AN 83/84 APPROXIMATION FOR A CUBIC THREE-OCCURRENCE HORN CLASS

FRANK VEGA

Abstract. We study an unweighted MAX-HORN-2SAT class, which is also a class of MAX 2-SAT instances, in which every Boolean variable has exactly three literal occurrences. For a simple cubic graph $H = (V, E)$ the formula T_H satisfies

$$\mathrm{OPT}(T_H) = 4|V| + |E| + \alpha(H),$$

and every assignment of value A is repaired in polynomial time to an independent set of size at least $A - (4|V|+|E|)$. This gives a polynomial-time 83/84 approximation from a cubic independent-set approximation and a near-bipartite cut algorithm, an L-reduction from cubic MIS, and APX-completeness. For simple d-regular graphs the optimum is $(d + 1)|V| + |E| + \alpha(H)$ and the factors are: $(3d + 2)/(3d + 3)$ from the all-true assignment, which is at least 83/84 for $d \geq 27$; 83/84 from the Berman–Fujito guarantee for $d \geq 22$ and, with a semidefinite cut rounding, for $d = 21$; and $(21d + 20)/(21d + 21)$ from any 6/7-approximate independent set. For arbitrary simple graphs the optimum is $n + 3m + \alpha(H)$.

MAX-CUT on a d-regular graph G reduces to MAX 2-SAT with the clauses $(x_u \vee x_v)$ and $(\neg x_u \vee \neg x_v)$ per edge, of value $m + \mathrm{cut}(S)$. The conflict graph $H(G)$ is $(d + 1)$-regular and $\alpha(H(G)) = m + \mathrm{maxcut}(G)$. A Goemans–Williamson cut gives on $T_{H(G)}$ the factor $1 - \frac{1-\alpha_\varepsilon}{6d+12} \geq 83/84$ for every $d \geq 1$, and a γ-dependent analysis gives $1 - \frac{0.1053}{6d+11.5}$, where $\alpha_\varepsilon = \alpha_{\mathrm{GW}} - 10^{-5}$ and $\gamma = \mathrm{maxcut}(G)/m$. For arbitrary simple graphs without isolated vertices the factor exceeds 0.9932. A cut guarantee c_d on graphs of maximum degree d gives $1 - \frac{1-c_d}{6d+12}$, which exceeds 0.9977 in expectation for $d = 3$. A ratio on $T_{H(G)}$ and a cut guarantee are equivalent, so the cut recovered is the Goemans–Williamson or bounded-degree cut used, and the MAX-CUT guarantee is not improved.

The main consequence is that on the class $\{T_{H(G)}\}$ of three-occurrence Horn formulas obtained from MAX-CUT instances the proved factor 83/84 exceeds the Unique Games thresholds ≈ 0.94616 and ≈ 0.94017 that hold for all instances of weighted MAX-HORN-2SAT and of MAX 2-SAT.

For a restricted hard-split family with unit vertex rewards and edge-clause weights at least one we obtain 27/28, and 257/260 with edge weights four. The Unique Games thresholds of Brakensiek–Huang–Zwick (≈ 0.94616), Lewin–Livnat–Zwick (≈ 0.94017) for MAX 2-SAT, and Khot–Kindler–Mossel–O'Donnell (≈ 0.87856) for MAX-CUT quantify over all instances; the guarantees above concern classes with three occurrences per variable, exceed the first two thresholds, and are consistent with the conjecture. Max-Ones 2-SAT for negative (Horn) formulas equals Maximum Independent Set and Min-Ones 2-SAT for positive formulas equals Minimum Vertex Cover, with the same optima, so the published hardness of these two problems applies to them, and $\mathrm{OPT}(T_H) = 4|V| + |E| + \mathrm{MaxOnes}(N_H)$ for the cubic negative formula N_H. The Lean 4 development formalizes the formula T_H for every degree and every degree sequence, its score decomposition, repair and optimum identity, the arithmetic transfers, the MAX-CUT clause count and lifting step, and the hard-score optima $4m + \alpha$ and $m + \alpha$; the Goemans–Williamson, Berman–Fujito and Unique Games statements are explicit hypotheses taken from the literature.

hexagon:2611.00827v2

1. Introduction

MAX-HORN-2SAT maximizes the number of satisfied unary and binary Horn clauses. We use the polarity convention of MAX $\{1, 2\}$-HORN-SAT in which a clause does not contain two negative literals. A global complementation converts this to the conventional logical definition without changing objective values. Every clause has at most two literals, so each formula is also an instance of MAX 2-SAT, and every guarantee proved below for the Horn class is a guarantee for this class of MAX 2-SAT instances, for cubic graphs, d-regular graphs, and non-regular conflict graphs alike. On the class $\{T_{H(G)}\}$ of formulas obtained from the conflict graphs of d-regular MAX-CUT instances, Theorem 15 gives the factor $1 - \frac{1-\alpha_\varepsilon}{6d+12} \geq 83/84$ for every $d \geq 1$, and Theorem 21 extends this conflict-graph guarantee beyond regular base graphs G. The

Date: November 21, 2026.

2020 *Mathematics Subject Classification.* 68Q25, 68W25, 68Q15, 68Q17, 05C69.

Key words and phrases. approximate-computing, max-sat, independent-set-problem, lean4, np-hard, satisfiability, unique-games-conjecture.

construction below starts from cubic Maximum Independent Set (MIS), uses four clones for every graph vertex, and connects them by a directed cycle of implications. One clone participates in a unit clause and the other three in the three incident edge clauses. Each clone additionally occurs in its incoming and outgoing implication. Consequently every formula variable occurs exactly three times. The reduction from cubic MIS to MAX-HORN-2SAT composes unit, edge, and consistency clauses; the quantitative contribution is the exact optimum identity and the explicit cardinality bound of the repair, together with their Lean formalization. The contribution is quantitative rather than merely satisfiability preserving. We prove an explicit repair of every inconsistent clone assignment, obtain an exact affine optimum identity, and transfer the Berman–Fujito cubic-MIS guarantee [11] to the Horn class. Cubic MIS is APX-complete [1], and the same repair proves APX-completeness of the image class. The central result is Theorem 22: the factor 83/84 on the three-occurrence class $\{T_{H(G)}\}$ lies above the Unique Games thresholds of Brakensiek–Huang–Zwick and Lewin–Livnat–Zwick, which quantify over all instances. Section 6 defines the forced-equality occurrence split with hard implication cycles, which Sections 7 and 9 use.

2. The cubic three-occurrence construction

Let $H = (V, E)$ be a finite simple cubic graph, with $n = |V|$ and $m = |E|$. Thus $m = 3n/2$. At each vertex v, label its incident edges with the three ports $1, 2, 3$. Introduce four Boolean variables

$$z_{v,0}, z_{v,1}, z_{v,2}, z_{v,3}.$$

The formula T_H contains the following clauses.

(1) For every vertex v, one unit clause $(\neg z_{v,0})$.
(2) For every edge $e = uv$, whose endpoint ports are p at u and q at v, one edge clause $(z_{u,p} \vee z_{v,q})$.
(3) For every vertex v and $i \in \mathbb{Z}/4\mathbb{Z}$, one consistency implication $(\neg z_{v,i} \vee z_{v,i+1})$.

Every clause has at most one negative literal. Every clone occurs in exactly one source clause, one outgoing implication, and one incoming implication; hence it occurs exactly three times. There are $4n$ variables and $n + m + 4n = 5n + m$ clauses. For an assignment z, let

$$\begin{aligned} U(z) &= |\{v : z_{v,0} = 0\}|, \\ F(z) &= |\{uv \in E : z_{u,p} = z_{v,q} = 0\}|, \\ T(z) &= |\{(v, i) : z_{v,i} = 1,\ z_{v,i+1} = 0\}|. \end{aligned}$$

These are respectively the satisfied unit clauses, falsified edge clauses, and falsified consistency implications. Therefore

$$\mathrm{val}_{T_H}(z) = U(z) + (m - F(z)) + (4n - T(z)). \tag{1}$$

3. Constructive repair and exact optimum

Call a vertex selected when $z_{v,0} = 0$, and clean when it is selected and its four-bit cycle has no 1-to-0 transition.

Lemma 1 (Cycle lemma)**.** *If $z_{v,0} = 0$ and the cycle at v has no 1-to-0 transition, then $z_{v,i} = 0$ for all i.*

Proof. A nonconstant cyclic binary word has both a 0-to-1 and a 1-to-0 transition. Since the word begins with zero, absence of a 1-to-0 transition forces it to be constantly zero. □

Let $C(z)$ be the clean selected vertices. Delete the left endpoint of every edge induced by $C(z)$, ignoring repetitions, and call the remaining set $R(z)$.

Lemma 2 (Quantitative repair)**.** *$R(z)$ is independent and*

$$|R(z)| \geq \mathrm{val}_{T_H}(z) - (4n + m).$$

The set is computable in polynomial time.

Proof. Every edge induced by $C(z)$ loses an endpoint, so $R(z)$ is independent. At most $T(z)$ selected vertices are nonclean, because each such vertex contributes at least one transition. By Lemma 1, both endpoint clones of every edge induced by $C(z)$ are zero; hence every such edge is counted by $F(z)$. The number of distinct removed endpoints is at most the number of edges induced by $C(z)$, which is at most $F(z)$. Thus

$$|R(z)| \geq U(z) - T(z) - F(z).$$

Substituting (1) gives the claimed inequality. □

Theorem 3 (Exact identity)**.** *For every simple cubic graph H,*

$$\mathrm{OPT}(T_H) = 4n + m + \alpha(H).$$

Moreover, an assignment of value A decodes in polynomial time to an independent set of size at least $A - (4n + m)$.

Proof. The upper bound follows from Lemma 2. For the lower bound, let I be an independent set and set all four clones of v to zero exactly when $v \in I$. Every consistency implication and every edge clause is satisfied, while precisely $|I|$ unit clauses are satisfied. The value is $4n+m+|I|$. Taking I maximum proves equality. □

4. The 83/84 approximation

For a simple cubic graph put $B = 4n+m = 11n/2$. Every independent set sends all its incident edges to its complement, so $\alpha \leq n/2$. If $|I| \geq 6\alpha/7$, the ratio $(B+|I|)/(B+\alpha)$ is at least $(11+6/7)/12 = 83/84$. We give an endpoint argument that only requires the published cubic MIS factor arbitrarily close to 7/6 [11].

