)$. At a child, writing $t^\star = \nu/\delta$, the candidate $-(b + s_v t^\star)/(2a)$ is a quotient of integers whose absolute values are at most $2^{O(\kappa)} \max\{|\nu|, \delta\}$ and $2^{O(\kappa)}\delta$, so its h exceeds $h(t^\star)$ by $O(\kappa)$. Hence $h(x_v^\star) = O(n\kappa)$ for every v, and $\langle x^\star \rangle$ is polynomial in $\mathcal{L}$, and so are the values compared when selecting each $x_v^\star$, which are obtained from the candidates, $t^\star$ and the arc coefficients by $O(1)$ operations. ☐

We are now ready to prove Theorem 1.

Proof of Theorem 1. We first bound the number of operations. Process vertices leaf-to-root, and let $d_v := |\mathrm{ch}(v)|$ denote the number of children of v. Let m_v denote the total number of breakpoints of the μ_j, $j \in \mathrm{ch}(v)$; by Proposition 1, $m_v \le \sum_{j\in \mathrm{ch}(v)} 2|T_j| = 2(|T_v| - 1)$. By Lemma 18 the arcs of each child value function μ_j are produced in the order of their breakpoints, up to a global reversal according to the sign of s_j, so that the d_v breakpoint lists are available in increasing order, after reversing some of them, at a cost of $O(m_v)$. Forming φ_v amounts to merging them, discarding the breakpoints outside $(0,1)$; each comparison of two breakpoints costs $O(1)$ operations by Lemma 16. We merge by a single scan that repeatedly selects the smallest among the d_v current heads, at a cost of $O(d_v m_v)$ comparisons, while maintaining the sum of the arc coefficients incrementally: crossing a breakpoint of a child μ_j updates the running sum by subtracting the coefficients of the arc of μ_j that ends there and adding those of the arc that begins, at a cost of $O(1)$ operations, the initial sum costing $O(d_v)$. Several children may have a breakpoint at the same point, in which case all their updates are performed there; and consecutive intervals of the merged list carrying the same polynomial are coalesced, so that the arcs of φ_v are maximal. The merge thus costs $O(d_v m_v + d_v)$ operations and comparisons, including the $O(m_v)$ for reversing. For $v \neq r$, computing μ_v then costs $O(K_v)$ operations by Lemma 18 applied to $(\varphi_v|_{[0,1]}, s_v)$; its hypotheses hold because $\varphi_v|_{[0,1]} \in \mathcal{Q}$ (Proposition 1), its arcs have rational coefficients and its breakpoints are breakpoints of the μ_j (Lemma 12), whose representations are provided by Lemma 18, and $s_v = 2q_{p(v)v} \in \mathbb{Q} \setminus \{0\}$. By

Proposition 1, $K_v \le 2|T_v| - 1$, and

$$\sum_{v\in V} |T_v| = \sum_{u\in V} \bigl(\mathrm{depth}(u) + 1\bigr) \le n^2,$$

because $\mathrm{depth}(u) \le n-1$ for every u; hence $\sum_{v\in V} K_v = O(n^2)$. The merging is equally cheap: since $m_v \le 2(|T_v| - 1) \le 2n$ and $\sum_{v\in V} d_v = n-1$, we have $\sum_{v\in V}(d_v m_v + d_v) \le (2n+1)(n-1) = O(n^2)$. Adding the $O(\sum_v K_v)$ operations of the backtracking of Lemma 19, the total is $O(n^2)$ arithmetic operations and comparisons. Lemma 19 then returns an optimal solution.

We now bound the lengths of the numbers involved. The following rationals have length polynomial in $\mathcal{L}$: the entries of Q and c; the arc coefficients of the φ_v and μ_v, the coefficients of their defining quadratics and the rationals of the representations of their breakpoints, by Lemma 15; the partial sums of at most n arc coefficients formed during a merge, as sums of at most n such rationals; and the coordinates of $x^\star$ and the values compared when selecting them, by Lemma 19. Every other number computed by the algorithm is obtained from these by $O(1)$ arithmetic operations: this holds for the rationals representing the auxiliary nested radicals of Lemma 17, since every endpoint there is computed from arc coefficients and breakpoints alone, and for the rationals computed within a sign determination of Lemma 16. Hence every number computed has length polynomial in $\mathcal{L}$.

Finally, the operation count depends on n alone and every intermediate number has length polynomial in the input length, so the algorithm is strongly polynomial. $\square$

2.5 Relation to the literature

Kolmogorov, Pock and Rolinek [25] consider problems over a tree of the form

$$\min\Bigl\{\sum_{v\in V} g_v(x_v) + \sum_{v\ne r} h_v\bigl(x_v - x_{p(v)}\bigr) \ : \ x \in \mathbb{R}^V\Bigr\},$$

and solve them by dynamic programming when the terms h_v are (truncated) total variation terms; Kuric, Ahmetspahic and Pock [27] extend this type of algorithm to piecewise quadratic g_v and h_v, motivated by total generalized variation models. Up to the fact that in [27] these terms are finite on $\mathbb{R}$, Problem (QP) is of this form when its interaction graph is a tree: by the identity $s_v x_{p(v)} x_v = \frac{1}{2}s_v(x_{p(v)}^2 + x_v^2) - \frac{1}{2}s_v(x_v - x_{p(v)})^2$, one can take $h_v(d) = -\frac{1}{2}s_v d^2$, and $g_v(x) = \bigl(q_{vv} + \sum_{u:\{u,v\}\in E} q_{uv}\bigr)x^2 + c_v x$ on $[0,1]$ and $g_v(x) = +\infty$ elsewhere. The message sent from v to $p(v)$ by the dynamic program of [27] is then $\mu_v(t) - q_{p(v)v}t^2$. For nonconvex terms, the algorithm of [27] has exponential worst-case time and memory complexity, whereas, for the terms arising from Problem (QP), the value functions have linearly many arcs (Proposition 1), and Theorem 1 gives a strongly polynomial algorithm performing $O(n^2)$ operations.

A second, more closely related, line of work concerns *convex quadratic optimization with indicator variables,* that is, the problem

$$\min\Bigl\{\tfrac{1}{2}x^\top Qx + c^\top x + \lambda^\top z \ : \ x_i(1 - z_i) = 0 \text{ for } i \in [n],\ x \in \mathbb{R}^n,\ z \in \{0,1\}^n\Bigr\},$$

where $Q \succ 0$ and $\lambda > 0$. When the interaction graph of Q is a tree, Bhathena, Fattahi, Gómez and Küçükyavuz [1] give an $O(n^2)$ algorithm based on a parametric analysis of the value functions of a dynamic program over the tree. For positive definite Q whose interaction graph has bounded treewidth, the same authors give a parametric dynamic program over a tree decomposition [2], whose complexity depends on structural and numerical parameters, including treewidth, volume growth, conditioning and a margin parameter, and is linear in n when the relevant parameters are bounded appropriately; Problem (QP), instead, is strongly $\mathcal{NP}$-hard already when the interaction graph has treewidth two (Theorem 3). The technique of [1] is closely related to ours.

Similarities. Both algorithms process the tree from the leaves to the root. Their *parametric cost* at v, the optimal value of the subproblem on the subtree rooted at v as a function of x_v, plays the role of our φ_v, and their recursion is the analogue of (3): for $v \neq r$, minimizing the parametric cost at v plus $q_{p(v)v}\, x_{p(v)} x_v$ over x_v yields a function of $x_{p(v)}$, expressed in [1] through a Fenchel conjugate, which corresponds to our value function μ_v. Both analyses isolate a class of piecewise quadratic functions preserved by the recursion, their *consistent* functions, up to an indicator term, and our class $\mathcal{Q}$, and share three ingredients: conjugation adds at most two arcs ([1, Proposition 1] and Lemma 9); the conjugate is computed by a linear-time sweep ([1, Proposition 2, Theorem 3] and Lemmas 17 and 18); and every arc has a closed form ([1, Lemma 1] and Lemma 14). In both, the number of arcs adds at vertices with several children, which yields $O(n^2)$ operations.

Differences. The two works differ in the problem solved, in the role of convexity, and in the model of computation; the last two are where the main difficulties of our analysis lie.

(*i*) *Problem structure.* Neither problem is a special case of the other. The problem of [1] is mixed-integer, with unbounded continuous variables whose support is chosen by the binary variables and penalized by $\lambda^\top z$; Problem (QP) is continuous, with variables confined to a box. Accordingly, the coordinates that are not free in Lemma 14 sit at either endpoint of their range, whereas in [1] they are all equal to zero.

(*ii*) *Nonconvexity.* In [1], $Q \succ 0$, so every arc of every parametric cost, up to the indicator term, is a strongly convex quadratic [1, Lemma 1], and this is why consistency is preserved along the recursion. We make no assumption on Q: an arc of φ_v may be strictly convex, affine or strictly concave, and consistency fails. What survives is the weaker property that every kink is concave (Proposition 1), which defines $\mathcal{Q}$ and holds for the restriction to $[0, 1]$ of every consistent function. Contact must then be excluded from the interior of strictly concave arcs (Lemma 7), and the sweep must replace every strictly concave arc by its chord (Lemma 11); both steps are vacuous in the convex setting. The count of at most one arc of μ_v per arc of φ_v (Lemmas 7 and 8) is the analogue of the monotonicity of their indexing function [1, Lemma 5], and holds here for arcs of any curvature. Finally, the positive definiteness of the free block becomes a conclusion instead of a hypothesis: in Lemma 14, $Q_S \succ 0$ follows from Lemma 13, which makes the length bounds of Lemma 15 available.

(*iii*) *Model of computation.* The bound $O(n^2)$ of [1] counts arithmetic operations on real numbers and does not estimate the lengths of the numbers; already the normalization $q_{ii} = 1$ made there, by $x_i \leftarrow x_i/\sqrt{q_{ii}}$, takes the data out of $\mathbb{Q}$. We work in the Turing model. In both settings the breakpoints are irrational already for rational data, for instance slopes of common tangents to two parabolas; we compute and compare them exactly, with $O(1)$ operations on rational numbers and no root extraction (Lemma 16), whereas [1, Section 4] reports that in floating-point arithmetic their computation becomes unstable on large trees and gives a correction step. That the numbers produced along the recursion have polynomial length is not evident from the recursion itself; we obtain it from the closed form of Lemma 14, which exhibits every arc as the value of an unconstrained quadratic minimization with a positive definite matrix (Lemma 15).

We believe that our techniques can be applied to prove that the algorithm of [1] is in fact strongly polynomial, once we remove the normalization $q_{ii} = 1$. Namely, by [1, Lemma 1], every arc of their parametric costs is also the value of an unconstrained quadratic minimization with a positive definite matrix, so the argument of Lemma 15 bounds the lengths of its coefficients, and their breakpoints can be computed and compared exactly as in Lemma 16.

2.6 Higher degree box-constrained polynomial optimization

Theorem 1 shows that a box-constrained quadratic program can be solved in strongly polynomial-time when the interaction graph is a forest. It is therefore natural to ask whether a similar tractability result continues to hold for polynomials of higher degree under the same sparsity assumption. The next result shows that this is not the case: already for quartic objectives, the problem becomes strongly $\mathcal{NP}$-hard even when the interaction graph is a path. Thus, the quadratic structure of Problem (QP), and not only the acyclicity of its interaction graph, is essential to our tractability result.

Let $f(x) = \sum_{\alpha \in \mathbb{Z}^n_{\geq 0}} c_\alpha x^\alpha$ be a polynomial, where $x^\alpha := \prod_{i=1}^n x_i^{\alpha_i}$. We define the *interaction graph* of f as the graph $G = (V, E)$ with vertex set $V = [n]$, containing one vertex for each variable x_i, and with $\{i, j\} \in E$ for some $i \neq j \in [n]$ if and only if there exists α with $c_\alpha \neq 0$ such that $\alpha_i > 0$ and $\alpha_j > 0$. Namely, two distinct vertices are adjacent if and only if the corresponding variables appear together in a monomial with nonzero coefficient.

