$\hbar$-PERFECT GRAPHS: EXACT CERTIFICATES FOR SIMULTANEOUS PAULI-VARIANCE VIOLATIONS, AND THREE ERRATA

KOREA SUPERINTELLIGENCE LABS

ABSTRACT. Let $S_1, \dots, S_n$ be Pauli strings on m qubits and let G be their frustration graph, whose edges record the anticommuting pairs. Xu, Wang, Ye, Koßmann, Schwonnek and Winter attach to G and a weight $w \in \mathbb{R}^n_+$ the *weighted beta number* $\beta(G,w) = \sup_\rho \sum_i w_i \langle S_i \rangle^2_\rho$, call G *$\hbar$-perfect* when $\beta(G,w) = \alpha(G,w)$ for every $w \geq 0$, and classify all graphs on at most nine vertices up to 78 undetermined ones. The only upper bound on an anti-cycle beta number available there, or in the companion paper of Xu, Schwonnek and Winter, is $\beta(\overline{C}_n, \mathbf{1}) \leq \vartheta(\overline{C}_n) = 1 + 1/\cos(\pi/n)$ for odd n; the sharper value $\beta(\overline{C}_7, \mathbf{1}) = (9 + 4\sqrt{2})/7 = 2.0938363\dots$ that the companion records is established there only from below, by one explicit state. We prove the exact rational upper bound $\beta(\overline{C}_7, \mathbf{1}) \leq \frac{421}{200} = 2.105$, strictly below $\vartheta(\overline{C}_7) = 2.1099162642\dots$, together with $\beta(\overline{C}_5, \mathbf{1}) \leq \frac{1029}{500}$; combined with an explicit Gaussian-integer state this brackets $2.0937171893\dots \leq \beta(\overline{C}_7, \mathbf{1}) \leq 2.105$ by machine. The certificate is the second level of a sum-of-squares hierarchy, made computable by an exact symmetry reduction: its 448×448 Gram matrix is shown to equal $\hat{O}(Z \otimes \mathbf{1})\hat{O}$ for an explicit orthogonal signed permutation $\hat{O}$ and a 28×28 Hermitian Z, so only 28×28 is ever factored. We also give exact rational witnesses of $\hbar$-imperfectness for six graphs, four of which the source determines only numerically, two of them at weighted facets that no uniform-weight search can see; and we record three errata in the source's tables and appendix. Every numbered statement below is machine-checked in Lean 4, apart from the two elementary hand steps that §3.3 isolates.

1. INTRODUCTION

1.1. **Frustration graphs and the beta number.** A *Pauli string* on m qubits is a tensor product of m factors drawn from $\{\mathbf{1}, X, Y, Z\}$, normalised so that it is a Hermitian involution on $\mathbb{C}^d$, $d = 2^m$. Any two Pauli strings either commute or anticommute, and a set $\{S_1, \dots, S_n\}$ of them therefore determines a graph G on $\{1, \dots, n\}$ — its *frustration graph* — with $ij \in E(G)$ exactly when $S_i S_j = -S_j S_i$. Conversely every graph arises this way.

Xu, Wang, Ye, Koßmann, Schwonnek and Winter [1] ask how much the frustration graph alone controls the simultaneously attainable squared expectation values of the strings. Their central invariant is the *weighted beta number*, their Definition 1, which reads verbatim

> "For a given frustration graph G w.r.t. $\{S_i\}$ and a vector $w \in \mathbb{R}^n_+$, we define the weighted beta number as $\beta(G,w) := \sup_{\rho \in \mathcal{S}(\mathcal{H})} \sum_i w_i \langle S_i \rangle^2_\rho$,"

the supremum running over all density matrices on $\mathbb{C}^d$. Two facts make this a graph invariant worth the name: it does not depend on the chosen realisation of G by Pauli strings [2], and it always dominates the weighted independence number $\alpha(G,w) = \max\{\sum_{i \in I} w_i : I \subseteq V(G) \text{ independent}\}$. The class singled out in [1] is where the two agree; its Definition 3 reads, verbatim,

> "A graph G is called $\hbar$-perfect if $\mathrm{STAB}(G) = \mathrm{BETA}(G)$, or equivalently, for any $w \in \mathbb{R}^n_+$, $\beta(G,w) = \alpha(G,w)$."

Only the second form is used below, so we never need the beta body $\mathrm{BETA}(G)$; $\mathrm{STAB}(G)$ is as usual the convex hull of the incidence vectors of the independent sets of G. Perfect graphs and h-perfect graphs are $\hbar$-perfect [1], and the class is strictly larger; the smallest graph outside it is the anti-heptagon $\overline{C}_7$, the complement of the seven-cycle, for which [1] writes in its introduction

> "this conjecture is refuted, with the smallest counterexample of $\hbar$-imperfectness provided by the anti-heptagon"

and refers to [2]. Every statement, table, figure and appendix section of [1] cited below is cited by the number it carries in the arXiv v1 build of its own LaTeX source, which is the only version that exists; the theorem counter of that version restarts in the appendix, so appendix material is always located by its appendix section as well.

1.2. **What the source leaves open.** Two open ends of [1] are the subject of this paper.

The first is its classification. It determines $\hbar$-perfectness for every graph on at most nine vertices except for 78 of them, and reports the count for $n = 9$ as an interval rather than a number: the caption of its Table I reads "the number of $\hbar$-perfect graphs is in the range $[259583, 259661]$ since there are 78 undetermined graphs". The determinations that were made are in part numerical. Its Appendix B 3 states, verbatim, of the nine-vertex sweep,

> "There are only 2963 graphs which are neither reducible nor induced subgraphs of G_{15}, for which we apply the numerical tools and find 2880 ones to be $\hbar$-perfect (up to a $< 1 \times 10^{-5}$ gap between the lower and upper bounds), $\bar{C}_9$ and other 4 graphs as illustrated in Fig. [11](b-e) to be $\hbar$-imperfect, and 78 undetermined ones. For those 78 undetermined ones, the gap between the upper bound of the weighted beta number and the corresponding weighted independence number is less than 3×10^{-3}."

— the bracketed figure number replacing a symbolic cross-reference in the source, here and in the next quotation — and, earlier in the same section, of the eight-vertex sweep,

> "Then there are 100 undetermined graphs, for which we apply the numerical tools and find them to be $\hbar$-perfect (up to a $< 1 \times 10^{-5}$ gap between the lower and upper bounds), except one graph as in Fig. [11](a) which is named as G_8."

So five nine-vertex graphs and one eight-vertex graph are declared $\hbar$-imperfect on the strength of a floating-point semidefinite program and a see-saw iteration agreeing to 10^{-5}, computed with Mosek 10.2 (its Appendix B 1). We write G_8 for the eight-vertex graph of Figure 11(a) — the name [1] itself gives it in the sentence just quoted — and F_b, F_c, F_d, F_e for the four nine-vertex graphs of its parts (b)–(e).