Theorem 4. *On the class $\{T_H : H$ is simple and cubic$\}$ there is a polynomial-time 83/84 approximation.*

Proof. Return the better lift of two independent sets. First use Berman–Fujito with fixed parameters giving solution fraction $r = 6/7-10^{-6}$. If $2\alpha/n \leq 9999/10000$, then $(83-84r)\alpha \leq (11+84\cdot 10^{-6})(9999/20000)n \leq 11n/2 = B$, which gives the required ratio. For the second set, apply the near-satisfiable Max Cut guarantee of Trevisan [10]. For fixed $\epsilon = 10^{-4}$, a graph having a cut of fraction at least $1-\epsilon$ admits a polynomially computable cut of fraction at least 0.95 for all sufficiently large n: the stated bound is $1 - 4\sqrt{\epsilon} + 8\epsilon - o(1) > 0.95$. Fixed small orders may be solved exhaustively, including the empty graph. The independent set of size α and its complement cut 3α edges, a fraction $2\alpha/n$. Thus the cut guarantee applies in the remaining case. From each side delete one endpoint of each internal edge, taking the union of the deletions on that side. The larger remaining independent set has size at least $(n-0.05m)/2 = 37n/80 > 3n/7 \geq 6\alpha/7$. The endpoint calculation above now applies. The algorithm runs both branches without knowing α. □

4.1. **Generalization to degree $d \geq 3$.** The same construction and ratio transfer extend to every fixed degree $d \geq 3$. Let $H = (V, E)$ be simple and d-regular, so $m = dn/2$. At each vertex introduce $d+1$ Boolean clones together with one unit clause, d edge clauses (via a port labelling of the incident edges), and a directed cycle of $d+1$ consistency implications. Every clone still occurs in exactly three places. The score identity becomes

$$\mathrm{OPT}(T_H) = (d+1)n + m + \alpha(H).$$

The quantitative repair of Lemma 2 is independent of the cycle length, and the same deletion argument yields an independent set of size at least $A - ((d+1)n + m)$. Write

$$B = (d+1)n + m = n\Big(\tfrac{3d}{2} + 1\Big), \qquad \beta = \tfrac{3d}{2} + 1.$$

Any independent set cuts exactly $d\alpha$ edges, hence $\alpha \leq n/2$. Consequently any independent set of size at least $6\alpha/7$ already produces the ratio

$$\frac{B + |I|}{B + \alpha} \geq \frac{\beta + 3/7}{\beta + 1/2}.$$

The right-hand side is strictly increasing in β for $\beta > 0$ and equals 83/84 when $d = 3$. Therefore the ratio is at least 83/84 for every $d \geq 3$.

Proposition 5. *Let H be simple and d-regular, $d \geq 3$, and $B = (d+1)n + m$.*

(1) *The assignment with every clone true has value B. It is a $\frac{3d+2}{3d+3}$-approximation on $\{T_H\}$.*
(2) *If $d \geq 27$, then $\frac{3d+2}{3d+3} \geq \frac{83}{84}$.*
(3) *If $d \geq 22$ and an independent set I satisfies $5\alpha \leq (d+3)|I|$, the Berman–Fujito guarantee* [11]*, its lift has value at least $\frac{83}{84}\,\mathrm{OPT}(T_H)$.*
(4) *If an independent set I satisfies $7|I| \geq 6\alpha$, its lift has value at least $\frac{21d+20}{21d+21}\,\mathrm{OPT}(T_H)$.*
(5) *If $d = 21$, there is a polynomial-time $83/84$ approximation on $\{T_H : H \text{ is simple and } 21\text{-regular}\}$.*
(6) *If an independent set I satisfies $28|I| + d\alpha \geq 27\alpha$, that is $|I| \geq \rho(d)\alpha$ with $\rho(d) = \frac{27-d}{28}$, its lift has value at least $\frac{83}{84}\,\mathrm{OPT}(T_H)$.*
(7) *For every fixed $\varepsilon > 0$ there is a polynomial-time algorithm on $\{T_H : H \text{ is simple and } d\text{-regular}\}$ whose value is at least $\left(1 - \frac{d-2+\varepsilon}{(3d+3)(d+3+\varepsilon)}\right)\mathrm{OPT}(T_H)$. This exceeds $\frac{3d+2}{3d+3}$ for every $d \geq 3$, and the limit $\varepsilon \to 0$ is $1 - \frac{d-2}{3(d+1)(d+3)}$.*

Proof. With every clone true, the n unit clauses fail and the m edge clauses and all $(d+1)n$ consistency implications hold, so the value is B. Since $\alpha \leq n/2$ and $2B = (3d+2)n$, we get $B \geq (3d+2)\alpha$, which gives the first claim. For the second, $3d + 2 \geq 83$. For the Berman–Fujito claim, multiply the required inequality $B + 84|I| \geq 83\alpha$ by $d+3$ and use $(d+3)|I| \geq 5\alpha$, $(d+3)B \geq (d+3)(3d+2)\alpha$ and $3d^2 - 72d + 177 \geq 0$ for $d \geq 22$. The fourth claim is the endpoint computation above: using $7|I| \geq 6\alpha$, $(21d+21)(B+|I|) - (21d+20)(B+\alpha) \geq B - (3d+2)\alpha \geq 0$. For $d = 21$ put $a = \alpha/n$, so $B = 32.5\,n$ and $m = 10.5\,n$. The algorithm returns the better lift of two independent sets, without knowing a. The first is a Berman–Fujito set with $\varepsilon = 1/50$, which has size at least $50\alpha/241$ since the guarantee for degree 21 is $24/5 + \varepsilon = 241/50$. The second is obtained from a cut: solve the Goemans–Williamson semidefinite relaxation of Max Cut on H to additive error 10^{-3}, round it, and derandomize the rounding by the method of Mahajan and Ramesh [17, 19]; then delete one endpoint of every edge inside each side, and keep the larger side. If $a \leq 99/200$, the first set suffices, because

$$84 \cdot \frac{50}{241}\alpha \geq 83\alpha - B \quad\Longleftrightarrow\quad \frac{15803}{241}\,a \leq 32.5,$$

and $\frac{15803}{241} \cdot \frac{99}{200} = \frac{1564497}{48200} < 32.5$. If $a > 99/200$, the independent set of size α cuts 21α of the m edges, a fraction $2a > 0.99$, so the relaxation has value at least $c = 0.989$. By the analysis of Goemans and Williamson, rounding yields expected cut fraction at least $\arccos(1-2c)/\pi$ whenever $c \geq 0.845$; this holds because $\arccos(1-2x)/\pi$ is convex on $[1/2, 1]$ and its convex minorant coincides with it for $x \geq 0.845$. Here $\arccos(-0.978)/\pi > 0.93$. Deleting one endpoint of each uncut edge leaves an independent set of size at least $(n - 0.07m)/2 = (1 - 0.735)n/2 > 0.12\,n$ on the larger side. Then $B + 84|I| \geq 32.5\,n + 10.08\,n > 41.5\,n \geq 83\alpha$, using $\alpha \leq n/2$. In both cases $B + 84|I| \geq 83\alpha$, which is the required ratio. For the sixth claim, $B \geq (3d+2)\alpha$ by the argument of the first claim, and the hypothesis gives $84|I| \geq 3(27-d)\alpha$. Adding the two gives $B + 84|I| \geq 83\alpha$. For the seventh claim, the Berman–Fujito set I with $5\alpha \leq (d+3+\varepsilon)|I|$ satisfies $\alpha - |I| \leq \frac{d-2+\varepsilon}{d+3+\varepsilon}\alpha$, and $B + \alpha \geq (3d+3)\alpha$, so $\frac{B+|I|}{B+\alpha} = 1 - \frac{\alpha - |I|}{B+\alpha} \geq 1 - \frac{d-2+\varepsilon}{(3d+3)(d+3+\varepsilon)}$; the loss is below $\frac{1}{3d+3}$ because $d - 2 + \varepsilon < d + 3 + \varepsilon$. □

The threshold $\rho(d) = (27-d)/28$ equals $6/7$ for $d = 3$. Selected values:

d	3	4	5	6	8	10	12	16
$\rho(d)$	6/7	23/28	11/14	3/4	19/28	17/28	15/28	11/28

The value $83/84$ for $d = 3$ uses the Berman–Fujito branch and the near-bipartite branch of Theorem 4. The $6/7$ hypothesis in the fourth claim is the hypothesis whose cubic instance is supplied by that theorem. For $\varepsilon \to 0$ the ratio of the seventh claim is $1 - \frac{d-2}{3(d+1)(d+3)}$. Selected values:

d	4	5	6	20
$1 - \frac{d-2}{3(d+1)(d+3)}$	$\frac{103}{105}$	$\frac{47}{48}$	$\frac{185}{189}$	$\frac{159}{161}$
$\frac{3d+2}{3d+3}$	$\frac{14}{15}$	$\frac{17}{18}$	$\frac{20}{21}$	$\frac{62}{63}$

Explicitly, the endpoint ratio obtained from a 6/7-approximate independent set is

$$\frac{21d+20}{21d+21}.$$

Selected values:

d	3	4	5	6	8	10	12
ratio	83/84	104/105	125/126	146/147	188/189	230/231	272/273
	≈ 0.9881	≈ 0.9905	≈ 0.9921	≈ 0.9932	≈ 0.9947	≈ 0.9957	≈ 0.9963

As $d \to \infty$ both $\frac{3d+2}{3d+3}$ and $\frac{21d+20}{21d+21}$ tend to 1.