Theorem 2. *Minimizing a polynomial of degree four with integer coefficients over the unit hypercube is strongly $\mathcal{NP}$-hard even when its interaction graph is a path.*

Proof. Consider the decision version of the problem: given a degree four polynomial $f(x)$ with integer coefficients and $\gamma \in \mathbb{Q}$, decide whether

$$\min\{f(x) : x \in [0,1]^n\} \leq \gamma.$$

To prove the statement, it suffices to show that this decision problem is strongly $\mathcal{NP}$-hard. We give a polynomial-time reduction from the strongly $\mathcal{NP}$-complete PRODUCT PARTITION problem [29]. Given positive integers $a_1, \dots, a_n$, the PRODUCT PARTITION problem asks whether there exists a subset $S \subseteq [n]$ such that

$$\prod_{i \in S} a_i = \prod_{i \in [n] \setminus S} a_i. \tag{9}$$

Define $P := \prod_{i=1}^n a_i$, introduce the variables $x_1, \dots, x_{2n} \in [0, 1]$, and let $x_0 := 1$. Consider the polynomial

$$\begin{aligned} f(x) := &\sum_{i=1}^n (a_i x_i - x_{i-1})^2 && (10)\\ &+ \sum_{i=1}^n (x_{n+i} - x_{n+i-1})^2 (x_{n+i} - a_i^2 x_{n+i-1})^2 && (11)\\ &+ (x_{2n} - 1)^2. && (12) \end{aligned}$$

Clearly, $f(x)$ is a polynomial of degree four with integer coefficients. We first prove that the minimum of $f(x)$ over $[0,1]^{2n}$ is zero if and only if the given PRODUCT PARTITION instance is feasible. First, suppose that $x \in [0,1]^{2n}$ satisfies $f(x) = 0$. Since every term in (10)–(12) is nonnegative, every squared term must vanish. From (10) we obtain $a_i x_i = x_{i-1}$ for all $i \in [n]$. Using $x_0 = 1$ recursively gives

$$x_i = \frac{1}{\prod_{j=1}^i a_j}, \quad \forall i \in [n]. \tag{13}$$

and in particular

$$x_n = \frac{1}{P}. \tag{14}$$

Next, from (11), for every $i \in [n]$ we obtain

$$(x_{n+i} - x_{n+i-1})(x_{n+i} - a_i^2 x_{n+i-1}) = 0.$$

Hence, for each $i \in [n]$, either $x_{n+i} = x_{n+i-1}$ or $x_{n+i} = a_i^2 x_{n+i-1}$. Let $S \subseteq [n]$ be the set of indices for which $x_{n+i} = a_i^2 x_{n+i-1}$. Using (14), we obtain

$$x_{2n} = \frac{1}{P} \prod_{i \in S} a_i^2. \tag{15}$$

Moreover, from (12) we obtain $x_{2n} = 1$, implying that

$$\prod_{i \in S} a_i^2 = P = \prod_{i \in [n]} a_i = \Big(\prod_{i \in S} a_i\Big)\Big(\prod_{i \in [n] \setminus S} a_i\Big).$$

Since all a_is are positive, we deduce that $\prod_{i \in S} a_i = \prod_{i \in [n] \setminus S} a_i$. Thus S solves the Product Partition instance.

Conversely, suppose that $S \subseteq [n]$ satisfies (9). Define

$$x_i := \frac{1}{\prod_{j=1}^{i} a_j}, \quad \forall i \in [n]. \tag{16}$$

Since $a_j \geq 1$ for all $j \in [n]$, we have $x_i \in (0, 1]$ for all $i \in [n]$. For each $i \in [n]$, define

$$x_{n+i} := \frac{1}{P} \prod_{\substack{j \in S \\ j \leq i}} a_j^2. \tag{17}$$

Since $a_j \geq 1$, we have $\prod_{j \in S,\, j \leq i} a_j^2 \leq \prod_{j \in S} a_j^2$. From (9) it follows that $\prod_{j \in S} a_j^2 = P$. Therefore, $0 < x_{n+i} \leq 1$, for all $i \in [n]$. We then have $x \in [0, 1]^{2n}$. From (16) it follows that

$$a_i x_i - x_{i-1} = 0 \quad \forall i \in [n],$$

implying that every term in (10) vanishes. Furthermore, from (17) it follows that

$$x_{n+i} = \begin{cases} a_i^2 x_{n+i-1}, & i \in S, \\ x_{n+i-1}, & i \notin S, \end{cases} \quad \forall i \in [n].$$

Therefore, every term in (11) also vanishes. Finally, by construction we have $x_{2n} = 1$, so (12) vanishes as well. Therefore $f(x) = 0$.

We next examine the interaction graph of $f(x)$. Expanding the objective function gives

$$\begin{aligned} f(x) = &\sum_{i=1}^{n} \left(a_i^2 x_i^2 - 2a_i x_{i-1} x_i + x_{i-1}^2\right) \\ &+ \sum_{i=1}^{n} \left(x_{n+i}^4 - 2(a_i^2+1) x_{n+i-1} x_{n+i}^3 + (a_i^4 + 4a_i^2 + 1) x_{n+i-1}^2 x_{n+i}^2 - 2a_i^2(a_i^2+1) x_{n+i-1}^3 x_{n+i} + a_i^4 x_{n+i-1}^4\right) \\ &+ x_{2n}^2 - 2x_{2n} + 1, \end{aligned}$$

where $x_0 := 1$. It is then simple to see that the interaction graph of $f(x)$ is the path $x_1 - x_2 - \cdots - x_{2n}$.

Finally, we verify strong $\mathcal{NP}$-hardness. The polynomial $f(x)$ has $2n$ variables and $O(n)$ monomials after expansion. Let $a_{\max} := \max_{i \in [n]} a_i$. It follows from the above expansion that every coefficient of $f(x)$ has absolute value at most $6a_{\max}^4$. Since PRODUCT PARTITION is strongly $\mathcal{NP}$-complete [29], we may restrict attention to instances for which $a_{\max}$ is polynomially bounded in n. Hence all numerical coefficients of the constructed instance are polynomially bounded in its number of variables. The box bounds are 0 and 1, and we set the decision threshold to $\gamma := 0$. The construction is clearly polynomial in the encoding length of the PRODUCT PARTITION instance. Therefore, the decision problem is strongly $\mathcal{NP}$-hard. $\square$

3 Treewidth two: Strong $\mathcal{NP}$-hardness

In this section, we show that Problem (QP) is strongly $\mathcal{NP}$-hard even when the interaction graph has treewidth two. This result in turn implies that Theorem 1 is sharp with respect to the treewidth. Let us first give a formal definition of treewidth. Given a graph $G = (V, E)$, a *tree decomposition* of G is a pair $(\mathcal{X}, T)$, where $\mathcal{X} = \{X_1, \ldots, X_m\}$ is a family of subsets of V, called *bags*, and T is a tree with m vertices, the i-th vertex being associated with the bag X_i, such that:

1. $V = \bigcup_{i \in [m]} X_i$.
2. For every edge $\{u, v\} \in E$, there is a bag X_i, $i \in [m]$, with $u, v \in X_i$.
3. For each vertex $v \in V$, the set of all bags containing v induces a connected subtree of T.

The *width* of a tree decomposition $(\mathcal{X}, T)$ is the size of its largest bag X_i minus one. The *treewidth* $\mathrm{tw}(G)$ of a graph G is the minimum width among all possible tree decompositions of G.

We are now ready to present our $\mathcal{NP}$-hardness result.

Theorem 3. *Problem* (QP) *is strongly $\mathcal{NP}$-hard even when the interaction graph has treewidth two. Moreover, the result holds for instances satisfying*

$$Q \in \mathbb{Z}^{n \times n}, \quad c \in \mathbb{Z}^n, \quad \|Q\|_{\max} \le 5, \quad \|c\|_\infty \le 4,$$

where $\|Q\|_{\max} := \max_{p,q} |Q_{pq}|$ and n is the number of variables of the instance.

It is well-known that quadratic programming is in $\mathcal{NP}$ [40]. Therefore, Theorem 3 implies that Problem (QP) is strongly $\mathcal{NP}$-complete, even if the interaction graph has treewidth two.

Proof. Consider the decision version of Problem (QP): given symmetric $Q \in \mathbb{Q}^{n \times n}$, $c \in \mathbb{Q}^n$, and $\gamma \in \mathbb{Q}$, decide whether

$$\min \left\{ x^\top Q x + c^\top x : x \in [0,1]^n \right\} \le \gamma.$$

To prove the statement, it suffices to show that this decision problem is strongly $\mathcal{NP}$-hard. We give a polynomial-time reduction from SUBSET SUM [18]. Throughout this proof, n denotes the number of items of the SUBSET SUM instance, and N the number of variables of the constructed instance of Problem (QP). Given integers $a_1, \ldots, a_n \in \mathbb{Z}_{\ge 0}$ and a target $T \in \mathbb{Z}_{\ge 0}$, the SUBSET SUM asks whether there exists $x \in \{0,1\}^n$ such that

$$\sum_{i=1}^{n} a_i x_i = T.$$

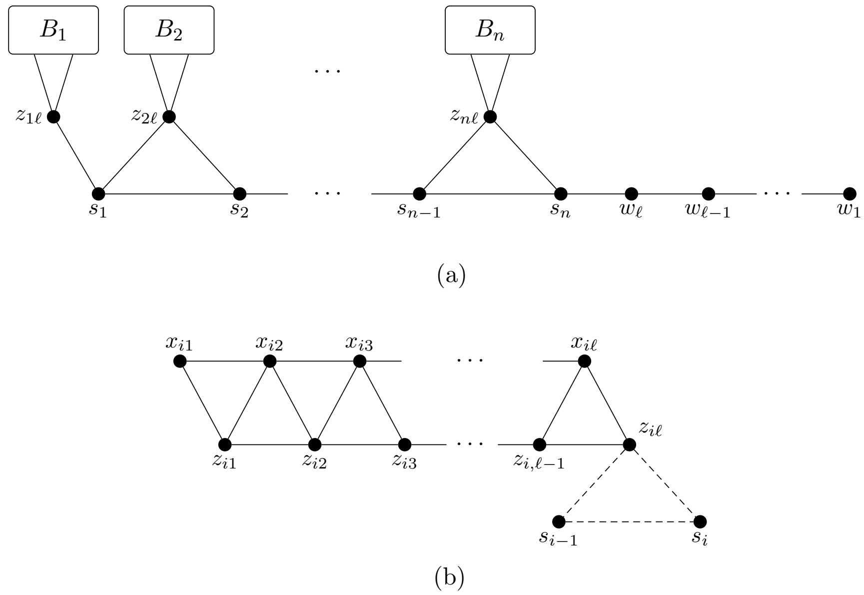

Figure 1: The interaction graph of the instance of Problem (QP) constructed in the proof of Theorem 3, drawn for $b_{ik} = 1$ for all $i \in [n]$ and $k \in [\ell]$; if $b_{ik} = 0$, the two edges of x_{ik} incident to $z_{i,k-1}$ and z_{ik} are absent. In (a), the vertices $s_1, \ldots, s_n$ form a path, the vertices $w_1, \ldots, w_\ell$ form a path attached to s_n, and, for $i \geq 2$, the vertex $z_{i\ell}$ forms a triangle with s_{i-1} and s_i, while $z_{1\ell}$ is adjacent to s_1 alone. In either case $z_{i\ell}$ is the unique vertex through which the block B_i, depicted in (b), is attached to the rest of the graph, and the two edges leaving $z_{i\ell}$ upwards lead to $z_{i,\ell-1}$ and $x_{i\ell}$.

Define

$$U := \sum_{i=1}^{n} a_i.$$

The cases with $T = 0$, $U = 0$, and $T > U$ can be decided trivially. We therefore assume that $1 \leq T \leq U$. Define

$$\ell := \lceil \log_2(U+1) \rceil \qquad \text{and} \qquad D := 2^\ell.$$

In particular, we have $D > U$. For each $i \in [n]$, write the binary expansion of a_i as

$$a_i = \sum_{k=1}^{\ell} b_{ik} 2^{k-1}, \qquad b_{ik} \in \{0, 1\}. \tag{18}$$

Similarly, write the binary expansion of T as

$$T = \sum_{k=1}^{\ell} t_k 2^{k-1}, \qquad t_k \in \{0, 1\}. \tag{19}$$

We construct an instance of Problem (QP) as follows. Introduce the following variables:

- $s_i \in [0, 1]$ for all $i \in [n]$,
- $w_k \in [0, 1]$ for all $k \in [\ell]$,
- $x_{ik}, z_{ik} \in [0, 1]$ for all $i \in [n]$ and all $k \in [\ell]$.

We further define $z_{i0} := 0$ for all $i \in [n]$, $s_0 := 0$, and $w_0 := 0$. Consider the following quadratic function:

$$\Psi(x, z, s, w) := \sum_{i=1}^{n} \Big(x_{i1}(1 - x_{i1}) + \sum_{k=2}^{\ell} (x_{ik} - x_{i,k-1})^2 \Big) \tag{20}$$

$$+ \sum_{i=1}^{n} \sum_{k=1}^{\ell} (2z_{ik} - z_{i,k-1} - b_{ik} x_{ik})^2 \tag{21}$$

$$+ \sum_{i=1}^{n} (s_i - s_{i-1} - z_{i\ell})^2 \tag{22}$$

$$+ \sum_{k=1}^{\ell} (2w_k - w_{k-1} - t_k)^2 \tag{23}$$

$$+ (s_n - w_\ell)^2. \tag{24}$$

See Figure 1 for the interaction graph of the resulting instance. Every term in (20)–(24) is nonnegative over the unit box. Consequently, $\Psi(x, z, s, w) \geq 0$ for every feasible point. We first show that the minimum of Ψ is zero if and only if the original Subset Sum instance is feasible. Suppose that $\Psi(x, z, s, w) = 0$. Since all the terms in (20)–(24) are nonnegative over the unit box, each of them vanishes. The first term in (20) implies that $x_{i1} \in \{0, 1\}$ for all $i \in [n]$. The remaining terms in (20) imply that

$$x_{ik} = x_{i,k-1} \quad \forall i \in [n],\ k \in \{2, \ldots, \ell\}.$$

Thus, for every i, there exists $\bar{x}_i \in \{0, 1\}$ such that

$$x_{ik} = \bar{x}_i \quad \forall k \in [\ell].$$

From (21), we get $z_{ik} = \frac{1}{2}(z_{i,k-1} + b_{ik}\bar{x}_i)$ for all $i \in [n]$ and all $k \in [\ell]$, which together with $z_{i0} = 0$ implies that

$$z_{ik} = \frac{\bar{x}_i}{2^k} \sum_{r=1}^{k} b_{ir} 2^{r-1} \quad \forall i \in [n],\ k \in [\ell].$$

In particular, from (18) it follows that

$$z_{i\ell} = \frac{a_i \bar{x}_i}{D}. \tag{25}$$

Similarly, from (23) we get $w_k = \frac{1}{2}(w_{k-1} + t_k)$ for all $k \in [\ell]$, which together with $w_0 = 0$ implies:

$$w_k = \frac{1}{2^k} \sum_{r=1}^{k} t_r 2^{r-1} \quad \forall k \in [\ell].$$

In particular, from (19) it follows that

$$w_\ell = \frac{T}{D}. \tag{26}$$

From (22) we get $s_i = s_{i-1} + z_{i\ell}$ for all $i \in [n]$. Using $s_0 = 0$ and (25), we obtain

$$s_n = \sum_{i=1}^{n} z_{i\ell} = \frac{1}{D} \sum_{i=1}^{n} a_i \bar{x}_i. \tag{27}$$

Finally, from (24) we deduce that

$$s_n = w_\ell. \tag{28}$$

Combining (26), (27), and (28) gives $\sum_{i=1}^n a_i \bar{x}_i = T$. Hence $\bar{x} \in \{0,1\}^n$ is a solution of the Subset Sum instance.

