U.C. Berkeley — CS278: Computational Complexity | Handout N17
Professor Luca Trevisan | 10/22/02
Scribe: Andrej Bogdanov

Notes for Lecture 17

Last time we showed that **PCP** theorem implies hardness of approximation for several **NP** complete problems. Today we introduce Håstad's **PCP** verifier, which yields stronger hardness results than the raw **PCP** theorem. In particular, we derive asymptotically optimal hardness of approximation results for Max E3LIN2 and Max 3SAT. We also show that Max 2SAT and Min Vertex Cover are hard to approximate within $21/22 + o(1)$ and $7/6 + o(1)$ of their optimum, respectively. Finally, we outline the connection between the **PCP** theorem and Håstad's verifier.

1 Håstad's verifier

To describe Håstad's verifier, we will need a more detailed parametrization of **PCP**. Say a language L is in $\mathbf{PCP}_{c,s}[r(n), q(n)]$ if there is a probabilistic polynomial time verifier V which, given input x and oracle access to a proof π, uses $r(|x|)$ random bits, reads $q(|x|)$ bits of π and such that

$$\begin{aligned} x \in L &\quad\Rightarrow\quad \exists \pi : \Pr[V^\pi(x) = 1] \geq c \\ x \notin L &\quad\Rightarrow\quad \forall \pi : \Pr[V^\pi(x) = 1] < s. \end{aligned}$$

The values c and s are called the *completeness* and the *soundness* of the verifier, respectively. If, in addition to these conditions, the verifier V makes its queries nonadaptively, we say $L \in \mathbf{naPCP}_{c,s}[r(n), q(n)]$.

Theorem 1 (Håstad) *For every constant $\epsilon > 0$, $\mathbf{NP} = \mathbf{naPCP}_{1-\epsilon,1/2+\epsilon}[O(\log n), 3]$. Moreover, the verifier for any $L \in \mathbf{NP}$ is simple: V^π generates three indices i, j, k and one bit b and accepts iff $\pi[i] \oplus \pi[j] \oplus \pi[k] = b$.*

We will sketch how to derive Håstad's theorem from the **PCP** theorem in the next few lectures. We remark that the theorem achieves the best possible values for its parameters in several ways. In particular, we have the following lower bounds:

1. Number of queries. (Folklore) For every $\epsilon > 0$, $\mathbf{P} = \mathbf{PCP}_{1,1-\epsilon}[O(\log n, 2)]$ and (Zwick) $\mathbf{P} = \mathbf{PCP}_{1-\epsilon,1-\epsilon^{1/3}}[O(\log n), 2]$.

2. Completeness vs. soundness. (Zwick) If $c/(1-s) \geq 2$, then $\mathbf{P} = \mathbf{PCP}_{c,1-s}[O(\log n), 3]$.

3. Completeness. (Zwick) $\mathbf{P} = \mathbf{naPCP}_{1,5/8}[O(\log n), 3]$.

However, perfect completeness can be bought by giving away soundness or nonadaptivity:

1. (Håstad) For all $\epsilon > 0$, $\mathbf{NP} = \mathbf{naPCP}_{1,3/4+\epsilon}[O(\log n), 3]$. It is not known whether this soundness constant is optimal for perfect completeness.

2. (Guruswami, Lewin, Sudan and Trevisan) For all $\epsilon > 0$, $\mathbf{NP} = \mathbf{PCP}_{1,1/2+\epsilon}[O(\log n), 3]$.

2 Applications to hardness of approximation

We show that Håstad's theorem implies the following hardness of approximation factors: $1/2 + o(1)$ for Max E3LIN2, $7/8 + o(1)$ for Max 3SAT, $21/22 + o(1)$ for Max 2SAT and $7/6 - o(1)$ for Min Vertex Cover. The proofs use approximation preserving reductions similar to the ones we saw last time.

