Approximating Sparse Quadratic Programs

Danny Hermelin*,[1], Leon Kellerhals[2], Rolf Niedermeier[2], and Rami Pugatch†,[1]

[1]Technische Universität Berlin, Chair of Algorithmics and Computational Complexity,
Berlin, Germany, `leon.kellerhals@tu-berlin.de, rolf.niedermeier@tu-berlin.de`
[2]Ben-Gurion University of the Negev, Department of Industrial Engineering and Management,
Beer Sheva, Israel, `hermelin@bgu.ac.il, rpugatch@bgu.ac.il`

Abstract. Given a matrix $A \in \mathbb{R}^{n \times n}$, we consider the problem of maximizing $x^T Ax$ subject to the constraint $x \in \{-1, 1\}^n$. This problem, called MaxQP by Charikar and Wirth [FOCS'04], generalizes MaxCut and has natural applications in data clustering and in the study of disordered magnetic phases of matter. Charikar and Wirth showed that the problem admits an $\Omega(1/\lg n)$ approximation via semidefinite programming, and Alon, Makarychev, Makarychev, and Naor [STOC'05] showed that the same approach yields an $\Omega(1)$ approximation when A corresponds to a graph of bounded chromatic number. Both these results rely on solving the semidefinite relaxation of MaxQP, whose currently best running time is $\tilde{O}(n^{1.5} \cdot \min\{N, n^{1.5}\})$, where N is the number of nonzero entries in A and $\tilde{O}$ ignores polylogarithmic factors.

In this sequel, we abandon the semidefinite approach and design purely combinatorial approximation algorithms for special cases of MaxQP where A is sparse (*i.e.*, has $O(n)$ nonzero entries). Our algorithms are superior to the semidefinite approach in terms of running time, yet are still competitive in terms of their approximation guarantees. More specifically, we show that:

- MaxQP admits a $(1/2\Delta)$-approximation in $O(n \lg n)$ time, where Δ is the maximum degree of the corresponding graph.
- Unit MaxQP, where $A \in \{-1, 0, 1\}^{n \times n}$, admits a $(1/2d)$-approximation in $O(n)$ time when the corresponding graph is d-degenerate, and a $(1/3\delta)$-approximation in $O(n^{1.5})$ time when the corresponding graph has δn edges.
- MaxQP admits a $(1 - \varepsilon)$-approximation in $O(n)$ time when the corresponding graph and each of its minors have bounded local treewidth.
- Unit MaxQP admits a $(1 - \varepsilon)$-approximation in $O(n^2)$ time when the corresponding graph is H-minor free.

1 Introduction

In this paper we are interested in the following (integer) quadratic problem which was coined MaxQP by Charikar and Wirth [13]. Given an $n \times n$ symmetric matrix with zero valued diagonal entries A, $a_{i,j} \in \mathbb{R}$ for all $i, j \in \{1, \ldots, n\}$, we want to maximize

$$\mathrm{val}_x(A) \;=\; \sum_{i=1}^{n} \sum_{j=1}^{n} a_{i,j} x_i x_j \qquad \text{s.t. } x_i \in \{-1, 1\} \text{ for all } i \in \{1, \ldots, n\}. \tag{1}$$

---

[*] The first author's work was supported by the Israel Science Foundation (Grant no. 1070/20).
[†] The fourth author's work was supported by the Israel Science Foundation (Grant no. 776/19).

Observe that the requirement that all diagonal values of A are zero is to avoid the term $\sum_i a_{i,i}$ which is constant for all choices of $x_{i,i}$ in (1). Furthermore, a non-symmetric matrix A can be replaced with an equivalent symmetric A' by setting $a'_{i,j} = a'_{j,i} = \frac{1}{2} \cdot (a_{i,j} + a_{j,i})$ without changing $\mathrm{val}_x(A)$, and so the requirement that A is symmetric is just for convenience's sake.

Our interest in MaxQP lies in the fact that it is a generic example of integer quadratic programming which naturally appears in different contexts. Below we review three examples:

- *Graph cuts:* Readers familiar with the standard quadratic program formulation of MaxCut [20] will notice the similarity to (1). Indeed, given a graph $G = (V, E)$ with vertex set $V = \{1, \dots, n\}$ and edge weights $a_{i,j} \geq 0$ for each $\{i, j\} \in E$, the corresponding MaxQP instance on $-1 \cdot A$ has an optimum solution of value $2k - \sum_{i,j} a_{i,j}$ if and only if G has a maximum cut of total weight k. Thus, MaxQP with only negative A entries can be used to solve MaxCut exactly, implying that even this special case is NP-hard. Furthermore, this special case translates to the closely related MaxCut Gain problem [13, 25].

- *Correlation clustering:* In correlation clustering [6, 12, 14, 31], we are provided with pairwise judgments of the similarity of n data items. In the simplest version of the problem there are three possible inputs for each pair: similar (*i.e.* positive), dissimilar (*i.e.* negative), or no judgment. In a given clustering of the n items, a pair of items is said to be in *agreement* (*disagreement*) if it is a positive (negative) pair within one cluster or a negative (positive) pair across two distinct clusters. In MaxCorr, the goal is to maximize the *correlation* of the clustering; that is, the absolute difference between the number of pairs in agreement and the number of pairs in disagreement, across all clusters. Note that when only two clusters are allowed, this directly corresponds to Unit MaxQP, the variant of MaxQP where $a_{i,j} \in \{-1, 0, 1\}$ for each entry $a_{i,j}$ of A.

- *Ising spin glass model:* Spin glass models are used to in physics to study disordered magnetic phases of matter. Such system are notoriously hard to solve, and various techniques to approximate the free energy were developed. In the Ising spin-glass model [7, 32], each node in the graph represents a single spin which can either point up (+1) or down (-1), and neighboring spins (i, j) may have either positive or negative coupling energy $a_{i,j}$ between them. The energy of this system (when there is no external field) is given by its Hamiltonian $H = -1 \cdot \sum_{i,j} a_{i,j} \alpha(i) \alpha(j)$, where $\alpha(i) \in \{-1, 1\}$ is the spin at site i. A famous problem in the physics of spin-glasses is the characterization of the ground state — the state that minimizes the energy of the system. This problem is precisely MaxQP.

It is convenient to view MaxQP in graph-theoretic terms. Let $G = (V, E)$ be the *graph associated with* A, where $V = \{1, \dots, n\}$ and $E = \{\{i, j\} : a_{i,j} \neq 0\}$. The first algorithmic result for MaxQP was due to Bieche *et al.* [10] and Barahona *et al.* [9] who studied the problem in the context of the Ising spin glass model. They showed that when G is restricted to be planar, the problem is polynomial-time solvable via a reduction to maximum-weight matching. At the same time, Barahona proved that the problem is NP-hard for three-dimensional grids [7] or apex graphs (graphs with a vertex whose removal leaves the graph planar) [8].

1.1 Approximation Algorithms

As MaxQP is NP-hard, even for restricted instances, our focus is naturally on polynomial-time approximation algorithms. We note that the fact that the values of A are allowed to be both positive and negative makes MaxQP quite unique in the context of approximation and presents several challenges. First of all, there is an immediate equivalence between MaxQP and the problem of minimizing (1), as maximizing $\mathrm{val}_x(A)$ is the same as minimizing $\mathrm{val}_x(-1 \cdot A)$. Furthermore, solutions might have negative values; that is, we might have $\mathrm{val}_x(A) < 0$ for certain solutions x. This poses an extra challenge since a solution with a non-positive value is not an $f(n)$-approximate solution, for any function f, in case the optimum is positive (which it always is whenever $A \neq 0$, see Charikar and Wirth [13] and our Lemma 11). In particular, a uniformly at random chosen solution x has $\mathrm{val}_x(A) = 0$ on expectation, and unlike MaxCut, such a solution is unlikely to be useful as any kind of approximation.