5. APX-completeness

Theorem 6. *The cubic three-occurrence Horn class is APX-complete under L-reductions. The reduction from cubic MIS has constants* 23 *and* 1*.*

Proof. A greedy maximal independent set in a graph of maximum degree three has size at least $n/4$, so $n \le 4\alpha(H)$. Since $m = 3n/2$,

$$\mathrm{OPT}(T_H) = \frac{11n}{2} + \alpha(H) \le 22\alpha(H) + \alpha(H) = 23\alpha(H).$$

For an assignment of value A, Lemma 2 returns R with $|R| \ge A - (4n+m)$. Hence

$$\alpha(H) - |R| \le (4n + m + \alpha(H)) - A = \mathrm{OPT}(T_H) - A.$$

These are the two L-reduction inequalities. Cubic MIS is APX-complete and Theorem 4 places the image class in APX. □

5.1. **Hardness for larger degree.** The same argument yields APX-completeness for every fixed $d \ge 3$. Maximum Independent Set remains APX-complete on simple d-regular graphs (and on graphs of maximum degree d) for each fixed $d \ge 3$ [1, 11]. The construction of T_H from a d-regular graph is an L-reduction whose constants depend only on d. A greedy maximal independent set in a degree-d graph has size at least $n/(d+1)$, so $n \le (d+1)\alpha$. With $m = dn/2$ one obtains

$$\mathrm{OPT}(T_H) = (d+1)n + m + \alpha \le (d+1)^2\Big(\tfrac{3d}{2} + 1\Big)\alpha + \alpha.$$

Thus the first L-reduction constant is $O(d^2)$. The second constant remains exactly 1 by the quantitative repair: every assignment of value A produces an independent set R satisfying $\alpha - |R| \le \mathrm{OPT}(T_H) - A$. Consequently the image class

$$\{T_H : H \text{ is simple and } d\text{-regular}\}$$

is APX-complete for every fixed $d \ge 3$. By Proposition 5(1) the image class is approximable to the factor $1 - \frac{1}{3d+3}$ for every d, so it lies in APX and is APX-complete. Under the Unique Games Conjecture the unrestricted weighted MAX-HORN-2SAT problem is hard to approximate above the constant $\alpha^* \approx 0.94616$ of Brakensiek–Huang–Zwick. The factor $\frac{3d+2}{3d+3}$ exceeds α^* for every $d \ge 6$, and Theorem 4 gives $83/84 > \alpha^*$ for $d = 3$. Both statements concern bounded-occurrence classes and are consistent with the bounded-occurrence improvements of Ghoshal et al. [16] and with the Unique Games Conjecture.

6. The forced-equality occurrence split

Let F be a weighted MAX-HORN-2SAT instance. Its source clauses have nonnegative rational weights. Give every literal occurrence a distinct clone, and replace that occurrence in its source clause by the clone without changing the clause weight. If the occurrences of an original variable x receive clones $x_0, \dots, x_{d-1}$, add the directed cycle

$$(\neg x_i \vee x_{i+1}) \qquad (i \bmod d).$$

These implications are *hard*: every feasible assignment must satisfy them. They contribute zero to the maximized objective. This is a partial weighted MAX-HORN-2SAT instance, denoted by $P(F)$. Note that the construction remains faithful even when a source clause contains several occurrences of the same underlying variable, or both polarities of a variable: by Lemma 7, feasibility forces all clones of a variable to take the same value, so the split clause is satisfied exactly when the source clause is. Every clone appears

once in a weighted source clause, once negatively in its outgoing hard implication, and once positively in its incoming hard implication. Thus every clone, viewed as a variable of the split formula $P(F)$, has exactly three occurrences.

Lemma 7 (Hard-cycle consistency)**.** *A clone assignment satisfies every hard implication on one directed cycle if and only if all clones on that cycle have the same value.*

Proof. A constant cycle satisfies every implication. Conversely, a nonconstant cyclic binary word has a 1-to-0 transition, and the implication at that transition is false. □

Theorem 8 (Exact hard-split identity)**.** *The transformation $F \mapsto P(F)$ is polynomial, and*

$$\mathrm{OPT}_{\mathrm{partial}}(P(F)) = \mathrm{OPT}_{\mathrm{weighted}}(F).$$

More strongly, feasible split assignments and source assignments have exactly the same attainable objective values under encoding and decoding.

Proof. Lift a source assignment by giving every occurrence clone the value of its original variable. The result satisfies all hard cycles and preserves every source clause's truth value and weight. Conversely, by Lemma 7, every feasible split assignment is constant on each original-variable cycle. Decode that common value. Again every weighted source clause has the same truth value. Encoding and decoding therefore preserve the objective pointwise, proving both claims. □

Corollary 9 (Ratio preservation)**.** *For every $0 \le \rho \le 1$, a polynomial-time ρ-approximation for the hard three-occurrence image $\{P(F)\}$ yields a polynomial-time ρ-approximation for weighted MAX-HORN-2SAT.*

Proof. Apply the image-class algorithm and decode its feasible assignment. Both its value and the optimum are unchanged by Theorem 8. □

7. A hard-split family attaining the target ratio

Let H be simple cubic. In the source $F_H(w)$ put a unit clause $(\neg x_v)$ of weight one at each vertex and a clause $(x_u \vee x_v)$ of rational weight $w_e \ge 1$ on each edge. Write $W = \sum_e w_e$. Apply the ordinary occurrence split to this source, with zero-score hard cycles. The four ports identify the three edge occurrences and the unit. The source graph can also be recovered by contracting each hard cycle and reading the binary source clauses.

Theorem 10. *The feasible optimum of $P(F_H(w))$ is $W+\alpha(H)$. For all $w_e \ge 1$ there is a polynomial-time $27/28$ approximation. If $W \ge 6n$, this class admits a polynomial-time $257/260 > 83/84$ approximation. In particular this holds if every edge clause weighs four.*

Proof. For the set S of false source variables the score is $W+|S|-\sum_{e\in E(H[S])} w_e$. Delete from S one endpoint of each internal edge, counting repeated deletions only once. The independent set left has size at least $|S| - |E(H[S])|$, which bounds the excess over W from above since $w_e \ge 1$. Conversely every independent set realizes $W + |I|$ by setting its variables false and copying this choice to all ports. This proves the identity and the constructive repair. For $n > 0$ we have $W \ge m = 3n/2 \ge 3\alpha$. Even a single cubic MIS run with $20|I| \ge 17\alpha$ gives $(W + |I|)/(W + \alpha) \ge (3 + 17/20)/4 = 77/80 > 0.94615981$. To attain $27/28$, use both candidates from Theorem 4. In its near-bipartite case $|I| \ge 6\alpha/7$, giving $(3 + 6/7)/4 = 27/28$. In its other case $|I| \ge r\alpha$ and $(27 - 28r)\alpha \le (3 + 28 \cdot 10^{-6})(9999/20000)n \le 3n/2 \le W$, which is exactly the sufficient inequality. Since $20/17 > 7/6$, a fixed-parameter Berman–Fujito run gives $20|I| \ge 17\alpha$. Here $W \ge 6n \ge 12\alpha$, so $260(W + |I|) - 257(W + \alpha) \ge 3W - 36\alpha \ge 0$. The empty instance satisfies the division-free bound as well. □

In Lean, the hard optima $4m+\alpha$ and $m+\alpha$ and the $257/260$ and $77/80$ transfers from a $17/20$ independent set are proved. The $27/28$ statement is the arithmetic implication from the numerical hypothesis $6\alpha \le 7s$, or from $20000\alpha \le 9999n$ together with $5999993\alpha \le 7000000s$, to $28(m+s) \ge 27(m+\alpha)$ for an independent set of size s. The algorithm that supplies this hypothesis is the Berman–Fujito and Trevisan branch proved in the text above and is not formalized in Lean. A second Lean statement makes the two branch guarantees explicit: if $20000\alpha \le 9999n$ implies $5999993\alpha \le 7000000s$ (Berman–Fujito branch) and $20000\alpha > 9999n$ implies $37n \le 80s$ (near-bipartite branch), then $28(m + s) \ge 27(m + \alpha)$.

Why hardening can change the ratio. In the unit-weight version, take four disjoint copies of the cube Q_3. Each has independence number four by its parity classes and a perfect matching. Take a parity class in each of three copies and an antipodal pair in the fourth. The pair is maximal because every other vertex neighbors an endpoint. Thus $|I| = 14$, $\alpha = 16$, $n = 32$, and $m = 48$. The same feasible lift meets $7|I| \geq 6\alpha$ and scores 190/192 relative to optimum on T_H, but only 62/64 on $P(F_H(1))$. The 128 satisfied implications have ceased to earn points. This refutes a transfer based solely on the MIS guarantee and this decoding, even for maximal sets with perfectly consistent ports.

8. A signed graph for general weighted inputs

The unsigned construction can fail even on K_4. Start with two indexed copies of $(x \vee x)$, assign their occurrences the opposite pairs of a four-cycle, and join each source pair. The resulting graph is simple cubic. Each singleton is a maximum independent set and gives a nonconstant cyclic word. For an example with distinct variables within each clause, three indexed copies of $(x \vee y)$ give two triangles joined by a matching. An independent set taking a vertex in each triangle at different indices is maximum and violates both cycles under membership decoding. In the unsigned graph a pair edge forbids only simultaneous membership. Complementation reverses a directed implication but does not turn it into a prohibition on simultaneous membership. The paper uses the at-most-one-negative convention.
Unary terminals. For the path a–w–b, the terminal states 00, 01, 10, and 11 give maximum total cardinalities $1, 1, 1, 2$. It realizes a one-unit bonus for 11 when another constraint synchronizes the ends. Both mixed states extend to independent sets. The following signed construction instead gives a unit clause one reward vertex attached to its falsifying port.

Theorem 11. *For any weighted source F of width at most two, let D be its number of literal occurrences, W its total nonnegative clause weight, and $M = W + 1$. One can construct a simple vertex-weighted graph J of maximum degree three such that $\alpha_w(J) = DM + \mathrm{OPT}(F)$. Every independent set of weight A decodes in polynomial time to a feasible split assignment of score at least $A - DM$.*

Proof. At occurrence i create adjacent ports T_i, F_i of weight M. Join T_i to F_{i+1} around each variable cycle. For a singleton cycle the pair edge already supplies this conflict; insert it just once. Discard variables with no occurrences. For each literal occurrence create a reward vertex of weight equal to its clause weight and join it to the port that falsifies that literal. Join the two reward vertices of each binary clause. Every port has at most two ring neighbors and one reward neighbor. Distinct literal occurrences have distinct vertices even for repeated literals. The graph has $3D$ vertices, linear combinatorial size and polynomial bit complexity with binary rational weights. At most one port per pair and one reward per clause can be selected. If fewer than D ports are selected, the weight is at most $(D-1)M + W < DM$. Any constant-per-variable encoding already reaches DM. Otherwise each pair has exactly one selected port. Choosing T_i forces T_{i+1}, so all pairs of a variable choose the same side. Each selected reward then witnesses a true literal. Reading the ports gives score at least $A - DM$. Conversely encode any assignment and select one true witness per satisfied clause. The weight is exactly DM plus its score. These operations prove both claims, including for empty clauses. □

This construction preserves multiplicative guarantees on the excess $A - DM$. A guarantee on the total weight alone can be almost perfect when the decoded score is zero: for k positive unit clauses on one variable, the all-false port selection has fraction $(k+1)/(k+2)$ and decodes to zero. Already $k = 82$ gives 83/84. Thus neither the signed repair nor a polynomial algorithm for unweighted cubic MIS supplies the unrestricted weighted ratio.