Linear equations modulo 2. Max E3LIN2 is the problem of finding a maximum assignment to a linear system of equations modulo 2, with three distinct variables per equation. The value of an assignment is the number of equations satisfied by it. The problem of deciding whether a satisfying assignment of an E3LIN2 instance exists is **P**—it can be solved by Gaussian elimination. Obtaining a 1/2-approximation for Max E3LIN2 is trivial: Consider the two assignments $x_i = 0$ for all i and $x_i = 1$ for all i, and choose the one that satisfies more equations. This assignment will satisfy at least half of the equations.

Surprisingly, a 1/2 approximation is asymptotically optimal, unless $\mathbf{P} = \mathbf{NP}$. To show this, we construct a reduction ψ from an arbitrary **NP** language L to E3LIN2. Consider Håstad's verifier V for L: On input z, V^π generates a random string r of length $R = O(\log |z|)$, computes three indices $i(r), j(r), k(r)$ and a bit $b(r)$, queries π at positions i, j, k and accepts iff $\pi[i(r)] \oplus \pi[j(r)] \oplus \pi[k(r)] = b(r)$. The reduction ψ maps z to the system of equations

$$\{x_{i(r)} \oplus x_{j(r)} \oplus x_{k(r)} = b(r) \mid r \in \{0,1\}^R\}.$$

This is a polynomial-time reduction, as the number of equations is $m = |\{0,1\}^R| = |z|^{O(1)}$. If $z \in L$, there exists a proof π such that $V^\pi(x)$ accepts with probability $1 - \epsilon$ over the choice of r. In this case, the assignment $x_i = \pi[i]$ satisfies $(1-\epsilon)m$ of the equations $\psi(z)$. If, on the other hand, some assignment $x_i = a_i$ satisfies $(1/2+\epsilon)m$ of the equations $\psi(z)$, then the proof π such that $\pi[i] = a_i$ makes $V^\pi(z)$ accept with probability $(1/2 + \epsilon)m$, implying $z \in L$. It follows that E3LIN2 is inapproximable within a factor of $\frac{(1/2+\epsilon)m}{(1-\epsilon)m} \leq 1/2 + 2\epsilon$, unless $\mathbf{P} = \mathbf{NP}$.

Max 3SAT. The $7/8 + \epsilon$ inapproximability of Max 3SAT follows by reduction from Max E3LIN2. The reduction ψ from E3LIN2 to 3SAT maps each equation $x \oplus y \oplus z = b$ into a conjunction of four clauses:

$$\begin{aligned} x \oplus y \oplus z = 0 \quad &\Leftrightarrow \quad (\bar{x} \vee \bar{y} \vee \bar{z}) \wedge (\bar{x} \vee y \vee z) \wedge (x \vee \bar{y} \vee z) \wedge (x \vee y \vee \bar{z}) \\ x \oplus y \oplus z = 1 \quad &\Leftrightarrow \quad (x \vee y \vee z) \wedge (x \vee \bar{y} \vee \bar{z}) \wedge (\bar{x} \vee y \vee \bar{z}) \wedge (\bar{x} \vee \bar{y} \vee z). \end{aligned}$$

Consider an instance I of E3LIN and an assignment a for I. If an equation in I is satisfied by a, all four corresponding clauses in $\psi(I)$ are satisfied by a. Otherwise, exactly three of these four clauses are satisfied by a. Therefore, if $(1-\epsilon)m$ equations of I are satisfiable, then $4(1-\epsilon)m + 3\epsilon m$ clauses of $\psi(I)$ are satisfiable (by the same assignment). On the other hand, if at most $(1/2+\epsilon)m$ equations of I are satisfiable, then at most $4(1/2+\epsilon)m + 3(1/2-\epsilon)m$ clauses of $\psi(I)$ are satisfiable. The ratio between these two values is less than $7/8 + \epsilon$.

The factor 7/8 is optimal: All instances $\psi(I)$ have exactly three literals per clause, so a random assignment is expected to satisfy $7m/8$ clauses. Even for arbitrary 3SAT instances, which may contain two literal and one literal clauses, there is a 7/8 approximation algorithm based on semidefinite programming.