Alon and Naor [2] were the first to show that these difficulties can be overcome by carefully rounding a semidefinite relaxation of MaxQP. In particular, they studied the problem when G is bipartite, and showed that using a rounding technique that relies on the famous Grothendieck inequality, one can obtain an approximation factor guarantee of ≈ 0.56 for the bipartite case. Later, together with Makarychev and Makarychev [1], they showed that the integrality gap of the semidefinite relaxation is $O(\lg \chi(G))$ and $\Omega(\lg \omega(G))$, where $\chi(G)$ and $\omega(G)$ are the chromatic and clique numbers of G, respectively. In particular, this gap is constant for several interesting graph classes such as d-degenerate graphs and H-minor free graphs, and it generalizes the previous result of Alon and Naor [2] as $\chi(G) \leq 2$ when G is bipartite.

Theorem 1 ([1, 2]). MaxQP *restricted to graphs of* $O(1)$ *chromatic number can be approximated within a factor of* $\Omega(1)$ *in polynomial time.*

Regarding the general version of the problem, where G can be an arbitrary graph, an integrality gap of $O(\lg n)$ for the semidefinite relaxation was first shown by Nesterov [29]. However, his proof was non-constructive. Charikar and Wirth [13] made his proof constructive, and provided a rounding procedure for the relaxation that guarantees $\Omega(1/\lg n)$-approximate solutions regardless of the structure of G.

Theorem 2 ([13, 29]). MaxQP *can be approximated within a factor of* $\Omega(1/\lg n)$ *in polynomial time.*

As for the time complexity of the algorithm in Theorems 1 and 2 above, Arora, Hazan, and Kale [4] provided improved running times for several semidefinite programs, including the relaxation of MaxQP. They showed that this relaxation can be solved (to within any constant factor) in $\tilde{O}(n^{1.5} \cdot \min\{N, n^{1.5}\})$ time, where N is the number of nonzero entries in A and $\tilde{O}$ ignores polylogarithmic factors. Thus, for general matrices A, this running time is $O(n^3)$, and for matrices with $O(n)$ nonzero entries this is $O(n^{2.5})$.

There has also been work on approximation lower bounds for MaxQP. Alon and Naor [2] showed that MaxQP restricted to bipartite graphs cannot be approximated within $16/17 + \varepsilon$ unless P=NP, while Charikar and Wirth [13] showed that, assuming P≠NP, the problem admits no $(11/13 + \varepsilon)$-approximation when G is an arbitrary graph. Both these results follow somewhat directly from the $16/17 + \varepsilon$ lower bound for MaxCut [23]. In contrast, Arora *et al.* [3] showed a much stronger lower bound by proving that there exists a constant $c > 0$ such that MaxQP cannot be approximated within $\Omega(1/\lg^c n)$, albeit under the weaker assumption that $\mathrm{NP} \not\subseteq \mathrm{DTime}(n^{\lg^{O(1)} n})$.

1.2 Our results

In this paper we focus on *sparse* graphs, *i.e.*, graphs where the number of edges m is $O(n)$. This corresponds to matrices A having $O(n)$ nonzero entries. Note that MaxQP remains APX-hard in this case as well (see Theorem 9 in Appendix B). Nevertheless, we show that one can abandon the semidefinite approach in favor of simpler, "purely combinatorial" algorithms, while still maintaining comparable performances. In particular, our algorithms are faster than than those obtained from the semidefinite approach whose fastest known implementation requires $O(n^{2.5})$ time [4]. Furthermore, most of them are quite easy to implement.

1.2.1 Generic classes of sparse graphs

We begin by considering three basic classes of sparse graphs. Our first result concerns bounded degree graphs. We show that a simple greedy algorithm for computing a matching in G can be used to obtain an $\Omega(1)$-approximate solution to the corresponding MaxQP instance.

Theorem 3. *Let* $\Delta \geq 1$*.* MaxQP *restricted to graphs of maximum degree* Δ *can be approximated within a factor of* $1/2\Delta$ *in* $O(n \lg n)$ *time.*

Next we consider d-degenerate graphs. Recall that a graph is *d-degenerate* if each of its subgraphs has a vertex of degree at most d. We show that by using a more elaborate structure than the matching used in the proof of Theorem 3, one can obtain an $\Omega(1)$ approximation in this setting for the unit weight case.

Theorem 4. *There is a* $1/2d$ *approximation algorithm for* Unit MaxQP *restricted to* d*-degenerate graphs with* $O(n)$ *running time.*

Finally, we consider *δ-dense* graphs, graphs which have at most δn edges for some fixed δ. By a slight modification of the idea used to prove both theorems above, we obtain the following generalization of Theorem 4 at the cost of a slight increase in the running time and decrease in the approximation factor guarantee.

Theorem 5. *There is a* $1/3\delta$ *approximation algorithm for* Unit MaxQP *restricted to* δ*-dense graphs without isolated vertices that runs in* $O(n^{1.5})$ *time.*

Observe that all the three results above improve on the running time of Theorem 1. Furthermore, while Theorem 1 provides an $\Omega(1)$ approximation for graphs of bounded degree and bounded degeneracy, this is not true for graphs of bounded density. For example, consider a graph consisting of a clique of size $\sqrt{n}$ together with a perfect matching on the remaining vertices. The result of Alon *et al.* [1] implies that the semidefinite relaxation has an integrality gap of $O(\lg n)$ on such a graph, while the algorithm in Theorem 5 provides an $\Omega(1)$ approximation.

1.2.2 H-minor free graphs

We next consider graph classes that exclude certain minors. A graph G is H-minor free for a fixed graph H if one cannot obtain in G an isomorphic copy of H by a series of edge contractions, edge deletions, and vertex deletions. We begin by considering the class of apex-minor free graphs (recall that a graph is *apex* if it contains one specific vertex whose deletion results in a planar graph). This graph class is better known as the class of minor-closed graphs with bounded local treewidth [18], and includes well-studied classes such as planar and bounded genus graphs. We tailor the approach of Eppstein [18] and Grohe [21] for designing approximation algorithms for apex-minor-free graphs to our setting to obtain the following result.

Theorem 6. *Let* $\varepsilon > 0$*. There is an* $O(n)$ *time* $(1-\varepsilon)$*-approximation algorithm for* MaxQP *restricted to apex-minor free graphs.*

We then show that Theorem 6 can be extended to general H-minor free graphs with unit weights, although this requires an extra $O(n)$ factor in the running time. This algorithm is obtained by using a partitioning algorithm of Demaine *et al.* [15] which is based on the graph minor decomposition of Robertson and Seymour [30].