Conversely, suppose that $\bar{x} \in \{0,1\}^n$ satisfies $\sum_{i=1}^n a_i \bar{x}_i = T$. Set $x_{ik} := \bar{x}_i$ for all $i \in [n]$ and all $k \in [\ell]$, and define

$$\begin{aligned} z_{ik} &:= \frac{\bar{x}_i}{2^k} \sum_{r=1}^k b_{ir} 2^{r-1} \quad \forall i \in [n], k \in [\ell] \\ w_k &:= \frac{1}{2^k} \sum_{r=1}^k t_r 2^{r-1} \quad \forall k \in [\ell] \\ s_i &:= \frac{1}{D} \sum_{j=1}^i a_j \bar{x}_j \quad \forall i \in [n]. \end{aligned}$$

It can be checked that substituting these values in (20)–(24), we get $\Psi = 0$. They also belong to the unit box. Indeed, for every i and k,

$$0 \le z_{ik} \le \frac{1}{2^k} \sum_{r=1}^k 2^{r-1} = 1 - \frac{1}{2^k} < 1,$$

and the same argument applies to w_k. Moreover, $0 \le s_i \le \frac{U}{D} < 1$. We have established that the minimum value of $\Psi(x, z, s, w)$ is zero if and only if the Subset Sum instance is feasible.

We next express the objective function in the form required by Problem (QP), which does not include an additive constant. The only constant terms produced by expanding Ψ occur in (23). Since $t_k \in \{0,1\}$, their sum is

$$C := \sum_{k=1}^\ell t_k^2 = \sum_{k=1}^\ell t_k.$$

Define $F(x, z, s, w) := \Psi(x, z, s, w) - C$. After expansion, F can be written as

$$F(y) = y^\top Q y + c^\top y$$

for a symmetric integral matrix Q and an integral vector c, where $y \in [0,1]^N$ collects all the introduced variables. Note that Q is integral because every product of two distinct variables arises from squaring an affine expression and therefore has an even coefficient in F. We then have

$$\min_{y \in [0,1]^N} F(y) = -C$$

if and only if the original Subset Sum instance is feasible. Thus, we define the decision threshold as $\gamma := -C$.

We now verify strong $\mathcal{NP}$-hardness. The constructed instance of Problem (QP) has $N = 2n\ell + n + \ell$ variables. Since

$$\ell = \lceil \log_2(U+1) \rceil = O\left(\log n + \max_{i=1,\dots,n} \log(a_i + 1) \right),$$

the constructed instance has length polynomial in the length of the Subset Sum instance. Moreover, all numerical coefficients in the resulting quadratic optimization problem are bounded independently of the Subset Sum instance. To see this, observe that each squared affine expression has coefficients in $\{-1, 0, 1, 2\}$, and every variable occurs in only a constant number of such expressions. A direct inspection of (20)–(24) gives $\|Q\|_{\max} \leq 5$ and $\|c\|_\infty \leq 4$. Furthermore, $|\gamma| = C \leq \ell \leq N$. Thus every number appearing in the constructed instance is bounded by a polynomial in the number of variables, and the decision version of Problem (QP) is therefore strongly $\mathcal{NP}$-hard [18].

It remains to establish the treewidth bound. We describe a tree decomposition for a graph that contains the interaction graph of Problem (QP) as a subgraph. Since deleting edges cannot increase treewidth, this suffices even if some quadratic coefficients cancel after expansion. For each $i \in [n]$, define the bag

$$A_i := \{s_{i-1}, s_i, z_{i\ell}\} \setminus \{s_0\}.$$

Arrange these bags in the path $A_1 - A_2 - \cdots - A_n$. For each $i \in [n]$ and $k \in [\ell]$, define the bag

$$P_{ik} := \{z_{i,k-1}, z_{ik}, x_{ik}\} \setminus \{z_{i0}\}.$$

Moreover, for each $i \in [n]$ and $k \in [\ell - 1]$, define the bag

$$R_{ik} := \{z_{ik}, x_{ik}, x_{i,k+1}\}.$$

For each fixed i, arrange these bags in the path $P_{i1} - R_{i1} - P_{i2} - R_{i2} - \cdots - R_{i,\ell-1} - P_{i\ell}$, and attach $P_{i\ell}$ to A_i. Define the bags

$$H_k := \{w_{k-1}, w_k\} \setminus \{w_0\}, \quad \forall k \in [\ell],$$

and arrange these bags in the path $H_1 - H_2 - \cdots - H_\ell$. Finally, introduce the bag $E := \{s_n, w_\ell\}$, and attach E to both A_n and H_ℓ. First, it is simple to check that every variable x, z, s, w appears in at least one of the bags defined above, implying that Property 1 of a tree decomposition is satisfied. Second, the variables corresponding to every quadratic term in Ψ are contained in at least one of these bags, implying that Property 2 of a tree decomposition is satisfied. It remains to verify Property 3. We do so by explicitly identifying, for each variable, all bags in which it occurs. The following cases arise:

(*i*) variables x_{ik}, $i \in [n]$, $k \in [\ell]$: if $k = 1$ and $\ell \geq 2$, then x_{i1} occurs in the bags P_{i1} and R_{i1}, which are adjacent along a path. For $2 \leq k \leq \ell - 1$, the variable x_{ik} occurs in $R_{i,k-1}$, P_{ik} and R_{ik}. These three bags occur consecutively along a path as $R_{i,k-1} - P_{ik} - R_{ik}$. Finally, $x_{i\ell}$ occurs exactly in $R_{i,\ell-1}$ and $P_{i\ell}$, which are adjacent. If $\ell = 1$, then x_{i1} occurs only in P_{i1}, and the claim is immediate.

(*ii*) variables z_{ik}, $i \in [n]$, $k \in [\ell]$: for $1 \leq k \leq \ell - 1$, z_{ik} occurs in the three bags P_{ik}, R_{ik}, and $P_{i,k+1}$, which appear consecutively as $P_{ik} - R_{ik} - P_{i,k+1}$. The variable $z_{i\ell}$ occurs in $P_{i\ell}$ and A_i. By construction, $P_{i\ell}$ is adjacent to A_i.

(*iii*) variables s_i, $i \in [n]$: for $i \in [n-1]$, the variable s_i occurs in the two bags A_i and A_{i+1}, which are adjacent along the path $A_1 - A_2 - \cdots - A_n$. The variable s_n occurs in A_n and E which are adjacent by construction.

(*iv*) variables w_k, $k \in [\ell]$: for $k \in [\ell-1]$, the variable w_k occurs in H_k and H_{k+1}, which are adjacent in the path $H_1 - H_2 - \cdots - H_\ell$. The variable w_ℓ occurs in H_ℓ and E which are adjacent by construction.

Hence property 3 of the tree decomposition holds. Every bag has cardinality at most three, and therefore the interaction graph has treewidth at most two and this completes the proof. □

Remark 2. The machinery in the proof of Theorem 3 is needed to establish *strong* $\mathcal{NP}$-hardness. If one were interested only in $\mathcal{NP}$-hardness, a substantially simpler construction would suffice. Indeed, given a Subset Sum instance $a_1, \dots, a_n, T$, one may set $U := \sum_{i=1}^n a_i$, introduce only variables $x_i, s_i \in [0,1]$, with $s_0 := 0$, and consider the problem of minimizing

$$(Us_n - T)^2 + \sum_{i=1}^n (Us_i - Us_{i-1} - a_i x_i)^2 + \sum_{i=1}^n x_i(1 - x_i).$$

The objective function has minimum value zero if and only if the Subset Sum instance is feasible, and its interaction graph has treewidth at most two, with bags $\{s_{i-1}, s_i, x_i\}$, $i \in [n]$. This construction, however, uses the integers a_i and U directly as coefficients and therefore establishes only weak $\mathcal{NP}$-hardness. The bit-serial construction used in the proof of Theorem 3 replaces these large numerical coefficients by polynomially many local relations with bounded coefficients, thereby transferring the information contained in the binary representations of the input integers from coefficient magnitude to the combinatorial structure of the instance. This idea is closely related to the bounded-coefficient encoding used by Eisenbrand et al. [16] in their study of integer programming (Section 5.2 of [16]).

Cifuentes and Parrilo [6] show that checking the feasibility of quadratic polynomial systems is $\mathcal{NP}$-hard even when the interaction graph is a path, using a reduction from Subset Sum (Example 1 in [6]). Their construction retains the potentially large Subset Sum integers as coefficients and therefore yields only weak $\mathcal{NP}$-hardness. Note that the technique used in the proof of Theorem 3 does not directly strengthen their result while preserving treewidth at most one: each bit-serial relation induces a triangle in the associated graph. Namely, the resulting construction has treewidth two. However, interestingly, the proof technique of Theorem 2 can be used to prove that checking the feasibility of quadratic polynomial systems is strongly $\mathcal{NP}$-hard even when the interaction graph is a path. There is also a close connection between our result and the construction of Bienstock and Muñoz [3]. Their construction yields weak $\mathcal{NP}$-hardness for a polynomial-feasibility problem whose interaction graph has treewidth two. In fact, replacing the large coefficients in the Bienstock–Muñoz construction by the bit-serial encoding used in Theorem 3, we deduce that their polynomial-feasibility problem is strongly $\mathcal{NP}$-hard.

4 Larger treewidths: A polynomial-time solvable class

In this section, we obtain sufficient conditions for the polynomial-time solvability of Problem (QP) for classes of instances whose interaction graphs may have unbounded treewidth. The key idea is that there always exists an optimal solution to Problem (QP) in which every variable with a nonpositive quadratic coefficient is binary. This allows us to separate the problem into a combinatorial part and a continuous part, and to exploit the tractability of the corresponding components together with suitably simple interactions between them. We first introduce some terminology and tools that will be used in the proofs of our results.

We make use of the following celebrated result in discrete optimization that is concerned with tractability of binary optimization problems whose interaction graph has bounded treewidth. In the following, by poly(x, y) we mean a polynomial function in x, y.

Proposition 2 ([8]). *Consider a binary optimization problem*

$$\min_{x \in \{0,1\}^V} \sum_{\ell=1}^m f_\ell(x_{S_\ell}),$$

where $S_\ell \subseteq V$ and $f_\ell : \{0,1\}^{S_\ell} \to \mathbb{Q}$ for every $\ell \in [m]$. Define the interaction graph $G = (V, E)$ by letting $\{i, j\} \in E$ whenever $i, j \in S_\ell$ for some $\ell \in [m]$. If $\mathrm{tw}(G) = \kappa$, then the problem can be solved by dynamic programming in $2^{O(\kappa)}\mathrm{poly}(|V|, m)$ operations, provided that the functions f_ℓ can be evaluated in polynomial time. In particular, if $\kappa \in O(\log \mathrm{poly}(|V|, m))$, then the problem can be solved in polynomial time.

4.1 Hyperplane arrangements and the essential rank

Hyperplane arrangements [38] are a useful geometric tool to partition the space into finitely many polyhedral regions. When the geometry of the hyperplanes is sufficiently low-dimensional, the number of such regions can be polynomial rather than exponential. We use this idea to reduce the number of cases that must be considered in our algorithm. In the following, we provide a brief overview of the results that we need to prove our result.

