


Computational Topology and the Unique Games Conjecture

Joshua A. Grochow¹

Department of Computer Science & Department of Mathematics
University of Colorado, Boulder, CO, USA
jgrochow@colorado.edu
 <https://orcid.org/0000-0002-6466-0476>

Jamie Tucker-Foltz²

Amherst College, Amherst, MA, USA
jtuckerfoltz19@amherst.edu

Abstract

Covering spaces of graphs have long been useful for studying expanders (as “graph lifts”) and unique games (as the “label-extended graph”). In this paper we advocate for the thesis that there is a much deeper relationship between computational topology and the Unique Games Conjecture. Our starting point is Linial’s 2005 observation that the only known problems whose inapproximability is equivalent to the Unique Games Conjecture – Unique Games and Max-2Lin – are instances of Maximum Section of a Covering Space on graphs. We then observe that the reduction between these two problems (Khot–Kindler–Mossel–O’Donnell, FOCS ’04; SICOMP ’07) gives a well-defined map of covering spaces. We further prove that inapproximability for Maximum Section of a Covering Space on (cell decompositions of) closed 2-manifolds is also equivalent to the Unique Games Conjecture. This gives the first new “Unique Games-complete” problem in over a decade.

Our results partially settle an open question of Chen and Freedman (SODA, 2010; Disc. Comput. Geom., 2011) from computational topology, by showing that their question is almost equivalent to the Unique Games Conjecture. (The main difference is that they ask for inapproximability over \mathbb{Z}_2 , and we show Unique Games-completeness over \mathbb{Z}_k for large k .) This equivalence comes from the fact that when the structure group G of the covering space is Abelian – or more generally for principal G -bundles – Maximum Section of a G -Covering Space is the same as the well-studied problem of 1-Homology Localization.

Although our most technically demanding result is an application of Unique Games to computational topology, we hope that our observations on the topological nature of the Unique Games Conjecture will lead to applications of algebraic topology to the Unique Games Conjecture in the future.

2012 ACM Subject Classification Theory of computation → Problems, reductions and completeness, Theory of computation → Computational geometry


Keywords and phrases Unique Games Conjecture, homology localization, inapproximability, computational topology, graph lift, covering graph, permutation voltage graph, cellular graph embedding

Digital Object Identifier 10.4230/LIPIcs.SoCG.2018.43

Related Version A full version of this paper is available as arXiv:1803.06800 [cs.CC].

¹ J. A. G. gratefully acknowledges support from NSF grants DMS-1750319 and DMS-1622390 during the course of this work.

² J. T.-F. gratefully acknowledges the support of his Schupf Scholarship at Amherst College.



Acknowledgements We thank Alex Kolla for a great talk on Unique Games that inspired us in this direction. We thank Ryan O’Donnell, Hsien-Chih Chang, and Jeff Erickson for useful discussions and feedback. We especially thank Thomas Church for helping us clarify some of the topological issues involved and suggesting how to modify the proof of Thm. 14 to extend it to the S_n case (Thm. 18).

1 Introduction

A unique game is a constraint satisfaction problem in which every constraint is between two variables, say x_i and x_j , and for each assignment to x_i , there is a unique assignment to x_j which satisfies the constraint, and vice versa; in particular, the domain of each variable must have the same size k . Khot [21] conjectured that, for any $\varepsilon, \delta > 0$, there is some k such that it is NP-hard to distinguish between instances of Unique Games in which at most a δ fraction of constraints can be satisfied from those in which at least a $1 - \varepsilon$ fraction of the constraints can be satisfied. The Unique Games Conjecture (UGC) rose to prominence in the past 15 years partly because it implies that our current best approximation algorithms for many problems are optimal assuming $P \neq NP$ (e. g., [21, 25, 11, 23, 26]), thus explaining the lack of further progress on these problems. It is also interesting because, unlike $P \neq NP$, the UGC is a more well-balanced conjecture, with little consensus in the community as to its truth or falsehood. This even-handedness, together with the progress made in the last 10 years (e. g., [7, 34, 35, 27, 30, 3, 6, 24]) suggests that the UGC might be closer to resolution than other major conjectures in complexity theory like P versus NP or VP versus VNP.

Khot, Kindler, Mossel, and O’Donnell [23] showed that the UGC is equivalent to its special case, Γ -Max-2Lin, in which every constraint is of the form $x_i - x_j = c$, treated as equations over \mathbb{Z}_k (the cyclic group of order k). This beautiful simplification might lead one to naively expect that the UGC is somehow primarily about linear algebra, but this is potentially misleading. Indeed, a key feature in the solution of linear systems of equations is the ability to perform Gaussian elimination by taking linear combinations of equations, but when the equations are not satisfiable, taking linear combinations of equations can significantly change the maximum fraction of equations that are satisfiable. This leads us to ask: is there a domain of classical mathematics – other than modern computer science – in which the UGC is naturally situated?

In this paper, we argue that (algebraic) topology is such a domain. The starting point for our investigation is Linial’s observation [29]³ that the only two known “UGC-complete” problems – UG itself and Γ -Max-2Lin – are in fact instances of finding a maximum section of a (G -)covering space over the underlying constraint graph of the CSP (topological terminology will be explained in §2; $G = S_k$ for UG and $G = \mathbb{Z}_k$ for Γ -Max-2Lin). In the case of Γ -Max-2Lin, we observe that this is naturally equivalent to the well-studied 1-Homology Localization problem from computational topology (see, e. g., [14, 9, 12, 16, 13, 10, 39, 15]). We also observe that the reduction from Γ -Max-2Lin to UG [23] gives a well-defined map of G -covering spaces.

³ In an earlier version of this paper we were unaware of Linial’s observation, which appears on slides 55–56 of [29]. Once we were made aware of this, for which we thank an anonymous reviewer and Hsien-Chih Chang, we checked Linial’s slides, and the first author remembered having *attended* the talk that Linial gave at MIT on 11 May 2005! This was, in fact, one of the first theory seminars the first author had ever attended, and at the time he certainly didn’t know what bundles were, nor the UGC; he also could not recall whether Linial actually made it to those slides that particular day. We believe that Linial was the first to make this observation.

To cement the topological nature of the UGC, we then show that Maximum Section of a G -Covering Space, or 1-Homology Localization, on (cell decompositions of) 2-manifolds, rather than graphs, is still UGC-complete. This gives the first new UGC-complete problem in over a decade. Of course, there is some subjectivity as to what counts as a UGC-complete problem being “distinct” from UG itself. In particular, Γ -Max-2Lin can be viewed as UG with certain additional hypotheses satisfied, but nonetheless “feels” different (this difference can be made a little more precise topologically, see Remark 3). In our proof, we’ll see that 1-Homology Localization on 2-manifolds can also be viewed as a special case of Γ -Max-2Lin satisfying certain additional hypotheses, but again, Homology Localization “feels different” to us. Regardless, our results draw what we believe is a new connection between UGC and computational topology.