Theorem 7. *For* $\varepsilon > 0$ *and any graph* H *there is an* $O(n^2)$ *time* $(1-\varepsilon)$*-approximation algorithm for* Unit MaxQP *restricted to* H*-minor free graphs.*

1.2.3 Maximum correlation

Finally, we note that our results have direct consequences for the MaxCorr problem: Charikar and Wirth [13] proved that an α-approximation algorithm for MaxQP implies an $\alpha/(2+\alpha)$-approximation algorithm for MaxCorr. Combining this with the results discussed above gives us the following:

Corollary 1. MaxCorr *can be approximated within a factor of*

- $1/(4d+1)$ *on* d*-degenerate graphs in* $O(n)$ *time;*
- $1/3-\varepsilon$ *on* H*-minor free graphs in* $O(n^2)$ *time;*
- $1/3-\varepsilon$ *on apex-minor free graphs in* $O(n)$ *time.*

2 Preliminaries

Throughout the paper we use $G=(V,E)$ to denote the graph associated with our input matrix A; that is, $V=\{1,\dots,n\}$ and $E=\{\{i,j\}:a_{i,j}\neq 0\}$. Thus, $n=|V|$ and we let $m=|E|$. We slightly abuse notation by allowing a solution x to denote either a vector in $\{-1,1\}^n$ indexed by V or a function $x:V\to\{-1,1\}$. For a solution x, we let $\mathrm{val}_x(G)=\sum_{\{i,j\}\in E}a_{i,j}x_ix_j$, and we let $\mathrm{opt}(G)=\max_x\mathrm{val}_x(G)$. We use $||A||$ to denote the sum of absolute values in A, *i.e.*, $||A||=\sum_{i,j}|a_{i,j}|$. Note that $\mathrm{opt}(G)\leq||A||$.

We use standard graph-theoretic terminology when dealing with the graph G, as in *e.g.* Diestel [16]. In particular, for a subset $V'\subseteq V$, we let $G[V']$ denote the subgraph of G *induced* by V'; *i.e.*, the subgraph with vertex set V' and edge set $\{\{u,v\}\in E:u,v\in V'\}$. We let $G-V'=G[V\setminus V']$, and for a subset of edges $E'\subseteq E$ we let $G-E'$ denote the graph (V,E') without isolated vertices. For a pair of disjoint subsets $V_1,V_2\subseteq V$, we let $E(V_1,V_2)=\{\{u,v\}\in E:u\in V_1,v\in V_2\}$. Finally, we use $N(v)=\{u:\{u,v\}\in E\}$ to denote the *neighborhood* of a vertex $v\in V$.

2.1 Useful observations

Note that for a uniformly chosen at random solution x, the value $a_{i,j}x_ix_j$ is zero in expectation for any edge $\{i,j\}\in E$. This implies that $\mathrm{opt}(G)\geq 0$. Moreover, a solution x with $\mathrm{val}_x(G)\geq 0$ can be computed in linear time:

Lemma 1. *One can compute in* $O(m+n)$ *time a solution* x *with* $\mathrm{val}_x(G)\geq 0$.

Proof. For each vertex $i\in V$, let $E(i)=\{\{i,j\}\in E:j<i\}$. Consider an arbitrary initial solution x, and let $z_i=\sum_{\{i,j\}\in E(i)}x_ix_ja_{i,j}$. Then z_i is the contribution of edges in $E(i)$ to $\mathrm{val}_x(G)$. We compute a solution x^* by scanning the vertices from 1 to n. For a given vertex i, we check whether $z_i<0$. If so, we set $x^*_i=-x_i$, and otherwise we set $x^*_i=x_i$. Note that $z^*_i=\sum_{\{i,j\}\in E(i)}x^*_ix^*_ja_{i,j}$ must now be positive. As the value of x^*_i does not change z^*_j for any $j<i$, when we finish our scan we have $z^*_i\geq 0$ for each $i\in\{1,\dots,n\}$. Thus, $\mathrm{val}_{x^*}(G)=\sum_i z^*_i\geq 0$. ☐

Lemma 2. *Let* $V_1,V_2\subseteq V$ *be two disjoint subsets of vertices, and let* $x^{(1)}$ *and* $x^{(2)}$ *be two solutions for* $G[V_1]$ *and* $G[V_2]$ *of value* z_1 *and* z_2 *respectively. Then at least one of the solutions* $x^{(1)}\cup x^{(2)}$ *and* $-x^{(1)}\cup x^{(2)}$ *has value* z_1+z_2 *for* $G[V_1\cup V_2]$.

Proof. Suppose $x^{(1)}\cup x^{(2)}$ has value less than z_1+z_2. This means that the total contribution of the edges in $E(V_1,V_2)$ is negative in this solution. Observe that in $-x^{(1)}\cup x^{(2)}$ each edge of $E(V_1,V_2)$ with negative contribution under $x^{(1)}\cup x^{(2)}$ now has positive contribution, and vice versa. The lemma thus follows. ☐

Combining Lemma 1 and Lemma 2 above, we get an important property of $\mathrm{val}(G)$, namely that it is monotone with respect to induced subgraphs.

Lemma 3. *Let H be an induced subgraph of G. Then given a solution x, one can compute in $O(n+m)$ time a solution x^* for G with $\mathrm{val}_{x^*}(G) \geq \mathrm{val}_x(H)$.*

Proof. Let $V_0 \subseteq V$ be the vertices of G which are not present in H. According to Lemma 1 we can compute a solution $x^{(0)}$ for $G[V_0]$ with value at least zero in linear time. According to Lemma 2 either $x^{(0)} \cup x$ or $x^{(0)} \cup -x$ has value at least $\mathrm{val}_{x^{(0)}}(G[V_0]) + \mathrm{val}_x(H) \geq \mathrm{val}_x(H)$. Thus, taking x^* to be the solution with higher value out of $x^{(0)} \cup x$ or $x^{(0)} \cup -x$ proves the lemma. ☐

3 Matching-Based Algorithms

In this section we present approximation algorithms for MaxQP using certain type of matchings that we compute for G. In particular, we provide proofs for Theorems 3, 4, and 5. The proof of each theorem is given in each subsection below.

3.1 Graphs of Bounded Degree

We begin with the case where G has bounded maximum degree Δ. Our algorithm for this case is quite simple: It greedily computes a matching and then outputs a solution corresponding to this matching. For an edge set $E' \subseteq E$, we let $w(E') = \sum_{\{i,j\} \in E'} |a_{i,j}|$ be the total absolute value of the edges in E'. To prove that our strategy gives a good approximation on the optimal value of the MaxQP instance, we first observe that, given a matching M, we can derive a solution with value at least $w(M)$.

Lemma 4. *Given a matching M, one can compute in $O(n)$ time a solution x with $\mathrm{val}_x(G) \geq w(M)$.*

Proof. Let $V_M = \{u \in V : \{u,v\} \in M\}$ be the set of vertices matched by M in G. By Lemma 3 it suffices to compute a solution x for the set of endpoints V_M in M such that $\mathrm{val}_x(G[V_M]) \geq w(M)$. We construct x by induction on $t = |M|$. For $t = 1$, let $\{u,v\} \in M$. If $a_{u,v} > 0$, we choose x_u and x_v to be of equal value, and otherwise we choose x_u and x_v to be of opposite value. Then $\mathrm{val}_x(G[\{u,v\}]) = |a_{i,j}| = w(M)$. Suppose now that $t > 1$ and let $M_0 = M \setminus \{e_t\}$. By induction, we have a solution $x^{(0)}$ for V_{M_0} with $\mathrm{val}_{x^{(0)}}(G[V_{M_0}]) \geq w(M_0)$. Let $e_t = \{u,v\}$ and let $x^{(t)} \colon \{u,v\} \to \{-1,1\}$ such that $\mathrm{val}_{x^{(t)}}(G[\{u,v\}]) = w(e_t)$ as in the case of $t = 1$. Then, by Lemma 2, either $x^{(0)} \cup x^{(t)}$ or $x^{(0)} \cup -x^{(t)}$ has value at least $w(M_0) + w(e_t) = w(M)$, and we are done. ☐

For our algorithm we choose a particular matching M^* which we obtain greedily as follows: We first sort the edges $\{u,v\} \in E$ in non-increasing order according to their absolute values $|a_{u,v}|$. We then iteratively choose the first edge among the remaining edges, and remove all edges that share an endpoint with this edge, and stop when no more edges are left. The matching M^* is the set of all edges selected in this process.