The second is quantitative. For the anti-cycles, [1] states, in its Appendix C 2, $2 = \alpha(\overline{C}_n) < \beta(\overline{C}_n) \le \vartheta(\overline{C}_n)$ with

(1)
$$\vartheta(\overline{C}_n) = 1 + \frac{1}{\cos(\pi/n)} \quad (n \text{ odd}), \qquad \vartheta(\overline{C}_5) = \sqrt{5} = 2.2360679775\ldots, \quad \vartheta(\overline{C}_7) = 2.1099162642\ldots.$$

Here ϑ is the Lovász theta function of [3]: [1] calls it "the Lovász number" and cites [3] for it, both at (1) and where it introduces the weighted form $\vartheta(G, w) \ge \beta(G, w)$; it works with the theta body $\mathrm{TH}(G)$, and it asks in its conclusion for an extension of the sandwich theorem relating α, ϑ and the clique covering number. The inequality $\beta \le \vartheta$ itself is quoted there as known — its Appendix C 2 opens a proof with the words "It is known that" followed by $\beta(G, w) \in [\alpha(G, w), \vartheta(G, w)]$ — and attributed to earlier work on uncertainty relations, so (1) is the best upper bound *available* in [1] rather than one proved there. Its companion [2] records the sharper number $\beta(\overline{C}_7, \mathbf{1}) = (9 + 4\sqrt{2})/7 = 2.0938363\ldots$, but what it proves is the lower half: it exhibits one state of an explicit realisation with $\sum_i \langle S_i \rangle^2 = (9 + 4\sqrt{2})/7$, and reports $\vartheta(\overline{C}_{2n+1}) > \beta(\overline{C}_{2n+1}) > 2$ for $n \le 10$ as a numerical observation, in a sentence that begins "Numerically, we have". So no exact upper bound below (1) is available in either paper, for $\overline{C}_7$ or for any other graph: as we recall in §2, a *lower* bound on a beta number is a single state and an *upper* bound is a positivity statement about a matrix-valued quadratic form, and the two are not comparable in cost.

1.3. **Results.** Our main result closes the second gap for the anti-heptagon.

Theorem 1.1. $\beta(\overline{C}_5, \mathbf{1}) \le \dfrac{1029}{500} = 2.058$ *and* $\beta(\overline{C}_7, \mathbf{1}) \le \dfrac{421}{200} = 2.105$. *In particular*

$$\beta(\overline{C}_7, \mathbf{1}) \;\le\; 2.105 \;<\; 2.1099162642\ldots = \vartheta(\overline{C}_7),$$

so the bound is strictly below $\vartheta(\overline{C}_7)$*, which is the only upper bound for this beta number available in* [1].

Both bounds are exact rationals with an explicit, machine-checkable certificate. What is new is the *upper* direction below (1): we have found no earlier proof of an upper bound on a beta number strictly below the theta bound. The search behind that sentence was a full reading of [1] and of [2] — the two papers that introduce and develop the beta number under that name — together with arXiv queries combining "frustration graph" with "semidefinite" and with "sum of squares", and "beta number" with "Pauli", and web searches on the same vocabulary; the last of these turned up [2]'s number $(9+4\sqrt{2})/7$, discussed above, whose proof there is one state and hence a lower bound. We claim no more than that. The source itself records that the quantity "[w]ithout an explicit name … has been considered before in the contexts of uncertainty relations … and algebraic relaxations to the Hamiltonian ground state problem", and that literature we did not survey; a reader of it may know of a bound we do not. For $n = 5$ the value is not new — $\overline{C}_5 \cong C_5$ is $\hbar$-perfect, so $\beta(\overline{C}_5, \mathbf{1}) = 2$, every graph on at most six vertices being $\hbar$-perfect by [1] (its Appendix A 2: "all the graphs with no more than 6 vertices are $\hbar$-perfect") — and only the certificate is. Together with an explicit state (Proposition 4.1) Theorem 1.1 brackets the unknown value:

Corollary 1.2. $$\frac{327090970}{156225001} = 2.0937171893\ldots \;\le\; \beta(\overline{C}_7, \mathbf{1}) \;\le\; \frac{421}{200} = 2.105.$$

The width of the bracket is 0.0113. Its lower end is the machine-checked one; the sharper lower bound $(9 + 4\sqrt{2})/7 = 2.0938363\ldots$ of [2] lies inside it, 0.0112 below the upper end, and would narrow the bracket to 0.0112 if one is content with a hand computation there. What the certificate is, and how a 448×448 semidefinite program becomes a 28×28 Hermitian one without any relaxation, is §3.

The complementary direction is §4: exact rational witnesses of $\hbar$-imperfectness for six graphs (Propositions 4.1 and 4.2 and Theorems 4.3–4.5). Four of the six are the graphs whose $\hbar$-imperfectness [1] establishes only numerically, in the sense of the passage quoted above; for these the mathematical claim is its, and what is new is that it is now a theorem with an exact integer witness. Two of the six, moreover, show no violation at the uniform weight and violate only at a weighted facet of $\mathrm{STAB}(G)$, so a uniform-weight search does not find them.

Section 5 records three errata in [1]: an independence number that does not match the graph drawn beside it in its nine-vertex table, two table entries drawn identically but given different numerical bounds, and one sentence of its appendix that is false as written, with $\overline{C}_9$ as an explicit counterexample. None of the three affects the classification — the third sentence is used only to justify a filter — but all three are in statements a reader would take at face value.

Section 6 collects, as clearly labelled computations outside the formal development, what we measured about the shape of the remaining problem: an exhaustive uniform-weight sweep of a natural class of nine-vertex graphs that produces no new $\hbar$-imperfect graph, the observation that each of the hardest graphs of the source's table is one tight inequality away from being decided, and the level-by-level map of the relaxation hierarchy that Theorem 1.1 sits in.

2. Preliminaries

2.1. **Realisations.** Fix a graph G on n vertices with adjacency matrix A. A *realisation* of G is a list $S_1, \ldots, S_n$ of Pauli strings on m qubits whose frustration graph is G; by Theorem 3 of [1], in its Appendix A 3, a shortest one has $m = \mathrm{rank}_{\mathbb{F}_2}(A)/2$, so $d = 2^m$ is 8 for the seven- and nine-vertex graphs treated here, except where noted. We use throughout the explicit form

$$P(a, b) = i^{\,|a \wedge b|}\, X^a Z^b, \qquad a, b \in \{0, 1, \ldots, d-1\}, \tag{2}$$

X^a and Z^b being the products of the single-qubit X's and Z's selected by the bit masks a and b, and $|a \wedge b|$ the number of common bits. Each $P(a, b)$ is a Hermitian involution with entries in $\{0, \pm 1, \pm i\}$, and $P(a, b)$ and $P(a', b')$ commute or anticommute according to the parity of the symplectic form $\langle a, b' \rangle + \langle a', b \rangle$; all realisations below are given by their mask pairs, and both the involutivity and the whole frustration pattern are recomputed entry by entry in the verification. The phase in (2) is fixed only up to a sign by the demand that $P(a, b)$ be a Hermitian involution: $i^{\,|a \wedge b|}$ and $i^{\,|a \wedge b| \bmod 2}$ both serve, and differ exactly when $|a \wedge b| \equiv 2 \pmod 4$. Nothing below depends on which is taken, every statement being about the squares $\langle S_i \rangle^2$; the verification of §4 and §5 uses $i^{\,|a \wedge b| \bmod 2}$ and the matrices

R_i of §3 are built with (2), and the two differ, for the mask data of Proposition 4.1, by a sign on one of the seven strings, which changes the sign of two of the R_i and is absorbed by a sign change of the single c variable that the two flipped R_i have in common in (†).

