Combinatorics, Probability and Computing (2026), **35**, pp. 83–88
doi:10.1017/S0963548325100187

CAMBRIDGE UNIVERSITY PRESS

ARTICLE

A note on the computational complexity of weak saturation

Martin Tancer* and Mykhaylo Tyomkyn*

Department of Applied Mathematics, Faculty of Mathematics and Physics, Charles University, Prague, Czech Republic
Corresponding author: Mykhaylo Tyomkyn; Email: tyomkyn@kam.mff.cuni.cz

(Received 21 January 2025; revised 15 August 2025; accepted 15 August 2025; first published online 16 October 2025)

Abstract
We prove that determining the weak saturation number of a host graph F with respect to a pattern graph H is computationally hard, even when H is the triangle. Our main tool establishes a connection between weak saturation and the shellability of simplicial complexes.

Keywords: Weak saturation; simplicial complex; shellability; NP-hardness
2020 MSC Codes: Primary: 05C35, 05E45; Secondary: 68Q17

1. Introduction

Weak saturation. Let F and H be graphs and let G be a spanning subgraph of F. We say that G is *weakly H-saturated in F*, if the edges of $E(F) \setminus E(G)$ can be ordered into a sequence $e_1, \dots, e_m$ in such a way that for every $i \in [m]$ the graph obtained from G by adding the edges $e_1, \dots, e_i$ contains a copy of H which contains e_i. The *weak saturation number* $\mathrm{wsat}(F, H)$ is defined as the minimum number of edges of a graph which is weakly H-saturated in F.

The concept of weak saturation was first introduced in 1968 by Bollobás [4] who considered the case when F and H are complete graphs and conjectured that $\mathrm{wsat}(K_n, K_t) = \binom{n}{2} - \binom{n-t+2}{2}$. This was confirmed by Frankl [10] and, independently, Kalai [14, 15] (a version for matroids was proven earlier by Lovász [16]) and extended by Alon [1] and Blokhuis [3]. Subsequently, weak saturation was studied for different classes of graphs and (in the analogous setting) hypergraphs; see [1, 2, 5, 6, 8, 9, 17, 18, 20, 21, 23, 27, 28, 25, 26]. As a common theme, upper bounds on $wsat(F, H)$ are usually established via simple constructions, while proving lower bounds tends to be much harder and typically requires methods from algebra or geometry.

In this note, we show that any kind of classification of weak saturation numbers in full generality is hopeless unless P = NP. More concretely, we show that determining the weak saturation number is already hard in a seemingly very simple case when H is the complete graph on three vertices K_3:

Theorem 1. *Given a graph F with n vertices as input, it is NP-hard to decide whether* $\mathrm{wsat}(F, K_3) = n - 1$.

Note that $\mathrm{wsat}(F, K_3) \geq n - 1$ for any connected n-vertex graph F, as any weakly K_3-saturated graph in F must be spanning. In our reduction, F will always be connected.

*Both authors were supported by GAČR grant 25-17377S.

© The Author(s), 2025. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives licence (https://creativecommons.org/licenses/by-nc-nd/4.0/), which permits non-commercial re-use, distribution, and reproduction in any medium, provided that no alterations are made and the original article is properly cited. The written permission of Cambridge University Press must be obtained prior to any commercial use and/or adaptation of the article.

https://doi.org/10.1017/S0963548325100187 Published online by Cambridge University Press

Our main tool to prove Theorem 1 is to establish a connection between weak saturation and shellability (and collapsibility) of simplicial complexes. We point out that a recent preprint [7] establishes a connection between weak saturation and d-collapsibility (closely related to collapsibility) in the context of fractional Helly-type theorems. Our setting, however, is quite different.

Simplicial complexes, shellability, and collapsibility. Let us recall that a(n abstract) *simplicial complex* is a set system K such that if $\sigma \in K$ and $\tau \subseteq \sigma$, then $\tau \in K$. We will consider only finite simplicial complexes. The set of *vertices* of K is the set $\bigcup K$. The elements of K are called *faces* of K and the *dimension* of a face $\sigma \in K$ is defined as $\dim \sigma := |\sigma| - 1$. The faces of dimension 1 are *edges* and the faces of dimension 2 are *triangles*. (A triangle in a simplicial complex should not be confused with a graph-theoretic triangle, a copy of K_3. We will avoid using the notion "triangle" in the latter context.) The *dimension* of the complex, $\dim K$ is defined as the maximum of the dimensions of faces in K. Given a non-negative integer k, the *k-skeleton* of a simplicial complex K is a subcomplex of K denoted $K^{(k)}$ consisting of faces of K of dimension at most k. From now on we regard graphs as (at most) 1-dimensional simplicial complexes. In particular, the 1-skeleton of a simplicial complex is a graph.