The UGC-completeness of this problem also partially settles a question of Chen and Freedman [14] on the complexity of the 1-Homology Localization problem on 2-manifolds. In particular, Chen and Freedman [14, p. 438, just before § 4.3] asked whether it was hard to approximate 1-Homology Localization with coefficients in \mathbb{Z}_2 on 2-manifolds; while some of the details are left unspecified, given the context in their paper we may conservatively infer (see our discussion in § 5) that they were asking for inapproximability to within all constant factors for triangulations of 2-manifolds. We show:

- Assuming UGC, for any constant $\alpha > 1$, there is a k such that 1-Homology Localization over \mathbb{Z}_k on cell decompositions of 2-manifolds cannot be efficiently approximated to within α . In particular, this problem is UGC-complete.
- Assuming UGC, for any $\varepsilon > 0$, there is a k such that 1-Homology Localization over \mathbb{Z}_k on *triangulations* of 2-manifolds cannot be efficiently approximated to within $7/6 - \varepsilon$.

Although the above are our most technically demanding results, which are applications of UGC to 1-Homology Localization – and, in the course of this, showing a new UGC-complete problem – we hope that the connections we have drawn between UGC and computational topology will lead to further progress on both of these topics in the future.

Related work. Linial [29] first observed that UG could be phrased in terms of Maximum Section of a Graph Lift (though we were unaware of this when we began our investigations, see Footnote 3). To our knowledge, since Linial’s observation there have been no other works relating approximation problems in computational topology with the Unique Games Conjecture, nor is there previous work on the problem of Maximum Section of G -Covering Spaces. In this paper we extend Linial’s observation by showing that the reduction of [23] gives a well-defined map of covering spaces, and we relate UGC to the well-studied problem of Homology Localization.

Here we briefly survey related work on approximation problems in computational topology, particularly those related to Homology Localization and the question of Chen and Freedman that we partially answer. Note that d -Homology Localization fixes the dimension of the homology considered, but allows the input to consist of d -homology classes on manifolds of arbitrarily large dimension. In our paper we consider 1-Homology Localization on graphs and 2-manifolds. For a more comprehensive overview of the area, as well as more direct motivations for the problem of Homology Localization, see [14, Sections 1 and 3].

1-Homology Localization with coefficients in \mathbb{Z}_2 is NP-hard to optimize exactly on simplicial complexes [12] and even on 2-manifolds [9]. Chen and Freedman showed it was NP-hard to approximate 1-Homology Localization on triangulations of 3-manifolds to within all constant factors, and that it was NP-hard to approximate d -Homology Localization on triangulations of manifolds for any $k \geq 2$. The best known algorithms for 1-Homology

Localization on a 2-manifold are given in several papers by Chambers, Erickson, and Nayyeri [9, 16, 10] (see also [39]). In particular, [9] solve the problem in polynomial time for fixed genus; however, for triangulations of 2-manifolds, the genus $g = \Theta(e/v)$ (follows from Euler's formula), and for instances of Unique Games to be difficult one must have the edge density e/v growing without bound, so their result seems not to solve the instances of 1-Homology Localization relevant for the UGC. Dey, Hirani, and Krishnamoorthy [15] showed that Homology Localization *in the 1-norm* over \mathbb{Z} can be done in polynomial time; in our paper we are primarily concerned with the 0-norm.

Organization. In § 2 we give preliminaries. § 3 contains the details of how to view Unique Games and Γ -Max-2Lin as instances of Maximum Section of a G -Covering Space, and the result that the KKMO reduction [23] gives a well-defined map of G -covering spaces. In § 4 we show that 1-Homology Localization on cell decompositions of 2-manifolds is UGC-complete. In § 5 we show how our techniques partially settle a question of Chen and Freedman, and in § 6 we discuss open questions. All omitted proofs are available in the full version, which also has a discussion of generalizations to non-Abelian groups G and arbitrary topological spaces X . Content that appears in both has the same numbers in both, and references to the full version are displayed as, e. g., Obs. 12^{FULL} or App. A^{FULL}.

2 Preliminaries

2.1 The Unique Games Conjecture and inapproximability

We refer to the textbooks [37, 5] for standard material on approximation algorithms and inapproximability, and to the survey [22] for more on the Unique Games Conjecture. Here we briefly spell out the needed definitions and one standard lemma that will be of use.

An instance of a constraint satisfaction problem (CSP) is specified by a set of variables x_1, \dots, x_n , for each variable x_i a domain D_i (which we will always take to be a finite set, and, in fact, we will have all D_i equal to one another), and a set of constraints. Each constraint is specified by a subset $\{x_{i_1}, \dots, x_{i_k}\}$ of k of the variables (each constraint may, in principle, have a different arity k), and a k -ary relation $R \subseteq D_{i_1} \times \dots \times D_{i_k}$. An assignment to the variables satisfies a given constraint if the assignment is an element of the associated R .

A CSP may be specified by restricting the arity and type of relations allowed in its instances, as well as the allowed domains for the variables. The value function associated to a CSP is $v(x, s) =$ the fraction of constraints in x satisfied by s , and we get the associated maximization problem. Given a CSP \mathcal{P} , the associated *gap problem* $\text{Gap}\mathcal{P}_{c,s}$ is the promise problem of deciding, given an instance x , whether $\text{OPT}(x) \leq s$ or $\text{OPT}(x) \geq c$. (An algorithm solving $\text{Gap}\mathcal{P}_{c,s}$ may make either output if x violates the promise, that is, if $s < \text{OPT}(x) < c$.) In general, the parameters c, s may depend on the problem size $|x|$.

If the optimization problem \mathcal{P} can be approximated to within a factor α by some algorithm, then essentially the same algorithm solves $\text{Gap}\mathcal{P}_{c,s}$ whenever $c/s > \alpha$. In the contrapositive, if $\text{Gap}\mathcal{P}_{c,s}$ is, for example, NP-hard, then so is approximating \mathcal{P} to within a factor c/s . The converse is false.

► **Problem (Unique Games).** The Unique Games problem with k colors, denoted $UG(k)$, is the CSP whose domains all have size exactly k , and where each constraint has arity 2 and is a bijection between the domains of its two variables.

The natural n -vertex graph associated to a UG instance – in which there is an edge (i, j) for each constraint on the pair (x_i, x_j) – is called its *constraint graph*. A $UG(k)$ instance is completely specified by its constraint graph, together with a permutation $\pi_{ij} \in S_k$ on each edge (i, j) , specifying the constraint that, for $i < j$, $x_i = \pi_{ij}(x_j)$.

► **Conjecture** (Khot [21], Unique Games Conjecture (UGC)). *For all $\varepsilon, \delta > 0$, there exists a $k \in \mathbb{N}$ such that $\text{GapUG}(k)_{1-\varepsilon, \delta}$ is NP-hard.*

Since the community is divided on this exact conjecture, the UGC is sometimes interpreted more liberally as saying that $\text{GapUG}(k)_{1-\varepsilon, \delta}$ is *somehow* hard, for example, not in P, BPP, or quasiP. All our results will work equally well under any of these interpretations, so we often just refer to “efficient approximation” or write “approximating ... is hard.”