2.2. **The two directions.** Fix $w \in \mathbb{R}^n_+$ and write $\mathbf{1}$ for the all-ones weight. Since $\beta(G,w)$ is a *supremum* over states, a single state ρ with $\sum_i w_i \langle S_i\rangle_\rho^2 > \alpha(G,w)$ is a complete proof of $\hbar$-imperfectness. For a pure state $\rho = |\psi\rangle\langle\psi|/\psi^\dagger\psi$ with $\psi \in \mathbb{Z}[i]^d$ this is the integer inequality

$$\sum_i w_i \big(\psi^\dagger S_i \psi\big)^2 \;>\; \alpha(G,w)\,\big(\psi^\dagger\psi\big)^2, \tag{3}$$

each $\psi^\dagger S_i \psi$ being a rational integer because S_i is Hermitian. That is the shape of §4.

The opposite direction has no finite witness of that kind. Write $H(c) = \sum_i c_i S_i$ and $q(c) = \sum_i c_i^2/w_i$. Fenchel duality applied to the square, $\sum_i w_i a_i^2 = \max_c [2c\cdot a - q(c)]$, together with homogeneity, gives the identity [1] uses for its own hand computation,

$$\beta(G,w) = \max_{c\neq 0} \frac{\lambda_{\max}\big(H(c)\big)^2}{q(c)}. \tag{4}$$

Because the S_i with $ij \in E(G)$ anticommute, their cross terms cancel and

$$H(c)^2 = \|c\|^2\,\mathbf{1} + 2 \sum_{i<j,\ ij\notin E(G)} c_i c_j\, S_i S_j, \tag{5}$$

so $\beta(G,w) \le \alpha$ is equivalent to the positive semidefiniteness, for every $c \in \mathbb{R}^n$, of

$$\big(\alpha\, q(c) - \|c\|^2\big)\mathbf{1} \;-\; 2 \sum_{i<j,\ ij\notin E(G)} c_i c_j\, S_i S_j \;\succeq\; 0. \tag{6}$$

This is a matrix-valued quadratic form in c: equivalently, block positivity of one explicit Hermitian matrix over $\mathbb{Z}[i]$ of size nd — 72×72 or 144×144 for the nine-vertex graphs — which is the entanglement-witness side of the separability duality and is co-NP-hard in general. Section 3 certifies an inequality of exactly the form (6), with α replaced by a rational $B > \alpha$, for the two anti-cycles $\overline{C}_5$ and $\overline{C}_7$.

2.3. **Facets.** Both STAB(G) and the beta body are down-closed convex corners in $\mathbb{R}^n_+$, so $\hbar$-perfectness is the equality $\beta(G,w) = \alpha(G,w)$ at each of the finitely many facet normals $w \ge 0$ of STAB(G) — a finite list, and the reduction [1] itself uses ("G_7 has only one facet whose norm vector contains no zero, and this norm vector is full of 1"). A facet normal with a zero entry reduces to the induced subgraph on its support. Two of the witnesses of §4 live at a *non*-uniform facet, and this is why they cannot be found by a uniform-weight search.

3. An exact upper bound below the theta bound

Throughout this section $G = \overline{C}_n$ with $n \in \{5,7\}$ odd, $w = \mathbf{1}$ and $\alpha = 2$. The non-edges of $\overline{C}_n$ are the edges of the cycle C_n, so (5) reads $H(c)^2 = \|c\|^2\mathbf{1} + 2\sum_{i\in\mathbb{Z}_n} c_i c_{i+1} P_i$ with $P_i = S_i S_{i+1}$, each P_i again a Hermitian involution.

3.1. **The polynomial inequality.** Realify: a vector $\psi = u + iv \in \mathbb{C}^d$ becomes $y = (u,v) \in \mathbb{R}^{2d}$, and a Hermitian $P = A + iB$ becomes the real symmetric $R = \left(\begin{smallmatrix} A & -B \\ B & A \end{smallmatrix}\right)$, so that $\psi^\dagger P \psi = y^{\mathsf{T}} R y$. Setting R_i for the realification of P_i and using (4) with $q(c) = \|c\|^2$, the bound $\beta(\overline{C}_n, \mathbf{1}) \le 1 + \rho$ becomes the assertion that

$$\rho\Big(\sum_i c_i^2\Big)^2 \Big(\sum_k y_k^2\Big) \;\ge\; 2\Big(\sum_i c_i^2\Big) \sum_{i\in\mathbb{Z}_n} c_i c_{i+1}\, y^{\mathsf{T}} R_i y \tag{†}$$

holds for every $(c,y) \in \mathbb{R}^{n+2d}$ — 13 real variables at $n = 5$, 23 at $n = 7$. The extra factor $\sum_i c_i^2$ on both sides is what makes (†) the *second* level of the natural sum-of-squares hierarchy in c: dropping it, and asking instead that $\rho\|c\|^2\mathbf{1} - 2\sum_i c_i c_{i+1} R_i$ be a sum-of-squares matrix, is the first level, and Remark 6.4 records where that level stops.

Certifying (†) as a sum of squares over the monomials $c_a c_b y_k$ requires a Gram matrix of size $\binom{n+1}{2}\cdot 2d$: 120 at $n=5$ and 448 at $n=7$. This is the obstacle. The exact $LDL^{\mathsf T}$ pivots of an $N\times N$ rational Gram matrix are ratios of consecutive leading principal minors, so their numerators grow like the Hadamard bound; measured on certificates of this problem, the common scale is $3\cdot 10^{38}$ already at $N=15$ and $1.2\cdot 10^{47}$ at $N=28$, which extrapolates to entries of about 10^3 digits at $N=448$ and a certificate of hundreds of megabytes.

3.2. **The symmetry reduction.** The way past it is that the 448×448 semidefinite program *is* a 28×28 Hermitian one, exactly, and not by relaxation.