Lemma 5. $w(M^*) \geq 1/2\Delta \cdot w(E)$.

Proof. Let $M^* = \{e_1, \dots, e_t\}$, with $w(e_1) \geq w(e_2) \geq \cdots \geq w(e_t)$, and let E_i be the set of edges that were removed after choosing e_i. Formally, $E_i = \{e \in E \colon e \cap e_i \neq \emptyset\} \setminus \bigcup_{j=1}^{i-1} E_j$. Let $e_i = \{u_i, v_i\}$. As all edges in E_i are incident to either u_i or v_i, we have that $|E_i|$ is at most the sum of the degrees of u_i and v_i, which is at most 2Δ. As e_i is the edge with the highest weight within E_i (otherwise we would have chosen a different edge for e_i), we have

$$w(\{e_i\}) \geq \frac{w(E_i)}{|E_i|} \geq \frac{w(E_i)}{2\Delta}.$$

Since $E_1, \ldots, E_t$ is a partition of all edges in G, we get

$$w(M) = \sum_{i=1}^{k} w(\{e_i\}) \geq \sum_{i=1}^{k} \frac{w(E_i)}{2\Delta} = \frac{w(E)}{2\Delta}. \qquad \square$$

Proof of Theorem 3. The matching M^* can be computed in $O(n \lg n)$ time. Due to Lemma 4 and Lemma 5, we obtain from M^* a solution x with

$$\mathrm{val}_x(G) \geq 1/2\Delta \cdot w(E) = 1/2\Delta \cdot ||A||$$

in $O(n)$ time. Since $\mathrm{opt}(G) \leq ||A||$, this solution is $1/2\Delta$-approximate. $\square$

3.2 Graphs of Bounded Degeneracy

We next present our approximation algorithm for Unit MaxQP restricted to d-degenerate graphs. Our algorithm extends the algorithm of the previous subsection by considering a structure more elaborate than a matching.

Let us begin with introducing some terminology. Let x be a solution for G. If for some edge $e = \{u, v\}$ we have $a_{u,v} x_u x_v > 0$, then we say that e is *good* in x, otherwise we say that it is *bad* in x. Let $\{u, v, w\} \subseteq V$ be the vertices of a triangle in G. Note that there exists a solution x for G where all all edges in this triangle are good iff $a_{u,v} \cdot a_{v,w} \cdot a_{w,u} = 1$. Hence, we call a triangle *good* if $a_{u,v} \cdot a_{v,w} \cdot a_{w,u} = 1$, and *bad* otherwise.

We next introduce a special type of subgraph that will be exploited by our algorithm. Recall that a graph is a *split graph* if its vertex set can be partitioned into a clique and an independent set. We call an n-vertex split graph *basic* if it is connected, and the class of vertices inducing a clique is of size 2 (note that this clique need not be maximal). A subgraph G' of G is *easy* if it is a basic split graph without any bad triangles. In this way, an easy subgraph of G contains one *center edge* whose endpoints are adjacent to all remaining *outside vertices*, and the outside vertices have no edges between them. Note that, possibly, an outside vertex can form a triangle with the center edge. In this case, the triangle must be good. An *easy packing* of G is a family $\mathcal{F} = \{F_1, \ldots F_t\}$ of pairwise disjoint subsets of vertices such that each F_i induces an easy subgraph in G. For each $F_i \in \mathcal{F}$, we use $m(F_i)$ to denote the number of edges in $G[F_i]$. Furthermore, we let $m(\mathcal{F}) = \sum_i m(F_i)$ and $V_{\mathcal{F}} = \bigcup_i F_i$.

Lemma 6. *Given an easy packing $\mathcal{F}$, one can compute in $O(n)$ time a solution x with $\mathrm{val}_x(G) \geq m(\mathcal{F})$.*

Proof. By Lemma 3, it suffices to compute a solution x for $V_{\mathcal{F}}$ of value at least $m(\mathcal{F})$. Let $\mathcal{F} = \{F_1, \ldots, F_t\}$. We construct such a solution by induction on t. For $t = 1$, let $\{u, v\}$ be the center edge of the easy subgraph $G[F_1]$. We assign to u an arbitrary value. To each neighbor w of u we assign the same value if $a_{u,v} = 1$ and the opposite value otherwise. Afterwards, we proceed analogously with the neighbors of v. This ensures that each edge incident to u or v contributes $+1$ to $\mathrm{val}_x(G[F_1])$. As each triangle in $G[F_1]$ is good and for each such triangle, two edges contribute $+1$ to the value, the third edge must contribute $+1$ as well. Thus, $\mathrm{val}_x(G[F_1]) = m(F_1) = m(\mathcal{F})$. Suppose then that $t > 1$ and let $\mathcal{F}_0 = \mathcal{F} \setminus \{F_t\}$. By induction, we have a solution $x^{(0)}$ for $V_{\mathcal{F}_0}$ with $\mathrm{val}_{x^{(0)}}(G[V_{\mathcal{F}_0}]) \geq m(\mathcal{F}_0)$. Let $x^{(t)}\colon F_t \to \{1, -1\}$ be such that $\mathrm{val}_{x^{(t)}}(G[F_t]) \geq |E(G[F_t])|$ as in the case of $t = 1$. Then, by Lemma 2, either $x^{(0)} \cup x^{(t)}$ or $x^{(0)} \cup -x^{(t)}$ have value at least $m(\mathcal{F}_0) + |E(G[F_t])| = m(\mathcal{F})$, and we are done. □

For our algorithm we need a particular easy packing which is inclusion-wise maximal. Moreover, we want to be able to compute this packing efficiently so as not to exceed the linear running time promised in Theorem 4. We next describe a procedure that achieves just that.

ALGORITHM EASYPACK:

1. Compute an inclusion-wise maximal matching M in G, and let I be the set of unmatched vertices.
2. Let $M^* = \emptyset$ and $I^* = I$.
3. For each $\{x, y\} \in M$, if there are two vertices $u, v \in I^*$ such that both u, x, y and v, x, y induce a triangle in G, then
   - $M^* = M^* \cup \{\{u, x\}, \{v, y\}\}$.
   - $I^* = I^* \setminus \{u, v\}$.

   Otherwise, $M^* = M^* \cup \{\{x, y\}\}$.
4. Initialize the easy packing with the matching M^*, that is, $\mathcal{F}^* = \{\{x, y\} : \{x, y\} \in M^*\}$.
5. For each $v \in I^*$, if there is an edge $\{x, y\} \in M^*$ such that v, x, y induce a path or a good triangle, then remove v from I^* and add it to the vertex set $F \in \mathcal{F}^*$ that contains x and y.

Lemma 7. *The algorithm above computes in $O(n)$ time an easy packing $\mathcal{F}^*$ with $m(\mathcal{F}^*) \geq |V_{\mathcal{F}^*}|/2$.*

Proof. Clearly, F^* is an easy packing. Steps 1, 2, and 4 can be performed in linear time in a straightforward manner. Step 3 can be executed in linear time by first adding a marker to every vertex in I and then, computing $N(x) \cap N(y)$ for each $\{x, y\} \in M$ in $O(\Delta^2) = O(1)$ time. If the intersection contains two marked vertices, we remove their marks, and add the corresponding edges to M^*. Finally, step 5 can be performed in linear time by storing for each matched vertex its partner and then iterating over the neighborhood of each $v \in I^*$. Thus, the entire algorithm can be executed in linear time. To complete the proof, observe that each $G[F_i]$, $F_i \in \mathcal{F}$, is connected and consists of at least one edge. This implies that $m(\mathcal{F}^*) \geq |V_{\mathcal{F}^*}|/2$, and so the lemma holds. □

Next we need to provide an upper bound on $\mathrm{opt}(G)$. For this we first make an observation on the matching M^* that we computed in Step 3 of our algorithm above.