9. Unique Games thresholds, MAX 2-SAT and MAX-CUT

Let α^* denote the constant of Brakensiek, Huang and Zwick for weighted MAX $\{1, 2\}$-HORN-SAT, approximately 0.94615981. Their Theorem 5.12 gives UG-hardness at $\alpha^* + \varepsilon$ for every $\varepsilon > 0$ [12]. In particular, $19/20 = 0.95 > \alpha^*$. A deterministic polynomial-time algorithm returning value at least $19\,\mathrm{OPT}(F)/20$ on *every* weighted source F would solve the UG-hard gap problem in polynomial time and, assuming UGC, imply P = NP. By Theorem 8 the same implication holds for a guarantee on the entire hard-split image. The worst-case bounds in Theorems 4 and 10 concern their specified cubic families. The constant

$\beta^\star \approx 0.94016567$ of Lewin, Livnat and Zwick is the UG-hard threshold for MAX 2-SAT [12]. Both UG-hardness statements quantify over all instances. They give no information about the class in which every variable has exactly three literal occurrences. For that class Theorem 4 gives the worst-case ratio 83/84, and

$$\frac{83}{84} > 0.94615981 > 0.94016567.$$

Ghoshal, Huang, Lee, Makarychev and Makarychev prove that MAX 2-SAT instances in which every variable appears in at most d constraints admit a $\beta^\star + \tilde{\Omega}(1/d^2)$-approximation [16]. Thus a guarantee above $\beta^\star$ on a bounded-occurrence class is established for every fixed d and is consistent with the Unique Games Conjecture. A guarantee above α^* or $\beta^\star$ on *all* instances would not be. For MAX 2-SAT instances in which every variable appears in at most two clauses, the problem can be solved exactly in linear time, while already three occurrences per variable give an NP-complete problem [21]. Thus the occurrence bound $d = 3$ is the first at which the problem is NP-hard, and for every fixed d the guarantee of [16] exceeds $\beta^\star$. The endpoint guarantee obtained from a 6/7-approximate independent set on a d-regular graph is

$$r(d) = \frac{21d + 20}{21d + 21},$$

which is strictly increasing for $d \geq 3$ since

$$r(d+1) - r(d) = \frac{21}{(21d + 21)(21d + 42)} > 0.$$

For $d = 3$ the independent-set hypothesis is supplied by Theorem 4, and $r(3) = 83/84$ exceeds $\beta^\star$ by ≈ 0.04793. For every $d \geq 6$ Proposition 5(1) gives the unconditional factor $\frac{3d+2}{3d+3} > 0.95 > \alpha^* > \beta^\star$ on the d-regular family. For $d \geq 3$ one has $\frac{d-2}{3(d+1)(d+3)} \leq \frac{4}{189} < 0.022$, because $63(d-2) \leq 4(d+1)(d+3)$ is equivalent to $(4d - 23)(d - 6) \geq 0$, which holds for every integer d. Hence for every $d \geq 4$ and every sufficiently small $\varepsilon > 0$ the factor of Proposition 5(7) exceeds $0.97 > \alpha^* > \beta^\star$ on the d-regular family, and together with $r(3) = 83/84$ every class $\{T_H : H \text{ is simple and } d\text{-regular}\}$ with $d \geq 3$ has a polynomial-time factor above α^*. These are statements about three-occurrence classes and do not concern all instances. The additive gain $\tilde{\Omega}(1/d^2)$ of [16] is $o(1)$ in d. For the hard-split family with $d = 3$ the proved ratios are those of Theorems 8 and 10. By Proposition 5(6), for each d with $4 \leq d \leq 20$ an independent-set algorithm of ratio $\rho(d) = (27 - d)/28$ on d-regular graphs would give 83/84 on the d-regular family, and $83/84 > \alpha^* > \beta^\star$. The thresholds of Brakensiek, Huang and Zwick quantify over all instances. OpenAI released on 6 October 2026 a manuscript claiming a proof of the Unique Games Conjecture, together with a Lean formalization; this claim has not been peer reviewed [20]. If it is confirmed, the UG-hardness statements above become NP-hardness statements.

9.1. **MAX-CUT as MAX 2-SAT.**

Proposition 12. *Let $G = (V, E)$ be a graph with n vertices and m edges. Let F_G be the formula with one variable x_v for every vertex and, for every edge uv, the two clauses $(x_u \vee x_v)$ and $(\neg x_u \vee \neg x_v)$. Then:*

1. *an assignment z satisfies exactly $m + \mathrm{cut}(S_z)$ clauses, where $S_z = \{v : z(x_v) = \textit{true}\}$ and $\mathrm{cut}(S)$ is the number of edges with exactly one endpoint in S; hence $\mathrm{OPT}(F_G) = m + \mathrm{maxcut}(G)$;*
2. *the variable x_v occurs in exactly $2\deg(v)$ clauses;*
3. *if an assignment satisfies at least $\rho\,\mathrm{OPT}(F_G)$ clauses, then the cut S_z has at least $\rho\,(m + \mathrm{maxcut}(G)) - m$ edges.*

Proof. For an edge uv whose endpoints receive different values both clauses hold. If both endpoints are true only $(x_u \vee x_v)$ holds, and if both are false only $(\neg x_u \vee \neg x_v)$ holds. Summing over the edges gives the first claim, and the second and third follow. □

The clause $(\neg x_u \vee \neg x_v)$ has two negative literals, so F_G is an instance of MAX 2-SAT that does not belong to the Horn class; no guarantee of this paper is claimed for it.

Proposition 13. *Let G be a simple d-regular graph with n vertices and $m = dn/2$ edges. Let H be the conflict graph of F_G: its $N = 4m$ vertices are the literal occurrences of F_G, and two vertices are adjacent if they lie in the same clause or are complementary literals. Then H is simple and $(d+1)$-regular, and*

$$\alpha(H) = \mathrm{OPT}(F_G) = m + \mathrm{maxcut}(G).$$

Consequently, with $D = d+1$, $M = DN/2$ and $B = (D+1)N + M = 2m(3d+5)$, the formula T_H satisfies $\mathrm{OPT}(T_H) = B + m + \mathrm{maxcut}(G)$, and an assignment of T_H of value A yields in polynomial time a cut of G with at least $A - B - m$ edges.

Proof. Every literal occurrence has exactly one neighbour in its clause. A positive occurrence of x_u is adjacent to the d negative occurrences of x_u, and symmetrically; an in-clause neighbour is a literal of another variable, so no edge is repeated. Hence H is simple and $(d+1)$-regular. An independent set contains at most one literal of each clause and no two complementary literals, so it is a consistent partial assignment satisfying as many clauses as the set has vertices; conversely, an assignment and one satisfied literal per satisfied clause give an independent set. Thus $\alpha(H) = \mathrm{OPT}(F_G)$, which equals $m + \mathrm{maxcut}(G)$ by Proposition 12. The identity for T_H is the degree-D exact identity. By the quantitative repair, an assignment of value A gives an independent set of size at least $A - B$, hence an assignment of F_G satisfying at least $A - B$ clauses, hence a cut with at least $A - B - m$ edges by Proposition 12. □

Corollary 14. *Let* $\mathrm{maxcut}(G) = \gamma m$ *and* $0 \le \rho \le 1$.

1. *If an assignment of F_G satisfies at least $\rho\,\mathrm{OPT}(F_G)$ clauses, its cut has at least $\big(\rho - \frac{1-\rho}{\gamma}\big)\,\mathrm{maxcut}(G)$ edges.*
2. *If G is d-regular and an assignment of T_H has value at least $\rho\,\mathrm{OPT}(T_H)$, the cut obtained from it has at least $\big(\rho - \frac{(1-\rho)(6d+11)}{\gamma}\big)\,\mathrm{maxcut}(G)$ edges.*

Proof. By Proposition 12, the cut has at least $\rho(m + \gamma m) - m$ edges; dividing by γm gives the first claim. By Proposition 13 it has at least $\rho(B + m + \gamma m) - B - m$ edges with $B + m = (6d+11)m$, which gives the second. □

Theorem 15. *For every $d \ge 1$ there is a polynomial-time algorithm that, given a simple d-regular graph G, returns an assignment of $T_{H(G)}$ of value at least*

$$\Big(1 - \frac{1-\alpha_\varepsilon}{6d+12}\Big)\,\mathrm{OPT}(T_{H(G)}) \;\ge\; \frac{83}{84}\,\mathrm{OPT}(T_{H(G)}),$$

where $H(G)$ is the conflict graph of Proposition 13 and $\alpha_\varepsilon \approx 0.87856$ Here $\alpha_\varepsilon = \alpha_{\mathrm{GW}} - 10^{-5} \ge 0.8785$, where $\alpha_{\mathrm{GW}} \approx 0.87856$ is the Goemans–Williamson constant.