The group. For $k\in\mathbb{Z}_n$ let $s_{ki}=+1$ if R_i and R_k commute and -1 otherwise, and suppose $\varepsilon^k\in\{\pm1\}^n$ satisfies $\varepsilon^k_i\varepsilon^k_{i+1}=s_{ki}$. Then $\sigma_k:(c,y)\mapsto(\varepsilon^k\odot c,\,R_k y)$ preserves both sides of (†): $\sum c_i^2$ is preserved because ε^k has entries ±1, $\sum y_k^2$ because R_k is a symmetric involution, and $y^{\mathsf T}R_kR_iR_ky=s_{ki}\,y^{\mathsf T}R_iy$. Such an ε^k exists exactly when $\prod_i s_{ki}=1$, i.e. when an even number of the R_i anticommute with R_k. At $n=5$ all pairs anticommute; at $n=7$ the commuting pairs are exactly those at cycle distance 3, so each R_k anticommutes with four others. Both patterns are verified.

Averaging. Each σ_k acts on the monomial basis by a signed permutation, so averaging a feasible Gram matrix over the group $\langle\sigma_k\rangle$ leaves it feasible and positive semidefinite. Writing $Q^{\alpha\beta}$ for the $2d\times 2d$ block of Q indexed by the degree-two c-monomials α,β, invariance says

$$R_kQ^{\alpha\beta}R_k=\chi_\alpha(k)\chi_\beta(k)\,Q^{\alpha\beta},\qquad \chi_\alpha(k)=\varepsilon^k_{\alpha_1}\varepsilon^k_{\alpha_2}.$$

Schur, concretely. Conjugation by the R_k acts on the Pauli basis of $M_d(\mathbb{C})$ by signs, and the sign pattern of P_u is $(\langle u,r_k\rangle)_k$ for the symplectic form, r_k being the mask of R_k. The masks of the R_k span the symplectic space, so at $d=8$ the 64 Pauli matrices have 64 distinct sign patterns and each pattern is realised by exactly one Pauli. Hence every $Q^{\alpha\beta}$ is a single Pauli matrix up to a complex scalar.

The change of basis. Let $u(\alpha)$ be the mask with sign pattern χ_α and O_α the realification of $P_{u(\alpha)}$ — a symmetric orthogonal signed permutation commuting with J, the realification of $i\mathbf{1}$. Then $Q^{\alpha\beta}=O_\alpha(X_{\alpha\beta}\mathbf{1}+Y_{\alpha\beta}J)O_\beta$, that is

$$Q=\hat O\,(X\otimes\mathbf{1}+Y\otimes J)\,\hat O,\qquad \hat O=\operatorname{blockdiag}(O_\alpha),\qquad Z:=X+iY \text{ Hermitian},\tag{7}$$

and $Q\succeq 0\iff Z\succeq 0$, because $\mathbb{R}^{\binom{n+1}{2}}\otimes\mathbb{R}^{2d}$ with complex structure $\mathbf{1}\otimes J$ is $\mathbb{C}^{\binom{n+1}{2}}\otimes\mathbb{C}^d$, on which $X\otimes\mathbf{1}+Y\otimes J$ acts as $(X+iY)\otimes\mathbf{1}$. So the 448×448 program is a 28×28 Hermitian one, and only Z is ever factored.

The certificate. From the exact $LDL^\dagger$ of Z over $\mathbb{Q}(i)$, writing $Z=\sum_s w_s\zeta_s\zeta_s^\dagger$ with $\zeta_s=p_s+iq_s$, the identity becomes

$$D\cdot\big[(\dagger)\text{ as a polynomial}\big]=\sum_{j=1}^{2d}\sum_s w_s\Big(\sum_\alpha c^\alpha\big[p_{s,\alpha}(O_\alpha y)_j-q_{s,\alpha}(JO_\alpha y)_j\big]\Big)^2,\tag{8}$$

one real square per component of each complex square. Since O_α and JO_α are signed permutations, each of the $2d$ linear maps in (8) is a sparse integer matrix of size $\binom{n+1}{2}\cdot 2d\times 2\binom{n+1}{2}$ with only $2\binom{n+1}{2}$ nonzero entries, and all $2d$ of them share one list of $\binom{n+1}{2}$ weighted integer vectors. The 10^3-digit growth predicted above is then confined to 29 scalars: the common denominator D of the $n=7$ certificate has 1052 decimal digits and the 28 weights w_s have between 915 and 1052, while no entry of the 28 integer vectors ζ_s exceeds 69 digits. That is what makes the file small — 1052 digits appear 29 times rather than in $448^2=200\,704$ Gram entries — and the emitted $n=7$ certificate is 53 kB. At $n=5$ the corresponding figures are 204 digits for D, 158 to 205 for the 15 weights, 24 digits for the vectors, and 8 kB.

3.3. **What is verified, and what is not.** Nothing in §3.2 is trusted: the symmetry argument explains where the data come from, and the verification re-checks the polynomial identity (8) from scratch, the only structural conditions being $D>0$ and $w_s\ge 0$. Any integer face matrix at all would be sound, since $\sum_s w_s(Vu_s)(Vu_s)^{\mathsf T}$ is positive semidefinite for every integer V and every $w_s\ge 0$; a wrong one can only make the identity fail. The general soundness statement behind this is Lemma 3.1 below.

Two steps between the verified statement and Theorem 1.1 are carried out by hand, in §2 and above, and are *not* part of the formal development. What is verified is (†) as a polynomial inequality about explicit integer matrices $R_0, \dots, R_{n-1}$, together with the facts that these are symmetric involutions whose commutation pattern is that of $\overline{C}_n$ — which is the only property of the realisation that §3.2 uses. The two unformalised steps are the identity (4) with (5), which is the source's own device, and the identification of the R_i with the realifications of S_iS_{i+1} for the explicit mask data. Both are elementary; neither is a computation. We state this plainly because the distinction matters for a reader who wants to know what a machine has and has not seen.

Lemma 3.1 (soundness of the certificate format)**.** *Let P be a system of polynomial inequality and equality constraints with integer coefficients, and let a certificate consist of an integer D, integer matrices $V^{(1)}, \dots, V^{(t)}$, and weighted integer vectors (w_s, u_s). If $D > 0$, all $w_s \geq 0$, and the polynomial identity $D \cdot p = \sum_j \sum_s w_s \big(\sum_a (V^{(j)}u_s)_a B_a\big)^2 +$ (multiplier terms) holds identically, then $p \geq 0$ at every real point of the set cut out by P.*

Lemma 3.1 is standard — it is Positivstellensatz bookkeeping — and is recorded here only because it is the statement the machine checks against; its proof is one rewrite per square. The technique it packages is not new either: facial and symmetry reduction of semidefinite and sum-of-squares programs is classical [5], and exact rational sum-of-squares by round-and-project is [6]. What is new in Theorem 1.1 is the number.