A polynomial-time *gap-preserving reduction* from $\text{Gap}\mathcal{P}_{c,s}$ to $\text{Gap}\mathcal{Q}_{c',s'}$ (say, both minimization or both maximization problems) is a polynomial-time function f such that $\text{OPT}_{\mathcal{P}}(x) \leq s \Rightarrow \text{OPT}_{\mathcal{Q}}(f(x)) \leq s'$ and $\text{OPT}_{\mathcal{P}}(x) \geq c \Rightarrow \text{OPT}_{\mathcal{Q}}(f(x)) \geq c'$. If \mathcal{P} is a maximization problem and \mathcal{Q} is a minimization problem, then a gap-preserving reduction is similarly an f such that $\text{OPT}_{\mathcal{P}}(x) \leq s \Rightarrow \text{OPT}_{\mathcal{Q}}(f(x)) \geq c'$ and $\text{OPT}_{\mathcal{P}}(x) \geq c \Rightarrow \text{OPT}_{\mathcal{Q}}(f(x)) \leq s'$.

We say informally that a problem \mathcal{P} is “UGC-complete” if there are gap-preserving reductions from $\text{Gap}\mathcal{P}_{\alpha, \beta}$ to $\text{GapUG}_{1-\varepsilon, \delta}$ and $\text{GapUG}_{1-\varepsilon, \delta}$ to $\text{Gap}\mathcal{P}_{\alpha, \beta}$ (where, in one direction, ε, δ may depend on α, β , and vice versa in the other direction) such that some UGC-like statement holds for \mathcal{P} – such as “For any $\alpha < 1, \beta > 0$ $\text{Gap}\mathcal{P}_{\alpha, \beta}$ is hard to approximate” – if and only if UGC holds. Prior to this paper, the only known UGC-complete problems were UG itself,⁴ and $\Gamma\text{-Max-2Lin}(q)$ [23]:

► **Problem** ($\text{Max-2Lin}(A)$ and $\Gamma\text{-Max-2Lin}(A)$). Let A be an Abelian group. $\text{Max-2Lin}(A)$, or $\text{Max-2Lin}(k)$ when $A = \mathbb{Z}_k$, consists of those instances of UG where every variable has A as its domain, and each constraint takes the form $ax_i + bx_j = c$ for some $a, b \in \mathbb{Z}$ and $c \in A$ (not necessarily the same a, b, c for all constraints). $\Gamma\text{-Max-2Lin}(A)$ is the same, except that all the constraints have the form $x_i - x_j = c$ for some $c \in A$ (not necessarily the same for all constraints).

We will use the following standard lemma, which allows one to add a small number of new constraints to a given graph in a way that preserves an inapproximability gap.

► **Lemma 1.** *For a class \mathcal{A} of graphs, let $UG_{\mathcal{A}}$ denote the Unique Games Problem on graphs from \mathcal{A} . Given two classes of graphs \mathcal{A}, \mathcal{B} , let $f : \mathcal{A} \rightarrow \mathcal{B}$ be a polynomial-time computable function such that for all $G \in \mathcal{A}$, $E(G) \subseteq E(f(G))$ and $|E(f(G)) \setminus E(G)| = O(v)$ where v is the number of vertices in G of degree ≥ 1 . If the number of edges added is at most av , then there is a gap-preserving reduction from $UG_{\mathcal{A}, 1-\varepsilon, \delta}$ to $UG_{\mathcal{B}, 1-\varepsilon_0, \delta_0}$ where $\varepsilon_0 = \varepsilon + \Delta$ and $\delta_0 = \delta + \Delta$, for any $1 > \Delta > 2\delta a / (1 + 2\delta a)$ (in particular, with $\Delta \rightarrow 0$ as $\delta \rightarrow 0$).*

In particular, if $UG_{\mathcal{A}}$ is UGC-hard, then so is $UG_{\mathcal{B}}$. The same holds with “UG” everywhere replaced by Max-2Lin or $\Gamma\text{-Max-2Lin}$.

The intuition here is that one can always satisfy a number of constraints linear in the number of vertices (just choose a spanning tree or forest), so adding another linear number of constraints will only affect the inapproximability gap by a constant, which is negligible as δ and ε get arbitrarily small. For completeness the full version contains its (easy) proof §2.1^{FULL}.

2.2 G -covering spaces of graphs

► **Definition 2** (Graph lifts, a.k.a covering graph). Let X be a graph. A *graph lift*, or *covering graph*, is another graph Y with a map $p : V(Y) \rightarrow V(X)$ that such that the restriction of p

⁴ And slight variants, for example UG on bipartite graphs [21], or a variant due to Khot and Regev [25] in which one tries to maximize the number of vertices in an induced subgraph all of whose constraints are satisfied, rather than simply maximizing the number of constraints satisfied.

to the neighborhood of each $v \in V(Y)$ is a bijection onto the neighborhood of $p(v) \in X$. If X is connected, then the number k of points in $p^{-1}(v)$ is independent of v , and we say Y is a k -sheeted cover of X . The set of vertices $p^{-1}(v)$ is called the *fiber* over v .

Graph lifts have found many uses in computer science and mathematics, particularly in the study of expanders (e. g., [8, 20, 31, 32, 2]) and, via the next example, Unique Games.

► **Example 3 (Label-extended graph).** Given an instance of Unique Games with constraint graph X on vertex set $[n]$ and domain size k , and permutations π_e on the directed edges $e \in E(X)$, its *label-extended graph* is a graph with vertex set $[n] \times [k]$, and with an edge from (v, i) to (w, j) iff $\pi_{v,w}(i) = j$. In particular, the label-extended graph is a k -sheeted covering graph of X .

In our setting, all of our covering graphs will come naturally with a group that acts on their fibers, and we would like to keep track of this group action, for reasons that will become clear in § 3. For example, the label-extended graph of a UG instance carries a natural action of S_k on each fiber (as would be the case with *any* k -sheeted covering graph), and the label-extended graph of a Max-2Lin(A) instance has a natural action of the Abelian group A on each fiber. From the point of view of approximation, keeping track of the group currently seems of little relevance, but it may be useful from the topological point of view, so we state our definitions and results carefully keeping track of the (monodromy) group.

► **Definition 4 (G -covering graph, see [18] and [1, Definition 12]⁵).** Let G be a group of permutations on a set of size k . A G -covering space of a graph X is a k -sheeted covering graph $Z = (V(X) \times [k], E)$ such that the permutations on each edge come from the action of the group G . Symbolically, for each edge $(u, v) \in X$, there is a group element $g_{u,v} \in G$ and Z contains an edge from (u, i) to (v, j) iff $g_{u,v}(i) = j$.

In topological terminology, this definition is equivalent to a “ G -bundle with finite fibers” or to a covering space of the graph whose monodromy group (the group generated by considering the permutations you get by going around cycles in the graph) is contained in G .

We consider a graph X as a 1-dimensional geometric simplicial complex in the natural way, in which each edge has length 1.

► **Definition 5 (Section of a covering graph).** Given a covering graph $p: Y \rightarrow X$, a *section* of p is a continuous map $s: X \rightarrow Y$ (of topological spaces, as above) such that $p(s(x)) = x$ for all x . That is, it is a choice, for each $x \in X$, of a unique point in $p^{-1}(x)$, in a way that varies continuously with x .