Let $\mathcal{H} = \{H_1, \ldots, H_M\}$ be a collection of hyperplanes in $\mathbb{R}^d$, where $H_i = \{x \in \mathbb{R}^d : a_i^\top x = b_i\}$ and $a_i \neq 0$. The *arrangement* generated by $\mathcal{H}$, denoted by $\mathcal{A}(\mathcal{H})$, is the collection of connected components of

$$\mathbb{R}^d \setminus \bigcup_{i=1}^{M} H_i.$$

The elements of $\mathcal{A}(\mathcal{H})$ are called the *cells* of the arrangement. For every cell $R \in \mathcal{A}(\mathcal{H})$, the sign of $a_i^\top x - b_i$ is constant on R for every $i \in [M]$. Consequently,

$$R = \bigcap_{i=1}^{M} \{x \in \mathbb{R}^d : \sigma_i(a_i^\top x - b_i) > 0\}$$

for suitable signs $\sigma_i \in \{-1, +1\}$. In particular, every cell is an open polyhedron. Moreover, by Lemma 2.5 in [14] the number of cells satisfies

$$|\mathcal{A}(\mathcal{H})| \leq \sum_{r=0}^{d} \binom{M}{r} \in O(M^d). \tag{29}$$

By Theorem 3.3 in [14], an arrangement $\mathcal{A}(\mathcal{H})$ of hyperplanes in $\mathbb{R}^d$ can be constructed in time $O(M^d)$. This implies that if $d \in O(1)$, then the cells of the arrangement can be computed in polynomial time. The upper bound on the number of cells given by (29) can be improved using the concept of essential rank. The *essential rank* of $\mathcal{H}$, denoted by ρ, is the dimension of the linear space spanned by the normals of the hyperplanes in the arrangement [38]. Define $W := \mathrm{span}\{a_i : i \in [M]\}$. We then have $\rho := \dim(W)$. An arrangement is essential when ρ is equal to the dimension of the ambient space. It then follows that $\mathcal{H}_W := \{H \cap W : H \in \mathcal{H}\}$ is an essential arrangement in W. Stanley [38] shows that the number of cells of $\mathcal{H}$ is equal to the number of cells of $\mathcal{H}_W$. This together with (29) implies that

$$|\mathcal{A}(\mathcal{H})| = |\mathcal{A}(\mathcal{H}_W)| \leq \sum_{r=0}^{\rho} \binom{M}{r} \in O(M^\rho).$$

4.2 A sufficient condition for polynomial-time solvability

Let $G = (V, E)$ be a graph. Recall that for a subset $S \subseteq V$, we denote by G_S the subgraph of G induced by S. Consider a vertex $v \in V$. We define the *neighborhood* of v in G as

$$N_G(v) := \big\{u \in V : \{u, v\} \in E\big\}. \tag{30}$$

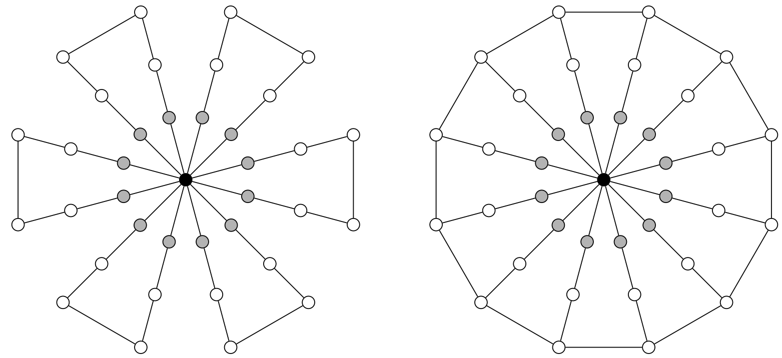

(a)

(b)

Figure 2: Two families of interaction graphs illustrating the application of Theorem 4. In both cases, V^+ consists of a single black vertex, the vertices in $N_G(V^+)$ are depicted in gray, and the remaining vertices are depicted in white. Moreover, we assume that $|N_G(V^+)| \in \Theta(|V|)$. In (a), each connected component D of G_{V^-} is a path. Therefore, $\mathrm{tw}(G_{V^-}) = 1$ and for each connected component D we have $|D \cap N_G(V^+)| = 2$. Thus all the assumptions of Theorem 4 are satisfied. In (b), G_{V^-} is connected with a unique connected component D, and therefore $|D \cap N_G(V^+)| = |N_G(V^+)| \in \Theta(|V|)$, implying that Theorem 4 does not apply. Part (a) also satisfies the assumptions of Corollary 3 because $|V^+| = 1$.

Similarly, for a subset $S \subseteq V$, we define the neighborhood of S in G as:

$$N_G(S) := \Big(\bigcup_{v \in S} N_G(v)\Big) \setminus S. \tag{31}$$

We are now ready to present the main result of this section. Figure 2 demonstrates the application of the theorem.

Theorem 4. *Let $G = (V, E)$ be the interaction graph of Problem* (QP)*. Define $V^+ := \{i \in V : q_{ii} > 0\}$, and $V^- := V \setminus V^+$. Suppose that the following conditions hold:*

1. $\mathrm{tw}(G_{V^-}) \in O(\log |V|)$ *and for every connected component D of G_{V^-}, we have*

$$|D \cap N_G(V^+)| \in O(\log |V|).$$

2. *For every connected component C of G_{V^+} and every $d \in \mathbb{Q}^C$, the optimization problem*

$$\min_{x_C \in [0,1]^C} \Big(x_C^\top Q_C x_C + d^\top x_C\Big)$$

   can be solved in time polynomial in $\mathcal{L} + \langle d \rangle$.
3. $\mathrm{rank}(Q_{V^+, N_G(V^+)}) \in O(1)$.

Then Problem (QP) *can be solved in time polynomial in the input length.*

Proof. It can be checked that there exists an optimal solution x^* of Problem (QP) with $x_i^* \in \{0, 1\}$ for all $i \in V^-$. Therefore, it suffices to solve the following mixed-binary quadratic program:

$$\begin{aligned} \min \quad & x^\top Q x + c^\top x \\ \text{s.t.} \quad & 0 \le x_i \le 1, \quad \forall i \in V^+ \\ & x_i \in \{0, 1\}, \quad \forall i \in V^-. \end{aligned} \tag{32}$$

Denote by $f(x)$ the objective function of Problem (32). Let $D_1, \dots, D_p$ denote the connected components of G_{V^-} and for each $k \in [p]$, define $A_k := D_k \cap N_G(V^+)$ and $F_k := D_k \setminus A_k$. Since $D_1, \dots, D_p$ are the connected components of G_{V^-}, there are no bilinear terms $x_i x_j$ in $f(x)$ with $i \in D_k$, $j \in D_\ell$ with $k \neq \ell \in [p]$. Hence the objective function can be written as:

$$f(x) = f_+(x_{V^+}) + \sum_{k=1}^{p} \Big(g_k(x_{D_k}) + 2 \sum_{i \in A_k} x_i \, Q_{V^+,i}^\top x_{V^+} \Big),$$

where

$$f_+(x_{V^+}) := x_{V^+}^\top Q_{V^+} x_{V^+} + c_{V^+}^\top x_{V^+},$$

and

$$g_k(x_{D_k}) := x_{D_k}^\top Q_{D_k} x_{D_k} + c_{D_k}^\top x_{D_k}.$$

Now fix $x_{V^+} \in [0,1]^{V^+}$. Then minimizing over the binary variables x_{V^-} decomposes over the connected components D_k, $k \in [p]$:

$$\min_{x_{V^-} \in \{0,1\}^{V^-}} f(x_{V^+}, x_{V^-}) = f_+(x_{V^+}) + \sum_{k=1}^{p} \theta_k(x_{V^+}),$$

where

$$\theta_k(x_{V^+}) := \min_{x_{D_k} \in \{0,1\}^{D_k}} \Big(g_k(x_{D_k}) + 2 \sum_{i \in A_k} x_i \, Q_{V^+,i}^\top x_{V^+} \Big).$$

Fix $k \in [p]$. For each $z \in \{0,1\}^{A_k}$, define

$$\mu_k(z) := \min_{w \in \{0,1\}^{D_k \setminus A_k}} g_k(z, w), \tag{33}$$

where $g_k(z,w)$ denotes the value of $g_k(x_{D_k})$ after letting $x_{A_k} = z$ and $x_{D_k \setminus A_k} = w$. More explicitly,

$$g_k(z,w) = w^\top Q_{F_k} w + (c_{F_k} + 2 Q_{F_k, A_k} z)^\top w + z^\top Q_{A_k} z + c_{A_k}^\top z.$$

Therefore, the interaction graph of the binary quadratic optimization problem in (33) is G_{F_k}. Since G_{F_k} is a subgraph of G_{D_k}, from assumption 1 it follows that $\mathrm{tw}(G_{F_k}) \leq \mathrm{tw}(G_{D_k}) \leq \mathrm{tw}(G_{V^-}) \in O(\log(|V|))$. Therefore, by Proposition 2, for every fixed $z \in \{0,1\}^{A_k}$, the value $\mu_k(z)$ can be computed in time polynomial in the input length. Moreover, since by assumption 1, we have $|A_k| \in O(\log |V|)$, all values of $\mu_k(z)$, $z \in \{0,1\}^{A_k}$ can be computed in time polynomial in $|V|$. Finally, since $p \leq |V|$, values of $\mu_k(z)$, $z \in \{0,1\}^{A_k}$ for all $k \in [p]$ can be computed in polynomial time.

For each $z \in \{0,1\}^{A_k}$, define

$$b_k(z) := 2 \sum_{i \in A_k} z_i Q_{V^+,i}. \tag{34}$$

Then we have

$$\theta_k(x_{V^+}) = \min_{z \in \{0,1\}^{A_k}} \left(\mu_k(z) + b_k(z)^\top x_{V^+} \right).$$

Therefore, Problem (32) is equivalent to the following optimization problem:

$$\min_{x_{V^+} \in [0,1]^{V^+}} \left(f_+(x_{V^+}) + \sum_{k=1}^{p} \min_{z \in \{0,1\}^{A_k}} \left(\mu_k(z) + b_k(z)^\top x_{V^+} \right) \right). \tag{35}$$

So far, for $k \in [p]$, we denoted by z a point in $\{0,1\}^{A_k}$. From now on, it will be useful to keep explicitly the dependence on k and we denote it instead by z_k.

Claim 1. *We can compute in time polynomial in $|V|$ a set*

$$Z \subseteq \{0,1\}^{A_1} \times \cdots \times \{0,1\}^{A_p}$$

of cardinality polynomial in $|V|$ such that Problem (35) *is equivalent to the following optimization problem:*

$$\min_{(z_1,\ldots,z_p)\in Z} \left(\min_{x_{V^+}\in[0,1]^{V^+}} \left(f_+(x_{V^+}) + \sum_{k=1}^{p} \left(\mu_k(z_k) + b_k(z_k)^\top x_{V^+} \right) \right) \right). \tag{36}$$

Proof of claim. The function $\theta_k(x_{V^+})$ is the minimum of finitely many affine functions of x_{V^+}. We partition $\mathbb{R}^{V^+}$ into regions over which, for each $k \in [p]$, the minimizing affine function is fixed. Such changes can occur only at points where two affine functions become equal. Therefore, for each $k \in [p]$ and every distinct pair $z_k, z'_k \in \{0,1\}^{A_k}$, we define the hyperplane

$$H_{k,z_k,z'_k} := \left\{ y \in \mathbb{R}^{|V^+|} : \mu_k(z_k) + b_k(z_k)^\top y = \mu_k(z'_k) + b_k(z'_k)^\top y \right\}. \tag{37}$$

If $b_k(z_k) = b_k(z'_k)$ then H_{k,z_k,z'_k} either defines the empty set or is all of $\mathbb{R}^{V^+}$, and therefore does not create a cell boundary. Hence, in the arrangement, we keep only the pairs for which $b_k(z_k) \neq b_k(z'_k)$. Thus, let

$$\mathcal{H} := \left\{ H_{k,z_k,z'_k} : k \in [p],\ z_k, z'_k \in \{0,1\}^{A_k},\ z_k \neq z'_k,\ b_k(z_k) \neq b_k(z'_k) \right\}. \tag{38}$$

Define

$$L := \operatorname{span}\left\{ Q_{V^+,i} : i \in N_G(V^+) \right\} \subseteq \mathbb{R}^{V^+}.$$

Then

$$\dim(L) = \operatorname{rank}(Q_{V^+,N_G(V^+)}) =: r.$$

For any $z_k \neq z'_k \in \{0,1\}^{A_k}$, from (34) it follows that

$$b_k(z_k) - b_k(z'_k) = 2 \sum_{i\in A_k} (z_i - z'_i) Q_{V^+,i}.$$

Since $A_k \subseteq N_G(V^+)$, it follows that $b_k(z_k) - b_k(z'_k) \in L$. Thus, every hyperplane normal in the arrangement belongs to L, implying that the essential rank ρ of $\mathcal{H}$ is upper bounded by r. Therefore, by Lemma 2.5 in [14] and Section 1.1 of [38] it follows that

$$|\mathcal{A}(\mathcal{H})| \leq \sum_{i=0}^{r} \binom{|\mathcal{H}|}{i} \in O(|\mathcal{H}|^r).$$

From (37) and (38), it follows that $|\mathcal{H}| \leq \sum_{k=1}^{p} \binom{2^{|A_k|}}{2}$. Since by assumption 1, we have $|A_k| \in O(\log|V|)$ for all k, and $p \leq |V|$, we deduce that $|\mathcal{H}| \in |V|^{O(1)}$. Moreover, by assumption 3, we have $r \in O(1)$. Therefore, the number of cells of $\mathcal{A}(\mathcal{H})$ is polynomial in $|V|$. Let W denote the linear space spanned by the normals of the hyperplanes in $\mathcal{H}$. Then $\dim(W) = \rho \leq r$. By the essentialization argument given above, the cells of $\mathcal{A}(\mathcal{H})$ are in one-to-one correspondence with the cells of the essentialized arrangement $\mathcal{H}_W := \{H \cap W : H \in \mathcal{H}\}$. Since $\rho \in O(1)$ and $|\mathcal{H}| \in |V|^{O(1)}$, by Theorem 3.3 in [14], applied to the arrangement $\mathcal{H}_W$ in the ρ-dimensional space W, these cells can be computed in time polynomial in $|V|$.