Proof of Theorem 1.1 and Corollary 1.2. For $n = 5$ take $\rho = \frac{529}{500}$ and for $n = 7$ take $\rho = \frac{221}{200}$, so that $1+\rho$ is $\frac{1029}{500}$ and $\frac{421}{200}$ respectively. The certificates are produced as in §3.2: solve the reduced $\binom{n+1}{2} \times \binom{n+1}{2}$ Hermitian program numerically, round to a rational matrix, factor it exactly over $\mathbb{Q}(i)$, and push the factors through (7). The identity (8) is then a finite check — 448 squares of 56 terms each at $n = 7$, about $1.4 \cdot 10^6$ monomial operations, 32 seconds — and it is this check, not the search that produced the data, that carries the proof, by Lemma 3.1. With (†) in hand, (4) and (5) give $\beta(\overline{C}_n, \mathbf{1}) \leq 1 + \rho$ as in §3.3. The numerical comparison $\frac{421}{200} = 2.105 < 2.1099162642\dots$ is (1). Corollary 1.2 combines this with the state of Proposition 4.1, whose exact value is $\frac{327090970}{156225001}$. □

3.4. **Sharpness and controls.** The bound of Theorem 1.1 is not tight: 2.105 against the value $(9 + 4\sqrt{2})/7 = 2.0938363\dots$ that [2] gives for $\beta(\overline{C}_7, \mathbf{1})$. How far from tight is itself certified, from the other side of the same polynomial.

Proposition 3.2. *The inequality (†) at $\rho = \frac{109}{100}$ — that is, the bound $\beta(\overline{C}_7, \mathbf{1}) \leq 2.09$ — is false: at the explicit integer point obtained by realifying the state of Proposition 4.1 the two sides of (†) stand in the ratio $\frac{447155038123}{408831003403} = 1.0937405294\dots$, so no $\rho \leq \frac{109}{100}$ is admissible.*

So the certified bracket for the level-two relaxation of $\overline{C}_7$ is $(2.0937405\dots, 2.105]$, and the witness of Proposition 4.1 reappears as the refuting point. Two further controls are verified: the certificate of Theorem 1.1 is rejected when it is offered for a strictly stronger target, at both $n = 5$ and $n = 7$, so it is not an artefact of a checker that accepts anything.

Remark 3.3 (earlier and weaker certificates; outside the numbered development). A cheaper certificate of the same kind, using a Positivstellensatz multiplier on $\sum_k y_k^2$ instead of the second level, gives $\beta(\overline{C}_5, \mathbf{1}) \leq \frac{103}{50}$ and $\beta(\overline{C}_7, \mathbf{1}) \leq \frac{14}{5}$; its Gram matrices are 15×15 and 28×28 and it needs no symmetry reduction. At $n = 5$ that already beats $\vartheta(\overline{C}_5)$, so the $n = 5$ half of Theorem 1.1 is an improvement of a bound rather than the first of its kind; at $n = 7$ it is worth 2.8, far above $\vartheta(\overline{C}_7)$, and it is Theorem 1.1 that first goes below the theta bound. The obstruction to sharpening the cheaper route is a Motzkin-type gap: it is admissible exactly when $\kappa(\sum c_i^2)^2 - \sum_i c_i^2 c_{i+1}^2$ is a sum of squares, whose true threshold is $\kappa = \frac{1}{4}$ but whose sum-of-squares threshold is $0.27636\dots$ at $n = 5$ and $0.26303\dots$ at $n = 7$ (bisected numerically).

4. Exact witnesses of $\hbar$-imperfectness

Each statement of this section is an instance of (3) together with two exact side conditions: that the listed mask pairs realise the stated graph (every $P(a, b)$ is a Hermitian involution and the pairwise

anticommutation pattern is the graph's, recomputed entry by entry over $\mathbb{C}^d$), and that the weighted independence number is the stated value (a sweep of all 2^n subsets). All arithmetic is over $\mathbb{Z}[i]$; no floating-point number occurs anywhere in these proofs.

Proposition 4.1 (the anti-heptagon)**.** *The seven three-qubit Pauli strings with mask pairs* $(4,0)$, $(2,0)$, $(0,4)$, $(0,6)$, $(7,6)$, $(4,7)$, $(5,3)$ *in the notation* (2) *realise* $\overline{C}_7$*, whose independence number is* 2*, and the explicit Gaussian-integer state* $\psi \in \mathbb{Z}[i]^8$ *recorded in the development satisfies* $\sum_{i=1}^{7}(\psi^\dagger S_i\psi)^2 > 2(\psi^\dagger\psi)^2$*. Its exact value is*

$$\frac{\sum_i(\psi^\dagger S_i\psi)^2}{(\psi^\dagger\psi)^2} = \frac{327090970}{156225001} = 2.0937171893\ldots,$$

so $\beta(\overline{C}_7, \mathbf{1}) \geq \frac{327090970}{156225001} > 2$.

That $\overline{C}_7$ is $\hbar$-imperfect is stated in [1] and attributed there to [2], which exhibits a realisation by seven three-qubit strings and a state with $\sum_i\langle S_i\rangle^2 = (9+4\sqrt{2})/7 = 2.0938363\ldots$; what Proposition 4.1 adds is a Gaussian-integer witness, checkable in the kernel over $\mathbb{Z}[i]$ with no algebraic number in sight, which is what Corollary 1.2 uses. Its seven integers $\psi^\dagger S_i\psi$ are 6914, 6876, 6695, −6695, 6876, 6914, −6876 — individual signs depending on the phase convention of (2), their squares not — and $\psi^\dagger\psi = 12499$.

Proposition 4.2 (the anti-nonagon)**.** *Nine explicit three-qubit Pauli strings realise* $\overline{C}_9$*, whose independence number is* 2*, and an explicit Gaussian-integer state gives* $\beta(\overline{C}_9, \mathbf{1}) \geq \frac{2641054}{1283689} = 2.0573939\ldots > 2$.

Again the claim itself is the source's — it asserts $2 = \alpha(\overline{C}_n) < \beta(\overline{C}_n) \leq \vartheta(\overline{C}_n)$ for every anti-cycle, and $\frac{2641054}{1283689} \leq \vartheta(\overline{C}_9) = 2.0641777\ldots$ as it must — and only the witness is ours. The next four are the graphs [1] determines numerically.

Theorem 4.3 (the two nine-vertex graphs at the uniform weight)**.** *Let* F_b *and* F_c *be as above, with* 25 *and* 26 *edges respectively. Each has independence number* 2*, each is realised by nine explicit three-qubit Pauli strings, and explicit Gaussian-integer states give*

$$\beta(F_b, \mathbf{1}) \;\geq\; \frac{45101177}{22227852} = 2.0290389\ldots, \qquad \beta(F_c, \mathbf{1}) \;\geq\; \frac{1044367973}{514699969} = 2.0290810\ldots,$$

so both are $\hbar$*-imperfect.*

The two graphs are non-isomorphic, having different numbers of edges, yet the see-saw of §6 returns beta numbers for them that agree to six decimal places, both near 2.0291773; we neither explain that coincidence nor claim it.