► **Example 6.** Consider a covering graph $p: Y \rightarrow X$ where X is a triangle ($V(X) = \{0, 1, 2\}$, $E(X) = \{\{1, 2\}, \{0, 1\}, \{0, 2\}\}$), Y is a 6-cycle with vertex set $\{0, \dots, 5\}$, in its natural ordering (edge set $\{\{i, i + 1 \pmod{6}\} : i \in \{0, \dots, 5\}\}$), and $p(i) = i \pmod{3}$. This covering

⁵ Note that here we consider G as a permutation group – that is, technically, an abstract group together with an action on a set of size k , as is done in Definition 12 of the preprint [1]. In [2, Definition 1] (and [1, Definition 1]) they define a “ G -lift” for an abstract group G as a G -covering graph where the action of G is the regular action on itself by left translations. To translate between this terminology, that of bundles, and that of voltage graphs [18], we have:

Action	Covering graph	Bundle	Lift	Voltage Graph
regular	regular covering space	principal bundle	G - G -lift	ordinary voltage graph
general	G -covering graph (not nec. regular)	general bundle with finite fibers	(G, S, \cdot) -lift	permutation voltage graph

graph has no section: For, without loss of generality, we may suppose it has a section s and $s(0) = 0$. Then by continuity (imagine dragging a point $x \in Z$ from the point 0 across the edges of G), $s(1) = 1$ and $s(2) = 2$. But then as we continue varying our point in X across the edge $\{2, 0\}$, we find that we must also have $s(0) = 3$, a contradiction. If we think of Y as “lying over” X in the manner specified by p , we see that it is the label-extended graph of the UG instance on X with domain size 2, and π_e being the unique transposition for every edge e . The fact that there is no section here corresponds precisely to the fact that the UG instance is not completely satisfiable; see Obs. 11.

Given two covering graphs $p_i: Y_i \rightarrow X$ ($i = 1, 2$) of the same graph X , a *homomorphism* between covering graphs is a continuous map $f: Y_1 \rightarrow Y_2$ such that $p_2 \circ f = p_1$. In particular, this means that the points in Y_1 in the fiber over $x \in X$ (that is, in $p_1^{-1}(x)$) are mapped to points in Y_2 that also lie in the fiber over the same point x .

► **Example 7** (Isomorphism of label-extended graphs). Given two instances of UG on the same constraint graph X , if their label-extended graphs are isomorphic as covering graphs of X , then there is a natural bijection between assignments to the variables in the two instances which precisely preserves the number of satisfied constraints. Indeed, such an isomorphism is nothing more than re-labeling the elements of the domain of each variable.

► **Observation 8.** Given two G -covering graphs $p_\ell: Y_\ell \rightarrow X$ ($\ell = 1, 2$) with edge permutations $\pi_{ij}^{(\ell)}$, any isomorphism of covering graphs between them has the following form: for each $i \in V(X)$ there is a permutation π_i (not necessarily in G) such that $\pi_{ij}^{(2)} = \pi_i^{-1} \pi_{ij}^{(1)} \pi_j$.

Conversely, given a G -covering space $p_1: Y_1 \rightarrow X$ with edge permutations π_{ij} , and an element $g_i \in G$ for each $i \in V(X)$, the G -covering space Y_2 defined by $\pi_{ij}^{(2)} = g_i \pi_{ij}^{(1)} g_j^{-1}$ is isomorphic to Y_1 .

► **Definition 9.** An *isomorphism of G -covering graphs* $p_\ell: Y_\ell \rightarrow X$ ($\ell = 1, 2$) is an isomorphism of covering graphs such that the π_i (notation from Obs. 8) can be chosen to lie in G . When we say two G -covering spaces are “isomorphic”, we mean isomorphic as G -covering spaces (not just as covering spaces).

All these notions generalize from graphs to topological spaces (see App. A^{FULL}).

2.3 Homology and cohomology in dimensions ≤ 2

Let us briefly recall the problem of 1-Homology Localization, specialized to our context. Given a simplicial complex, or more generally a combinatorial CW complex X (see §2.3^{FULL}) of dimension 2, the group of d -cycles ($d = 0, 1, 2$) with coefficients in an Abelian group A , denoted $C_d(X, A)$ is isomorphic to the group A^{n_d} , where n_d is the number of d -simplices in X ($d = 0$: vertices, $d = 1$: edges, $d = 2$: triangles or 2-cells). We identify the coordinates of such a vector with an assignment of an element of A to each d -simplex of X . The support of a d -chain $a \in C_d(X, A)$ is the set of d -simplices that appear in a with nonzero coefficient. The boundary of a 1-simplex $[i, j]$ is the 0-chain $\partial_1([i, j]) := [i] - [j]$, and this operator ∂_1 is extended to a function $C_1(X, A) \rightarrow C_0(X, A)$ by A -linearity. Similarly, the boundary of a 2-cell $[i_1, i_2, \dots, i_\ell]$ is the 1-cycle $\partial_2[i_1, i_2, \dots, i_\ell] := [i_1, i_2] + [i_2, i_3] + [i_3, i_4] + \dots + [i_{\ell-1}, i_\ell] - [i_1, i_\ell]$, and we extend this to a map $C_2(X, A) \rightarrow C_1(X, A)$ by A -linearity. When no confusion may arise, we may refer to both of these maps simply as ∂ . The image of the boundary map ∂_d is a subgroup of $C_{d-1}(X, A)$, denoted $B_{d-1}(X, A)$.

A d -cycle is a d -chain $a \in C_d(X, A)$ such that $\partial a = 0$. For example, if X is a graph, a 1-cycle is just a union of cycles, in the usual sense of cycles in a graph; if X is a 2-manifold,

the only 2-cycles are A -scalar multiples of the entire manifold; all vertices are 0-cycles. The d -cycles form a subgroup of $C_d(X, A)$ denoted $Z_d(X, A)$.

Two d -cycles that differ by the boundary of a $(d+1)$ -cycle are *homologous*. The d -homology classes form the quotient group $H_d(X, A) := Z_d(X, A)/B_d(X, A)$. $H_0(X, A) \cong A^c$ where c is the number of connected components, and if X is a closed 2-manifold then $H_2(X, A) = A$. For closed 2-manifolds, thus the main interest is in $H_1(X, A)$.

► **Problem (1-Homology Localization, 1-HomLoc)**. Given a simplicial complex (resp., combinatorial CW complex) X and a 1-cycle $a \in Z_1(X, A)$, determine the sparsest homologous representative of a , that is, a 1-cycle a' homologous to a with minimum support among all 1-cycles homologous to a .

Cohomology is, in a sense, dual to homology. A d -cochain on X with coefficients in an Abelian group A is a homomorphism $C_d(X, \mathbb{Z}) \rightarrow A$; equivalently, it is determined by its values (from A) on the d -simplices (or d -cells) of X . The d -cochains form a group $C^d(X, A) \cong A^{n_d}$, where n_d is the number of d -simplices or d -cells. Given a d -cochain $f: X_d \rightarrow A$ (X_d being the d -simplices or d -cells of X), its *coboundary* is the function $(\delta f): X_{d+1} \rightarrow A$ defined by $(\delta f)(\Delta) = f(\partial\Delta)$ for any $(d+1)$ -cell Δ , and then extended A -linearly. Thus $\delta f \in C^{d+1}(X, A)$. A *d-cocycle* is a d -cochain whose coboundary is zero, equivalently, a d -cochain that evaluates to 0 on the boundary of any $(d+1)$ -chain. The d -cocycles form a subgroup $Z^d(X, A) \leq C^d(X, A)$. A d -coboundary is the coboundary of some $(d-1)$ -cochain; these form a subgroup $B^d(X, A) \leq Z^d(X, A)$. Two d -cochains that differ by a d -coboundary are said to be *cohomologous*, and the cohomology classes form a group $H^d(X, A) := Z^d(X, A)/B^d(X, A)$. As with homology, for 2-manifolds the main cohomological interest is in H^1 . The support of a d -cocycle is the number of d -simplices to which it assigns a nonzero value.