Lemma 8. *Each edge in M^* forms at most one triangle with a vertex from I^*.*

Proof. Let $\{x, y\} \in M^*$. If this edge is also in M, then it forms a triangle with at most one vertex from I due to Step 3 of the algorithm above. Otherwise, one of its endpoints, say x, was originally unmatched, that is $x \in I$. Assume that there is a vertex $v \in I^* \subseteq I$ such that v, x, y form a triangle. Then $\{v, x\} \notin M$ but both v and x were in I, a contradiction to the maximality of M (see Step 1). □

Now we are ready to bound $\mathrm{opt}(G)$ from above. Recall that $V_{\mathcal{F}^*}$ is the set of vertices in the packing computed by our algorithm above.

Lemma 9. $\mathrm{opt}(G) \leq d \cdot |V_{\mathcal{F}^*}|$, *where d is the degeneracy of G.*

Proof. Let M^* be the matching and $\mathcal{F}^*$ be the easy packing computed above. Then $I^* = V \setminus V_{\mathcal{F}^*}$. Let T be the set of triangles that contain a vertex from I^*. We first show that the triangles in T are edge-disjoint. Observe that I^* is an independent set, as $I^* \subseteq I$ is a subset of the vertices that were left unmatched in step 1; an edge within this vertex set would contradict the maximality of the matching M computed in this step. So, if two triangles in T share an edge, then it must be an edge in M^*. But this is impossible according to Lemma 8, and so the triangles in T are indeed edge-disjoint.

Next observe that if a vertex v is in I^* at the end of the algorithm above, we know that either v is adjacent neither to x nor to y, or $\{v, x, y\}$ induces a bad triangle in G. Thus, T consists solely of bad triangles. Furthermore, for each $v \in I^*$, there are $\deg(v)/2$ bad triangles that contain v, where $\deg(v)$ is the degree of v in G. Thus, overall there are at least $|T| = \sum_{v \in I^*} \deg(v)/2$ bad triangles in G. As the triangles in T do not share edges, we obtain that there are at least $|T|$ bad edges in every solution for G. Consequently,

$$\mathrm{opt}(G) \leq (||A|| - |T|) - |T| = m - 2|T| = m - \sum_{v \in V'} \deg(v).$$

Now, as $G[V_{\mathcal{F}^*}]$ is a d-degenerate graph, it has at most $d \cdot |V_{\mathcal{F}^*}|$ edges. Thus,

$$\begin{aligned} 2\mathrm{opt}(G) &\leq 2m - 2\sum_{v \in I^*} \deg(v) = \sum_{v \in V} \deg(v) - 2\sum_{v \in I^*} \deg(v) \\ &\leq \sum_{v \in V_{\mathcal{F}^*}} \deg(v) \leq 2d \cdot |V_{\mathcal{F}^*}|, \end{aligned}$$

and the lemma follows. ☐

Proof of Theorem 4. By Lemma 6 and Lemma 7 we can compute in linear time an easy packing $\mathcal{F}^*$ and a solution x with $\mathrm{val}_x(G) \geq |V_{\mathcal{F}^*}|/2$. Using Lemma 9 we obtain

$$\mathrm{val}_x(G) \geq |V_{\mathcal{F}^*}|/2 \geq \mathrm{opt}(G)/2d,$$

and the theorem is proven. ☐

3.3 Graphs of Bounded Density

We next turn to handle graphs of bounded density δ without isolated vertices. Here we use a particular kind of easy packings described in section 3.2, were each vertex subset induces a star in G. A *star packing* of G is a family of pairwise disjoint subsets of vertices $\mathcal{F} = \{F_1, \dots, F_t\}$ such that each $G[F_i]$ is a star. Note that since a star is a basic split graph without bad triangles, a star packing is also an easy packing. We compute a specific star $\mathcal{F}^*$ packing inspired by the work of Erdős, Gyárfás, and Kohayakawa [19] concerning graph cuts.

To compute $\mathcal{F}^*$, we first compute a maximum matching M in G. Note that we require M to be of maximum cardinality, and not only maximal inclusion-wise. Let $V_M \subseteq V$ denote the vertices in G that are matched in M, and let $I = V \setminus V_M$. We initialize the packing $\mathcal{F}^*$ with the matching M, that is, $\mathcal{F}^* = \{\{x, y\} : \{x, y\} \in M\}$. We then iterate through vertices $v \in I$, and add v to any subset F such that $G[F \cup \{v\}]$ is a star. Let $V_{\mathcal{F}^*} = \bigcup_{F \in \mathcal{F}^*} F$ denote the set of vertices in the packing at the end of this process, and let $I^* = V \setminus V_{\mathcal{F}^*}$.

Lemma 10. *Let $F \in \mathcal{F}^*$. Then there is at most one vertex in $v \in I^*$ which is adjacent to any vertex in F.*

Proof. Let $\{x, y\} \subseteq F \cap V_M$ be the edge in M from which F was initialized. Then, as $F \setminus \{x, y\} \subseteq I$, any vertex in I^* can only be adjacent to x and y in F. Let $I^*(x) \subseteq I^*$ and $I^*(y) \subseteq I^*$ respectively denote the set of neighbors of x and y in I^*. If both $I^*(x) = \emptyset$ and $I^*(y) = \emptyset$ we are done. So assume w.l.o.g. that $I^*(x) \neq \emptyset$ and let $v \in I^*(x)$. Now if there exists a vertex $u \in I^*(y) \setminus I^*(x)$, then replacing $\{x, y\}$ with $\{x, v\}$ and $\{y, u\}$ in M gives us a matching M' with $|M'| > |M|$, contradicting the fact that M is a maximum matching. Thus, $I^*(y) \setminus I^*(x) = \emptyset$. Furthermore, if $v \notin I^*(y)$ then $G[F \cup \{v\}]$ is a star, and so v would have been added to F by the algorithm above. Thus, $I^*(x) = I^*(y)$. Finally, if there are two distinct vertices $u, v \in I^*(x)$, then $u, v \in I^*(y)$, and again we can replace $\{x, y\}$ with $\{x, v\}$ and $\{y, u\}$ in M. This implies that $|I^*(x) \cup I^*(y)| \leq 1$, and so the lemma is proven. ☐

Lemma 11. $m(\mathcal{F}^*) \geq m/3\delta$.

Proof. Since each $F \in \mathcal{F}^*$ induces a star in G we have $2m(\mathcal{F}^*) \geq |V_{\mathcal{F}^*}|$. Moreover, as by Lemma 10, each $F \in \mathcal{F}^*$ is adjacent to at most one vertex in I^*, and each vertex in I^* is adjacent to at least one $F \in \mathcal{F}^*$ as G has no isolated vertices, we have $|\mathcal{F}^*| \geq |I^*|$. Since each $F \in \mathcal{F}^*$ induces at least one edge, we have $m(\mathcal{F}^*) \geq |\mathcal{F}^*|$, and so $m(\mathcal{F}^*) \geq |I^*|$. Thus, combining all of the above, we get

$$3m(\mathcal{F}^*) \geq |V_{\mathcal{F}^*}| + |I^*| = n \geq m/\delta,$$

and the lemma is proven. ☐

Proof of Theorem 5. Due to Lemma 6 and 11, the packing $\mathcal{F}^*$ yields a solution x with $\mathrm{val}_x(G) \geq m/3\delta$. Since $\mathrm{opt}(G) \leq ||A|| = m$, this solution is $1/3\delta$-approximate. The running time for computing $\mathcal{F}^*$ is dominated by the computation of the maximum matching M for the initial packing, taking $O(m\sqrt{n}) = O(n^{1.5})$ time [28]; computing $\mathcal{F}^*$ from M can be done in $O(m+n) = O(n)$ time. ☐

4 Algorithms Based On Treewidth Partitions

In this section we present approximation algorithms for sparse MAXQP instances that exclude certain types of minors. Namely, we prove Theorems 6 and 7. Our algorithms all revolve around the Baker technique for planar graphs [5] and its generalizations [15, 17, 21], all using what we refer to here as a *treewidth partition* — a partition of the vertices of G into $V_0, \ldots, V_{k-1}$ such that $G - V_i$ has bounded treewidth for any subset V_i in the partition. As treewidth plays a central role here, we begin with formally defining this notion.

