Approaches to prove the Unique Games Conjecture

Given the sub-exponential time algorithm for unique games (Arora et al. [2015]), that under the exponential time hypothesis, the unique games conjecture implies the following conjecture:

> **Intermediate complexity conjecture:** There exist some $1 > c > s > 0$, $1 > \alpha > \beta > 0$ and a CSP $CSP_\Sigma(\mathcal{P})$ such that the c vs s problem for $CSP_\Sigma(\mathcal{P})$ *can* be solved in time $\exp(O(n^\alpha))$ but it *cannot* be solved in time faster than $\exp(\Omega(n^\beta))$.

This is a very interesting conjecture in its own right ,as it says that unlike the widely believed situation for *exact* computation, it is *not* the case that every CSP approximation problem either can be solved in polynomial time or requires $\exp(\Omega(n))$ time. Thus, if the Unique Games Conjecture is true, then the complexity landscape of approximation problems for CSP's is much richer (at least in this sense) than the one for exact computation. This issue of "intermediate complexity" also raises some obstacles for certain approaches for *proving* the unique games conjecture, as well as suggests certain directions for doing so.

Subexponential complexity and gadget reductions.

The popular approach to proving hardness of approximation for CSP's can be called the "label cover + gadget paradigm".

1. Definition. A *label cover predicate* is a predicate $LC\colon \Sigma' \times \Sigma' \to \{0,1\}$ such that there is some $|\Sigma'|/|\Sigma''|$-to-one functions π_1, π_2 such that $P(x,y) = 1$ iff $\pi_1(x) = \pi_2(y)$.

One canonical way to get a label cover instance is the "clause vs clause" construction. Suppose I is an instance of some CSP, say $3LIN(2)$ for concreteness, we can define a new label cover instance I' where for every equation $x_i + x_j + x_k = b$ we have a variable $X_{i,j,k}$ in the alphabet $[4]$ which we identify with the set of satisfying assignments to this equation. For every two equations that share a variable $x_i + x_j + x_k = b$, $x_i + x_j + x_k = b'$ we put in the constraint that this variable is identical in both, which corresponds to checking that two projections of $[4]$ to $\{0,1\}$ agree with one another.

One can relate the two instances as follows:

2. Exercise. Assume that the original $3LIN(2)$ instance was d regular and had m constraints (i.e., every variable participated in the same number of constraints).

- Prove that if there is an assignment $x \in \{0,1\}^n$ for the original instance satisfying $1-\epsilon$ fraction of the constraints then there is an assignment $y \in [4]^m$ satisfying at least $1-2\epsilon$ fraction of the constraints of the label cover instance.
- Prove that if there is an assignment $y \in [4]^m$ satisfying at least $1-\delta$ fraction of the constraints of the label cover instance then there is an assignment $x \in \{0,1\}^n$ satisfying at least $1-2\delta$ fraction of the constraints of the original instance.

Given a label cover instance, a canonical way to reduce it to a CSP instance is the following:

3. Definition. Let $\mathcal{LC}$ be a family of label cover predicates mapping Σ' to $\{0,1\}$ and $\mathcal{P}$ be a family of predicates mapping Σ^k to $\{0,1\}$ for some Σ, k. A $(c',s') \mapsto (c,s)$ *gadget reduction* from $CSP(\mathcal{LC})$ to $CSP(\mathcal{P})$ consists of an *encoding map* $E\colon \Sigma' \to \Sigma^t$ and a *gadget* map that takes a predicate $LC \in \mathcal{LC}$ to a $CSP(\mathcal{P})$ instance $\mathcal{G}_{LC}$ on $2t$ variables, such that for every n-variable instance I' of $CSP(\mathcal{LC})$, if we let I be the nt-variate instance in for every constraint $LC(x_i,x_j)=1$ we place the $\mathcal{G}_{LC}$ instance on the $2t$ variables of the i-th and j-th blocks then it holds that:

- If $x' \in {\Sigma'}^n$ satisfies at least a c' fraction of the constraints of I', then $x = (E(x'_1),\ldots,E(x'_n)) \in \Sigma^{nt}$ satisfies at least c fraction of the constraints of I.
- If $x \in \Sigma^{tn}$ satisfies at least s fraction of the constraints of I, then there exists some $x \in {\Sigma'}^n$ that satisfies at least s' fraction of the constraints of I'.

Note that for every t, a gadget reduction maps an instance I' of n variables and m constraints into an instance I of nt variables and at most $m(2t)^k|\mathcal{P}|$ constraints, which for $t,k,|\mathcal{P}|$ constant means that the size of I is linear in the size of I'. Hence in particular one can show the following:

4. Exercise. Prove that if there is a $(c',s') \mapsto (c,s)$ reduction from $CSP(\mathcal{LC})$ to $CSP(\mathcal{P})$ with parameter t, then if there is a $T(n)$ time algorithm for the c vs s problem for $CSP(\mathcal{P})$ then there is a $T(Cn)$ time algorithm for the c' vs s' problem for $CSP(\mathcal{LC})$ where C is a constant depending only on $|\mathcal{P}|,k,t$.

5. Exercise. Prove that under the assumptions above, if I' is a $CSP(\mathcal{LC})$ instance that has a degree d pseudo-distribution μ' such that $\tilde{\mathbb{E}}_{\mu'(x)} \frac{1}{I'} \sum_{f\in I'} f(x) \geq c'$ then there exists a degree d/C pseudo-distribution μ such that $\tilde{\mathbb{E}}_{\mu(x)} \frac{1}{|I|} \sum_{f\in I} f(x) \geq c$ where C is a constant depending only on $|\mathcal{P}|,k,t$.