Max 2SAT. Unlike for Max 3SAT, there is no representation of an E3LIN2 equation by a short 2CNF formula. The reduction from E3LIN2 to 2SAT uses auxiliary variables: Given an equation $x_1 \oplus x_2 \oplus x_3 = b$, there exists a 12 clause 2CNF ϕ over variables $x_1, x_2, x_3, y_1, y_2, y_3, y_4$ such that:

1. Every assignment of ϕ satisfies at most 11 clauses.
2. If $a_1 \oplus a_2 \oplus a_3 = b$, there is an assignment that satisfies 11 clauses of ϕ such that $x_1 = a_1$, $x_2 = a_2$ and $x_3 = a_3$.
3. If $a_1 \oplus a_2 \oplus a_3 \neq b$, every assignment of ϕ such that $x_1 = a_1$, $x_2 = a_2$ and $x_3 = a_3$ satisfies at most 10 clauses of ϕ.

We reduce an E3LIN2 instance I to a 2CNF instance $\psi(I)$ by applying this construction to every equation of I. We instantiate distinct copies of y_1, y_2, y_3 and y_4 for every equation of I. If $(1-\epsilon)m$ equations of I are satisfiable, then $11(1-\epsilon)m$ of their corresponding clauses in $\psi(I)$ are satisfiable (by property 2). If fewer than $(1/2+\epsilon)m$ equations of I are satisfiable, then at most $11(1/2+\epsilon)m + 10(1/2-\epsilon)m$ equations of $\psi(I)$ are satisfiable (by properties 1 and 3). The ratio between these two values is $21/22 + O(\epsilon) \approx 0.954 + O(\epsilon)$. Using the same approach, one can show that Max Cut cannot be approximated within $16/17 + \epsilon$, unless $\mathbf{P} = \mathbf{NP}$.

The best known approximation algorithm for Max 2SAT is based on semidefinite programming and guarantees a ≈ 0.93 approximation.

Min Vertex Cover. A *vertex cover* for a graph G is a set $C \subseteq V(G)$ such that every edge of G has at least one endpoint in C. Equivalently, C is a vertex cover of G if $V(G) - C$ is an independent set. We show that the minimum vertex cover of G cannot be approximated within 7/6, unless $\mathbf{P} = \mathbf{NP}$. Let I be an instance of E3LIN2. For every equation $x_1 \oplus x_2 \oplus x_3 = b$ in I, we produce four vertices for $\psi(I)$. Each of these corresponds to an assignment of x_1, x_2 and x_3 that is consistent with the equation. We draw edges between all pairs of vertices that originated from the same equation, as well as between all pairs of vertices that correspond to inconsistent assignments. If $(1-\epsilon)m$ equations of I are satisfiable, then $\psi(I)$ contains an independent set of size $(1-\epsilon)m$: The independent set is made of vertices corresponding to satisfied clauses consistent with the optimal satisfying assignment of I. The complement of this independent set is a vertex cover of size at most $(3+\epsilon)m$. If at most $(1/2+\epsilon)m$ equations of I are satisfiable, then every set of $(1/2+\epsilon)m + 1$ vertices of $\psi(I)$ contains an edge. Therefore the largest independent set of $V(G)$ has at most $(1/2+\epsilon)m$ vertices, so its smallest vertex cover has at least $(3\frac{1}{2} - \epsilon)m$ vertices. The ratio between these two values is $7/6 - O(\epsilon)$. The best known inapproximability ratio for this problem is $10\sqrt{5} - 21 \approx 1.36$. A trivial algorithm that picks all endpoints of an arbitrary maximal matching achieves a factor 2 approximation. No algorithm is known that improves this factor asymptotically.

3 From the PCP theorem to Raz's verifier

The **PCP** theorem says that for some constant $\delta > 0$, $\mathbf{NP} = \mathbf{PCP}_{1,\delta}[O(\log n), O(1)]$. Assuming this theorem, we sketch the proof of Theorem 1. The main ingredient is a proof system for **NP** of Raz that makes only two queries, but over proofs represented in a large alphabet. To prove Theorem 1, the first step is to show that the proof system of Raz is equivalent to $\mathbf{PCP}_{1,\delta}[O(\log n), O(1)]$. The second step is to construct a simulation of the two queries performed by Raz's verifier by three queries over an appropriate binary encoding of Raz's proof. Here we cover the first step of the argument.