Given a sequence $\vartheta_1, \dots, \vartheta_k$ of faces of K we denote by $K[\vartheta_1, \dots, \vartheta_k]$ the subcomplex of K *induced* by these faces, that is, the subcomplex formed by faces σ such that $\sigma \subseteq \vartheta_i$ for some $i \in [k]$. An inclusion maximal face of a simplicial complex is a *facet* and a simplicial complex is *pure* if all facets have the same dimension. A pure d-dimensional complex K is *shellable* if there is an ordering $\vartheta_1, \dots, \vartheta_m$ of all facets of K such that for every $i \in \{2, \dots, m\}$ the complex $K[\vartheta_i] \cap K[\vartheta_1, \dots, \vartheta_{i-1}]$ is pure and $(d-1)$-dimensional.[1]

A simplicial complex K' arises from K by an *elementary collapse* if there is a face τ of K contained in a single facet σ distinct from τ and K' is obtained from K by removing all faces containing τ. A simplicial complex K *collapses* to a subcomplex L, if there is a sequence $K = K_1, K_2, \dots, K_\ell = L$ of simplicial complexes such that K_{i+1} is obtained from K_i by an elementary collapse for $i \in [\ell - 1]$. A simplicial complex K is *collapsible* if it collapses to a point (an arbitrary vertex of K).

The *reduced Euler characteristic* of a complex K is defined as

$$\tilde{\chi}(K) = \sum_{i=-1}^{\dim K} (-1)^i f_i(K)$$

where $f_i(K)$ is the number of i-dimensional faces of K. (Note that the empty set has the dimension equal to -1.) Given a simplicial complex K, its *barycentric subdivision* $\mathrm{sd}K$ is a complex whose vertices are nonempty faces of K and whose faces are collections $\{\vartheta_1, \dots, \vartheta_k\}$ of faces of K with $\emptyset \neq \vartheta_1 \subsetneq \vartheta_2 \subsetneq \cdots \subsetneq \vartheta_k$.

Hardness of shellability. In [11], Goaoc, Paták, Patáková, Tancer, and Wagner proved that shellability is NP-hard. We state a corollary of the main technical proposition from [11] in a way convenient for us. In the statement, we use *3-CNF formulas*, that is, formulas in a conjunctive normal form where each clause contains three literals. We skip details, referring the reader to [11], as we use 3-CNF formulas only implicitly. We only need the fact that the decision problem of whether a 3-CNF formula is satisfiable is a well-known NP-hard problem, known as *3-satisfiability*.

Theorem 2 (Essentially Proposition 8 from [11]). *There is a polynomial time algorithm that produces from a given 3-CNF formula ϕ with t variables a pure connected 2-dimensional complex K_ϕ with $\tilde{\chi}(K_\phi) = t$ such that the following statements are equivalent:*

[1] In other words, ϑ_i meets the previous facets in a pure $(d-1)$-dimensional subcomplex.

https://doi.org/10.1017/S0963548325100187 Published online by Cambridge University Press

(i) The formula ϕ is satisfiable.

(ii) The second barycentric subdivision $\mathrm{sd}^2 K_\phi$ is shellable.

(ii') The third barycentric subdivision $\mathrm{sd}^3 K_\phi$ is shellable.

(ii") The forth barycentric subdivision $\mathrm{sd}^4 K_\phi$ is shellable.

(iii) The complex K_ϕ is collapsible after removing some t triangles.

(iii') The barycentric subdivision $\mathrm{sd} K_\phi$ is collapsible after removing some t triangles.

(iii") The second barycentric subdivision $\mathrm{sd}^2 K_\phi$ is collapsible after removing some t triangles.

We are really interested only in the items (i), (ii), and (iii"). The remaining items are auxiliary for explaining the proof.