Next, fix one cell $R \in \mathcal{A}(\mathcal{H})$. By definition of the arrangement, for every $k \in [p]$, the ordering of the affine functions $\mu_k(z_k) + b_k(z_k)^\top x_{V^+}$, $z \in \{0,1\}^{A_k}$, is constant on R. Hence, for every $k \in [p]$, there exists $\tilde{z}_k \in \{0,1\}^{A_k}$ such that

$$\min_{z \in \{0,1\}^{A_k}} \left(\mu_k(z_k) + b_k(z_k)^\top x_{V^+}\right) = \mu_k(\tilde{z}_k) + b_k(\tilde{z}_k)^\top x_{V^+} \qquad \forall x_{V^+} \in R.$$

For each cell $R \in \mathcal{A}(\mathcal{H})$ we then add $(\tilde{z}_1, \ldots \tilde{z}_p) \in \{0,1\}^{A_1} \times \cdots \times \{0,1\}^{A_p}$ to the set Z. Since $|Z| \leq |\mathcal{A}(\mathcal{H})|$, it follows that $|Z| \in |V|^{O(1)}$ and since $\rho \in O(1)$, the cells and hence Z can be computed in polynomial time. We now show that this construction of Z yields the equivalence between Problems (35) and (36). Fix $x_{V^+} \in [0,1]^{V^+}$. Define

$$\varphi(x_{V^+}) := f_+(x_{V^+}) + \sum_{k=1}^{p} \theta_k(x_{V^+}),$$

and

$$\psi_z(x_{V^+}) := f_+(x_{V^+}) + \sum_{k=1}^{p} \left(\mu_k(z_k) + b_k(z_k)^\top x_{V^+}\right),$$

for $z = (z_1, \ldots, z_p) \in Z$, so that φ is the objective value of Problem (35) and ψ_z, for $z \in Z$, is the inner objective value of Problem (36). For every $k \in [p]$ and every $z_k \in \{0,1\}^{A_k}$, by definition of θ_k we have

$$\mu_k(z_k) + b_k(z_k)^\top x_{V^+} \geq \theta_k(x_{V^+}).$$

Summing over k and adding $f_+(x_{V^+})$ gives $\psi_z(x_{V^+}) \geq \varphi(x_{V^+})$ for all $z \in Z$. Denote by $\overline{R}$ the closure of the cell R. Since

$$[0,1]^{V^+} \subseteq \bigcup_{R \in \mathcal{A}(\mathcal{H})} \overline{R},$$

there is a cell R with $x_{V^+} \in \overline{R}$. Let $\tilde{z} = (\tilde{z}_1, \ldots, \tilde{z}_p) \in Z$ be the tuple added to Z for this cell R, which by construction satisfies $\theta_k(y) = \mu_k(\tilde{z}_k) + b_k(\tilde{z}_k)^\top y$ for all $y \in R$ and for all $k \in [p]$. Both sides of this identity are continuous functions of y, so the identity extends from R to its closure $\overline{R}$, and in particular it holds at $y = x_{V^+}$. Hence $\varphi(x_{V^+}) = \psi_{\tilde{z}}(x_{V^+})$. Therefore,

$$\varphi(x_{V^+}) = \min_{z \in Z} \psi_z(x_{V^+}) \qquad \forall\, x_{V^+} \in [0,1]^{V^+}. \tag{39}$$

Taking the minimum with respect to $x_{V^+} \in [0,1]^{V^+}$ on both sides, and using that Z is finite so that the minimum over x_{V^+} and the minimum over $z \in Z$ can be interchanged, we obtain

$$\min_{x_{V^+} \in [0,1]^{V^+}} \varphi(x_{V^+}) = \min_{x_{V^+} \in [0,1]^{V^+}} \min_{z \in Z} \psi_z(x_{V^+}) = \min_{z \in Z} \min_{x_{V^+} \in [0,1]^{V^+}} \psi_z(x_{V^+}),$$

which is exactly the statement that Problems (35) and (36) are equivalent. $\diamond$

Therefore, for every $z = (z_1, \ldots, z_p) \in Z$, we solve the inner box-constrained quadratic program in (36), namely:

$$\min_{x_{V^+} \in [0,1]^{V^+}} \left(f_+(x_{V^+}) + \sum_{k=1}^{p} \left(\mu_k(z_k) + b_k(z_k)^\top x_{V^+}\right)\right), \tag{40}$$

and we denote by f_z^* its optimal value. For fixed $z \in Z$, define

$$d_z := c_{V^+} + \sum_{k=1}^{p} b_k(z_k), \qquad e_z := \sum_{k=1}^{p} \mu_k(z_k).$$

Then Problem (40) can be written as

$$e_z + \min_{x_{V^+} \in [0,1]^{V^+}} \left(x_{V^+}^\top Q_{V^+} x_{V^+} + d_z^\top x_{V^+} \right).$$

Let $C_1, \ldots, C_s$ denote the connected components of G_{V^+}. Since there are no bilinear terms $x_i x_{i'}$ with $i \in C_j$ and $i' \in C_\ell$ for $j \neq \ell$, we have

$$x_{V^+}^\top Q_{V^+} x_{V^+} + d_z^\top x_{V^+} = \sum_{j=1}^{s} \left(x_{C_j}^\top Q_{C_j} x_{C_j} + (d_z)_{C_j}^\top x_{C_j} \right).$$

Therefore, the optimal value f_z^* of Problem (40) is given by

$$f_z^* = e_z + \sum_{j=1}^{s} \min_{x_{C_j} \in [0,1]^{C_j}} \left(x_{C_j}^\top Q_{C_j} x_{C_j} + (d_z)_{C_j}^\top x_{C_j} \right).$$

The length of d_z is polynomial in the input length, since d_z is obtained by summing polynomially many rational vectors whose entries are coefficients of the objective function. Hence, by assumption 2, each of the above box-constrained quadratic programs can be solved in time polynomial in the input length. Since $s \leq |V|$, it follows that f_z^* can be computed in polynomial time for every $z \in Z$. The optimal value f^* of Problem (35) is then given by

$$f^* = \min_{z \in Z} f_z^*.$$

Since $|Z| \in |V|^{O(1)}$, we conclude that Problem (QP) can be solved in polynomial time in the input length. □

The three assumptions of Theorem 4 control three different sources of complexity in Problem (QP). Assumption 1 concerns the *combinatorial* part of the problem. The subgraph induced by the binary variables has logarithmic treewidth, and each of its connected components D interacts with the continuous variables through only $|D \cap N_G(V^+)| \in O(\log |V|)$ binary variables. Consequently, after fixing these interface variables, the remaining binary subproblem on each component can be solved in polynomial time. Assumption 2 concerns the *continuous* part. Once the binary choices are fixed, the problem over V^+ decomposes over the connected components C of G_{V^+}, and the box-constrained quadratic problem associated with each such component must remain polynomial-time solvable under an arbitrary perturbation of its linear term. Finally, assumption 3 controls the *interaction* between the binary and continuous parts. Although the binary variables may interact with many variables in V^+, the vectors describing these interactions span a space of constant dimension. This low-dimensional coupling is what allows the exponentially many possible combinations of binary choices to be reduced, through the hyperplane-arrangement argument, to polynomially many relevant combinations. Assumptions 1 and 3 of Theorem 4 are structural conditions that can be verified directly from the interaction graph and the matrix Q, respectively. In contrast, assumption 2 is less explicit, as it requires polynomial-time solvability of a family of box-constrained quadratic programs for every choice of the linear term. A simple sufficient condition for assumption 2 is $Q_C \succeq 0$ for every connected component C of G_{V^+}, in which case each of these problems is convex. We next identify additional sufficient conditions under which assumption 2 holds. To this end, we use the following result from [22].

Lemma 20 ([22])**.** *Consider Problem* (QP) *and let* $\kappa \in \{0, \dots, n-1\}$ *be the smallest integer such that every* $(n-\kappa) \times (n-\kappa)$ *principal submatrix of* Q *is positive semidefinite. If no such integer exists, set* $\kappa := n$*. Then Problem* (QP) *can be solved in time*

$$O\Big((\kappa+1)\Big(\frac{2en}{\kappa}\Big)^{\kappa}\mathrm{poly}(\mathcal{L})\Big) \tag{41}$$

if $\kappa \geq 1$*, and in time* $\mathrm{poly}(\mathcal{L})$ *if* $\kappa = 0$*.*

Notice that $\kappa = 0$ in Lemma 20 corresponds to the convex case, while if $\kappa \geq 1$, then Problem (QP) is not convex. Lemma 20 implies that if $\kappa \in O(1)$, then Problem (QP) can be solved in polynomial time. We then obtain the following corollary of Theorem 4.

Corollary 1. *Let* $G = (V, E)$ *be the interaction graph of Problem* (QP)*. Define* $V^+ := \{i \in V : q_{ii} > 0\}$*. Let* $C_1, \dots, C_s$ *denote the connected components of* G_{V^+}*. For each* $j \in [s]$*, let* κ_j *be the parameter defined in Lemma 20 for the matrix* Q_{C_j}*. Suppose that assumptions 1 and 3 of Theorem 4 are satisfied. Moreover, assume that for each* $j \in [s]$*, at least one of the following conditions holds:*

(a) $\mathrm{rank}(Q_{C_j}) \in O(1)$
(b) G_{C_j} *is a tree*
(c) $\kappa_j \log\left(1 + \frac{|C_j|}{\kappa_j}\right) \in O(\log|V|)$*, where we define* $0\log(1+\frac{|C_j|}{0}) := 0$*.*

Then Problem (QP) *can be solved in time polynomial in the input length.*

Proof. To prove the statement, it suffices to verify assumption 2 of Theorem 4. Fix $j \in [s]$ and an arbitrary $d \in \mathbb{Q}^{C_j}$, and consider the problem

$$\min_{x_{C_j} \in [0,1]^{C_j}} \left(x_{C_j}^\top Q_{C_j} x_{C_j} + d^\top x_{C_j}\right). \tag{42}$$

If part (a) holds; i.e., $\mathrm{rank}(Q_{C_j}) \in O(1)$, then by Theorem 1 in [20], Problem (42) can be solved in polynomial time. If part (b) holds; i.e., G_{C_j} is a tree, then by Theorem 1, Problem (42) can be solved in polynomial time. Suppose now that condition (c) holds. Notice that κ_j depends only on the quadratic matrix Q_{C_j} and is therefore unchanged when the linear coefficient is replaced by the arbitrary vector d. If $\kappa_j = 0$, Lemma 20 implies that Problem (42) can be solved in time polynomial in $\mathcal{L} + \langle d\rangle$. Suppose that $\kappa_j \geq 1$. By Lemma 20, Problem (42) can be solved in time

$$O\left((\kappa_j+1)\left(\frac{2e|C_j|}{\kappa_j}\right)^{\kappa_j}\mathrm{poly}(\mathcal{L} + \langle d\rangle)\right).$$

Since $\kappa_j \leq |C_j|$, we have $\kappa_j + 1 \leq 2^{\kappa_j} \leq \left(1 + \frac{|C_j|}{\kappa_j}\right)^{\kappa_j}$. Moreover, setting $t := |C_j|/\kappa_j \geq 1$, we have $2et \leq 8t \leq (1+t)^3$, and hence $\left(\frac{2e|C_j|}{\kappa_j}\right)^{\kappa_j} \leq \left(1 + \frac{|C_j|}{\kappa_j}\right)^{3\kappa_j}$. Therefore,

$$(\kappa_j+1)\left(\frac{2e|C_j|}{\kappa_j}\right)^{\kappa_j} \leq \left(1 + \frac{|C_j|}{\kappa_j}\right)^{4\kappa_j}.$$

By condition (c), $\left(1 + \frac{|C_j|}{\kappa_j}\right)^{4\kappa_j} \in |V|^{O(1)}$. It then follows that for every $j \in [s]$ and every $d \in \mathbb{Q}^{C_j}$, Problem (42) can be solved in time polynomial in $\mathcal{L} + \langle d\rangle$, implying that assumption 2 of Theorem 4 is satisfied and this completes the proof. $\square$

Consider condition (c) of Corollary 1. As detailed in [22], this condition has a natural interpretation in terms of the value of κ_j relative to the size of C_j. In particular, condition (c) is automatically satisfied if $|C_j| \in O(\log |V|)$, regardless of the value of κ_j, and it is also automatically satisfied if $\kappa_j \in O(1)$ regardless of the size of $|C_j|$.