A tree decomposition is a pair $(\mathcal{T}, \mathcal{X})$ where $\mathcal{X}$ is a family of vertex subsets of G, called *bags*, and $\mathcal{T}$ is a tree with $\mathcal{X}$ as its node set. The decomposition is required to satisfy (i) $\{X \in \mathcal{X} : v \in X\}$ is connected in $\mathcal{T}$ for each $v \in V$, and (ii) for each $\{u, v\} \in E$ there is a bag $X \in \mathcal{X}$ that contains both u and v. The *width* of a tree decomposition $(\mathcal{T}, \mathcal{X})$ is $\max_{X \in \mathcal{X}} |X| - 1$, and the *treewidth* of G is the smallest width amongst all its tree decompositions. The proof of the following lemma, relying on standard dynamic programming techniques, is deferred to Appendix A.

Lemma 12. MAXQP *restricted to graphs of treewidth at most k can be solved in $2^{O(k)} \cdot n$ time.*

4.1 Apex-minor free instances

Our starting point for the $(1 - \varepsilon)$-approximation for MAXQP on H-minor free graphs with H being an apex graph (Theorem 6) is a layer decomposition $L_0, \ldots, L_\ell \subseteq V$ of G. Here $L_0 = \{v\}$ for some arbitrary vertex $v \in V$, $L_i = \{u : d(v, u) = i\}$ are all vertices at distance i from v, for each $i \in \{1, \dots, \ell\}$. This is the standard starting point of all Baker-type algorithms, and can be computed in linear time via breadth-first search from v. Note that $L_0, \dots, L_\ell$ form a partition of V, and that for each $i \in \{0, \dots, \ell\}$, each vertex in L_i has neighbors only in $L_{i-1} \cup L_i \cup L_{i+1}$ (here and elsewhere in this section we set $L_{-1} = L_{\ell+1} = \emptyset$ when necessary).

Given $0 < \varepsilon \leq 1$, we let k be the smallest integer such that $4/k \leq \varepsilon$. For each $i \in \{0, \dots, k-1\}$, let $\mathcal{L}_i$ denote the union of all vertices in layers with index equal to $i (\mathrm{mod}\, k)$; that is, $\mathcal{L}_i = \bigcup_{j \equiv i (\mathrm{mod}\, k)} L_j$. We define two subgraphs of G: The graph G_i is the graph induced by $V - \mathcal{L}_i$, and the graph H_i is the graph induced by $N[\mathcal{L}_i]$. Note that there is some overlap between the vertices of G_i and H_i, but each edge of G appears in exactly one of these subgraphs. Also note that since there is an apex graph H that G does not contain as a minor, G and each of its minors have bounded local treewidth [18]; thus both G_i and H_i are bounded treewidth graphs [21].

Our algorithm computes k different solutions for G, and selects the best one (*i.e.*, the solution x with highest $\mathrm{val}_x(G)$) as its solution. For $i \in \{0, \dots, k-1\}$, we first compute an optimal solution for G_i in linear time using the algorithm given in Lemma 12. We then extend this solution to a solution $x^{(i)}$ for G as is done in Lemma 3. In this way we obtain in linear time k solutions $x^{(0)}, \dots, x^{(k-1)}$ with $\mathrm{val}_{x^{(t)}}(G) \geq \mathrm{opt}(G_i)$ for each $i \in \{0, \dots, k-1\}$. In Lemma 13 we argue that the solution of maximum objective value is $(1-\varepsilon)$-approximate to the optimum of G; the proof of Theorem 6 will then follow as a direct corollary.

Lemma 13. *There is a solution* $x \in \{x^{(0)}, \dots, x^{(k-1)}\}$ *with* $\mathrm{val}_x(G) \geq (1-\varepsilon) \cdot \mathrm{opt}(G)$.

Proof. Let x^* denote the optimal solution for G. Then, as the edge set of G is partitioned into the edges of G_i and H_i, we have

$$\mathrm{opt}(G_i) + \mathrm{opt}(H_i) \geq \mathrm{val}_{x^*}(G_i) + \mathrm{val}_{x^*}(H_i) = \mathrm{val}_{x^*}(G) = \mathrm{opt}(G)$$

for each $i \in \{0, \dots, k-1\}$. Next observe that any two subgraphs H_{i_1} and H_{i_2} with $|i_1 - i_2| \geq 4$ do not have vertices in common, nor are there any edges between these two subgraphs in G. It follows that for any $j \in \{0, 1, 2, 3\}$, the graph $\bigcup_{i \equiv j (\mathrm{mod}\, 4)} H_i$ is an induced subgraph in G, and so $\mathrm{opt}(G) \geq \sum_{i \equiv j (\mathrm{mod}\, 4)} \mathrm{opt}(H_i)$ by Lemma 3. Thus, we have

$$4 \cdot \mathrm{opt}(G) \geq \sum_j \sum_{i \equiv j (\mathrm{mod}\, 4)} \mathrm{opt}(H_i) = \sum_i \mathrm{opt}(H_i).$$

Combining the two inequalities above we get

$$\begin{aligned} \sum_{i=0}^{k-1} \mathrm{val}_{x^{(i)}}(G) \geq \sum_{i=0}^{k-1} \mathrm{opt}(G_i) &\geq k \cdot \mathrm{opt}(G) - \sum_{i=0}^{k-1} \mathrm{opt}(H_i) \\ &\geq k \cdot \mathrm{opt}(G) - 4 \cdot \mathrm{opt}(G) = (k-4)\mathrm{opt}(G). \end{aligned}$$

It follows that the best solution out of $x^{(0)}, \dots, x^{(k-1)}$ has value at least $(1-4/k) \cdot \mathrm{val}(G)$, which is at least $(1-\varepsilon) \cdot \mathrm{val}(G)$ since $4/k \leq \varepsilon$. ☐

4.2 General H-minor free instances

To obtain the $(1-\varepsilon)$-approximation for Unit MaxQP on H-minor free graphs for any fixed graph H (Theorem 7), we make use of the lower bound obtained in Section 3.3, our algorithm for MaxQP restricted to bounded-treewidth graphs, and the following theorem by Demaine *et al.* [15, Theorem 3.1]:

Theorem 8 ([15]). *For a fixed graph H, there is a constant c_H such that, for any integer $k \geq 1$ and for every H-minor free graph G, the vertices of G can be partitioned into k sets such that the graph obtained by taking the union of any $k-1$ of these sets has treewidth at most $c_H \cdot k$. Furthermore, such a partition can be found in polynomial time.*

Note that this theorem gives a partition similar to the one used in the previous subsection, albeit slightly weaker. In particular, there is no restriction on the edges connecting vertices in different subsets of the partition as was the case in the previous subsection. It is for this reason that arbitrary weights are difficult to handle, and we need to resort to the lower bound of Lemma 11. Fortunately, for the unweighted case, we can use the fact that there exists some constant h depending only on H such that G has at most hn edges (see *e.g.* [16]). In particular, it can be shown that $h = O(n'\sqrt{\lg n'})$ [27], where n' is the number of vertices of H. Combining this fact with Lemma 11, we get:

Lemma 14. $\operatorname{opt}(G) \geq m/3h$.