Raz's proof system. A *two-proof system* for L with soundness γ is a probabilistic polynomial-time verifier V with oracle access to two proofs, π_1 and π_2 over an alphabet of size $1/\gamma^{O(1)}$ such that on input x, V performs exactly one query into each of π_1 and π_2, and

$$\begin{aligned} x \in L &\Rightarrow \exists \pi_1, \pi_2 : \Pr[V^{\pi_1,\pi_2}(x) = 1] = 1 \\ x \notin L &\Rightarrow \forall \pi_1, \pi_2 : \Pr[V^{\pi_1,\pi_2}(x) = 1] < \gamma. \end{aligned}$$

We show that the **PCP** theorem implies the existence of a two-proof system with soundness δ for some $\delta > 0$ for any language in **NP**. First, using a canonical reduction, we can restate the **PCP** theorem in the following form: For every $L \in \mathbf{NP}$, there is a polynomial-time reduction ψ that maps instances of L into $3CNF$ formulas such that

$$\begin{aligned} x \in L &\Rightarrow \psi(x) \text{ is satisfiable} \\ x \notin L &\Rightarrow \text{At most } (1-\delta)m \text{ clauses of } \psi(x) \text{ are satisfiable.} \end{aligned}$$

We now construct a two-proof system for 3SAT. The proofs π_1 and π_2 are supposed to represent different encodings of an assignment for the input ϕ. For each variable x_i, the value $\pi_1[i] \in \{0, 1\}$ represents an assignment of x_i. For each clause C_j, the value $\pi_2[j] \in \{1, \ldots, 7\}$ represents one of the seven satisfying assignments for C_j. The verifier V^{π_1,π_2} chooses a random clause C_j, a random variable x_i that appears in C_j and accepts iff the assignments $\pi_1[i]$ and $\pi_2[j]$ are consistent.

If ϕ is satisfiable, the proofs π_1 and π_2 encoding a satisfying assignment of ϕ make V^{π_1,π_2} accept with probability 1. Now suppose that at most $(1-\delta)m$ clauses of ϕ are satisfiable. Then at least a δ fraction of entries in π_2 are inconsistent with π_1. If one of these entries in π_2 is selected by the verifier, at least one of its three literals is assigned inconsistently in π_1. Therefore an inconsistency is detected with probability $\delta/3$.

We will need a two-proof system for **NP** with much better soundness. We expect to increase soundness by "repeating" the two-proof protocol for 3SAT several times. A k-repetition two-proof system for 3SAT is specified as follows: The proof entries $\pi_1[i_1, \ldots, i_k] \in \{0, 1\}^k$ and $\pi_2[j_1, \ldots, j_k] \in \{1, \ldots, 7\}^k$ represent partial assignments to the k-tuples of variables $x_{i_1}, \ldots, x_{i_k}$ and k-tuples of clauses $C_{j_1}, \ldots, C_{j_k}$, respectively. The verifier V^{π_1,π_2} chooses a random set of k clauses, a random variable in each clause and accepts iff the corresponding assignments in π_1 and π_2 are consistent. We might expect that the k-repetition of a two-proof system with soundness $1-\gamma$ has soundness $(1-\gamma)^k$. Surprisingly, this is not the case; in the next lecture we will see a two-proof system with soundness $1/2$ such that its 2-repetition has soundness only 5/8. Fortunately, repetition does amplify soundness at an exponential rate:

Theorem 2 (Raz) *If V is a two-proof verifier with soundness γ, the k-repetition of V has soundness $2^{-\Omega(k)}$, where the constant in the $\Omega()$ notation depends only on γ and on the alphabet size of the proof system.*

Applying this theorem (which we won't prove), we conclude that for every $\gamma > 0$ there exists a 2-proof system for 3SAT with soundness γ with proofs over alphabets of size $1/\gamma^{O(1)}$. In the next three lectures we will show that for suitably small $\gamma = \gamma(\epsilon)$, this 2-proof system can be simulated by $\mathbf{naPCP}_{1-\epsilon,1/2+\epsilon}[O(\log n), 3]$.