► **Problem (1-Cohomology Localization, 1-CohoLoc(G))**. Let G be a group (see App. A.3^{FULL} for the definitions in the non-Abelian case). Given a simplicial complex (resp. combinatorial CW complex) X and a 1-cocycle $a \in Z^1(X, G)$, find the sparsest cohomologous representative.

On closed surfaces, we have the following equivalence between these problems:

► **Observation 10** (J. Erickson, personal communication). *1-Cohomology Localization on CW complexes that are closed surfaces is equivalent to 1-Homology Localization on CW complexes that are closed surfaces, and dually (swapping the order of homology and cohomology).*

The proof essentially follows from (the proof of) Poincaré Duality; see §2.4^{FULL}.

3 The only known UGC-complete problems are Maximum Section of a G -Covering Graph

In this section we carefully write out the proof of Linial's observation [29] that UG (and Γ -Max-2Lin) is a special case of the following topological problem. We further observe that the reduction between these two problems [23] preserves the topological covering space structure of these problems.

► **Problem (Maximum Section of a (G -)Covering Graph)**. Let G be a group of permutations. Given a graph X and a G -covering graph $p: Y \rightarrow X$, find the partial section of p that is defined on as many edges of X as possible.

► **Observation 11** (Linial [29, pp. 55–56]). *UG(k) is the same as the Maximum Section of a S_k -Covering Graph Problem. Furthermore, given two isomorphic S_k -covering spaces of the same constraint graph X , there is a bijection between assignments to the two instances of UG that exactly preserves the number of constraints satisfied.*

Proof. Given an instance of $UG(k)$ on constraint graph X , we have seen in Example 3 that the label-extended graph Y of this instance is a covering graph of X ; let $p: Y \rightarrow X$ be the natural projection. As every constraint is a permutation in S_k , (p, Y) is clearly an S_k -covering space of X . Now suppose $u: X \rightarrow [k]$ is an assignment to the variables of X . We claim that this assignment can be extended to a section of p over the subset of $E(X)$ consisting precisely of those edges of X corresponding to constraints satisfied by u . For suppose $(i, j) \in E(X)$ and the constraint on (i, j) is satisfied by u , that is, $\pi_{ij}(u(i)) = u(j)$. Then in Y , there is an edge from $(i, u(i))$ to $(j, u(j))$ by construction. We may thus extend u to send points of the edge $(i, j) \in E(X)$ to the points of the edge $((i, u(i)), (j, u(j))) \in E(Y)$ bijectively and continuously, and thus extend u to a section of p that is defined over any edge satisfied by u .

Conversely, suppose $s: X' \rightarrow Y$ is a section of the restriction of p to $X' \subseteq X$, that is, $p|_{X'}: p^{-1}(X') \rightarrow X'$. We may use s to define a partial assignment to the variables of X . Namely, for any $i \in V(X')$ (that is, $i \in V(X)$ and $i \in X'$), define $u(i)$ by the equation $s(i) = (i, u(i)) \in V(Y)$. We claim that any edge of X contained entirely in X' is satisfied by this assignment u . Indeed, suppose the edge $(i, j) \in E(X)$ is contained entirely in X' . As s assigns $u(i)$ to i and $u(j)$ to j , and is continuous over all of X' , there must be an edge from $s(i) = (i, u(i))$ to $s(j) = (j, u(j))$ in Y . But this is the same as saying that $\pi_{ij}(u(i)) = u(j)$, and thus the constraint on this edge is satisfied. Therefore, maximizing the cardinality of the number of edges over which a section of p exists is the same as maximizing the cardinality of the number of constraints satisfied.

Finally, it is a folklore result that there is a bijection between assignments to two instances of $UG(k)$ on the same constraint graph X whose label-extended graphs are isomorphic graph lifts. Indeed, such an isomorphism corresponds simply to re-labeling the domain of each variable. As S_k is the maximal permutation group on a set of size k , Obs. 8 says that two isomorphic G -covering spaces of X are isomorphic graph lifts. ◀

► **Observation 12** (cf. Linial [29]). *Let A be an Abelian group. The Γ -Max-2Lin(A) Problem is the same as the Maximum Section of an A -Covering Graph Problem, where we view A as a permutation group acting on itself by translations. Furthermore, given two isomorphic A -covering spaces of the same constraint graph X , there is a bijection between assignments to the two instances of UG that exactly preserves the number of constraints satisfied.*

The proof is essentially the same as above; see Obs. 12^{FULL} for a little more detail.

Khot, Kindler, Mossel, and O’Donnell [23] prove that Γ -Max-2Lin(q) is UG-hard by giving a gap-preserving reduction from UG. Our next result is that this reduction sends isomorphic S_k -covering spaces to isomorphic \mathbb{Z}_q -covering spaces. See Prop. 13^{FULL} for the proof.

► **Proposition 13.** *The reduction [23] from $UG(k)$ to Γ -Max-2Lin(q) gives a well-defined map {isomorphism classes of S_k -covering spaces} \rightarrow {isomorphism classes of \mathbb{Z}_q -covering spaces}.*

► **Remark.** It is well-understood that one difference between $UG(k)$ and Γ -Max-2Lin(k) is that if an instance of Γ -Max-2Lin(k) is satisfiable, then one can choose an arbitrary vertex, assign an arbitrary value in \mathbb{Z}_k to this vertex, and propagate this value across the entire graph

using the constraints, and this will always be a solution. In contrast, even if an instance of $UG(k)$ is satisfiable, starting from a given vertex there may be (in the worst case) only one value that can be assigned to that vertex in such a way that the propagated assignment is actually satisfying.

In topological language, the above translates nearly exactly as follows (see Footnote 5): A principal bundle is trivial (a direct product) if and only if it has a section and if so, such a section can be found by making an arbitrary choice at one vertex and propagating. We note that when A is Abelian and the action is faithful, every A -covering space is a principal A -bundle, whereas this need not be the case for non-Abelian groups.

4 1-Homology Localization on cell decompositions of 2-manifolds is UGC-complete

► **Theorem 14.** *1-Homology Localization on cell decompositions of closed orientable surfaces is UGC-complete. More precisely, the Unique Games Conjecture holds if and only if for any $\varepsilon, \delta > 0$, there is some $k = k(\varepsilon, \delta)$ such that $\text{Gap1-HomLoc}_{1-\varepsilon, \delta}$ on cell decompositions of closed orientable surfaces with coefficients in \mathbb{Z}_k is NP-hard.*

To make the equivalence here a bit cleaner, we introduce a small variant of $\Gamma\text{-Max-2Lin}$ on surfaces instead of graphs:

► **Problem ($\Gamma\text{-Max-2Lin}(A)$ on surfaces).** Let A be an Abelian group. Given a cell decomposition of a closed surface X , and a 1-cocycle on X with coefficients in A treat the 1-cocycle as defining an instance of $\Gamma\text{-Max-2Lin}(A)$. In other words, this problem is the same as $\Gamma\text{-Max-2Lin}(A)$ on the 1-skeleton X_1 of X , except that we only consider instances of $\Gamma\text{-Max-2Lin}(A)$ in which the sum of the constraints along each cycle of X_1 that is the boundary of a 2-cell of X is zero.