Proof. Let $c = \mathrm{maxcut}(G)$. Solving the Goemans–Williamson relaxation to additive error 10^{-5} and rounding with derandomization [17, 19] returns in polynomial time a set S with $\mathrm{cut}(S) \ge \alpha_\varepsilon\, c$ (the relaxation error 10^{-5} is absorbed in α_ε). Set x_v true exactly for $v \in S$ and choose one true literal in every satisfied clause of F_G. These $m + \mathrm{cut}(S)$ literal occurrences are pairwise non-adjacent in $H(G)$, because they lie in distinct clauses and no two are complementary. So they form an independent set I with $|I| = m + \mathrm{cut}(S)$, while $\alpha(H(G)) = m + c$ by Proposition 13. Since $c \le m$,

$$|I| \ge m + \alpha_\varepsilon\, c \ge \frac{1+\alpha_\varepsilon}{2}\,(m + c) = \frac{1+\alpha_\varepsilon}{2}\,\alpha(H(G)),$$

because $(1-\alpha_\varepsilon)(m-c) \ge 0$. Let $D = d+1$, so $H(G)$ is D-regular and $B = (D+1)N + M \ge (3D+2)\alpha(H(G))$ by $\alpha(H(G)) \le N/2$. The lift of I has value $B + |I|$, hence

$$\frac{B + |I|}{B + \alpha(H(G))} \ge 1 - \frac{\frac{1-\alpha_\varepsilon}{2}\,\alpha(H(G))}{B + \alpha(H(G))} \ge 1 - \frac{1-\alpha_\varepsilon}{6d+12}.$$

Finally $\frac{1-\alpha_\varepsilon}{6d+12} \le \frac{1-\alpha_\varepsilon}{18} < \frac{1}{84}$ for $d \ge 1$, since $1 - \alpha_\varepsilon \le 0.1215 < 18/84 \approx 0.214$. □

Proposition 16. *Let G be d-regular with $\mathrm{maxcut}(G) = \gamma m$, $\frac{1}{2} \le \gamma \le 1$. Let h be the largest convex function on $[0,1]$ with $h(x) \le \frac{1}{\pi}\arccos(1-2x)$; thus $h(x) = \alpha_{\mathrm{GW}} x$ for $x \le x^* \approx 0.8446$ and $h(x) = \frac{1}{\pi}\arccos(1-2x)$ for $x \ge x^*$. The Goemans–Williamson algorithm returns in polynomial time a cut S with*

$\mathrm{cut}(S) \geq (h(\gamma) - 10^{-5})m$, *and the lift of the corresponding independent set of* $H(G)$ *has on* $T_{H(G)}$ *the ratio at least*

$$1 - \frac{\gamma - h(\gamma) + 10^{-5}}{6d + 11 + \gamma} \;\geq\; 1 - \frac{0.1053}{6d + 11.5}.$$

The cut fraction that Corollary 14(2) recovers from this ratio equals $(h(\gamma) - 10^{-5})/\gamma$.

Proof. Let $x_e \in [0,1]$ be the semidefinite value of the edge e and $s = \sum_e x_e \geq \gamma m$ (up to the additive error $10^{-5}m$). The expected cut of hyperplane rounding is $\sum_e \frac{1}{\pi}\arccos(1-2x_e) \geq \sum_e h(x_e) \geq m\, h(s/m) \geq m\, h(\gamma)$ by convexity and monotonicity of h, and the derandomization of [19] preserves it. By Proposition 13 $\alpha(H(G)) = m + \gamma m$, and $B = (6d+10)m$, so the value of the lift is $B + m + \mathrm{cut}(S)$ against the optimum $B + m + \gamma m = (6d + 11 + \gamma)m$, which gives the first bound. On $[\frac{1}{2}, 1]$ the function $\gamma - h(\gamma)$ is maximal where $\gamma(1-\gamma) = \pi^{-2}$, that is at $\gamma^* = \frac{1}{2}(1 + \sqrt{1 - 4/\pi^2}) \approx 0.8856$, with $\gamma^* - h(\gamma^*) \approx 0.10526$. Since $6d + 11 + \gamma \geq 6d + 11.5$, the second bound follows. For the last claim, with $1 - \rho = \frac{\gamma - \hat{h}}{6d+11+\gamma}$ and $\hat{h} = h(\gamma) - 10^{-5}$,

$$\rho - \frac{(1-\rho)(6d+11)}{\gamma} = 1 - \frac{(1-\rho)(6d+11+\gamma)}{\gamma} = 1 - \frac{\gamma - \hat{h}}{\gamma} = \frac{\hat{h}}{\gamma}. \qquad \square$$

Remark 17. *With the linear bound* $\mathrm{cut}(S) \geq \alpha_\varepsilon \gamma m$ *the lift has on* $T_{H(G)}$ *the ratio* $\rho_d(\gamma) = \frac{6d+11+\alpha_\varepsilon\gamma}{6d+11+\gamma} = 1 - \frac{(1-\alpha_\varepsilon)\gamma}{6d+11+\gamma}$, *and Corollary 14(2) applied to* $\rho_d(\gamma)$ *gives the cut fraction* $1 - \frac{(1-\alpha_\varepsilon)\gamma}{6d+11+\gamma} - \frac{(1-\alpha_\varepsilon)(6d+11)}{6d+11+\gamma} = \alpha_\varepsilon$ *for every* $\gamma \in (0,1]$. *The algorithm does not use* γ. *Since* $\sup_{0<\gamma\leq 1} \frac{(1-\alpha_\varepsilon)\gamma}{6d+11+\gamma} = \frac{1-\alpha_\varepsilon}{6d+12}$, *the ratio* $1 - \frac{1-\alpha_\varepsilon}{6d+12}$ *holds for every* G, *while the instance ratio* $\rho_d(\gamma)$ *recovers the constant* α_ε *exactly. Moreover* $\rho_d(\gamma) - \alpha_\varepsilon = \frac{(1-\alpha_\varepsilon)(6d+11)}{6d+11+\gamma} > 0$, *so* $\rho_d(\gamma) > \alpha_\varepsilon$ *for every* $\gamma \in (0,1]$. *Proposition 16 replaces* $\alpha_\varepsilon\gamma$ *by* $h(\gamma) - 10^{-5} \geq \alpha_\varepsilon\gamma$.

Proposition 18. *Let* $c_d \in (0,1]$ *and suppose a polynomial-time algorithm returns on every graph of maximum degree* d *a cut with at least* $c_d\,\mathrm{maxcut}(G)$ *edges; if the algorithm is randomized and the guarantee is for the expected cut, every statement below is for the expected value. Then on* $\{T_{H(G)} : G \text{ is } d\text{-regular}\}$ *a polynomial-time algorithm has the ratio at least* $1 - \frac{1-c_d}{6d+12}$.

Proof. As in the proof of Theorem 15 with 0.8785 replaced by c_d: $|I| \geq \frac{1+c_d}{2}\alpha(H(G))$ and $B + \alpha(H(G)) \geq (3D+3)\alpha(H(G))$ with $D = d+1$, so the loss is at most $\frac{1-c_d}{2(3d+6)}$. $\square$

For graphs of maximum degree d, Feige, Karpinski and Langberg give a guarantee $\alpha_{\mathrm{GW}} + \epsilon_d$ with $\epsilon_d > 0$, and 0.921 for $d = 3$ [13]; Halperin, Livnat and Zwick give an algorithm whose expected approximation ratio on graphs of maximum degree 3 is at least 0.9326 [14]; Hsieh and Kothari give $\alpha_{\mathrm{GW}} + \tilde{\Omega}(1/d^2)$ [15]. With $c_3 = 0.9326$ as an expected guarantee, Proposition 18 gives the expected ratio $1 - \frac{0.0674}{30} > 0.9977$ on $\{T_{H(G)} : G \text{ cubic}\}$. These bounded-degree guarantees exceed α_{GW} on a class of graphs and are consistent with the threshold of [18], which quantifies over all graphs.

Remark 19. *Theorem 15 applies to every degree* $D = d+1$ *with* $4 \leq D \leq 20$, *that is* $3 \leq d \leq 19$, *on the class* $\{T_{H(G)} : G \text{ is } (D-1)\text{-regular}\}$. *In that range* $|I| \geq \frac{1+\alpha_\varepsilon}{2}\,\alpha(H(G))$ *exceeds* $\rho(D) = \frac{27-D}{28} \leq \frac{23}{28}$, *so Proposition 5(6) gives the ratio* 83/84 *without any assumption on an independent-set algorithm. Since* $\frac{1+\alpha_\varepsilon}{2} > 6/7$, *Proposition 5(4) gives the ratio* $\frac{21D+20}{21D+21}$, *which equals* $\frac{104}{105}$ *for* $D = 4$ *and* $\frac{440}{441}$ *for* $D = 20$. *For these* G *the constant of Corollary 14 is* $B + m = (6D+5)m$; *for* $D = 4$ *it equals* $29m$, *and an assignment of* $T_{H(G)}$ *of ratio* $\frac{104}{105}$ *yields a cut with at least* $\frac{104}{105}\,\mathrm{maxcut}(G) - \frac{29}{105}m$ *edges, which is* $\frac{49}{105}m$ *when* $\mathrm{maxcut}(G) = \frac{3}{4}m$. *The statement concerns the conflict-graph class; for arbitrary* D*-regular graphs* H *with* $4 \leq D \leq 20$ *the ratio* 83/84 *is given by Proposition 5(6) under the hypothesis stated there.*

9.2. **Generalization to arbitrary simple graphs.** The three-occurrence construction extends without change of method to any finite simple graph $H = (V, E)$ with $n = |V|$ and $m = |E|$, even when H is not regular. At each vertex $v \in V$ of degree $d_v = \deg_H(v)$, label its incident edges with the ports $1, \ldots, d_v$ and introduce $d_v + 1$ Boolean variables

$$z_{v,0}, z_{v,1}, \ldots, z_{v,d_v}.$$

The formula T_H contains:

(1) for every vertex $v \in V$, one unit clause $(\neg z_{v,0})$;
(2) for every edge $e = uv \in E$ with endpoint ports p at u and q at v, one edge clause $(z_{u,p} \vee z_{v,q})$;
(3) for every vertex $v \in V$ and $i \in \mathbb{Z}/(d_v+1)\mathbb{Z}$, one consistency implication $(\neg z_{v,i} \vee z_{v,i+1})$.

Every clause has at most one negative literal. Each clone $z_{v,i}$ occurs once in a source clause (the unit clause at v for $i = 0$, or the edge clause at port i for $1 \le i \le d_v$), once negatively in $(\neg z_{v,i} \vee z_{v,i+1})$, and once positively in $(\neg z_{v,i-1} \vee z_{v,i})$. Thus every variable of T_H still has exactly three literal occurrences. Using $\sum_{v\in V} d_v = 2m$, the total number of variables is

$$\sum_{v\in V}(d_v+1) = n + 2m,$$

and the total number of clauses is $n + m + (n + 2m) = 2n + 3m$. Put

$$B(H) = \sum_{v\in V}(d_v+1) + m = n + 3m.$$

Lemma 20 (General quantitative repair and exact identity)**.** *For every finite simple graph $H = (V, E)$ with $n = |V|$ and $m = |E|$,*

$$\mathrm{OPT}(T_H) = n + 3m + \alpha(H).$$

Moreover, every assignment z of value A decodes in polynomial time to an independent set $R(z) \subseteq V$ of size at least $A - (n + 3m)$.