Our algorithm proceeds as follows. Fix $k \geq 6h/\varepsilon$, and let $V_0, \ldots, V_{k-1}$ denote the partition of V computed by the algorithm from Theorem 8. For each $i \in \{0, \ldots, k-1\}$, let E_i denote the set of edges $E(V_i, \bigcup_{j \neq i} V_j)$, and let $m_i = |E_i|$. Furthermore, let $G_i = G - E_i$. As both $G[V_i]$ and $G[V \setminus V_i]$ have bounded treewidth, we can compute an optimal solution for each of these subgraphs (and therefore also for G_i) using the algorithm in Lemma 12. Using Lemma 2, we can extend the optimal solutions for $G[V_i]$ and $G[V \setminus V_i]$ to a solution $x^{(i)}$ for G with value

$$\operatorname{val}_{x^{(i)}}(G) \geq \operatorname{opt}(G_i).$$

On the other hand, the optimal solution of G cannot do better than

$$\operatorname{opt}(G_i) + m_i \geq \operatorname{opt}(G).$$

Combining the two inequalities above, we can bound the sum of the objective values obtained by all our solutions by

$$\begin{aligned}\sum_{i=0}^{k-1} \operatorname{val}_{x^{(i)}}(G) &\geq \sum_{i=0}^{k-1} \operatorname{opt}(G_i) \geq \sum_{i=0}^{k-1} (\operatorname{opt}(G) - m_i) \\ &= \sum_{i=0}^{k-1} \operatorname{opt}(G) - 2m \geq (k - 6h) \cdot \operatorname{opt}(G),\end{aligned}$$

where the last inequality follows from Lemma 14. Thus at least one of these solutions has value at least $(k-6h)/k \cdot \operatorname{opt}(G)$, which is greater than $(1-\varepsilon)\operatorname{opt}(G)$ by our selection of parameter k.

To analyze the time complexity of our algorithm, observe that computing each solution $x^{(i)}$ requires $O(n)$ time according to Lemma 12 and Lemma 2. Thus, the time complexity of the algorithm is dominated by the time required to compute the partition promised by Theorem 8. Demaine *et al.* [15] showed that this partition can be computed in linear time given the graph decomposition promised by Robertson and Seymour's graph minor theory [30]. In turn, Grohe *et al.* [22] presented an $O(n^2)$-time algorithm for this decomposition, improving earlier constructions [15, 24]. Thus, the total running time of our algorithm can also be bounded by $O(n^2)$. This completes the proof of Theorem 7.

5 Conclusion

We presented efficient combinatorial approximation algorithms for sparse instances of MaxQP without resorting to the semidefinite relaxation, as done by Alon and Naor [2] and Charikar and

Wirth [13]. From a theoretical perspective, we still leave open whether there is a fast algorithm for d-degenerate MaxQP instances which obtains an $\Omega(1)$ approximation factor guarantee. Even more interesting is to design a purely combinatorial algorithm for general MaxQP instances with an approximation guarantee of $\Omega(1/\lg n)$. Finally, the simplicity of our algorithms compels the study of their usability in practice, especially for characterizations of ground states of spin glass models.

References

[1] Noga Alon, Konstantin Makarychev, Yury Makarychev, and Assaf Naor. Quadratic forms on graphs. In *Proceedings of the 37th Annual ACM Symposium on Theory Of Computing (STOC)*, pages 486–493, 2005. 3, 4

[2] Noga Alon and Assaf Naor. Approximating the cut-norm via Grothendieck's inequality. *SIAM Journal on Computing*, 35(4):787–803, 2006. 3, 12, 16

[3] Sanjeev Arora, Eli Berger, Elad Hazan, Guy Kindler, and Muli Safra. On non-approximability for quadratic programs. In *Proceedings of the 46th Annual IEEE Symposium on Foundations Of Computer Science (FOCS)*, pages 206–215, 2005. 3

[4] Sanjeev Arora, Elad Hazan, and Satyen Kale. Fast algorithms for approximate semidefinite programming using the multiplicative weights update method. In *46th Annual IEEE symposium on Foundations Of Computer Science (FOCS)*, pages 339–348, 2005. 3, 4

[5] Brenda S. Baker. Approximation algorithms for NP-complete problems on planar graphs. *Journal of the ACM*, 41(1):153–180, 1994. 10

[6] Nikhil Bansal, Avrim Blum, and Shuchi Chawla. Correlation clustering. *Machine Learning*, 56(1-3):89–113, 2004. 2

[7] Francisco Barahona. On the computational complexity of ising spin glass models. *Journal of Physics A: Mathematical, Nuclear and General*, 15:3241–3253, 1982. 2

[8] Francisco Barahona. The max-cut problem on graphs not contractible to K_5. *Operations Research Letters*, 2(3):107–111, 1983. 2

[9] Francisco Barahona, Roger Maynard, Rammal Rammal, and Jean-Pierre Uhry. Morphology of ground states of two-dimensional frustration model. *Journal of Physics A: Mathematical, Nuclear and General*, 15:673–699, 1982. 2

[10] Isabelle Bieche, Roger Maynard, Rammal Rammal, and Jean-Pierre Uhry. On the ground states of the frustration model of a spin glass by a matching method of graph theory. *Journal of Physics A: Mathematical, Nuclear, and General*, 13:2553–2576, 1980. 2

[11] Hans L. Bodlaender, Pål Grønås Drange, Markus S. Dregi, Fedor V. Fomin, Daniel Lokshtanov, and Michal Pilipczuk. A $c^k n$ 5-approximation algorithm for treewidth. *SIAM Journal on Computing*, 45(2):317–378, 2016. 16

[12] Moses Charikar, Venkatesan Guruswami, and Anthony Wirth. Clustering with qualitative information. *Journal of Computer and System Sciences*, 71(3):360–383, 2005. 2

[13] Moses Charikar and Anthony Wirth. Maximizing quadratic programs: extending Grothendieck's inequality. In *Proceedings of the 45th annual IEEE symposium on Foundations Of Computer Science (FOCS)*, pages 54–60, 2004. 1, 2, 3, 5, 13

[14] Erik D. Demaine, Dotan Emanuel, Amos Fiat, and Nicole Immorlica. Correlation clustering in general weighted graphs. *Theoretical Computer Science*, 361(2-3):172–187, 2006. 2

[15] Erik D. Demaine, Mohammad Taghi Hajiaghayi, and Ken ichi Kawarabayashi. Algorithmic graph minor theory: Decomposition, approximation, and coloring. In *Proceedings of the 46th Annual IEEE symposium on Foundations Of Computer Science (FOCS)*, pages 637–646, 2005. 4, 10, 11, 12

[16] Reinhard Diestel. *Graph Theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, 5th edition, 2016. 5, 12

[17] David Eppstein. Subgraph isomorphism in planar graphs and related problems. *Journal of Graph Algorithms & Applications*, 3(3):1–27, 1999. 10

[18] David Eppstein. Diameter and treewidth in minor-closed graph families. *Algorithmica*, 27(3):275–291, 2000. 4, 11

[19] Paul Erdős, András Gyárfás, and Yoshiharu Kohayakawa. The size of the largest bipartite subgraphs. *Discrete Math*, 177(1-3):267–271, 1997. 9

[20] Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, 1995. 2

[21] Martin Grohe. Local tree-width, excluded minors, and approximation algorithms. *Combinatorica*, 23(4):613–632, 2003. 4, 10, 11