Because Proposition 8 from [11] is not formulated exactly this way, we briefly explain how Theorem 2 follows from [11]: The construction of K_ϕ is according to [11, Proposition 8]. The fact that the number of variables of ϕ equals $\tilde{\chi}(K_\phi)$ is the content of [11, Proposition 12]. Then the statements (i), (ii), (ii'), (iii), and (iii') of Theorem 2 are explicitly stated as equivalent statements in the (joint) proof of Theorems 4 and 5 in [11]. It remains to argue that (ii") and (iii") are equivalent as well. The proof of Theorems 4 and 5 in [11] contains in particular implications (ii) $\Rightarrow$ (ii') $\Rightarrow$ (iii') $\Rightarrow$ (i). The implications (ii') $\Rightarrow$ (ii") $\Rightarrow$ (iii") $\Rightarrow$ (i) work with the exactly same reasoning, which proves the equivalence.[2] Finally, it is possible to check that K_ϕ is connected directly from the construction in [11]. Alternatively, Skotnica and Tancer proved [24, Appendix A] that K_ϕ from exactly this construction is homotopy equivalent to the wedge of spheres. This also implies that K_ϕ is connected.

In our proof of Theorem 1, we use Theorem 2 with $L_\phi = \mathrm{sd}^2 K_\phi$, extending it to the following setting.

Theorem 3. *There is a polynomial time algorithm that produces from a given 3-CNF formula ϕ with t variables a pure 2-dimensional connected complex L_ϕ with $\tilde{\chi}(L_\phi) = t$ such that the following statements are equivalent:*

(i) The formula ϕ is satisfiable.

(ii) The complex L_ϕ is shellable.

(iii) The complex L_ϕ is collapsible after removing some t triangles.

(iv) We have $\mathrm{wsat}(L_\phi^{(1)}, K_3) = n - 1$ where n is the number of vertices of L_ϕ.

(v) The complex L_ϕ is collapsible after removing some number of triangles.

The proof of Theorem 3 is given in the next section. Theorem 1 follows immediately from the equivalence of (i) and (iv) in Theorem 3.

Proof of Theorem 1. The equivalence of (i) and (iv) in Theorem 3 provides a polynomial time reduction from 3-satisfiability to determining whether $\mathrm{wsat}(F, K_3) = n - 1$. (Note that the graph $L_\phi^{(1)}$ can be constructed from L_ϕ in polynomial time.) Given that 3-satisfiability is NP-hard, it follows that the latter problem is NP-hard as well. □

[2] For implications (ii) $\Rightarrow$ (ii') and (ii') $\Rightarrow$ (ii"), it is not hard to check directly, at least for 2-complexes, that if a complex is shellable, then so is its barycentric subdivision. In [11], there is a detour via vertex-decomposability needed for other results of the paper, but it can be avoided. The implications (ii') $\Rightarrow$ (iii') and (ii") $\Rightarrow$ (iii') come from a result of Hachimori [12] stating that, assuming a certain link condition, shellability of the second barycentric subdivision is equivalent to the collapsibility of the complex after removing a suitable number of triangles. Finally, the implications (iii') $\Rightarrow$ (i) and (iii") $\Rightarrow$ (i) follow directly from Proposition 8(iii) of [11] where the exact subdivision considered is irrelevant.

https://doi.org/10.1017/S0963548325100187 Published online by Cambridge University Press

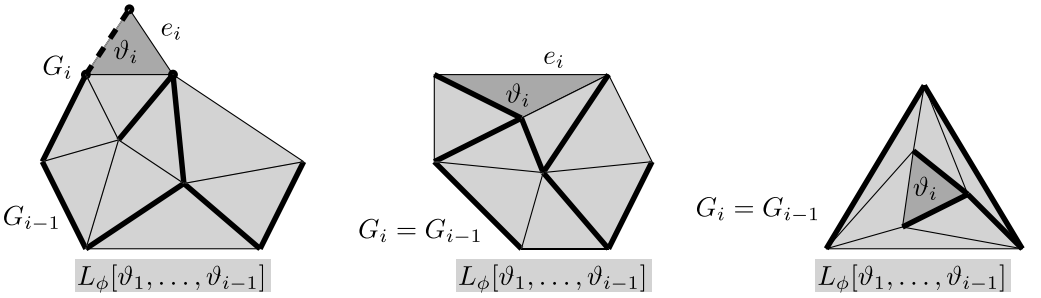


Figure 1. The figure displays three options of how ϑ_i may meet $L_\phi[\vartheta_1, \ldots, \vartheta_{i-1}]$ according to the number of shared edges. (The third displayed case is not fully realistic globally because the displayed $L_\phi[\vartheta_1, \ldots, \vartheta_{i-1}]$ is not shellable. However, it becomes realistic if we assume that the outer face is also part of the complex and it is actually ϑ_1.)