Theorem 4.4 (an eight-vertex graph, at a weighted facet)**.** *Let* G_8 *be as above, with* 15 *edges, and let* $\omega = (1,1,1,1,2,1,2,1)$*. Then* $\alpha(G_8, \omega) = 3$*, and an explicit Gaussian-integer state on eight three-qubit Pauli strings realising* G_8 *gives*

$$\beta(G_8, \omega) \;\geq\; \frac{167391146}{54982225} = 3.0444592\ldots \;>\; 3,$$

so G_8 *is* $\hbar$*-imperfect. The same state gives no violation at the uniform weight.*

Theorem 4.5 (a nine-vertex graph, at a weighted facet)**.** *Let* F_e *be as above, with* 18 *edges, and let* $\omega = (1,1,1,2,1,1,1,1,1)$*. Then* $\alpha(F_e, \omega) = 3$ *and an explicit Gaussian-integer state gives* $\beta(F_e, \omega) \geq \frac{315096058}{103489929} = 3.0447026\ldots > 3$.

The weights in Theorems 4.4 and 4.5 are the point. Both graphs also have $\alpha(G, \mathbf{1}) = 3$, and both states are checked in the kernel to give no violation at $w = \mathbf{1}$; the violated inequality is a non-uniform facet of STAB(G), and locating it required enumerating the facet normals with nonnegative support (235 of them for G_8, 12 for F_e) rather than searching states at the uniform weight. This is also why the sweep of Remark 6.1, which is a uniform-weight sweep, does not contain them.

Two controls pin what the witnesses do and do not prove.

Proposition 4.6 (negative controls)**.** *The state of Proposition 4.1 certifies* $\beta(\overline{C}_7, \mathbf{1}) > \frac{209}{100}$ *but not* $\beta(\overline{C}_7, \mathbf{1}) > \frac{21}{10}$*; the state of Proposition 4.2 does not certify* $\beta(\overline{C}_9, \mathbf{1}) > \frac{206}{100}$*; toggling a single edge of* $\overline{C}_7$ *makes the realisation check fail; and* $\alpha(\overline{C}_7)$ *is neither* 1 *nor* 3*.*

The upper controls are calibrated against (1): $\vartheta(\overline{C}_7) = 2.1099\ldots$ and $\vartheta(\overline{C}_9) = 2.0642\ldots$, so a witness claiming more would be contradicting [1] rather than confirming it.

Proposition 4.7 (a positive control at the boundary)**.** *Let G_7 be the seven-vertex graph whose $\hbar$-perfectness* [1] *proves by hand. Seven explicit three-qubit Pauli strings realise it, $\alpha(G_7, \mathbf{1}) = 2$, and the best state produced by the same search certifies $\sum_i \langle S_i\rangle^2 > \frac{199}{100}$ and does not cross* 2*.*

The name G_7 is [1]'s own: its Appendix A 2, headed "$\hbar$-perfectness of G_7 and G_{15}", introduces the graph of its Figure 7 as "named as G_7 here" and closes an unnumbered proof — the statement carries no number in that version — with "we conclude that $\beta(G_7) = 2$". We took the graph from that figure, whose vertices are labelled by an explicit realisation; it is isomorphic to the seven-vertex graph used here, and $\beta(G_7) = 2$ is also a theorem of [2], proved there by embedding G_7 in the lexicographic product $C_5[K_2]$. Proposition 4.7 is the control that makes the rest of the section meaningful: a pipeline that could not tell G_7 from $\overline{C}_7$ would prove nothing at all. Numerically the same search returns $\beta(G_7, \mathbf{1}) = 2.000000000000$ over 1500 restarts, in agreement with those proofs.

Remark 4.8 (one graph not certified; outside the numbered development)*.* The remaining graph of the source's figure, F_d, needs a four-qubit realisation, $d = 16$. Its certificate is computed and exact — $\omega = (1, 1, 1, 1, 1, 2, 1, 2, 1)$, $\alpha = 3$, $\beta \geq \frac{555850571}{184334929} = 3.0154381\ldots$ — but the realisation check at $d = 16$ costs eight times the integer arithmetic of a $d = 8$ graph, a measured 19 GB in a single statement, and it was therefore not carried out in the kernel. It is recorded here as a computation, not as a theorem.

Remark 4.9 (where the figure graphs come from)*.* Figure 11 of [1] is included in its e-print as a vector PDF and carries no machine-readable edge list. The five graphs used above were therefore read off the content stream of that file: the straight strokes are the edges and their endpoints are the vertices, giving $8, 9, 9, 9, 9$ vertices and $15, 25, 26, 20, 18$ edges for parts (a)–(e). The extraction was carried out twice, independently, and the two results agree; the drawn graphs are isomorphic to the five used here, and the $\mathbb{F}_2$-ranks of their adjacency matrices are $6, 6, 6, 8, 6$, which is why part (d) alone needs four qubits. A reader with the authors' code release can check the same lists against it.

5. Three errata

The three statements of this section are facts about data transcribed from the drawings and tables of [1]; each is a finite check. The table concerned is its Table II, "The top 12 graphs with the largest gaps", all on nine vertices; its twelve entries are drawn as `tikz` pictures in the e-print source, and the adjacency used below is the one those pictures draw. Eleven of the twelve entries are distinct as drawn.

Theorem 5.1 (independence number in the nine-vertex table)**.** *The second graph of Table II of* [1]*, as drawn, has $\alpha(G, \omega) = 3$ for the weight $\omega = (1, 1, 1, 1, 1, 1, 1, 1, 2)$ printed beside it, not the printed value* 2*.*

Proof. In the drawing, vertex 9 — the one drawn black, carrying weight 2 — is adjacent to vertices $1, 3, 4, 5, 6, 7, 8$ and not to vertex 2. Hence $\{2, 9\}$ is independent of weight 3, and a sweep of all 2^9 subsets shows that no independent set weighs more. □

The third entry of the same table is the second one plus the two edges $\{2, 9\}$ and $\{5, 8\}$, which makes two dropped coordinate pairs in the second picture the likely cause. Independently, and exactly: the printed ω is not a facet normal of the drawn graph at all. Its face $\{x \in \mathrm{STAB}(G) : \omega \cdot x = 3\}$ is the single vertex $e_2 + e_9$, since $\{2, 9\}$ is the only independent set of ω-weight 3, whereas a facet of $\mathrm{STAB}(G) \subset \mathbb{R}^9$ is 8-dimensional; the *uniform* weight, by contrast, does cut out a facet there, its 13 tight independent sets spanning an 8-dimensional face, with $\alpha = 2$.