From the second characterization in the definition (“in other words...”), it is clear that $\Gamma\text{-Max-2Lin}(A)$ on cell decompositions of surfaces is potentially easier than $\Gamma\text{-Max-2Lin}(A)$ on graph. We show that $\Gamma\text{-Max-2Lin}(A)$ on surfaces nonetheless remains UGC-complete.

Proof of Thm. 14. The proof proceeds as follows, and will take up the remainder of this section: $\Gamma\text{-Max-2Lin}(A)$ on graphs $\leq \Gamma\text{-Max-2Lin}(A)$ on surfaces $\cong 1\text{-CohoLoc}$ on surfaces $\cong 1\text{-HomLoc}$ on surfaces, where all reductions here are gap-preserving reductions. The first reduction is the technically tricky part, embodied in Prop. 17 below. The second equivalence is Obs. 15, which we do first since it will inform some of our subsequent discussion. The final equivalence is Obs. 10. ◀

► **Observation 15.** *$\Gamma\text{-Max-2Lin}(A)$ on a cell decomposition of a surface X is equivalent (under gap-preserving reductions) to $1\text{-CohoLoc}(A)$ on the same cell decomposition of the same surface.*

Proof. Given an instance of $\Gamma\text{-Max-2Lin}(A)$ specified by constants $a_{ij} \in A$ (that is, with constraints $x_i - x_j = a_{ij}$ for all edges in the cell decomposition of X), by the definition of $\Gamma\text{-Max-2Lin}(A)$ on surfaces the function $a: (i, j) \mapsto a_{ij}$ is a 1-cocycle. We claim that the following is a natural bijection between assignments to the variables and cohomologous 1-cocycles such that the set of constraints satisfied by an assignment is precisely the complement of the support of the associated 1-cocycle: Given an assignment α_i to the variables x_i , treat α as a 0-cochain, and consider the 1-cocycle $a - \delta\alpha$. Note that α satisfies some edge (i, j) iff $(a - \delta\alpha)(i, j) = 0$, for $(a - \delta\alpha)(i, j) = a_{ij} - \alpha_i + \alpha_j$.

Conversely, given a cohomologous 1-cocycle a' , the difference $a - a'$ is the coboundary of some 0-cochain: $a - a' = \delta\alpha$; treat α as an assignment to the variables. Using the same equation as before, we see that the support of a' is precisely the complement of the set of constraints satisfied by α .

Thus maximizing the number of satisfied constraints is equivalent to minimizing the support of a cohomologous cocycle. All that remains to check is that the equivalence above is indeed gap-preserving, noting that $\Gamma\text{-Max-2Lin}(A)$ is a maximization problem while $1\text{-CohoLoc}(A)$ is a minimization problem. Here we take the value of a cocycle to be the *fraction* of nonzero edges. The above shows that if the maximum fraction of satisfiable constraints in the $\Gamma\text{-Max-2Lin}(A)$ instance is ρ , then the minimum fraction of edges in the support of a cohomologous cocycle is $1 - \rho$ (since the number of edges in the same in both instances). So if a $\geq 1 - \varepsilon$ fraction of the constraints are satisfiable, then there is a cohomologous cocycle with support consisting of $\leq \varepsilon$ fraction of the edges; and if a $\leq \delta$ fraction of the constraints are satisfiable, then every cohomologous cocycle contains a $\geq 1 - \delta$ fraction of the edges. ◀

Our strategy for the reduction from $\Gamma\text{-Max-2Lin}(A)$ on graphs to $\Gamma\text{-Max-2Lin}(A)$ on surfaces will be to take an arbitrary graph and embed it as the 1-skeleton of a closed surface in polynomial time, in a gap-preserving way. The complication is that we must be careful about adding 2-cells. Given an instance of $\Gamma\text{-Max-2Lin}(A)$ on a graph X , and some cycle in X such that the sum of the constraints around the cycle is nonzero, we cannot simply “fill in” the cycle with a 2-cell. If we added such a cell to the complex, then the instance would no longer correspond to a 1-cocycle. Thus, we may only add 2-cells to cycles that are satisfiable in the given instance.

Furst, Gross, and McGeoch [17] give a polynomial-time algorithm to find the *maximal genus embedding* of a graph, that is, an embedding on an orientable closed surface such that the complement of the graph decomposes into a disjoint union of disks, in such a way as to maximize the genus of the surface. In a cellular embedding, Euler’s polyhedral formula applies: $V - E + F = 2 - 2g$, where V , E , and F are the numbers of vertices, edges, and faces of the complex, and g is the genus of the surface. Holding V and E fixed, we see that the problem of maximizing the genus is equivalent to minimizing the number of faces – and hence minimizing the additional properties an instance of $\Gamma\text{-Max-2Lin}(A)$ on a surface must satisfy compared to being on a graph. In the extreme case, their algorithm may find an embedding with only one face, which “wraps around” the surface touching every edge twice, once on each side. Because the orientations of the region on either side oppose each other, the boundary of this region will always be zero. Therefore, if such a region is added to the complex, *any 1-cochain will still be a 1-cocycle*.

The algorithm of Furst, Gross, and McGeoch [17] is based on the work of Xuong [38], who gave a complete characterization of how many regions one needs to embed a graph. We restate the special case of his main result in which only one region is needed.

▶ **Theorem 16** (Xuong [38]). *A connected graph G has a one-face cellular embedding into a closed orientable surface if and only if there exists a spanning tree T such that every connected component of $G \setminus T$ has an even number of edges.*

See [17, Lemmas 3.1 & 3.3] for a proof of this special case.

We are now ready to formally describe our reduction and prove its correctness.

▶ **Proposition 17.** *There is a reduction of the form in Lemma 1 from $\Gamma\text{-Max-2Lin}(A)$ on graphs to $\Gamma\text{-Max-2Lin}(A)$ on cell decompositions of surfaces, and therefore $\Gamma\text{-Max-2Lin}(A)$ on surfaces is UGC-complete.*

Proof. Suppose we are given a 1-cocycle on a graph X (every 1-cochain on a graph is a 1-cocycle, since there are no 2-cells). We describe an algorithm to transform X into a graph X' that admits a cellular embedding with one region. First, as X is the constraint graph of a Γ -Max-2Lin(A) instance, we may assume without loss of generality that X has no isolated vertices. Now, if X is disconnected, connect the components together, using one edge for each extra component. If the total number of edges is now odd, append a leaf attached by an edge to any vertex. Finally, add a “universal” vertex u , with an edge to every other vertex. Each new edge added in the preceding steps can be labeled with any constraint; it will not matter, but say $x_i - x_j = 0$ for concreteness. Let X' be the resulting graph. Letting T be the star spanning tree centered at u – that is, consisting of precisely the edges incident on u – we see that $X' \setminus T$ has only one component, with an even number of edges, so X' can be embedded in polynomial time with one region. The fact that this reduction preserves the inapproximability gap is the content of Lemma 1: If Γ -Max-2Lin(A) was hard on arbitrary graphs, it will still be hard on graphs with one-face embeddings, as this reduction added at most $(\frac{v}{2} - 1) + 1 + (v + 1) = O(v)$ edges. ◀

This completes the proof of Thm. 14. We would like to further reduce the instance so that it is a simplicial complex, rather than just a cell decomposition, which would prove that the simplicial version of the problem is UGC-hard as well. However, to break a 1-region embedding into triangles seems to require so many edges that it seems impossible to preserve the inapproximability gap completely. In the next section, we manage to preserve a 7/6 gap.