Finally, we remark that the items (ii), (iii) and (v) in Theorem 3 are again only auxiliary in order to show conveniently the equivalence of (i) and (iv).

2. The proof of Theorem 3

The aim of this section is to prove Theorem 3, completing the proof of Theorem 1.

As stated earlier, given a 3-CNF formula ϕ we take K_ϕ from Theorem 2 and set $L_\phi := \mathrm{sd}^2 K_\phi$. Given that K_ϕ is 2-dimensional, the complexity of L_ϕ grows only by a constant factor when compared with K_ϕ, and so L_ϕ can be constructed in polynomial time in the size of K_ϕ, hence in polynomial time in the size of ϕ. The complex L_ϕ is connected because K_ϕ is connected. We also remark that $\tilde{\chi}(L_\phi) = \tilde{\chi}(K_\phi) = t$ because a complex and its barycentric subdivision have the same reduced Euler characteristic.[3] The items (i), (ii), and (iii) of Theorem 3 are equivalent due to Theorem 2. We will now show implications (ii) $\Rightarrow$ (iv) $\Rightarrow$ (v) $\Rightarrow$ (iii), completing the proof.

Proof of (ii) $\Rightarrow$ (iv). Consider a *shelling* $\vartheta_1, \ldots, \vartheta_m$ of L_ϕ, that is, a sequence of all facets of L_ϕ witnessing that L_ϕ is shellable. Given that L_ϕ is pure 2-dimensional, all facets are triangles.

We construct a spanning tree G in the 1-skeleton $L_\phi^{(1)}$ in the following way. First we set G_1 inside ϑ_1 to contain all three vertices and two arbitrarily chosen edges. Next, for $i \in \{2, \ldots, m\}$, we inductively assume that we have a spanning tree G_{i-1} of $L_\phi[\vartheta_1, \ldots, \vartheta_{i-1}]^{(1)}$ and we construct a spanning tree G_i of $L_\phi[\vartheta_1, \ldots, \vartheta_i]^{(1)}$; see Fig. 1 for an illustration. We distinguish two cases. If ϑ_i meets the preceding triangles in two or three edges we set $G_i := G_{i-1}$. Note that $L_\phi[\vartheta_1, \ldots, \vartheta_{i-1}]$ and $L_\phi[\vartheta_1, \ldots, \vartheta_i]$ have the same sets of vertices in this case. Thus G_i is indeed a spanning tree. If ϑ_i meets the preceding triangles in a single edge, then we add to G_{i-1} the new vertex of ϑ_i (i.e., the vertex not contained in $L_\phi[\vartheta_1, \ldots, \vartheta_{i-1}]$) and one edge inside ϑ_i containing this vertex. Other cases are not possible because $\vartheta_1, \ldots, \vartheta_m$ is a shelling. We set $G := G_m$.

In order to finish the proof, we claim that G is weakly K_3-saturated in $L_\phi^{(1)}$. Indeed, the saturating sequence follows the shelling. The first edge e_1 is the unique edge of ϑ_1 not contained in G_1. This edge completes a copy of K_3 inside ϑ_1 (more precisely $L_\phi[\vartheta_1]^{(1)}$). Next, for $i \in \{2, \ldots, m\}$, we observe that $L_\phi[\vartheta_1, \ldots, \vartheta_i]^{(1)}$ contains at most one more edge not contained in G_i than $L_\phi[\vartheta_1, \ldots, \vartheta_{i-1}]^{(1)}$. It is exactly one edge e_i in the cases that ϑ_i meets the preceding triangles in one or two edges; see again Fig. 1 for an example. (With a slight abuse of notation we denote

[3] A complex and its barycentric subdivision are homeomorphic (see, for example [19, §15]), and the (reduced) Euler characteristic is an invariant under homeomorphism, even under homotopy [13, Theorem 2.44]. Alternatively, in the case of 2-complexes, one can check directly from the definition of the barycentric subdivision that given a 2-complex K with n vertices, m edges and t triangles, $\mathrm{sd}K$ contains $n + m + t$ vertices, $2m + 6t$ edges and $6t$ triangles.