Theorem 5.2 (two identical entries with different bounds)**.** *The third and fourth graphs of Table II are the same graph — their drawings agree edge for edge, and the two `\draw` lists that produce them are identical character for character — yet they are given different numerical upper bounds γ_{ub}, namely* 2.00175 *and* 2.00118*. Both entries print the same weight* $(1, \ldots, 1, 2)$*, the same $\alpha = 2$ and the same see-saw value* 2.00000*.*

One of the two drawings, or one of the two numbers, is therefore wrong; we do not know which. Both entries also carry 25 edges and a three-qubit realisation.

Theorem 5.3 (a false sentence in the appendix)**.** *$\overline{C}_9$ is ħ-imperfect and contains neither $\overline{C}_7$ nor G_8 as an induced subgraph. Hence the sentence of* [1]*, in its Appendix B 3,*

> *"Otherwise, G cannot have other ħ-imperfect graphs as induced subgraphs, since all ħ-imperfect graphs with no more than 9 vertices have either $\overline{C}_7$ or G_8 as an induced subgraph"*

is false as written.

Proof. ħ-imperfectness of $\overline{C}_9$ is Proposition 4.2. For the induced subgraphs no isomorphism test is needed, only two invariants, each checked over all subsets: $\overline{C}_7$ is 4-regular on seven vertices and no seven-vertex induced subgraph of $\overline{C}_9$ is 4-regular, while G_8 has 15 edges on eight vertices and every eight-vertex induced subgraph of $\overline{C}_9$ has 21. □

The graphs $F_b, \ldots, F_e$, which the next paragraph of [1] names as ħ-imperfect, are counterexamples to the same sentence. The intended bound is presumably eight rather than nine vertices, and the surrounding procedure is unaffected, the sentence being used only to justify a filter.

6. What is finite, and what is not

The results above are all finite checks, but the problem they sit in is not, and it is worth being exact about where the boundary runs. Everything in this section is a floating-point or convex-hull computation, reported as such; none of it is part of the verified development, and none of it is claimed as a theorem.

Remark 6.1 (an exhaustive uniform-weight sweep). A graph G on nine vertices has $\alpha(G) = 2$ exactly when $\overline{G}$ is triangle-free and has an edge, so that class sits inside the complements of the 1897 triangle-free nine-vertex graphs (OEIS `A006785` [8], whose $a(9) = 1897$ was reproduced here with `nauty` [7] as an enumeration control), missing only the edgeless $\overline{G}$, where $\alpha(G) = 1$. Computing $\beta(G, \mathbf{1})$ for all 1897 of them by the see-saw iteration of (4), to about 10^{-10}, gives exactly 69 graphs with $\beta(G, \mathbf{1}) > 2$: 66 contain an induced $\overline{C}_7$, and the remaining three are $\overline{C}_9$ and the graphs F_b, F_c of Theorem 4.3. **No new ħ-imperfect graph.** The detector was calibrated on those three known positives at a quarter of the budget used for the sweep; total cost about one core-hour. The scope of this negative is narrow and worth stating: it covers the uniform weight only, and only the graphs with $\alpha(G) = 2$. It says nothing about the undetermined graphs with $\alpha(G) \geq 3$, nor about non-uniform facets — which is precisely where the graphs of Theorems 4.4 and 4.5 hide.

Remark 6.2 (the undetermined graphs are one inequality each). For each of the eleven distinct graphs of the source's Table II, STAB(G) has between 356 and 976 facets, of which exactly *one* has full support. A facet normal with a zero entry reduces to the induced subgraph on its support, and [1] has determined every graph on at most eight vertices; so deciding one of these graphs is one tight instance of (6), not a facet enumeration. At that facet $\beta = \alpha$ to 10^{-14} in every case, consistent with all eleven being ħ-perfect; were that true of all 78 undetermined graphs, the interval $[259583, 259661]$ would collapse to its upper end. For ten of the eleven the unique full-support normal is the ω printed in the source; the exception is the entry of Theorem 5.1.

Remark 6.3 (why a subdivision certificate cannot work). Deciding one undetermined graph is block positivity of an explicit 72×72 or 144×144 Hermitian matrix over $\mathbb{Z}[i]$, *at the boundary of the cone*, since ħ-perfectness means $\beta = \alpha$ exactly. The cheap sufficient condition — positive semidefiniteness of the full matrix, the degree-$(1,1)$ certificate — fails everywhere it was tried, including on the ħ-perfect G_7 (smallest eigenvalue -1.199, and -1.000 for $\overline{C}_5, \overline{C}_7, \overline{C}_9$), and the fractional clique cover gives $2.5, 2.333, 2.25, 2.5$ for $C_5, \overline{C}_7, \overline{C}_9, G_7$ against a target of 2. A Bernstein-coefficient certificate on a box is worse than expensive: the leading minor route needs $33^9 \approx 4.6 \cdot 10^{13}$ root coefficients, and, the bound being attained, no finite subdivision can reach it at any budget. An exact rational sum-of-squares certificate — the route of §3 — is the honest attack.

Remark 6.4 (the hierarchy, level by level). Bisecting the threshold B at which each relaxation of (†) first becomes feasible, by alternating projections, gives the following map (numerical; brackets are the bisection intervals). The plain block-matrix test $\rho\mathbf{1} - C \succeq 0$ is worth $B = 3$ at every n. The first level, sum-of-squares over the monomials $c_i y_k$, has threshold in $(2.23605957, 2.23606873)$ at $n = 5$ and $(2.1099121, 2.1100586)$ at $n = 7$ — each interval containing $\vartheta(\overline{C}_n)$ from (1) and nothing else of interest, so *the theta bound appears to be exactly the first-level bound here*, and the second level is the first one that can improve on it. This is a measurement, not a proof; it is what one expects, since that relaxation is $\max 2\sum_i v_i^{\mathsf{T}} R_i v_{i+1}$ over $\sum_i \|v_i\|^2 = 1$, the "one y per coordinate" relaxation of β. The second level, over $c_a c_b y_k$, has threshold in $(2.05474396, 2.05476379)$ at $n = 5$ and $(2.10097656, 2.10126953)$ at $n = 7$; the exact bounds of Theorem 1.1 sit just above these. One methodological warning came out of the same computation: solving the $n = 5$ second level naively in the full 120-monomial basis reports a threshold of 2.1524, which is simply wrong, since 2.058 is certified exactly in Theorem 1.1. Alternating projections onto a subspace and a cone converge like $k^{-1/2}$, so a threshold read off a large-basis run is an upper bound on the threshold and never the threshold.