With G -covering spaces suitably defined, we show essentially the same result for non-Abelian G . We use G -covering spaces rather than 1-CohoLoc because 1-CohoLoc only corresponds to principal G -covering spaces, in which G acts on itself by translations. Using the following theorem, we could have shown UGC-completeness of Maximum Section of a G -Covering Space on cell decompositions of 2-manifolds without going through Max-2Lin; we chose the above route as the concepts with Abelian coefficient groups are simpler and more well-known. For the needed non-Abelian definitions and for the proof of this next result, see App. A^{FULL}. The hypothesis of the following result is satisfied by the symmetric groups S_k and all finite simple groups [33, 28].

► **Theorem 18.** *Let G be a group such that every product of commutators of G is equal to a commutator. Then Maximum Section of G -Covering Spaces on graphs reduces to Maximum Section of G -Covering Spaces on cell decompositions of 2-manifolds. In particular, the latter problem for (S_k) -covering spaces of cell decompositions of 2-manifolds is UGC-complete.*

5 On a question of Chen and Freedman

“This raises the open question [of] whether localizing a one-dimensional class of a 2-manifold is NP-hard to approximate...” –Chen and Freedman [14, p. 438]

Note that Chen and Freedman showed that d -Homology Localization, for any fixed $d \geq 2$ is indeed NP-hard to approximate to within any constant factor for triangulations of manifolds (of unbounded dimension), as well as 1-HomLoc for triangulations of 3-manifolds. One might thus infer from the above quote that they were asking the same question – inapproximability to within any constant factor – for triangulations of 2-manifolds. We note that although a greedy algorithm gives a k -approximation to Unique Games over \mathbb{Z}_k [36, Appendix], when we translate this maximization problem to the minimization problem of 1-HomLoc, we have essentially no control over the approximation ratio. We nonetheless show some inapproximability, by a different method (see §5^{FULL} for the proof).

► **Theorem 19.** *Assuming UGC, for any $\varepsilon > 0$, there is some $k = k(\varepsilon)$ such that it is hard to approximate 1-HomLoc over \mathbb{Z}_k on triangulations of surfaces to within a factor of $7/6 - \varepsilon$.*

Elsewhere in their paper they consider cell decompositions. If we relax their question to cell decompositions rather than triangulations, and we allow the coefficient group to be \mathbb{Z}_k (rather than just \mathbb{Z}_2), then Thm. 14 states that their question becomes equivalent to UGC.

6 Future directions

Although our most technically demanding results were applying UGC to 1-Homology Localization – and, in the course of this, showing a new UGC-complete problem – we hope that the connections we have drawn between UGC and computational topology will lead to further progress on both topics in the future. Here we highlight a few specific questions suggested by our investigations.

Inapproximability of 1-Homology Localization for triangulations of 2-manifolds? Although our results partially settle a question of Chen and Freedman [14, p. 438], we leave open the following questions:

► **Open Question 20.** Show (unconditionally) that there is a $c > 1$ such that it is NP-hard to c -approximate 1-Homology Localization over \mathbb{Z}_2 on triangulations of 2-manifolds.

► **Open Question 21.** Does UGC imply that for all $c > 1$, it is hard to c -approximate 1-Homology Localization on triangulations of 2-manifolds (over \mathbb{Z}_k for $k = k(c)$)?

We note that our reduction from Thm. 19 does not provide a strong enough gap for these problems to be immediately answered by the known NP-hardness results for Max-2Lin [19, 4]. In particular, Håstad shows that $\text{GapMax-2Lin}(2)_{\frac{12}{16}, \frac{11}{16}}$ is NP-hard (up to an additive arbitrary δ in the soundness and completeness), but if we plug these values into ε_0, δ_0 from our proof, we get a ratio of $\frac{92}{129} < 1$. A quick calculation shows that, to get any NP-hardness result (without UGC) from the proof of Thm. 19, we would need an inapproximability ratio for Max-2Lin of strictly greater than $\frac{12}{5}$, which is not provided by [19] nor [4]. Given the best upper bounds on approximating Max-2Lin(p) [4], which are just slightly less than p , such a ratio is not possible for $p = 2, 3$, but is possible already for $p = 5$.

G -covering spaces for other families of groups and group actions. The viewpoint of G -covering spaces suggests it might be fruitful to consider instances of UG that correspond to S_n -covering spaces for other actions of S_n . For example, one might consider the action of S_n on unordered k -tuples $\binom{[n]}{k}$, or even on n -vertex graphs $2^{\binom{[n]}{2}}$ (here, each variable in the UG would have the set of n -vertex graphs as its domain). We note that although much of the structure of any such instance is governed by the permutation constraints, the approximability properties may change significantly by varying the action. The general linear groups $\text{GL}_n(\mathbb{F}_q)$, acting on the vector space \mathbb{F}_q^n , as well as Schur functors thereof, or other representations of $\text{GL}_n(\mathbb{F}_q)$, strike us as leading to other possibly interesting approximation problems deserving further study. If one is looking for hard instances of Unique Games, one must construct such covering spaces so that they are not (sufficiently good) expanders [7]; see [2] for results on the expansion of G -covering spaces.