Proof. For an assignment z, let $U(z) = |\{v \in V : z_{v,0} = 0\}|$, let $F(z)$ be the number of falsified edge clauses, and let $T(z) = \sum_{v\in V} |\{i \in \mathbb{Z}/(d_v+1)\mathbb{Z} : z_{v,i} = 1,\ z_{v,i+1} = 0\}|$. Since there are $n + 2m$ consistency implications in total,

$$\mathrm{val}_{T_H}(z) = U(z) + (m - F(z)) + (n + 2m - T(z)) = n + 3m + U(z) - F(z) - T(z).$$

Call a vertex $v \in V$ selected when $z_{v,0} = 0$, and clean when it is selected and its $(d_v + 1)$-bit cycle has no 1-to-0 transition. By Lemma 1 (which holds for any cycle length $d_v + 1 \ge 1$), every clone $z_{v,i}$ of a clean selected vertex v is zero. Let $C(z)$ be the set of clean selected vertices, and form $R(z)$ by deleting one endpoint of each edge induced by $C(z)$, ignoring repeated deletions. Then $R(z)$ is independent in H. Each nonclean selected vertex contributes at least one 1-to-0 transition, so at most $T(z)$ selected vertices are nonclean. Every edge induced by $C(z)$ has both endpoint clones equal to zero and is therefore counted by $F(z)$, so at most $F(z)$ vertices are removed from $C(z)$. Consequently,

$$|R(z)| \ge |C(z)| - F(z) \ge U(z) - T(z) - F(z) = \mathrm{val}_{T_H}(z) - (n + 3m).$$

This gives the constructive repair and the upper bound $\mathrm{OPT}(T_H) \le n + 3m + \alpha(H)$. Conversely, for any independent set $I \subseteq V$, setting $z_{v,i} = 0$ for all $v \in I$ and $0 \le i \le d_v$, and $z_{v,i} = 1$ otherwise, satisfies all $n + 2m$ consistency implications, all m edge clauses, and $|I|$ unit clauses, realizing value $n + 3m + |I|$. Taking I maximum yields equality. □

Theorem 21 (Non-regular conflict-graph family)**.** *Let $G = (V_G, E_G)$ be any finite simple graph with $m = |E_G| \ge 1$ edges and no isolated vertices, and let $H(G)$ be the conflict graph of F_G from Proposition 13. Then every variable of $T_{H(G)}$ has exactly three literal occurrences, and there is a polynomial-time algorithm returning an assignment of $T_{H(G)}$ of value at least*

$$\Big(1 - \frac{1 - \alpha_\varepsilon}{18}\Big)\,\mathrm{OPT}(T_{H(G)}) > 0.9932\ \mathrm{OPT}(T_{H(G)}).$$

Proof. Each vertex $v \in V_G$ of degree $d_v = \deg_G(v) \ge 1$ contributes d_v positive literal occurrences and d_v negative literal occurrences to F_G. Thus $H(G)$ has $N = 2\sum_{v\in V_G} d_v = 4m$ vertices. Every literal occurrence in F_G has 1 in-clause neighbor and d_v complementary neighbors in $H(G)$ (with no overlap, as in Proposition 13), so those $2d_v$ vertices each have degree $d_v + 1$ in $H(G)$. Hence the number of edges of $H(G)$ is

$$M = |E(H(G))| = \frac{1}{2}\sum_{v\in V_G} 2d_v(d_v+1) = \sum_{v\in V_G}(d_v^2 + d_v) = 2m + \sum_{v\in V_G} d_v^2.$$

Since $d_v \geq 1$ for every $v \in V_G$, we have $d_v^2 \geq d_v$, so $\sum_{v \in V_G} d_v^2 \geq \sum_{v \in V_G} d_v = 2m$ and therefore $M \geq 4m$. By Lemma 20,

$$\mathrm{OPT}(T_{H(G)}) = B(H(G)) + \alpha(H(G)), \qquad B(H(G)) = N + 3M \geq 4m + 12m = 16m.$$

On the other hand, as in Proposition 13, an independent set in $H(G)$ selects at most one literal per clause of F_G and no complementary pair, so $\alpha(H(G)) = \mathrm{OPT}(F_G) = m + \mathrm{maxcut}(G) \leq 2m$. Consequently, $B(H(G)) \geq 16m \geq 8\,\alpha(H(G))$. Now run the derandomized Goemans–Williamson algorithm [17, 19] on G to obtain in polynomial time a cut S with $\mathrm{cut}(S) \geq \alpha_\varepsilon \,\mathrm{maxcut}(G)$. Selecting one true literal per satisfied clause of F_G under S yields an independent set I in $H(G)$ of size $|I| = m + \mathrm{cut}(S)$. Because $\mathrm{maxcut}(G) \leq m$, we have $(1 - \alpha_\varepsilon)(m - \mathrm{maxcut}(G)) \geq 0$, and hence

$$|I| \geq m + \alpha_\varepsilon \,\mathrm{maxcut}(G) \geq \frac{1 + \alpha_\varepsilon}{2}\big(m + \mathrm{maxcut}(G)\big) = \frac{1 + \alpha_\varepsilon}{2}\,\alpha(H(G)).$$

Lifting I to $T_{H(G)}$ gives an assignment of value $B(H(G)) + |I|$, whose approximation ratio is

$$\begin{aligned}
\frac{B(H(G)) + |I|}{B(H(G)) + \alpha(H(G))} &= 1 - \frac{\alpha(H(G)) - |I|}{B(H(G)) + \alpha(H(G))} \\
&\geq 1 - \frac{\frac{1-\alpha_\varepsilon}{2}\,\alpha(H(G))}{9\,\alpha(H(G))} \\
&= 1 - \frac{1 - \alpha_\varepsilon}{18} \\
&> 0.9932.
\end{aligned}$$

□

9.3. **Consequences of the MAX-CUT reduction for the Unique Games Conjecture.** Khot, Kindler, Mossel and O'Donnell show that, assuming the Unique Games Conjecture, MAX-CUT cannot be approximated in polynomial time within any constant above $\alpha_{\mathrm{GW}} = \min_{0<x\leq 1} \frac{\arccos(1-2x)}{\pi x} \approx 0.87856$ [18], the guarantee of the Goemans–Williamson algorithm. This threshold quantifies over all graphs. Theorem 15 and Theorem 21 concern the formulas $T_{H(G)}$ obtained from MAX-CUT instances G through Proposition 12 and Proposition 13. Every variable of such a formula occurs exactly three times, and the guarantees are 83/84 for d-regular G and 0.9932 for G without isolated vertices, both above α^* and $\beta^\star$. These are statements about a three-occurrence class and are therefore consistent with the hardness statements of this section.

Theorem 22 (Ratio above the Unique Games thresholds)**.** *For every $d \geq 1$ there is a polynomial-time algorithm that, given a simple d-regular graph G, returns an assignment of the three-occurrence Horn formula $T_{H(G)}$ of value at least $\frac{83}{84}\,\mathrm{OPT}(T_{H(G)})$, and*

$$\frac{83}{84} > \alpha^* \approx 0.94616 > \beta^\star \approx 0.94017.$$

Under the Unique Games Conjecture no polynomial-time algorithm has a factor above $\alpha^* + \varepsilon$ on all instances of weighted MAX $\{1,2\}$-HORN-SAT, or above $\beta^\star + \varepsilon$ on all instances of MAX 2-SAT, for any $\varepsilon > 0$. The class $\{T_{H(G)}\}$ is a class of Horn formulas, hence of MAX 2-SAT instances, in which every variable has exactly three literal occurrences, and it is consistent with the conjecture.

Proof. The algorithm and the factor 83/84 are Theorem 15. The numerical inequalities and the hardness statements are those quoted in the first paragraph of Section 9 and [12]. □

Take the ratio $\rho_d = 1 - \frac{1-\alpha_\varepsilon}{6d+12}$ of Theorem 15, which does not depend on γ. Then the cut fraction guaranteed by Corollary 14(2) is

$$1 - \frac{1 - \alpha_\varepsilon}{6d + 12}\Big(1 + \frac{6d + 11}{\gamma}\Big) \leq 1 - \frac{1 - \alpha_\varepsilon}{6d + 12}\,(6d + 12) = \alpha_\varepsilon,$$

with equality exactly for $\gamma = 1$. This is a lower bound valid for every $\gamma \in (0, 1]$ that uses only the value ρ_d; it is the worst case over γ of the corollary and is not a property of the algorithm. The algorithm does not obtain its cut from ρ_d. It first computes S with $\mathrm{cut}(S) \geq \alpha_\varepsilon \,\mathrm{maxcut}(G)$ and then lifts the independent set I of size $m + \mathrm{cut}(S)$ of Theorem 15. For this lift the instance ratio is $\rho_d(\gamma) = \frac{6d+11+\alpha_\varepsilon\gamma}{6d+11+\gamma} > \alpha_\varepsilon$ by

Remark 17, and Corollary 14(2) applied to $\rho_d(\gamma)$ returns exactly the cut fraction α_ε of the cut S. The MAX-CUT guarantee is thus the one of [17, 19], and the reduction transfers it to a guarantee on $T_{H(G)}$ that exceeds α^*. A bounded-degree cut guarantee c_d enters only through Proposition 18. Lean proves the clause-count identity $m + \mathrm{cut}(S)$ and the arithmetic of Corollary 14 (`CutTransfer`), the division-free inequality behind Proposition 16 (`CutTransfer`, `gap_ratio`), the strict inequality $\rho_d(\gamma) > \alpha_\varepsilon$ of Remark 17 in division-free form (`CutTransfer`, `lifted_ratio_gt_gw`), the lifting step with $18785\,\alpha(H) \le 20000\,|I|$, which follows from $|I| \ge \frac{1+\alpha_\varepsilon}{2}\alpha(H)$ (`CutLift`), and the degree-dependent identity and ratio step for non-regular graphs together with the arithmetic $8\alpha \le B$ of Theorem 21 (`GeneralDegreeFormula`) and the count $M \ge 4m$ (`ConflictDegree`), and the division-free form of Proposition 18 (`CitedResults`). The Goemans–Williamson guarantee and the Khot–Kindler–Mossel–O'Donnell threshold are taken from the literature (`CitedResults`) and are not formalized.