Clearly, assumption 3 of Theorem 4 depends on the matrix Q. We next obtain a sufficient condition that is independent of the coefficients of the objective function. Let H be the bipartite graph with bipartition V^+ and $N_G(V^+)$ and edge set

$$E(H) := \big\{\{i,j\} \in E : i \in V^+,\ j \in N_G(V^+)\big\}. \tag{43}$$

Thus, the edges of H correspond to the nonzero entries of $Q_{V^+,N_G(V^+)}$. Let $\nu(H)$ denote the cardinality of a maximum matching of H. By Theorem 1 in [15], we have

$$\operatorname{rank}(Q_{V^+,N_G(V^+)}) \leq \nu(H).$$

Since H is bipartite, König's theorem (see Theorem 2.1.1 in [13]) implies $\nu(H) = \tau(H)$, where $\tau(H)$ denotes the minimum cardinality of a vertex cover of H. Therefore, we obtain the following result:

Corollary 2. *Let $G = (V, E)$ be the interaction graph of Problem* (QP). *Define $V^+ := \{i \in V : q_{ii} > 0\}$. Let H be the bipartite graph with bipartition V^+ and $N_G(V^+)$ and with the edge set defined by* (43). *Suppose that assumptions 1 and 2 of Theorem 4 are satisfied. Moreover, suppose that the minimum cardinality $\tau(H)$ of a vertex cover of H satisfies $\tau(H) \in O(1)$. Then Problem* (QP) *can be solved in time polynomial in the input length.*

By construction we have

$$\operatorname{rank}(Q_{V^+,N_G(V^+)}) \leq |V^+|.$$

Therefore, in the special case where $|V^+| \in O(1)$, the assumptions of Theorem 4 can be simplified to obtain the following result:

Corollary 3. *Let $G = (V, E)$ be the interaction graph of Problem* (QP). *Define $V^+ := \{i \in V : q_{ii} > 0\}$. Suppose that $|V^+| \in O(1)$ and that assumption 1 of Theorem 4 is satisfied. Then Problem* (QP) *can be solved in time polynomial in the input length.*

Proof. It suffices to show that assumptions 2 and 3 of Theorem 4 are satisfied. Since $|V^+| \in O(1)$, every connected component C of G_{V^+} has constant size. Hence, by Lemma 20 the box-constrained quadratic problem associated with C can be solved in polynomial time for every linear coefficient, and assumption 2 of Theorem 4 is satisfied. Finally, $\operatorname{rank}(Q_{V^+,N_G(V^+)}) \leq |V^+| \in O(1)$, so assumption 3 is also satisfied. The result then follows. $\square$

The following results are concerned with special cases where the vertex sets V^+ and V^- are highly structured.

Corollary 4. *Let $G = (V, E)$ be the interaction graph of Problem* (QP). *Suppose that $V^+ := \{i \in V : q_{ii} > 0\}$ is a vertex cover of G. Suppose that assumptions 2 and 3 of Theorem 4 are satisfied. Then Problem* (QP) *can be solved in time polynomial in the input length.*

Proof. It suffices to show assumption 1 of Theorem 4 holds. Since V^+ is a vertex cover, $V^- := V \setminus V^+$ is a stable set. Hence, every connected component D of G_{V^-} consists of a single vertex, implying that $\operatorname{tw}(G_{V^-}) = 0$ and $|D \cap N_G(V^+)| \leq 1$ for every connected component D of G_{V^-}. $\square$

Corollary 5. *Let $G = (V, E)$ be the interaction graph of Problem* (QP)*. Suppose that $V^+ := \{i \in V : q_{ii} > 0\}$ is a stable set of G. Suppose that assumptions 1 and 3 of Theorem 4 are satisfied. Then Problem* (QP) *can be solved in time polynomial in the input length.*

Proof. It suffices to show that assumption 2 of Theorem 4 is satisfied. Since V^+ is a stable set, every connected component C of G_{V^+} consists of a single vertex. The result then follows. $\square$

Consider a bipartite graph G with bipartition V^+, V^-. Then both V^+ and V^- are stable sets, and every connected component of both G_{V^+} and G_{V^-} consists of a single vertex. Therefore, the following result is immediate from Corollary 1.

Corollary 6. *Let $G = (V, E)$ be the interaction graph of Problem* (QP)*. Define $V^+ := \{i \in V : q_{ii} > 0\}$ and $V^- := V \setminus V^+$. Suppose that G is bipartite with bipartition V^+, V^- and* $\mathrm{rank}(Q_{V^+, N_G(V^+)}) \in O(1)$*. Then Problem* (QP) *can be solved in time polynomial in the input length.*

The next result considers the case where every vertex in V^- has degree at most two in G.

Corollary 7. *Let $G = (V, E)$ be the interaction graph of Problem* (QP)*. Define $V^+ := \{i \in V : q_{ii} > 0\}$ and $V^- := V \setminus V^+$. Suppose that every vertex in V^- has degree at most two in G. Moreover, assume that assumptions 2 and 3 of Theorem 4 are satisfied. Then Problem* (QP) *can be solved in time polynomial in the input length.*

Proof. It suffices to show that assumption 1 of Theorem 4 is satisfied. Let $D_1, \ldots, D_p$ denote the connected components of G_{V^-} and as before, for each $k \in [p]$, define $A_k := D_k \cap N_G(V^+)$. Since every vertex in V^- has degree at most two in G, every connected component D_k of G_{V^-} is either a path, a cycle, or an isolated vertex. If D_k is a cycle, then every vertex of D_k has two neighbors in D_k. Therefore, no vertex of D_k can be adjacent to a vertex in V^+, since otherwise its degree in G would be at least three. Hence, in this case $A_k = \emptyset$. If D_k is a path, then only the endpoints of the path can be adjacent to vertices in V^+. Indeed, every internal vertex of the path already has two neighbors in D_k and therefore cannot have an additional neighbor in V^+. Hence, in this case $|A_k| \leq 2$. Finally, if D_k is an isolated vertex, then $|A_k| \leq 1$. Therefore, $|A_k| \leq 2$ for all $k \in [p]$. Moreover, since every D_k is a path, a cycle, or an isolated vertex, we have $\mathrm{tw}(G_{V^-}) \leq 2$. Therefore, assumption 1 of Theorem 4 is satisfied and this completes the proof. $\square$

The following example demonstrated that our sufficient conditions establish polynomial-time solvability of some class of nonconvex box-constrained quadratic programs for which the Hessian is full-row rank and the treewidth of the interaction graph is unbounded.

Example 1. For each $m \geq 1$, consider an instance of Problem (QP) with

$$Q = \begin{pmatrix} I_m & J_m \\ J_m & -I_m \end{pmatrix},$$

where I_m is an $m \times m$ identity matrix, J_m is an $m \times m$ matrix of all ones, and c is an arbitrary rational vector. It can be checked that Q is indefinite and has full-row rank $2m$. The interaction graph in this case is the complete bipartite graph $K_{m,m}$ with bipartition V^+ and V^-. We then have $\mathrm{tw}(K_{m,m}) = m$. Moreover, $Q_{V^+, V^-} = J_m$, implying that $\mathrm{rank}(Q_{V^+, V^-}) = 1$. Therefore, all assumptions of Corollary 6 are satisfied and hence, Problem (QP) can be solved in polynomial time in the input length.

4.3 A more general scheme

In [22], the author obtains a sufficient condition under which Problem (QP) can be solved in polynomial time. Next we state this result in a slightly more general form. The proof however follows directly from the proof of Theorem 1 in [22].

Theorem 5 ([22])**.** *Let $G = (V, E)$ be the interaction graph of Problem* (QP)*. Define $V^+ := \{i \in V : q_{ii} > 0\}$ and $V^- := V \setminus V^+$. Let $C_1, \dots, C_s$ denote the connected components of G_{V^+}. Let $\bar{G}$ be the graph on V^- obtained from G_{V^-} by making $N_G(C_j)$ a clique for every $j \in [s]$. Suppose that the following conditions hold:*

1. $\mathrm{tw}(\bar{G}) \in O(\log |V|)$.
2. *For each $j \in [s]$ and every $d \in \mathbb{Q}^{C_j}$, problem* (42) *can be solved in time polynomial in $\mathcal{L} + \langle d \rangle$.*

Then Problem (QP) *can be solved in time polynomial in the input length.*

Remark 3. Theorems 4 and 5 are incomparable. First, assumption 2 of Theorem 4 and that of Theorem 5 are identical. Second, assumption 1 of Theorem 5 implies that $|N_G(C_j)| \in O(\log |V|)$ for all $j \in [s]$. To see this, observe that each $N_G(C_j)$ is a clique in $\bar{G}$ and hence we have $|N_G(C_j)| \leq \mathrm{tw}(\bar{G}){+}1 \in O(\log |V|)$. That is, all continuous components have *small neighborhoods*, an assumption that does not need to hold in Theorem 4. To see this, for each $m \geq 1$, let $V^+ = \{u\}$ and $V^- = \{v_1, \dots, v_m\}$. Assume that V^- is a stable set and u is adjacent to every vertex in V^-. The connected components of G_{V^-} are the singletons $\{v_i\}$, and hence $\mathrm{tw}(G_{V^-}) = 0$ and $|D \cap N_G(V^+)| = 1$ for every connected component D of G_{V^-}. The unique connected component of G_{V^+} is the singleton $\{u\}$, and $\mathrm{rank}(Q_{V^+, N_G(V^+)}) \leq 1$. Thus, all assumptions of Theorem 4 are satisfied. Now consider Theorem 5. We have $N_G(V^+) = N_G(u) = V^-$, implying that $\bar{G} = K_m$. Therefore, $\mathrm{tw}(\bar{G}) = m - 1$, so assumption 1 of Theorem 5 is violated. Consequently, Theorem 4 does not follow from Theorem 5. Conversely, for each $m \geq 1$, let $V^+ := \{p_1, \dots, p_m\}$ and $V^- := \{z_1, \dots, z_{m+1}\}$. Let G_{V^-} be the path $z_1 - z_2 - \cdots - z_{m+1}$, let V^+ be a stable set, and let p_i be adjacent to z_i and z_{i+1} for every $i \in [m]$. The connected components of G_{V^+} are the singletons $C_i := \{p_i\}$, and $N_G(C_i) = \{z_i, z_{i+1}\}$. Since $\{z_i, z_{i+1}\}$ is already an edge of G_{V^-}, making $N_G(C_i)$ a clique adds no edge to G_{V^-}. Hence $\bar{G} = G_{V^-}$, and therefore $\mathrm{tw}(\bar{G}) = 1$. Moreover, each C_i is a singleton, so assumption 2 of Theorem 5 is satisfied. Thus, Theorem 5 applies. On the other hand, G_{V^-} is connected, so it has a single connected component $D = V^-$. Moreover, every vertex of V^- is adjacent to a vertex in V^+, and hence $D \cap N_G(V^+) = V^-$. Therefore, $|D \cap N_G(V^+)| = m + 1 = \Theta(|V|)$, and assumption 1 of Theorem 4 is violated. Consequently, Theorem 5 does not follow from Theorem 4.

The proof techniques of Theorems 4 and 5 are fundamentally different. To prove Theorem 4, we first eliminate each binary component $D_k \setminus A_k$, $k \in [p]$ and subsequently add a clique on A_k. This is a polynomial-time operation thanks to the two key facts in assumption 1. In contrast, to prove Theorem 5, the author first eliminates each continuous component C_j, $j \in [s]$ and subsequently adds a clique $N_G(C_j)$ to graph G. This step can be performed in polynomial time, because from assumption 1 of Theorem 5, it follows that $|N_G(C_j)| \in O(\log |V|)$ for all $j \in [s]$. The incomparability of Theorems 4 and 5 suggests combining the two elimination schemes. Namely, for an instance of Problem (QP) that does not satisfy the assumptions of either theorem, we first eliminate a subset of continuous components C_j of G_{V^+} with $|N_G(C_j)| \in O(\log |V|)$ using the technique of Theorem 5 and then apply Theorem 4 to the remaining graph.

Theorem 6. *Let $G = (V, E)$ be the interaction graph of Problem* (QP)*. Define $V^+ := \{i \in V : q_{ii} > 0\}$ and $V^- := V \setminus V^+$. Let $C_1, \dots, C_s$ denote the connected components of G_{V^+}. Suppose that the following conditions hold:*

1. *There exists a subset $J \subseteq [s]$ such that, letting $R := V^+ \setminus \bigcup_{j \in J} C_j$ and letting $\widehat{G}$ be the graph on V^- obtained from G_{V^-} by making $N_G(C_j)$ a clique for every $j \in J$, we have:*

   (i) $\text{tw}(\widehat{G}) \in O(\log |V|)$*, and* $|D \cap N_G(R)| \in O(\log |V|)$ *for every connected component D of $\widehat{G}$*

   (ii) $\text{rank}\big(Q_{R,N_G(R)}\big) \in O(1)$*.*

2. *For every $j \in [s]$ and every $d \in \mathbb{Q}^{C_j}$, the optimization problem* (42) *can be solved in time polynomial in $\mathcal{L} + \langle d \rangle$.*

Then Problem (QP) *can be solved in time polynomial in the input length.*

Proof. We first eliminate the continuous components C_j, $j \in J$, using the proof technique of Theorem 1 in [22]. Since $N_G(C_j)$ is a clique in $\widehat{G}$, assumption 1 implies that $|N_G(C_j)| \in O(\log |V|)$ for every $j \in J$. Therefore, by assumption 2, all such components can be eliminated in polynomial time. We are then left with an optimization problem of the form

$$\min_{\substack{x_R \in [0,1]^R \\ x_{V^-} \in \{0,1\}^{V^-}}} \Big\{ x_R^\top Q_R x_R + c_R^\top x_R + g(x_{V^-}) + 2 \sum_{i \in N_G(R)} x_i Q_{R,i}^\top x_R \Big\},$$

where g is a binary objective function whose interaction graph is $\widehat{G}$. We now apply the proof technique of Theorem 4, with R in place of V^+. The only difference is that the binary objective $g(x_{V^-})$ need not be quadratic. However, when the variables in $D \cap N_G(R)$ are fixed, for a connected component D of $\widehat{G}$, the remaining binary subproblem has interaction graph contained in $\widehat{G}_D$. Hence, by assumption 1 and Proposition 2, each such subproblem can still be solved in polynomial time. Since $|D \cap N_G(R)| \in O(\log |V|)$, all required values can be computed in polynomial time. The remainder of the proof of Theorem 4 applies unchanged: condition 1(ii) yields the required constant-rank bound for the hyperplane-arrangement argument, while assumption 2 guarantees polynomial-time solvability of the remaining continuous components under arbitrary linear perturbations. Therefore, Problem (QP) can be solved in polynomial time. □

Theorems 4 and 5 are special cases of Theorem 6. Namely, letting $J = \emptyset$, we get $\widehat{G} = G_{V^-}$ and $R = V^+$; therefore, Theorem 6 reduces to Theorem 4. Alternatively, letting $J = [s]$, we get $R = \emptyset$ and $\widehat{G} = \bar{G}$. So assumption (i) reduces to $\text{tw}(\bar{G}) \in O(\log |V|)$ and assumption (ii) becomes vacuous. We thus recover Theorem 5. Notice that the key to the power of Theorem 6 is a careful selection of the subset J. A simple rule is to let J consist of the index set of all continuous components with a logarithmic-size neighborhood. The following example shows that this rule is not sufficient. Namely, in general, among the continuous components having logarithmic-size neighborhoods, it may be necessary to eliminate only a proper subset in the first phase.