In particular this means that, if we assume that the original label cover instance could not be solved in time $\exp(n^{1-\epsilon})$ then the same holds for the resulting instance, and if we had an $\Omega(n)$ lower bound on the sos degree for the original instance then that lower bound carries over to the resulting instance. If the unique games conjecture is NP hard, then (assuming the ETH), the corresponding computational problem cannot be solved in time $\exp(n^{o(1)})$, while we know that it *can* be solved by an $\exp(n^{\epsilon})$ time algorithm for some small $\epsilon > 0$, and in fact by the degree n^{ϵ} sos program. This means that if want to establish the UGC via a gadget reduction, we'd better start with a label cover instance that has intermediate complexity, in both the time and the sos degree senses.

On label cover instances with intermediate complexity.

Some of the approaches to *prove* the unique games conjecture involve gadget reduction on top of certain label cover instances. Thus these approaches attempt to first prove (variants of) the "intermediate complexity conjecture" and then use that to derive the unique games conjecture. This raises the question of what properties of label cover instances could yield to them having intermediate complexity in certain approximation regimes. Assuming the unique games conjecture then having 1 to 1 projections (or even $O(1)$ to 1) is one such property, but is it easier to show this for other properties? Can we use sos to get some intuition on whether we expect this to be true?

The original way to manufacture label cover instances that are very hard to approximate was to start with a label cover instance over alphabet Σ with say 1 vs $1-\epsilon$ hardness (e.g., by starting from 3SAT) and then transform it into an instance over alphabet $\Sigma^{\otimes t}$ with gap, say, 1 vs $(1-\epsilon^{O(1)})^{\Omega(t)}$ using an amplification result such as the parallel repetition theorem (Raz [1995]). For example, if we started with the label cover corresponding to a 3XOR instance, we would get a label cover instance of alphabet $\Sigma = [4]^t$ where the projection maps Σ to an alphabet of size $2^t = \sqrt{|\Sigma|}$ and the hardness of approximation would be 1 vs $|\Sigma|^{-\epsilon}$ for some $\epsilon > 0$.

The parallel repetition theorem blows up an instance of size n to size $N = n^t$, and so one could a priori conjecture that the label cover problem with gap of $1-\epsilon$ vs ϵ has intermediate complexity, in the sense that it is NP hard but has an algorithm that runs in time $2^{N^{1/t}}$ or so where N is the instance size. However, this turns out to be false. Moshkovitz and Raz [2010] showed an alternative construction to get 1 vs $1-\epsilon$ hardness for label cover using only quasilinear blowup

from, say, 3SAT. In the sos world, this is even easier. If we consider the $3LIN(GF(2^t))$ problem, where the equations $x_i + x_j + x_k = b$ are taken in the field $GF(2^t)$ then the same proof as Grigoriev shows that a random instance (where one would not be able to satisfy more than $2^{-t} + o(1)$ fraction of the constraints) has a pseudo-distribution that pretends to be completely satisfiable. Here to the corresponding label cover would involve a projection of 2^{2t} to 2^t or an alphabet Σ to alphabet of size $\sqrt{|\Sigma|}$.

Another amplification construction is the "match/confuse games" of Feige and Kilian [1994] . In this construction, one takes a basic instance (such as the $3LIN(2)$ instance), and transforms it into a label cover where each variable corresponds to t tuples of constraints, and we put a constraint between pairs of tuples where $t - t'$ of the constraints are *identical* (for some $t' \ll t$) and the rest share a variable. One can show that this again amplifies the gap to something like 1 vs $2^{-\Omega(t')}$, but now the projections are "smoother" or closer to being 1 to 1 in the sense that they map the alphabet $\Sigma = [4]^t$ (in the case when the underlying CSP is $3LIN(2)$) to an alphabet of size $[4]^{t-t'}2^{t'} = |\Sigma|^{1-o(1)}$, since when two constraints are identical we require their projection to be identical too. Label cover instances of this type are sometimes known as *smooth label cover*. One can also think of a smooth label cover as a CSP over a family of k-ary predicates $\mathcal{P} \subseteq \{0,1\}^{\Sigma^k}$ that satisfy that for every $P \in \mathcal{P}$, $P^{-1}(1)$ is an error correcting code of the maximum distance of $k-1$ (i.e., every two distinct vectors agree on at most a single coordinate). It is an interesting open question whether there are sos instances of smooth label cover that require linear degree to obtain a 1 vs ϵ approximation, a positive answer can be interpreted as an obstacle to various approaches to proving the UGC.[1]

[1] The right notion of approximation here seems to be *strong soundness* where in the soundness case, not only every assignment x satisfies at most an ϵ fraction of the constraints, but even that the average fractional Hamming distance of the projection of x to a clause and $P^{-1}(1)$ is at least $1 - \epsilon$.

References

Sanjeev Arora, Boaz Barak, and David Steurer. Subexponential algorithms for unique games and related problems. *J. ACM*, 62(5): 42, 2015.

Uriel Feige and Joe Kilian. Two prover protocols: low error at affordable rates. In *STOC*, pages 172–183. ACM, 1994.

Dana Moshkovitz and Ran Raz. Sub-constant error probabilistically checkable proof of almost-linear size. *Computational Complexity*, 19 (3):367–422, 2010.

Ran Raz. A parallel repetition theorem. In *STOC*, pages 447–456. ACM, 1995.