10. Max-Ones and Min-Ones 2-SAT

In Max-Ones (respectively Min-Ones) 2-SAT one is given a 2-CNF formula and must find a satisfying assignment with the largest (respectively smallest) number of true variables. Write $\mathrm{ones}(a)$ for the number of true variables of an assignment a.

10.1. **Negative formulas and independent sets.** A *negative* 2-CNF formula has only clauses $(\neg x_u \vee \neg x_v)$ and $(\neg x_u)$; it is a Horn formula. Let Z be the set of variables with a unit clause and let G_φ be the graph on the variables outside Z with an edge uv for every clause $(\neg x_u \vee \neg x_v)$ with $u, v \notin Z$. Clauses meeting Z are satisfied by every assignment that is false on Z.

Lemma 23. *An assignment satisfies a negative 2-CNF formula φ if and only if it is false on Z and its set of true variables is an independent set of G_φ. Hence* $\mathrm{MaxOnes}(\varphi) = \alpha(G_\varphi)$*, and a satisfying assignment with k true variables is the same object as an independent set of size k.*

Proof. A clause $(\neg x_u \vee \neg x_v)$ with $u, v \notin Z$ is falsified exactly when both x_u and x_v are true, that is, when u, v are adjacent true variables. A unit clause $(\neg x_u)$ is falsified exactly when x_u is true. The remaining clauses contain a variable of Z and are satisfied when that variable is false. □

Every simple graph G is G_φ for the negative formula with the clause $(\neg x_u \vee \neg x_v)$ for each edge, and in that formula x_u occurs $\deg(u)$ times. Thus Max-Ones for negative 2-CNF formulas is Maximum Independent Set with the same optimum and the same approximation ratio, and negative formulas in which every variable occurs at most Δ times correspond to graphs of maximum degree at most Δ.

For a simple cubic graph H let N_H be the negative formula with $(\neg x_u \vee \neg x_v)$ for every edge uv. Every variable of N_H occurs exactly three times, $\mathrm{MaxOnes}(N_H) = \alpha(H)$ by Lemma 23, and Theorem 3 reads

$$\mathrm{OPT}(T_H) = 4|V| + |E| + \mathrm{MaxOnes}(N_H).$$

For a d-regular graph H the same holds with $(d+1)|V| + |E|$ in place of $4|V| + |E|$, and for arbitrary simple graphs with $|V| + 3|E|$.

10.2. **Positive formulas and vertex covers.** A *positive* 2-CNF formula ψ has only clauses $(x_u \vee x_v)$ and (x_u); it is not a Horn formula. Let F be the set of variables with a unit clause and G_ψ the graph on the variables outside F with an edge uv for every clause $(x_u \vee x_v)$ with $u, v \notin F$.

Lemma 24. *An assignment satisfies ψ if and only if it is true on F and its set of true variables outside F is a vertex cover of G_ψ. Hence* $\mathrm{MinOnes}(\psi) = |F| + \tau(G_\psi)$*, where τ is the vertex cover number, and $\tau(G_\psi) = |V(G_\psi)| - \alpha(G_\psi)$.*

Proof. A unit clause (x_u) is falsified exactly when x_u is false. A clause $(x_u \vee x_v)$ with $u, v \notin F$ is falsified exactly when both are false, so it is satisfied exactly when the true set contains u or v. Clauses meeting F are satisfied by the units. The last identity is Gallai's identity for a graph. □

Every simple graph G is G_ψ for the positive formula with $F = \emptyset$ and the clause $(x_u \vee x_v)$ for each edge; then $\mathrm{MinOnes}(\psi) = \tau(G)$ and x_u occurs $\deg(u)$ times. Min-Ones for positive 2-CNF formulas is therefore Minimum Vertex Cover with the same optimum and the same approximation ratio. Formulas that contain both a positive and a negative literal in some clause, such as the implications $(\neg x \vee y)$ of T_H, are not covered by Lemmas 23 and 24.

10.3. **Hardness of Maximum Independent Set and Vertex Cover.** By Lemmas 23 and 24 and the remarks after them, the following published results apply verbatim to Max-Ones for negative 2-CNF formulas and to Min-Ones for positive 2-CNF formulas, with N the number of variables and Δ the maximum number of occurrences of a variable.

(1) Maximum Independent Set is NP-complete [2]; it cannot be approximated within $N^{1-\epsilon}$ for any $\epsilon > 0$ unless NP = ZPP [3], and this holds under P $\neq$ NP [4].
(2) Minimum Vertex Cover cannot be approximated within $2-\epsilon$ for any $\epsilon > 0$ under the Unique Games Conjecture [5, 6].
(3) Maximum Independent Set on graphs of maximum degree Δ cannot be approximated within $\Delta/\log^2\Delta$ under the Unique Games Conjecture [7] and within $\Delta/\log^4\Delta$ unless P = NP [8], up to constant factors; it can be approximated within $O(\Delta\log\log\Delta/\log\Delta)$ [9].
(4) Maximum Independent Set and Minimum Vertex Cover are APX-complete on cubic graphs [1], so Max-Ones for negative 2-CNF formulas and Min-Ones for positive 2-CNF formulas are APX-complete when every variable occurs exactly three times.

By Lemma 23 the cubic negative formula N_H and the formula T_H of Section 2 have optima that differ by the constant $4|V| + |E|$, and Theorem 3 converts an assignment of T_H of value A into a satisfying assignment of N_H with at least $A - (4|V| + |E|)$ true variables.

11. Lean and computational verification

The module `PalomaMax2SAT.CubicThreeOccurrence` represents the cubic graph with three ports at every vertex. It formalizes the transition count, false edge clauses, explicit deletion repair, independence of the repaired set, its cardinality bound, the matching assignment from an independent set, and the exact optimum identity. It also proves the division-free inequality

$$84(4n + m + |I|) \geq 83(4n + m + \alpha)$$

from $2m = 3n$, $2\alpha \leq n$, and $7|I| \geq 6\alpha$, together with the general arithmetic transfer

$$84\big((d+1)n + m + |I|\big) \geq 83\big((d+1)n + m + \alpha\big)$$

valid for every $d \geq 3$ under the analogous hypotheses $2m = dn$ and $2\alpha \leq n$. It also proves the rational identity $11+6/7 = (83/84)\cdot 12$ and the monotonicity statement that the endpoint ratio $(\beta+3/7)/(\beta+1/2)$ is at least $83/84$ for all $\beta \geq 11/2$. GitHub Actions builds the module and runs the public theorem comparison. The module `PalomaMax2SAT.SemanticHorn` defines T_H explicitly as a family of clauses, each a list of literals over the clone variables, indexed by the vertices (unit clauses), the edges (edge clauses) and the pairs (v, i) (consistency implications). It proves that the number of clauses of T_H satisfied by z equals $U(z) + (m - F(z)) + (4n - T(z))$, that the maximum of this number over all assignments is exactly $4n + m + \alpha$, that every clause has at most one negative literal, that T_H has $5n + m$ clauses, and that every variable occurs exactly three times. The Berman–Fujito guarantee, the Trevisan near-bipartite branch, the Goemans–Williamson rounding, the APX-completeness of cubic Maximum Independent Set, the analytic part of Proposition 16 and the discussion of the Unique Games Conjecture are proved in the text or taken from the literature and are not formalized in Lean. The module `PalomaMax2SAT.DegreeTransfer` formalizes the arithmetic of Proposition 5. From $2m = dn$ and $2\alpha \leq n$, the value $(d + 1)n + m$ satisfies $(3d+2)(B+\alpha) \leq (3d+3)B$ with $B = (d+1)n+m$. For $d \geq 27$ it satisfies $83(B+\alpha) \leq 84B$. For $d \geq 22$ and $5\alpha \leq (d+3)s$ one gets $83(B+\alpha) \leq 84(B+s)$. From $6\alpha \leq 7s$ one gets $(21d+20)(B+\alpha) \leq (21d+21)(B+s)$, and from $27\alpha \leq 28s + d\alpha$ one gets $83(B + \alpha) \leq 84(B + s)$. For $d = 21$ it proves the two-branch arithmetic: from $2m = 21n$, $2\alpha \leq n$, $200\alpha \leq 99n \Rightarrow 50\alpha \leq 241s$ and $99n < 200\alpha \Rightarrow 3n \leq 25s$ one gets $83(22n+m+\alpha) \leq 84(22n+m+s)$. It also proves the two branches of the $d = 21$ argument separately, the inequality $\frac{15803}{241} \cdot \frac{99}{200} < \frac{65}{2}$, the fact that $200\alpha > 99n$ gives a cut fraction above 0.99, the step from at most $7m/100$ uncut edges to $3n \leq 25|I|$, and the tabulated values of $\rho(d)$ and of $(21d + 20)/(21d + 21)$. The independent-set guarantees are hypotheses; Lean does not formalize the semidefinite rounding. The module `PalomaMax2SAT.DegreeFormula` defines T_H for every d as an explicit list of Horn clauses over the $d+1$ clones of each vertex. It proves that the number of satisfied clauses equals $U(z) + (m - F(z)) + ((d+1)n - T(z))$, that the optimum is exactly $(d + 1)n + m + \alpha$, that every clause is Horn, that T_H has $(d + 2)n + m$ clauses and that every variable occurs exactly three times. It derives $2m = dn$ from the port structure