Example 2. Let $G = (V, E)$ be the interaction graph of Problem (QP). For each integer $m \geq 4$, let $\ell := \lceil \log_2 m \rceil$. Let $W := \{w_1, \ldots, w_m\}$ and suppose that G_W is the path $w_1 - w_2 - \cdots - w_m$. Moreover, for each $i \in [\ell]$, let $D_i := \{v_{i1}, \ldots, v_{i\ell}\}$ and suppose that G_{D_i} is the path $v_{i1} - v_{i2} - \cdots - v_{i\ell}$. Assume that $W, D_1, \ldots, D_\ell$ are pairwise disjoint and define $V^- := V \setminus V^+ = W \cup D_1 \cup \cdots \cup D_\ell$. Moreover, suppose that $W, D_1, \ldots, D_\ell$ are the connected components of G_{V^-}. Define $V^+ := \{u, c\} \cup \{b_1, \ldots, b_m\}$, and suppose that V^+ is a stable set. In addition to the edges of G_{V^-} specified above, suppose that G contains exactly the following edges: $\{b_j, w_j\}$ for every $j \in [m]$, $\{u, w_1\}$, $\{u, v_{ij}\}$ for every $i, j \in [\ell]$, and $\{c, v_{i1}\}$ for every $i \in [\ell]$. We then have $|V| = 2m + \ell^2 + 2 = \Theta(m)$, implying that $\ell = \Theta(\log |V|)$, while $m \notin O(\log |V|)$. We first show that neither Theorem 4 nor Theorem 5 applies to this instance. Consider first Theorem 4. Recall that $W = \{w_1, \ldots, w_m\}$ is a connected component of G_{V^-} and every w_j is adjacent to $b_j \in V^+$. Hence $W \cap N_G(V^+) = W$, and therefore

$|W \cap N_G(V^+)| = m \notin O(\log|V|)$. Thus assumption 1 of Theorem 4 is violated. Next consider Theorem 5; By construction, $N_G(u) = \{w_1\} \cup D_1 \cup \cdots \cup D_\ell$, and hence $|N_G(u)| = 1 + \ell^2$. In the graph $\bar{G}$, the set $N_G(u)$ is a clique. Therefore, $\mathrm{tw}(\bar{G}) \geq |N_G(u)| - 1 = \ell^2$. Since $\ell = \Theta(\log|V|)$, we have $\ell^2 \notin O(\log|V|)$, and hence Theorem 5 does not apply.

We now consider the strategy of first eliminating every connected component C of G_{V^+} satisfying $|N_G(C)| \in O(\log|V|)$ and then applying Theorem 4. Since V^+ is a stable set, its connected components are $\{u\}, \{c\}, \{b_1\}, \ldots, \{b_m\}$. For every $j \in [m]$, $N_G(b_j) = \{w_j\}$, while $N_G(c) = \{v_{11}, v_{21}, \ldots, v_{\ell 1}\}$. Hence $|N_G(b_j)| = 1$ and $|N_G(c)| = \ell = O(\log|V|)$. On the other hand, $|N_G(u)| = 1 + \ell^2 \notin O(\log|V|)$. Thus this strategy eliminates $\{b_j\}$, $j \in [m]$, together with $\{c\}$, while retaining $\{u\}$. Eliminating $\{b_j\}$ creates no interaction between distinct binary variables, since $N_G(b_j) = \{w_j\}$. We next consider the elimination of c. In this case, $\{v_{11}, \ldots, v_{\ell 1}\}$ becomes a clique in the resulting graph $\bar{G}$ and therefore the components $D_1, \ldots, D_\ell$ are joined into a single connected component $D^* := D_1 \cup \cdots \cup D_\ell$. Since u is adjacent to every vertex in D^*, we have $|D^* \cap N_G(u)| = |D^*| = \ell^2 = \Theta((\log|V|)^2)$. Thus the resulting instance still violates assumption 1 of Theorem 4.

Finally, consider eliminating only the collection $\mathcal{E} := \{\{b_j\} : j \in [m]\}$. Since each b_j has the single neighbor w_j, eliminating these components creates no interaction between distinct binary variables. Hence the connected components of the resulting binary interaction graph remain $W, D_1, \ldots, D_\ell$. We then have $|W \cap N_G(\{u, c\})| = 1$ and $|D_i \cap N_G(\{u, c\})| = |D_i| = \ell = O(\log|V|)$ for every $i \in [\ell]$. Moreover, G_W and G_{D_i}, $i \in [\ell]$ are paths, and hence $\mathrm{tw}(G_{V^-}) = 1$. Thus assumption 1 of Theorem 4 is satisfied. The retained positive components $\{u\}$ and $\{c\}$ are singletons, so assumption 2 of Theorem 4 is satisfied. Finally, $\mathrm{rank}\left(Q_{\{u,c\}, N_G(\{u,c\})}\right) \leq 2$, and hence assumption 3 is satisfied as well. Thus, eliminating all continuous components having logarithmic-size neighborhoods does not suffice, whereas eliminating the proper subset $\mathcal{E}$ and then applying Theorem 4 yields a polynomial-time algorithm for solving Problem (QP).

The next proposition indicates that the subset J whose existence is required in condition 1 of Theorem 6 can be constructed in polynomial time.

Proposition 3. *Under the notation of Theorem 6, whenever condition 1 of that theorem holds, a subset $J \subseteq [s]$ satisfying condition 1 can be found in time polynomial in the input length.*

Proof. Let $n := |V|$. Since condition 1 of Theorem 6 holds, there exist constants $\alpha, \beta > 0$ and $r \in \mathbb{Z}_{\geq 0}$, independent of n, and a subset $J^* \subseteq [s]$ such that, with $R^* := V^+ \setminus \cup_{j \in J^*} C_j$ and with $\widehat{G}^*$ defined as in Theorem 6, we have $\mathrm{tw}(\widehat{G}^*) \leq \alpha \log n$, $|D \cap N_G(R^*)| \leq \beta \log n$ for every connected component D of $\widehat{G}^*$, and $\mathrm{rank}(Q_{R^*, N_G(R^*)}) \leq r$. Set $k := \lceil \alpha \log n \rceil$ and $h := \lceil \beta \log n \rceil$. For every $j \in [s]$, define

$$L_j := \mathrm{span}\{Q_{v,V^-} : v \in C_j\}.$$

Moreover, for every $I \subseteq [s]$ with $|I| \leq r$, define

$$L_I := \mathrm{span}\Big\{Q_{v,V^-} : v \in \bigcup_{i \in I} C_i\Big\},$$

where we define $L_\emptyset := \{0\}$. We consider only subsets I for which $\dim(L_I) \leq r$. For each such I, define

$$J_I := \{j \in [s] : L_j \not\subseteq L_I\},$$

and let

$$R_I := V^+ \setminus \bigcup_{j \in J_I} C_j.$$

Let $\widehat{G}_I$ be the graph on V^- obtained from G_{V^-} by making $N_G(C_j)$ a clique for every $j \in J_I$. By the definition of J_I, if $C_j \subseteq R_I$, then $L_j \subseteq L_I$. Hence $\operatorname{span}\{Q_{v,V^-} : v \in R_I\} \subseteq L_I$. Conversely, for every $i \in I$ we have $L_i \subseteq L_I$, and therefore $i \notin J_I$. Thus $C_i \subseteq R_I$ for every $i \in I$, which implies $L_I \subseteq \operatorname{span}\{Q_{v,V^-} : v \in R_I\}$. Therefore,

$$\operatorname{span}\{Q_{v,V^-} : v \in R_I\} = L_I. \tag{44}$$

Since R_I is a union of connected components of G_{V^+}, there are no edges between R_I and $V^+ \setminus R_I$. Consequently, $N_G(R_I) \subseteq V^-$. Moreover, every column of Q_{R_I,V^-} indexed by $V^- \setminus N_G(R_I)$ is zero. Hence, using (44), we get:

$$\operatorname{rank}\big(Q_{R_I,N_G(R_I)}\big) = \operatorname{rank}(Q_{R_I,V^-}) = \dim(L_I) \leq r. \tag{45}$$

The algorithm considers every set $I \subseteq [s]$ with $|I| \leq r$ and $\dim(L_I) \leq r$. For each such I, it computes the connected components of $\widehat{G}_I$. The candidate is discarded if $|D \cap N_G(R_I)| > h$ for some connected component D of $\widehat{G}_I$. For every remaining candidate, it applies the treewidth algorithm of [26] to $\widehat{G}_I$ with parameter k. That algorithm runs in time $2^{O(k)}n$ and either returns a tree decomposition of width at most $2k+1$ or determines that $\operatorname{tw}(\widehat{G}_I) > k$. If a tree decomposition is returned, the algorithm outputs J_I. By construction every subset returned in this way satisfies the requirements of condition 1 of Theorem 6.

It remains to show that, whenever condition 1 of Theorem 6 holds, the algorithm returns a subset $J \subseteq [s]$. Recall the subset J^* whose existence follows from condition 1, and set

$$L^* := \operatorname{span}\{Q_{v,V^-} : v \in R^*\}.$$

Since the columns of Q_{R^*,V^-} outside $N_G(R^*)$ are zero, the rank assumption on J^* gives $\dim(L^*) = \operatorname{rank}\big(Q_{R^*,N_G(R^*)}\big) \leq r$. Choose a basis of L^* consisting of rows of Q_{R^*,V^-}, and let I^* be the set of indices of the components containing these basis rows; take $I^* = \emptyset$ if $L^* = \{0\}$. Then $|I^*| \leq r$. Since R^* is a union of components C_j, every C_i with $i \in I^*$ is contained in R^*. Hence all rows associated with these components belong to L^*, whereas the chosen basis rows span L^*. Consequently,

$$L_{I^*} = L^*. \tag{46}$$

We next claim that

$$J_{I^*} \subseteq J^*. \tag{47}$$

Indeed, if $j \notin J^*$, then $C_j \subseteq R^*$, and therefore $L_j \subseteq L^* = L_{I^*}$ by (46). Hence $j \notin J_{I^*}$, proving (47). By (44) and (46),

$$\operatorname{span}\{Q_{v,V^-} : v \in R_{I^*}\} = \operatorname{span}\{Q_{v,V^-} : v \in R^*\}.$$

Fix $u \in V^-$. By the definition of the interaction graph, $u \in N_G(R_{I^*})$ if and only if $Q_{vu} \neq 0$ for some $v \in R_{I^*}$. By the equality of the two spans, this holds if and only if the u-th coordinate is nonzero in some vector of $\operatorname{span}\{Q_{v,V^-} : v \in R^*\}$, which is equivalent to $u \in N_G(R^*)$. Therefore,

$$N_G(R_{I^*}) = N_G(R^*). \tag{48}$$

By (47), every clique added to construct $\widehat{G}_{I^*}$ is also added to construct $\widehat{G}^*$. Thus $\widehat{G}_{I^*}$ is a subgraph of $\widehat{G}^*$, and hence

$$\operatorname{tw}(\widehat{G}_{I^*}) \leq \operatorname{tw}(\widehat{G}^*) \leq k.$$

Furthermore, every connected component D_I of $\widehat{G}_{I^*}$ is contained in some connected component D^* of $\widehat{G}^*$. By (48),

$$|D_I \cap N_G(R_{I^*})| \leq |D^* \cap N_G(R^*)| \leq h.$$

Therefore the candidate I^* satisfies $|D_I \cap N_G(R_{I^*})| \le h$ for every connected component D_I of $\widehat{G}_{I^*}$. Moreover, $\mathrm{tw}(\widehat{G}_{I^*}) \le k$. Hence the treewidth algorithm returns a decomposition for $\widehat{G}_{I^*}$, and the algorithm outputs J_{I^*}. Finally, the number of sets I considered is at most

$$\sum_{\ell=0}^{\min\{r,s\}} \binom{s}{\ell} \le \sum_{\ell=0}^{r} s^\ell = s^{O(r)}.$$

Since r is fixed, this is polynomial in n. For each candidate, all required linear-algebra and graph operations take polynomial time in $\mathcal{L}$, while the treewidth algorithm takes $2^{O(k)}n$ time. Since $k = O(\log n)$, the total running time is polynomial in the input length. $\square$

Proposition 3 shows that the subset J in Theorem 6 need not be given explicitly, since a suitable choice can be found in polynomial time whenever condition 1 holds. For certain graph classes, however, a suitable choice of J can be identified directly from the graph structure. The following corollary gives such a class. Recall that a graph is a *cactus graph* (resp. *block graph*) if each of its maximal biconnected components (blocks) is an edge or a cycle (resp. clique). Moreover, a graph is a *block-cactus graph* if each of its maximal biconnected components is either a clique or a cycle.