References

- 1 Naman Agarwal, Karthekeyan Chandrasekaran, Alexandra Kolla, and Vivek Madan. On the expansion of group-based lifts. arXiv:1311.3268 [cs.DM], 2016.
- 2 Naman Agarwal, Karthekeyan Chandrasekaran, Alexandra Kolla, and Vivek Madan. On the expansion of group-based lifts. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2017, August 16-18, 2017, Berkeley, CA, USA*, pages 24:1–24:13, 2017. doi:10.4230/LIPIcs.APPROX-RANDOM.2017.24.
- 3 Naman Agarwal, Guy Kindler, Alexandra Kolla, and Luca Trevisan. Unique games on the hypercube. *Chic. J. Theoret. Comput. Sci.*, pages Article 1, 20, 2015. doi:10.4086/cjtc.2015.001.
- 4 Gunnar Andersson, Lars Engebretsen, and Johan Håstad. A new way of using semidefinite programming with applications to linear equations mod p . *J. Algorithms*, 39(2):162–204, 2001. Originally appeared in SODA 1999. doi:10.1006/jagm.2000.1154.
- 5 Sanjeev Arora and Boaz Barak. *Computational complexity*. Cambridge University Press, Cambridge, 2009. A modern approach. doi:10.1017/CBO9780511804090.
- 6 Sanjeev Arora, Boaz Barak, and David Steurer. Subexponential algorithms for unique games and related problems. *J. ACM*, 62(5):Art. 42, 25, 2015. Originally appeared in FOCS 2010. doi:10.1145/2775105.
- 7 Sanjeev Arora, Subhash A. Khot, Alexandra Kolla, David Steurer, Madhur Tulsiani, and Nisheeth K. Vishnoi. Unique games on expanding constraint graphs are easy: Extended abstract. In *STOC '08: 40th Annual ACM Symposium on Theory of Computing*, pages 21–28, 2008. doi:10.1145/1374376.1374380.
- 8 Yonatan Bilu and Nathan Linial. Lifts, discrepancy and nearly optimal spectral gap. *Combinatorica*, 26(5):495–519, 2006. doi:10.1007/s00493-006-0029-7.
- 9 Erin W. Chambers, Jeff Erickson, and Amir Nayyeri. Minimum cuts and shortest homologous cycles. In *SoCG '09: Proceedings of the 25th Annual Symposium on Computational Geometry*, pages 377–385, 2009. doi:10.1145/1542362.1542426.
- 10 Erin W. Chambers, Jeff Erickson, and Amir Nayyeri. Homology flows, cohomology cuts. *SIAM J. Comput.*, 41(6):1605–1634, 2012. Originally appeared in STOC 2009. doi:10.1137/090766863.
- 11 Shuchi Chawla, Robert Krauthgamer, Ravi Kumar, Yuval Rabani, and D. Sivakumar. On the hardness of approximating multicut and sparsest-cut. *computational complexity*, 15(2):94–114, Jun 2006. Originally appeared in CCC 2005. doi:10.1007/s00037-006-0210-9.
- 12 Chao Chen and Daniel Freedman. Quantifying homology classes II: Localization and stability. arXiv:07092512 [cs.CG], 2007.
- 13 Chao Chen and Daniel Freedman. Measuring and computing natural generators for homology groups. *Comput. Geom.*, 43(2):169–181, 2010. doi:10.1016/j.comgeo.2009.06.004.
- 14 Chao Chen and Daniel Freedman. Hardness results for homology localization. *Discrete Comput. Geom.*, 45(3):425–448, 2011. Originally appeared in SODA 2010. doi:10.1007/s00454-010-9322-8.
- 15 Tamal K. Dey, Anil N. Hirani, and Bala Krishnamoorthy. Optimal homologous cycles, total unimodularity, and linear programming. *SIAM J. Comput.*, 40(4):1026–1044, 2011. Originally appeared in STOC 2010. doi:10.1137/100800245.
- 16 Jeff Erickson and Amir Nayyeri. Minimum cuts and shortest non-separating cycles via homology covers. In *SODA '11: Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1166–1176. SIAM, Philadelphia, PA, 2011.
- 17 Merrick L. Furst, Jonathan L. Gross, and Lyle A. McGeoch. Finding a maximum-genus graph imbedding. *J. ACM*, 35(3):523–534, 1988. doi:10.1145/44483.44485.

- 18 Jonathan L. Gross and Thomas W. Tucker. Generating all graph coverings by permutation voltage assignments. *Discrete Math.*, 18(3):273–283, 1977. doi:10.1016/0012-365X(77)90131-5.
- 19 Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001. Originally appeared in STOC 1997. doi:10.1145/502090.502098.
- 20 Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bull. Amer. Math. Soc. (N.S.)*, 43(4):439–561, 2006. doi:10.1090/S0273-0979-06-01126-8.
- 21 Subhash Khot. On the power of unique 2-prover 1-round games. In *STOC '02: 34th Annual ACM Symposium on Theory of Computing*, pages 767–775. ACM, 2002. doi:10.1145/509907.510017.
- 22 Subhash Khot. On the unique games conjecture (invited survey). In *CCC '10: 25th IEEE Conference on Computational Complexity*, pages 99–121, June 2010. doi:10.1109/CCC.2010.19.
- 23 Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *SIAM J. Comput.*, 37(1):319–357, 2007. Originally appeared in FOCS 2004. doi:10.1137/S0097539705447372.
- 24 Subhash Khot and Dana Moshkovitz. Candidate hard unique game. In *STOC '16: 48th Annual ACM Symposium on Theory of Computing*, pages 63–76, 2016. doi:10.1145/2897518.2897531.
- 25 Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within $2 - \epsilon$. *J. Comput. System Sci.*, 74(3):335–349, 2008. Originally appeared in CCC 2003. doi:10.1016/j.jcss.2007.06.019.
- 26 Subhash A. Khot and Nisheeth K. Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative-type metrics into ℓ_1 . *J. ACM*, 62(1):8:1–8:39, 2015. Originally appeared in FOCS 2005. doi:10.1145/2629614.
- 27 Alexandra Kolla. Spectral algorithms for unique games. *Comput. Complexity*, 20(2):177–206, 2011. Originally appeared in CCC 2010. doi:10.1007/s00037-011-0011-7.
- 28 Martin W. Liebeck, E. A. O'Brien, Aner Shalev, and Pham Huu Tiep. The Ore conjecture. *J. Eur. Math. Soc. (JEMS)*, 12(4):939–1008, 2010. doi:10.4171/JEMS/220.
- 29 Nati Linial. Lifts of graphs. Slides of presentation, available at http://www.cs.huji.ac.il/~nati/PAPERS/lifts_talk.pdf, 2005.
- 30 Konstantin Makarychev and Yury Makarychev. How to play unique games on expanders. In Klaus Jansen and Roberto Solis-Oba, editors, *Approximation and Online Algorithms: 8th International Workshop, WAOA 2010, Liverpool, UK, September 9-10, 2010. Revised Papers*, pages 190–200, 2011. doi:10.1007/978-3-642-18318-8_17.
- 31 Adam W. Marcus, Daniel A. Spielman, and Nikhil Srivastava. Interlacing families I: Bipartite Ramanujan graphs of all degrees. *Ann. of Math. (2)*, 182(1):307–325, 2015. doi:10.4007/annals.2015.182.1.7.
- 32 Adam W. Marcus, Daniel A. Spielman, and Nikhil Srivastava. Interlacing families II: Mixed characteristic polynomials and the Kadison-Singer problem. *Ann. of Math. (2)*, 182(1):327–350, 2015. doi:10.4007/annals.2015.182.1.8.
- 33 Oystein Ore. Some remarks on commutators. *Proc. Amer. Math. Soc.*, 2:307–314, 1951. doi:10.2307/2032506.
- 34 Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *STOC '08: 40th Annual ACM Symposium on Theory of Computing*, pages 245–254, 2008. doi:10.1145/1374376.1374414.
- 35 Prasad Raghavendra and David Steurer. Graph expansion and the unique games conjecture. In *STOC '10: 42nd Annual ACM Symposium on Theory of Computing*, pages 755–764, 2010. doi:10.1145/1806689.1806792.

43:16 Computational Topology and the Unique Games Conjecture

- 36 Sartaj Sahni and Teofilo Gonzalez. P-complete approximation problems. *J. ACM*, 23(3):555–565, 1976. doi:10.1145/321958.321975.
- 37 Vijay V. Vazirani. *Approximation algorithms*. Springer-Verlag, Berlin, 2001.
- 38 Nguyen Huy Xuong. How to determine the maximum genus of a graph. *J. Combin. Theory Ser. B*, 26(2):217–225, 1979. doi:10.1016/0095-8956(79)90058-3.
- 39 Afra Zomorodian and Gunnar Carlsson. Localized homology. *Comput. Geom.*, 41(3):126–148, 2008. doi:10.1016/j.comgeo.2008.02.003.