Remark 6.5 (a sufficient criterion, proved but of small reach; not formalised). If $\overline{G}$ is triangle-free and every two disjoint edges of $\overline{G}$ span an odd number of $\overline{G}$-edges, then $\beta(G, \mathbf{1}) = \alpha(G, \mathbf{1}) = 2$. Indeed the hypothesis says the involutions $P_{ij} = S_i S_j$ over the non-edges pairwise anticommute, so $(\sum_e t_e P_e)^2 = \|t\|^2 \mathbf{1}$ with $t_{ij} = c_i c_j$; and $\sum_{ij \in E(\overline{G})} x_i x_j \le \frac{1}{2}(1 - 1/\omega(\overline{G})) = \frac{1}{4}$ on the simplex is Motzkin–Straus [4] for triangle-free $\overline{G}$, so (5) gives $\lambda_{\max}(H(c)^2) \le 2\|c\|^2$. The criterion is not among the graph-operation closure properties or the perfect ⇒ h-perfect ⇒ $\hbar$-perfect chain of [1], which are different mechanisms. Its reach, however, is tiny: enumerated with `nauty` [7], the hypothesis holds for $2, 4, 7, 10, 13, 17, 21$ of the $3, 7, 14, 38, 107, 410, 1897$ triangle-free graphs on $3, \dots, 9$ vertices, the edgeless one excluded in each case since there $\alpha(G) = 1$. At nine vertices every one of the 21 has a pair of vertices that is a copy or a split in the sense of [1] — two vertices with the same neighbourhood, adjacent or not — and 17 of the 21 have $\overline{G}$ disconnected as well, so G is a fully connected join; either way each is already reducible there, and the criterion settles none of the 78 undetermined graphs. We record it without a formal proof and without a claim of novelty: the underlying identity for anticommuting observables is close to standard uncertainty-relation folklore, and the literature [1] cites for that circle of ideas was not read here.

7. Verification

Every numbered statement above — Theorems 1.1, 4.3, 4.4, 4.5, 5.1, 5.2, 5.3, Corollary 1.2, Lemma 3.1, and Propositions 3.2, 4.1, 4.2, 4.6, 4.7 — is checked in the form stated by the Lean 4 proof assistant [9] over the Mathlib library [10], subject to the two hand steps that §3.3 isolates for Theorem 1.1 and Corollary 1.2. The development is four graph-specific modules and a general-purpose certificate layer; the graph-specific arithmetic is over $\mathbb{Z}[i]$ with Pauli strings stored as mask pairs and matrices built entry by entry from (2), in the phase convention noted in §2, so what is checked is the definition and not a numerical surrogate. It contains no `sorry` and declares no axiom of its own. The axiom sets divide in two. The 36 statements behind §4 and §5 — the realisation checks, the independence numbers over all 2^n subsets, the integer inequalities (3), the controls, and the induced-subgraph invariants — return `propext` alone, and one of them no axioms at all; there is no compiled evaluation, no external solver and no floating-point arithmetic in any of them. The two certificates of Theorem 1.1 return `propext`, `Classical.choice`, `Quot.sound` and, for the single evaluation of the polynomial identity (8), one compiled-evaluation axiom each; the soundness statement Lemma 3.1 and the whole certificate layer it belongs to use no compiled evaluation at all. For $n = 5$ the compiled evaluation is removable: the same check runs in the kernel in 98 seconds, at 12.1 GB of memory rather than 7.2 GB, and the compiled route was kept for the smaller footprint. Elaboration of the two certificates together takes 32 seconds of processor time; the witness modules take three minutes and 4.3 GB, the dominant cost being the entry-by-entry matrix products of the realisation checks. The source of the development is currently held in a private repository: repository reference to be added at submission.

On the reference list. References [1] and [2] were both read here in full from their arXiv LaTeX e-prints, every passage quoted above was re-read in those sources, and the statement, table, figure and appendix numbers cited above were taken from a build of each e-print rather than from a secondary listing. Both were re-checked on the arXiv on 3 September 2026: [1] is still v1 with no journal reference, and [2] is at v2 with the journal reference given below, which agrees with Crossref. The bibliographic data of [3], [4], [5], [6], [7], [9] and [10] were each verified at Crossref on the same day; [3] and [4] were also checked at zbMATH, and the volume number of [9] at dblp. Reference [8] was consulted online, and its $a(9) = 1897$ agrees with the `nauty` enumeration of Remark 6.1.

References

[1] Z.-P. Xu, J. Wang, Q. Ye, G. Koßmann, R. Schwonnek and A. Winter, *Simultaneous variances of Pauli strings, weighted independence numbers, and a new kind of perfection of graphs*, arXiv:2511.13531v1 [quant-ph], 17 November 2025. Read here in full from the arXiv LaTeX e-print of v1; all quotations above are from that source and are pinned to v1. Its accompanying code is at `https://github.com/wangjie212/BetaNumber`, which was not used here.

[2] Z.-P. Xu, R. Schwonnek and A. Winter, *Bounding the joint numerical range of Pauli strings by graph parameters*, PRX Quantum **5** (2024), no. 2, 020318, doi:10.1103/PRXQuantum.5.020318; preprint arXiv:2308.00753v2 [quant-ph]. Read here in full from the arXiv LaTeX e-print of v2; the anti-heptagon state, the value $(9 + 4\sqrt{2})/7$, the sentence beginning "Numerically, we have" and the theorem $\beta(G_7) = 2$ are all in its section *The beta parameter*, and are pinned to that version.

[3] L. Lovász, *On the Shannon capacity of a graph*, IEEE Transactions on Information Theory **25** (1979), no. 1, 1–7, doi:10.1109/TIT.1979.1055985. Cited for attribution only; it is also the reference [1] gives for its ϑ.

[4] T. S. Motzkin and E. G. Straus, *Maxima for graphs and a new proof of a theorem of Turán*, Canad. J. Math. **17** (1965), 533–540, doi:10.4153/CJM-1965-053-6. Cited for attribution only.

[5] K. Gatermann and P. A. Parrilo, *Symmetry groups, semidefinite programs, and sums of squares*, J. Pure Appl. Algebra **192** (2004), no. 1–3, 95–128, doi:10.1016/j.jpaa.2003.12.011. Cited for attribution only.

[6] H. Peyrl and P. A. Parrilo, *Computing sum of squares decompositions with rational coefficients*, Theoret. Comput. Sci. **409** (2008), no. 2, 269–281, doi:10.1016/j.tcs.2008.09.025. Cited for attribution only.

[7] B. D. McKay and A. Piperno, *Practical graph isomorphism, II*, J. Symbolic Comput. **60** (2014), 94–112, doi:10.1016/j.jsc.2013.09.003. Used as the enumeration and canonical-form tool of Remark 6.1.

[8] OEIS Foundation Inc., *The On-Line Encyclopedia of Integer Sequences*, sequence `A006785`, "Number of triangle-free graphs on n vertices", consulted 3 September 2026.

[9] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: Automated Deduction — CADE 28, Lecture Notes in Computer Science **12699**, Springer, 2021, 625–635, doi:10.1007/978-3-030-79876-5_37.

[10] The mathlib Community, *The Lean mathematical library*, in: Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020), ACM, 2020, 367–381, doi:10.1145/3372885.3373824.