Corollary 8. *Let $G = (V, E)$ be the interaction graph of Problem* (QP). *Define $V^+ := \{i \in V : q_{ii} > 0\}$. Suppose that G is a block-cactus graph with the clique number $\omega(G) \in O(\log |V|)$. Moreover, suppose that G_{V^+} is a forest. Let R be the union of the connected components of G_{V^+} containing at least one cut vertex of G. Suppose that the following conditions hold:*

1. *For every connected component D of $G_{V \setminus R}$, we have $|D \cap N_G(R)| \in O(\log |V|)$.*
2. $\mathrm{rank}(Q_{R,N_G(R)}) \in O(1)$.

Then Problem (QP) *can be solved in time polynomial in the input length.*

Proof. Let $C_1, \dots, C_s$ denote the connected components of G_{V^+}, and define

$$J := \{j \in [s] : C_j \text{ contains no cut vertex of } G\}.$$

By the definition of R, we have $R = V^+ \setminus \bigcup_{j \in J} C_j$. We apply Theorem 6 with this choice of J. Fix $j \in J$. Since no vertex of C_j is a cut vertex of G, all vertices of C_j belong to the same block B of G. Moreover, every neighbor of C_j also belongs to B. Indeed, a non-cut vertex of a graph belongs to a unique block, and the connectedness of C_j therefore forces all vertices of C_j and all edges incident to them to belong to the same block. Since G is a block-cactus graph, B is either a clique or a cycle. If B is a clique, then $N_G(C_j)$ is already a clique. If B is a cycle, then C_j is a path, since G_{V^+} is a forest, and therefore $|N_G(C_j)| \le 2$. The latter follows since every non-cut vertex of a cycle block has degree two. Therefore, making $N_G(C_j)$ a clique either adds no edge or adds the edge joining the two ends of the path through C_j.

We first verify that assumption 1 of Theorem 6 holds. Let $\widehat{G}$ be the graph on V^- obtained from G_{V^-} by making $N_G(C_j)$ a clique for every $j \in J$. In a clique block, this operation adds no edge. In a cycle block, the added edge, if any, can be obtained by contracting the path through C_j. Consequently, $\widehat{G}$ is a minor of $G_{V \setminus R}$ and hence also a minor of G. The treewidth of a block-cactus graph is the maximum of the treewidths of its blocks. A clique block has treewidth one less than its cardinality, while a cycle block has treewidth at most two. Therefore, $\mathrm{tw}(G) \le \max\{\omega(G) - 1, 2\} \in O(\log |V|)$, and hence $\mathrm{tw}(\widehat{G}) \in O(\log |V|)$. Since R is a union of connected components of G_{V^+}, $N_G(R) \subseteq V^-$. For every connected component D of $G_{V \setminus R}$, replacing each C_j, $j \in J$, by a clique on $N_G(C_j)$ preserves connectivity among the vertices of $D \cap V^-$. Consequently, the connected components of $\widehat{G}$ are precisely the nonempty sets $D \cap V^-$, where D ranges over the connected components of $G_{V \setminus R}$. Since $N_G(R) \subseteq V^-$, we have $(D \cap V^-) \cap N_G(R) = D \cap N_G(R)$. It follows from assumption 1 that every connected component $\widehat{D}$ of $\widehat{G}$ satisfies $|\widehat{D} \cap N_G(R)| \in O(\log |V|)$. Thus assumption 1 of Theorem 6 is satisfied. Finally, since G_{V^+} is a forest, every connected component C_j of G_{V^+} is a tree. Hence, by Theorem 1, for every $j \in [s]$ and every $d \in \mathbb{Q}^{C_j}$, the optimization problem (42) can be solved in polynomial time. Thus assumption 2 of Theorem 6 is satisfied. Finally, assumption 2 of the corollary is precisely condition 1(ii) of Theorem 6, and this completes the proof. □

Remark 4. Corollary 8 is not implied by either Theorem 4 or Theorem 5. To see this, consider the following family of instances. For each $m \geq 3$, let $V^+ := \{r, b_1, \ldots, b_m\}$ and $V^- := \{w_1, \ldots, w_m\} \cup \{z_1, \ldots, z_m\}$. Let $w_1, \ldots, w_m$ induce a cycle, let b_j be adjacent only to w_j for every $j \in [m]$, and let r be adjacent to w_1 and to every z_i, $i \in [m]$. There are no other edges. It can be checked that the graph G is a cactus, and V^+ is a stable set. The only cut vertex of G belonging to V^+ is r. Hence, with the notation of Corollary 8, we have $R = \{r\}$. The connected components of $G_{V \setminus R}$ consist of the component $\{w_1, \ldots, w_m\} \cup \{b_1, \ldots, b_m\}$, and the isolated vertices $z_1, \ldots, z_m$. Each of these components intersects $N_G(r)$ in exactly one vertex. Therefore, assumption 1 of Corollary 8 is satisfied. Moreover, $\text{rank}(Q_{R,N_G(R)}) \leq 1$, so Corollary 8 applies. However, Theorem 4 does not apply. Indeed, the cycle $D := \{w_1, \ldots, w_m\}$ is a connected component of G_{V^-} and, since w_j is adjacent to $b_j \in V^+$ for every $j \in [m]$, we have $D \cap N_G(V^+) = D$. Thus $|D \cap N_G(V^+)| = m = \Theta(|V|)$, violating assumption 1 of Theorem 4. Finally, Theorem 5 does not apply because $N_G(r) = \{w_1, z_1, \ldots, z_m\}$. In the graph obtained by making the neighborhood of every continuous component a clique, $N_G(r)$ therefore induces a clique of cardinality $m + 1$. Hence the resulting graph has treewidth at least m, and assumption 1 of Theorem 5 is violated. Thus, Corollary 8 contains families of instances that are covered by neither Theorem 4 nor Theorem 5.

Funding: A. Del Pia is partially funded by AFOSR grant FA9550-23-1-0433 and ONR grant N00014-25-1-2490. A. Khajavirad is in part supported by AFOSR grant FA9550-23-1-0123 and by ONR grant N00014-25-1-2491. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the Air Force Office of Scientific Research or the Office of Naval Research.

References

[1] A. Bhathena, S. Fattahi, A. Gómez, and S. Küçükyavuz. A parametric approach for solving convex quadratic optimization with indicators over trees. *Mathematical Programming*, 218:291–336, 2026.

[2] A. Bhathena, S. Fattahi, A. Gómez, and S. Küçükyavuz. Solving convex quadratic optimization with indicators over structured graphs. *arXiv preprint arXiv:2603.02103*, 2026.

[3] D. Bienstock and G. Muñoz. LP formulations for polynomial optimization problems. *SIAM Journal on Optimization*, 28(2):1121–1150, 2018.

[4] E. Boros and P.L. Hammer. Pseudo-Boolean optimization. *Discrete applied mathematics*, 123(1):155–225, 2002.

[5] V. Chandrasekaran, N. Srebro, and P. Harsha. Complexity of inference in graphical models. In *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence UAI'08*, 2008.

[6] D. Cifuentes and P. A. Parrilo. Exploiting chordal structure in polynomial ideals: A gröbner bases approach. *SIAM Journal on Discrete Mathematics*, 30(3):1534–1570, 2016.

[7] Y. Crama. Linearization techniques and concave extensions in nonlinear $0-1$ optimization. Technical report, Rutgers University, 1988.

[8] Y. Crama, P. Hansen, and B. Jaumard. The basic algorithm for pseudo-Boolean programming revisited. *Discrete Applied Mathematics*, 29(2–3):171–185, 1990.

[9] A. Del Pia. Factorized binary polynomial optimization. *Mathematical Programming, Series A*, 219:799–828, 2026.

[10] A. Del Pia, S. S. Dey, and M. Molinaro. Mixed-integer quadratic programming is in NP. *Mathematical Programming, Series A*, 162(1):225–240, 2017.

[11] A. Del Pia and A. Khajavirad. Beyond hypergraph acyclicity: limits of tractability for pseudo-boolean optimization. *arXiv preprint arxiv:2410.23045*, 2024.

[12] S. S. Dey and A. Khajavirad. A second-order cone representable class of nonconvex quadratic programs. *Mathematical Programming*, 2026.

[13] R. Diestel. *Graph Theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, 5 edition, 2017.

[14] H. Edelsbrunner, J. O'Rourke, and R. Seidel. Constructing arrangements of lines and hyperplanes with applications. *SIAM Journal on Computing*, 15(2):341–363, 1986.

[15] J. Edmonds. Systems of distinct representatives and linear algebra. *Journal of Research of the National Bureau of Standards, Section B: Mathematics and Mathematical Physics*, 71B(4):241–245, 1967.

[16] F. Eisenbrand, C. Hunkenschröder, K.-M. Klein, M. Kouteckỳ, A. Levin, and S. Onn. An algorithmic theory of integer programming. *arXiv preprint arXiv:1904.01361*, 2019.

[17] B. Gardiner and Y. Lucet. Convex hull algorithms for piecewise linear-quadratic functions in computational convex analysis. *Set-Valued and Variational Analysis*, 18:467–482, 2010.

[18] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, San Francisco, 1979.

[19] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer-Verlag, Berlin, 1988.

[20] M. Hladík, M. Černỳ, and M. Rada. A new polynomially solvable class of quadratic optimization problems with box constraints. *Optimization Letters*, 15(6):2331–2341, 2021.

[21] R. Horst and H. Tuy. *Global Optimization, Deterministic Approaches*. Springer, 1996.

[22] A. Khajavirad. A polynomial-time solvable class of sparse box-constrained polynomial optimization problems. *arXiv preprint arXiv:2604.25033*, 2026.

[23] A. Khajavirad. Tight semidefinite programming relaxations for sparse box-constrained quadratic programs. *arXiv preprint arXiv:2601.18545*, 2026.

[24] S. Kim and M. Kojima. Exact solutions of some nonconvex quadratic optimization problems via SDP and SOCP relaxations. *Computational optimization and applications*, 26(2):143–154, 2003.

[25] V. Kolmogorov, T. Pock, and M. Rolinek. Total variation on a tree. *SIAM Journal on Imaging Sciences*, 9(2):605–636, 2016.

[26] T. Korhonen. A single-exponential time 2-approximation algorithm for treewidth. In *Proceedings of the 62nd IEEE Annual Symposium on Foundations of Computer Science (FOCS 2021)*, pages 184–192. IEEE, 2021.

[27] M. Kuric, J. Ahmetspahic, and T. Pock. Total generalized variation on a tree. *SIAM Journal on Imaging Sciences*, 17(2):1040–1077, 2024.

[28] C. Michini. Tight cycle relaxations for the cut polytope. *SIAM Journal on Discrete Mathematics*, 35(4):2908–2921, 2021.

[29] C. T. Ng, MS Barketau, T.C.E. Cheng, and M. Y. Kovalyov. "product partition" and related problems of scheduling and systems reliability: Computational complexity and approximation. *European Journal of Operational Research*, 207(2):601–604, 2010.

[30] M. Padberg. The Boolean quadric polytope: Some characteristics, facets and relatives. *Mathematical Programming*, 45(1–3):139–172, 1989.

[31] J.-S. Pang and S. Han. Some strongly polynomially solvable convex quadratic programs with bounded variables. *SIAM Journal on Optimization*, 33(2):899–920, 2023.

[32] A. P. Punnen. The bipartite QUBO. In *The Quadratic Unconstrained Binary Optimization Problem*, chapter 10, pages 261–300. Springer, 2022.

[33] Y. Qiu and E. A. Yıldırım. On exact and inexact RLT and SDP-RLT relaxations of quadratic programs with box constraints. *Journal of Global Optimization*, 90:293–322, 2024.

[34] R.T. Rockafellar. *Convex Analysis.* Princeton University Press, Princeton, 1970.

[35] A. Schrijver. *Theory of Linear and Integer Programming.* Wiley, Chichester, 1986.

[36] A. Schrijver. A combinatorial algorithm minimizing submodular functions in strongly polynomial time. *Journal of Combinatorial Theory, Series B*, 80(2):346–355, 2000.

[37] S. Sojoudi and J. Lavaei. Exactness of semidefinite relaxations for nonlinear optimization problems with underlying graph structure. *SIAM Journal on Optimization*, 24(4):1746–1778, 2014.

[38] R. P. Stanley. An introduction to hyperplane arrangements. In Ezra Miller, Victor Reiner, and Bernd Sturmfels, editors, *Geometric Combinatorics*, volume 13 of *IAS/Park City Mathematics Series*, pages 389–496. American Mathematical Society, 2007.

[39] S. P. Tarasov and L. G. Khachiyan. Polynomial solvability of convex quadratic programming. In *Doklady akademii nauk*, volume 248, pages 1049–1051. Russian Academy of Sciences, 1979.

[40] S. A. Vavasis. Quadratic programming is in NP. *Information Processing Letters*, 36:73–77, 1990.