[22] Martin Grohe, Ken-ichi Kawarabayashi, and Bruce A. Reed. A simple algorithm for the graph minor decomposition—Logic meets structural graph theory. In *Proceedings of the 24th annual ACM-SIAM Symposium On Discrete Algorithms (SODA)*, pages 414–431, 2013. 12

[23] Johan Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, 2001. 3, 16

[24] Ken ichi Kawarabayashi and Paul Wollan. A simpler algorithm and shorter proof for the graph minor decomposition. In *Proceedings of the 43rd ACM Symposium on Theory Of Computing (STOC)*, pages 451–458, 2011. 12

[25] Subhash Khot and Ryan O'Donnell. SDP gaps and UGC-hardness for max-cut-gain. *Theory of Computing*, 5(1):83–117, 2009. 2

[26] Ton Kloks. *Treewidth, Computations and Approximations*, volume 842 of *Lecture Notes in Computer Science*. Springer, 1994. 15

[27] Alexandr V. Kostochka. Lower bound of the Hadwiger number of graphs by their average degree. *Combinatorica*, 4(4):307–316, 1984. 12

[28] Silvio Micali and Vijay V. Vazirani. An $o(v\sqrt{|V|}|e|)$ algorithm for finding maximum matching in general graphs. In *Proceedings of the 21st Annual Symposium on Foundations Of Computer Science (FOCS)*, pages 17–27, 1980. 10

[29] Yurii Nesterov. Global quadratic optimization via conic relaxation. CORE Discussion Papers 1998060, Université catholique de Louvain, Center for Operations Research and Econometrics (CORE), 1998. 3

[30] Neil Robertson and Paul D. Seymour. Graph minors. XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory, Series B*, 89(1):43–76, 2003. 4, 12

[31] Chaitanya Swamy. Correlation clustering: maximizing agreements via semidefinite programming. In *Proceedings of the 15th Annual ACM-SIAM Symposium On Discrete Algorithms (SODA)*, pages 526–527, 2004. 2

[32] Michel Talagrand. *Spin Glasses: A Challenge for Mathematicians*, volume 46 of *A Series of Modern Surveys in Mathematics*. Springer, 1st edition, 2003. 2

A An Exact Algorithm for Bounded Treewidth Graphs

We prove Lemma 12 by presenting an algorithm for MAXQP restricted to graphs of treewidth at most k running in $2^{O(k)} \cdot n$ time. For this we require the concept of nice tree decompositions [26].

A tree decomposition $(\mathcal{T}, \mathcal{X})$ is *rooted* if there is a designated bag $R \in \mathcal{X}$ being the root of $\mathcal{T}$. A rooted tree decomposition is *nice* if each bag $X \in \mathcal{X}$ is either (i) a leaf node (X contains exactly one vertex and has no children in $\mathcal{T}$), (ii) an introduce node (X has one child Y in $\mathcal{T}$ with $Y \subset X$ and $|X \setminus Y| = 1$), (iii) a forget node (X has one child in Y in $\mathcal{T}$ with $X \subset Y$ and $|Y \setminus X| = 1$), or (iv) a join node (X has two children Y, Z in $\mathcal{T}$ with $X = Y = Z$). Given a tree decomposition, one can compute a corresponding nice tree decomposition with the same width in linear time [26].

Our algorithm employs the standard dynamic programming technique on nice tree decompositions.

Proof of Lemma 12. Let $(\mathcal{T}, \mathcal{X})$ be a nice tree decomposition of G of width k with root bag R. For a node $X \in \mathcal{X}$ let $\mathcal{T}_X$ be the subtree of $\mathcal{T}$ rooted at X. Furthermore, let G_X be the subgraph of G induced by the vertices in the bags of $\mathcal{T}_X$ (while $G[X]$ is the subgraph of G induced only by the vertices in X). We describe a table in which we have an entry $D[X, x]$ for each bag $X \in \mathcal{X}$ and for each solution $x : X \to \{-1, 1\}$. The entry $D[X, x]$ contains the value of an optimum solution for G_X, where the values of the vertices in X are fixed by the solution x.

If X is a *leaf node*, then G_X contains no edges and so $D[X, x] = 0$. If X is an *introduce node*, then let $v \in X \setminus Y$ be the introduced vertex, where Y is the child of X in $\mathcal{T}$, and let $x \setminus x_v$ be the solution x restricted to the vertices of Y. Then $D[X, x]$ additionally contains the value of all edges incident to v, that is,

$$D[X, x] = D[Y, x \setminus x_v] + \sum_{u \in N(v)} x_u x_v a_{u,v}.$$

If X is a *forget node*, then let $v \in Y \setminus X$ be the forgotten vertex, where Y is the child of X in $\mathcal{T}$. Then, every value except for x_v is set in x, so we must choose it so that the value is maximized. Then

$$D[X, x] = \max_{x_v : v \to \{-1,1\}} D[Y, x \cup x_v].$$

Finally, if X is a *join node*, then let Y and Z be the children of X in $\mathcal{T}$. Note that $D[Y, x] + D[Z, x]$ contains the value of $G[X]$ twice, so

$$D[X, x] = D[Y, x] + D[Z, x] - \mathrm{val}_x(G[X]).$$

The tree decomposition contains $O(n)$ nodes, and for each node there are at most $O(2^k)$ solutions; thus we need to compute $O(2^k \cdot n)$ entries $D[\cdot, \cdot]$, each of which can be computed in

$O(k)$ time. The optimum value is the maximum over all $O(2^k)$ solutions for the root bag R. As we can compute a tree decomposition of width $O(k)$ in $2^{O(k)} \cdot n$ time [11], we can compute $\mathrm{opt}(A)$ in $O(2^{O(k)} \cdot n + 2^k \cdot k \cdot n) = 2^{O(k)} \cdot n$ time. ☐

B A Hardness Result

Alon and Naor [2] show that MaxQP restricted to bipartite graphs is not approximable in polynomial time with a ratio of $16/17 + \varepsilon$ unless P=NP. Using the same idea, we show that this approximation lower bound also holds for Unit MaxQP on 2-degenerate bipartite graphs.

Theorem 9. *If* Unit MaxQP *restricted to* 2*-degenerate bipartite graphs can be approximated within a factor of* $(16/17 + \varepsilon)$ *in polynomial time, then P=NP.*

Proof. We reduce from unweighted MaxCut which does not admit a $(16/17+\varepsilon)$-approximation unless P=NP [23]. Given an undirected unweighted graph $G = (V, E)$, we construct a graph $G' = (V \cup V', E')$ by subdividing each edge in E, that is, for every edge $\{u, w\} \in E$ we add a vertex v to V' and the edges $\{u, v\}, \{v, w\}$ to E'. One edge has weight 1 while the other edge has weight -1. Clearly, G' is bipartite; V' is one bipartition. As all vertices in V' have degree two, G' is 2-degenerate as well.

Let x be a solution for G'. Observe that for every vertex $v \in V'$ we can assume that at least one of its incident edges contributes positively to $\mathrm{val}_x(G')$; otherwise multiply x_v by -1. Furthermore, note that the cut in G corresponding to x (restricted to V) is of size $\mathrm{val}_x(G')/2$: If both edges incident to v contribute positively to $\mathrm{val}_x(G')$, then the edge in G subdivided by v is cut. Otherwise, the two edges contribute 0 to $\mathrm{val}_x(G')$, and the corresponding edge in G is not cut.

It follows that if there is a $(16/17 + \varepsilon)$-approximation for MaxQP, then there is one for MaxCut as well, implying P=NP by [23]. ☐