https://doi.org/10.1017/S0963548325100187 Published online by Cambridge University Press

the edge by e_i though it may be not the i-th edge in the order, if some preceding edges are missing.) The edge e_i again completes the copy of K_3 inside ϑ_i, thereby inside $L_\phi[\vartheta_1, \dots, \vartheta_i]^{(1)}$. □

Proof of (iv) ⇒ (v). The facts that L_ϕ is connected and $\mathrm{wsat}(L_\phi^{(1)}, K_3) = n - 1$ imply that there exists a spanning tree G weakly K_3-saturated in $L_\phi^{(1)}$. We will show that $L_\phi^{(1)}$ collapses to G. This implies that $L_\phi^{(1)}$ is collapsible as any tree is collapsible.

Let $e_1, \dots, e_m$ be a sequence of edges witnessing that G is weakly K_3-saturated in $L_\phi^{(1)}$. For every such edge e_i we fix a copy J_i of K_3 it creates. Now we crucially use that L_ϕ is a barycentric subdivision of another complex. It is well known and not hard to show (at least for 2-complexes) that every copy of K_3 induces a triangle in L_ϕ. (In general barycentric subdivisions are *flag*, that is, every clique in the 1-skeleton induces a full simplex in the complex.) By ϑ_i we denote the triangle induced by J_i. We remark that the triangles ϑ_i are distinct because for $i < j$, ϑ_j contains the edge e_j while ϑ_i does not contain it.

We set L to be L_ϕ after removing all triangles that do not appear as ϑ_i for some $i \in [m]$. Now we perform elementary collapses on L in the reverse order of $e_1, \dots, e_m$. That is, we first claim that e_m is in a unique triangle ϑ_m. This is indeed the case as the triangles ϑ_i with $i < m$ do not contain e_m. We perform an elementary collapse on L removing e_m and ϑ_m. After performing this collapse we claim that e_{m-1} is in a unique triangle ϑ_{m-1}. This is indeed the case as ϑ_m has been already removed and the triangles ϑ_i with $i < m - 1$ do not contain e_{m-1}. We perform an elementary collapse on the intermediate complex removing e_{m-1} and ϑ_{m-1}. We continue this until we have collapsed L to G as required. (Note that every edge of L outside G appears as some e_i because $e_1, \dots, e_m$ witnesses that G is weakly K_3 saturated in $L_\phi^{(1)}$ and $L_\phi^{(1)} = L^{(1)}$.) □

Proof of (v) ⇒ (iii). Let L be a collapsible complex obtained from L_ϕ by removing k triangles. We know that $\tilde{\chi}(L_\phi) = t$, so $\tilde{\chi}(L) = t - k$ follows immediately from the definition of the reduced Euler characteristic. Because collapses preserve the homotopy type (this is explained, e.g., in a slightly more general setting in [22, Chapter 3]) and the (reduced) Euler characteristic is an invariant of the homotopy type (see [13, Theorem 2.44]), we deduce $\tilde{\chi}(L) = \tilde{\chi}(pt)$ where pt stands for a point.[4] However, $\tilde{\chi}(pt) = 0$ (the empty set and the point itself contribute -1 and 1, respectively). Thus we deduce $k = t$, which proves (iii). □

Acknowledgements

We would like to thank Adam Rajský for helpful early discussions.

References

[1] Alon, N. (1985) An extremal problem for sets with applications to graph theory. *J. Combin. Theory Ser. A* **40**(1) 82–89.
[2] Balogh, J., Bollobás, B., Morris, R. and Riordan, O. (2012) Linear algebra and bootstrap percolation. *J. Combin. Theory Ser. A* **119**(6) 1328–1335.
[3] Blokhuis, A. (1990) Solution of an extremal problem for sets using resultants of polynomials. *Combinatorica* **10**(4) 393–396.
[4] Bollobás, B. (1968) Weakly k–saturated graphs. In *Beiträge zur Graphentheorie (Kolloquium, Manebach, 1967)*. Leipzig: Teubner, pp. 25–31.
[5] Balogh, J. and Pete, G. (1998) Random disease on the square grid. In *Proceedings of the Eighth International Conference "Random Structures and Algorithms" (Poznan, 1997), Volume 13*, pp. 409–422.
[6] Bulavka, D., Tancer, M. and Tyomkyn, M. (2023) Weak saturation of multipartite hypergraphs. *Combinatorica* **43** 1081–1102.
[7] Chakraborti, D., Cho, M., Kim, J. and Kim, M. (2024) Colorful fractional Helly theorem via weak saturation, arXiv preprint arXiv:2408.15093.