and $2\alpha \le n$ for $d > 0$, proves that the lift of an independent set I has value $(d+1)n + m + |I|$, and combines these with the arithmetic of `DegreeTransfer` into statements about T_H itself: the baseline ratio $\frac{3d+2}{3d+3}$, the ratio 83/84 for $d \ge 27$, and the transfers of the Berman–Fujito, 6/7, threshold and $d = 21$ hypotheses. The module `PalomaMax2SAT.GeneralDegreeFormula` proves, for a per-vertex degree function d_v and cycles of length $d_v + 1$, the cycle lemma, the repair inequality $\mathrm{val}(z) \le B + |R(z)|$ with $R(z)$ independent, the value $B + |I|$ of the lift of an independent set, and hence the identity $\mathrm{OPT} = B + \alpha$ with $B = \sum_v (d_v + 1) + m$, which equals $n + 3m$ when $\sum_v d_v = 2m$ (Lemma 20); it also proves the division-free step $8\alpha \le B$, $18785\,\alpha \le 20000\,|I| \Rightarrow 9932(B+\alpha) \le 10000(B+|I|)$ used in Theorem 21. The explicit clause list and the three-occurrence count are formalized only for regular graphs, in `DegreeFormula`. The module `PalomaMax2SAT.GraphTransfer` formalizes the hard score of the four-port cycle formulation. It proves the exact optima $4m+\alpha$ (edge weight four) and $m+\alpha$ (unit weights), the repair of every feasible assignment, and the division-free transfers $260\,\mathrm{score} \ge 257(4m+\alpha)$ and $80\,\mathrm{score} \ge 77(m+\alpha)$ from an independent set of size s with $17\alpha \le 20s$. The statement $28(m+s) \ge 27(m+\alpha)$ is proved from the numerical hypothesis $6\alpha \le 7s$, or from $20000\alpha \le 9999n$ and $5999993\alpha \le 7000000s$; a bridge statement takes the two branch guarantees $20000\alpha \le 9999n \Rightarrow 5999993\alpha \le 7000000s$ and $20000\alpha > 9999n \Rightarrow 37n \le 80s$ as hypotheses. Lean does not prove that an algorithm attains either guarantee. The module `PalomaMax2SAT.HardOccurrenceSplit` contains the hard-cycle consistency lemma used by this module. The independent Python audit checks 16,314 small Horn formula shapes and verifies equality of the complete sets of attainable source and feasible-split objective values. The repository is available at [22]. An additional Python audit checks all independent sets on 258 small weighted formulas and enumerates labelled simple subcubic graphs through five vertices for the two-terminal path profile. The general signed-graph theorem and the use of the external approximation algorithms are proved in the text. For every degree $e+2$ it also proves the division-free form of the seventh claim of Proposition 5, with $\varepsilon = 1/c$ for a natural number c, both as arithmetic and for the lift of an independent set of the formula T_H. The arithmetic inequality $63d \le 4(d+1)(d+3) + 126$ used in the sweep over $d \ge 3$ is also proved in `DegreeTransfer`. The module `PalomaMax2SAT.CutTransfer` proves Proposition 12(1), that an assignment of F_G satisfies exactly $m + \mathrm{cut}(S)$ clauses, and the arithmetic of the two ratio transfers of Corollary 14; the conflict-graph identity of Proposition 13 is proved in `ConflictGraph`. The module `PalomaMax2SAT.CutLift` proves that an independent set I of H with $18785\,\alpha(H) \le 20000\,|I|$ lifts to the ratio 83/84 on T_H for every degree $D \ge 1$, which is the last step of Theorem 15; the construction of I from a cut is in `ConflictGraph` and the cut guarantee is taken from the literature. The module `PalomaMax2SAT.CitedResults` states the published guarantees used by the $d = 21$ argument as explicit hypotheses, namely the existence of a cut with at most $7m/100$ uncut edges [17, 19] and of a Berman–Fujito independent set [11]. It proves that such a cut yields an independent set I with $3n \le 25|I|$ and that, under the two hypotheses, some independent set lifts to a 83/84-approximation of T_H for $d = 21$. The module `PalomaMax2SAT.ConflictGraph` proves $\alpha(H(G)) = m + \mathrm{maxcut}(G)$ for every graph given by its edge list, and `PalomaMax2SAT.ConflictDegree` proves that the conflict graph is $(d+1)$-regular when G is d-regular, that x_v occurs $2\deg(v)$ times, that $H(G)$ has $4m$ vertices, and that $M \ge 4m$ for every graph (Theorem 21). The module `PalomaMax2SAT.LReduction` proves that a graph of maximum degree Δ has an independent set I with $n \le (\Delta+1)|I|$, the first L-reduction constant $\mathrm{OPT}(T_H) \le 23\,\alpha(H)$ for cubic graphs and $2\,\mathrm{OPT}(T_H) \le ((3d+2)(d+1)+2)\alpha(H)$ for d-regular graphs, and the second constant 1 (Theorem 6); the APX-completeness of cubic Maximum Independent Set is cited. The module `PalomaMax2SAT.PaperArithmetic` proves the comparison $\frac{83}{84} > 0.94615981 > 0.94016567$ of Theorem 22, the two-branch arithmetic of Theorem 4 with the two branch guarantees as hypotheses, the numerical constants of Theorems 15 and 21 and of Proposition 18, the bound $\gamma/(K+\gamma) \le 1/(K+1)$ of Remark 17, the cube example of Section 7, and the tabulated values. The module `PalomaMax2SAT.SignedGraph` proves Theorem 11: the independent sets of the signed graph are the pairs of a port set and a reward set satisfying the four edge conditions, their weight is at most $DM + \mathrm{OPT}(F)$, an independent set of weight A decodes to an assignment of score at least $A - DM$, and the bound is attained; the vertex count and the maximum degree of the graph are not formalized. The module `PalomaMax2SAT.OnesBridge` proves, for a cubic port graph H with $\alpha = \alpha(H)$, that $\mathrm{OPT}(T_H) = 4|V| + |E| + \alpha$, that $\mathrm{MaxOnes}(N_H) = \alpha$ and that the minimum number of true variables of the positive formula is $|V| - \alpha$. The module `PalomaMax2SAT.OnesOptimal` proves Lemmas 23 and 24 and Gallai's identity: an assignment satisfies the negative formula exactly when

its true set is an independent set avoiding the unit variables, so the maximum number of true variables is the independence number, and it satisfies the positive formula exactly when the unit variables are true and the remaining true variables cover the remaining edges, so the minimum number of true variables is $|F|$ plus the minimum such cover. The hardness results of Section 10 are not formalized. The semidefinite rounding, the Unique Games Conjecture and its hardness thresholds are not formalized; the notes in `doc/formalization.md` list each cited result with its reference.

References

[1] P. Alimonti and V. Kann, *Some APX-completeness results for cubic graphs*, Theoretical Computer Science 237 (2000), 123–134. https://doi.org/10.1016/S0304-3975(98)00158-3.

[2] R. M. Karp, *Reducibility among combinatorial problems*, in Complexity of Computer Computations, Plenum, 1972, 85–103.

[3] J. Håstad, *Clique is hard to approximate within* $n^{1-\epsilon}$, Acta Mathematica 182 (1999), 105–142.

[4] D. Zuckerman, *Linear degree extractors and the inapproximability of max clique and chromatic number*, Theory of Computing 3 (2007), 103–128.

[5] S. Khot, *On the power of unique 2-prover 1-round games*, STOC 2002, 767–775.

[6] S. Khot and O. Regev, *Vertex cover might be hard to approximate to within* $2-\epsilon$, Journal of Computer and System Sciences 74 (2008), 335–349.

[7] P. Austrin, S. Khot and M. Safra, *Inapproximability of vertex cover and independent set in bounded degree graphs*, Theory of Computing 7 (2011), 27–43.

[8] S. O. Chan, *Approximation resistance from pairwise-independent subgroups*, Journal of the ACM 63 (2016), Article 27. https://doi.org/10.1145/2873054.

[9] E. Halperin, *Improved approximation algorithms for the vertex cover problem in graphs and hypergraphs*, SIAM Journal on Computing 31 (2002), 1608–1623.

[10] L. Trevisan, *Max Cut and the smallest eigenvalue*, STOC 2009. https://theory.stanford.edu/~trevisan/pubs/maxcutstoc.pdf.

[11] P. Berman and T. Fujito, *On approximation properties of the independent set problem for low degree graphs*, Theory of Computing Systems 32 (1999), 115–132. https://doi.org/10.1007/s002240000113.

[12] J. Brakensiek, N. Huang, and U. Zwick, *Tight approximability of MAX 2-SAT and relatives, under UGC*, Proceedings of SODA 2024, 1328–1344. https://doi.org/10.1137/1.9781611977912.53.

[13] U. Feige, M. Karpinski, and M. Langberg, *Improved approximation of Max-Cut on graphs of bounded degree*, Journal of Algorithms 43 (2002), 201–219.

[14] E. Halperin, D. Livnat, and U. Zwick, *MAX CUT in cubic graphs*, Journal of Algorithms 53 (2004), 169–185. https://doi.org/10.1016/j.jalgor.2004.06.001.

[15] J.-T. Hsieh and P. K. Kothari, *Approximating Max-Cut on bounded degree graphs: tighter analysis of the FKL algorithm*, ICALP 2023, arXiv:2206.09204. https://arxiv.org/abs/2206.09204.

[16] S. Ghoshal, N. Huang, E. Lee, K. Makarychev, and Y. Makarychev, *Threshold rounding and bounded-degree Boolean MAX 2-CSP*, arXiv:2607.11050, 2026. https://arxiv.org/abs/2607.11050.

[17] M. X. Goemans and D. P. Williamson, *Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming*, Journal of the ACM 42 (1995), 1115–1145.

[18] S. Khot, G. Kindler, E. Mossel, and R. O'Donnell, *Optimal inapproximability results for MAX-CUT and other 2-variable CSPs?*, SIAM Journal on Computing 37 (2007), 319–357.

[19] S. Mahajan and H. Ramesh, *Derandomizing approximation algorithms based on semidefinite programming*, SIAM Journal on Computing 28 (1999), 1641–1663.

[20] OpenAI, *Proof of Khot's Unique Games Conjecture* (entry 102 of the OpenAI mathematics release, 6 October 2026), GitHub repository `openai/math`; not peer reviewed. https://github.com/openai/math.

[21] V. Raman, B. Ravikumar, and S. Srinivasa Rao, *A simplified NP-complete MAXSAT problem*, Information Processing Letters 65 (1998), 1–6.

[22] F. Vega, *Paloma: MAX-HORN-2SAT and MAX-CUT formalization*, GitHub, 2026. https://github.com/frankvegadelgado/paloma.

Information Physics Institute, 840 W 67th St, Hialeah, FL 33012, USA

Email address: vega.frank@gmail.com