[4] Alternatively, one can check directly from the definition of an elementary collapse that it removes the same number of faces of odd and of even size. Therefore elementary collapses preserve the (reduced) Euler characteristic.

https://doi.org/10.1017/S0963548325100187 Published online by Cambridge University Press

[8] Erdős, P., Füredi, Z. and Tuza, Z. (1991) Saturated r-uniform hypergraphs. *Discrete Math.* **98**(2) 95–104.
[9] Faudree, R. J. and Gould, R. J. (2014) Weak saturation numbers for multiple copies. *Discrete Math.* **336** 1–6.
[10] Frankl, P. (1982) An extremal problem for two families of sets. *Eur. J. Combin.* **3**(2) 125–127.
[11] Goaoc, X., Paták, P., Patáková, Z., Tancer, M. and Wagner, U. (2019) Shellability is NP-complete. *J. ACM* **66**(3) 18.
[12] Hachimori, M. (2008) Decompositions of two-dimensional simplicial complexes. *Discrete Math.* **308**(11) 2307–2312.
[13] Hatcher, A. (2002) *Algebraic Topology*. Cambridge University Press.
[14] Kalai, G. (1984) Weakly saturated graphs are rigid. In *Convexity and Graph Theory (Jerusalem 1981)* , Volume 87 of North-Holland Math. Stud, pp. 189–190.
[15] Kalai, G. (1985) Hyperconnectivity of graphs. *Graphs Combin.* **1**(1) 65–79.
[16] Lovász, L. (1977) Flats in matroids and geometric graphs. In Combinatorial Surveys (Proc. Sixth British Combinatorial Conf., Royal Holloway Coll., Egham, 1977), pp. 45–86.
[17] Morrison, N. and Noel, J. A. (2018) Extremal bounds for bootstrap percolation in the hypercube. *J. Combin. Theory Ser. A* **156** 61–84.
[18] Moshkovitz, G. and Shapira, A. (2015) Exact bounds for some hypergraph saturation problems. *J. Combin. Theory Ser. B* **111** 242–248.
[19] Munkres, J. R. (1984) *Elements of Algebraic Topology*. Addison - Wesley.
[20] Pikhurko, O. (2001) Uniform families and count matroids. *Graphs Combin.* **17**(4) 729–740.
[21] Pikhurko, O. (2001) Weakly saturated hypergraphs and exterior algebra. *Combin. Probab. Comput.* **10**(5) 435–451.
[22] Rourke, C. P. and Sanderson, B. J. (1982) *Introduction to Piecewise-Linear Topology*. Springer Study Edition. Springer-Verlag, Reprint.
[23] Shapira, A. and Tyomkyn, M. (2023) Weakly saturated hypergraphs and a conjecture of Tuza. *Proc. Am. Math. Soc.* **151** 2795–2805.
[24] Skotnica, M. and Tancer, M. (2023) NP-hardness of computing PL geometric category in dimension 2. *SIAM J. Discrete Math.* **37**(3) 2016–2029.
[25] Tuza, Z. (1988) Extremal problems on saturated graphs and hypergraphs. In *Eleventh British Combinatorial Conference (London, 1987)*, vol. 25, pp. 105–113.
[26] Tuza, Z. (1992) Asymptotic growth of sparse saturated structures is locally determined. *Discrete Math.* **108**(1-3) 397–402.
[27] Terekhov, N. and Zhukovskii, M. (2023) Weak saturation in graphs: a combinatorial approach, arXiv preprint arXiv:2305.11043.
[28] Terekhov, N. and Zhukovskii, M. (2024) Weak saturation rank: a failure of linear algebraic approach to weak saturation, arXiv preprint arXiv:2405.17857.

Cite this article: Tancer M and Tyomkyn M (2026). A note on the computational complexity of weak saturation. *Combinatorics, Probability and Computing* **35**, 83–88. https://doi.org/10.1017/S0963548325100187

https://doi.org/10.1017/S0963548325100187 Published online by Cambridge University Press