

Probabilistically Checkable Proofs + Hardness of Approximation.

Note Title

5/21/2013

Many hard problems (especially NP-hard problems) are optimization problems. (minimization or maximization)

e.g. Smallest tour, Smallest VC, largest I.S.

"opt" = value of the optimal solution.

Is approximating opt good enough?

- there are lots of heuristics.
- we want an approximation algorithm that guarantees some approximation ratio r .

on every input x

$$\frac{\text{opt}}{r} \leq \text{answer} \leq \text{opt} \quad (\text{for maximization})$$

$$\text{opt} \leq \text{answer} \leq r \cdot \text{opt} \quad (\text{for minimization})$$

Vertex Cover: Input: undirected graph $G = (V, E)$.

What is the smallest subset of the vertices that touches every edge?

Vertex Cover (VC) is NP-complete.

Here's an approximation algorithm:

pick an edge (x, y) , add $x + y$ to VC

discard all edges incident to x or y .

Continue until no edges remain.

approx ratio for this simple algorithm is 2:

For each edge considered, we add 2 vertices.

Edges considered do not share any endpoints

opt \geq 1 vertex per edge considered.

opt \geq (our VC) / 2.

There is a diverse array of ratios achievable.

Vertex Cover: 2.

MAX-3-SAT: $8/7$ (find assignment that satisfies the most clauses).

Set Cover: $\log n$.

Knapsack: $(1+\epsilon)$ for any $\epsilon > 0$.

Clique: $n/\log^2 n$.

Best Case:

Poly-time approximation scheme (PTAS)

for every $\epsilon > 0$

$(1+\epsilon)$ ratio can be achieved

the approx algorithm runs in $\text{poly}(n)$ where $n = \text{input size}$
but may depend exponentially on $1/\epsilon$.


How do we explain the failure to improve some of these?

(How to explain the failure to find a poly-time alg for SAT?)

Is it NP-hard to improve an approximation ratio?

In order to prove a statement like this, we need a
"gap-producing" reduction from L_1 to L_2 :

Minimization:



r -gap producing reduction: f : poly-time computable.

$x \in L_1 \implies \text{opt}(f(x)) \leq k$.

$x \notin L_1 \implies \text{opt}(f(x)) > rk$.

The target problem is not a language.
 It is a promise problem. (Yes \cup No instances are not all things)

The algorithm is promised that the input comes from the set of Yes or No instances.


Purpose: r -approximation alg for L_2 distinguishes between $f(\text{yes}) + f(\text{no})$. This can be used to decide L_1 .

Note that if it is NP-hard to compute with the promise, then it's also hard to compute without it.

Gap producing reductions are difficult.

Gap preserving reductions are easier.

many typical reductions are not gap preserving.



For example the standard reduction of k -SAT to 3-SAT is gap preserving.

Max- k -SAT w/ gap ϵ reduces to max-3-SAT w/ gap ϵ' .

If it's hard to approximate max- k -SAT to within ϵ , then it's hard to approximate max-3-SAT to within ϵ' .

MAX SNP (Papadimitriou + Yannakakis) a set of problems reducible to each other in this way.

A PTAS for a MAX-SNP-complete problem gives a PTAS for all the problems in MAX-SNP.

Missing piece: first gap-producing reduction.

Consider Max-k-SAT w/ promise gap ϵ .

input: instance ϕ k-CNF.

YES: Some assignment satisfies all the clauses.

NO: no assignment satisfies more than $(1-\epsilon)$ of the clauses.

Let's look at a Proof System view of this problem.

Suppose there is a reduction from an NP-hard problem to Max-k-SAT w/ promise gap ϵ .

Then the following protocol will solve the NP-hard problem:

Given x compute reduction to k-SAT ϕ_x

The verifier expects that the proof is a satisfying assignment to ϕ_x .

Verifier picks random clause ("local test")

and checks that it is satisfied by the assignment.

$$x \in L \Rightarrow \Pr[V \text{ accepts}] = 1$$

$$x \notin L \Rightarrow \Pr[V \text{ accepts}] < 1 - \epsilon.$$

This can be repeated $O(1/\epsilon)$ times for error $< 1/2$.

Note: Prover commits to the whole proof.

Verifier only looks at part of the proof.

\Rightarrow Verifier does only local test. An ϵ fraction of these tests will notice the invalidity.

\hookrightarrow looking for $\epsilon = 1/poly(n)$

PCP Probabilistically Checkable Proof

Novel way of verifying a proof.

- pick a random local test

- query proof in specified k locations (chosen w/ random bits)

\rightarrow then accept if the proof passes the test.

PCP $[r(n), q(n)]$ set of languages w/ p.p.t. verifier V that has (r, q) restricted access to the proof.

V tosses $O(r(n))$ coins.

V access the proof in $O(q(n))$ locations (bits).

(Completeness) $x \in L \Rightarrow \exists$ proof s.t.

$$\Pr [V(x, \text{proof}) \text{ accepts}] = 1.$$

(Soundness) $x \notin L \Rightarrow \forall$ proof*

$$\Pr [V(x, \text{proof}) \text{ accepts}] \leq 1/2.$$

Observations

$$\text{PCP} [1, \text{poly}(n)] = \text{NP}$$

verifier sees the entire poly-size proof.

$$\text{PCP} [\log n, 1] \subseteq \text{NP}$$

\rightarrow run through all $\log n$ possibilities.

witness y is the set of all bits examined for each random choice.

Can calculate prob of acceptance explicitly.

PCP Theorem

$$\text{PCP} [\log n, 1] = \text{NP}$$

Any problem in NP has a $[\log n, 1]$ prob. checkable proof.

How does this relate to hardness of approximations?

Corollary: Max-k-SAT is hard to approximate to within some constant ϵ .

Pf of cor: Use PCP $[\log n, 1]$ protocol for some NP-hard problem.

variables in ϕ :

$x_1 \dots x_m$
bits of the proof.

\Rightarrow Enumerate all $2^{O(\log n)} = \text{poly}(n)$ sets of queries.

\Rightarrow Construct a k -CNF ϕ_i for verifier's test

on each. \Rightarrow k -CNF function on only k bits.

$$m \leq 2^{O(\log n)} \cdot O(1)$$

query $i \in \{0, 1\}^{O(\log n)} \rightarrow \phi_i$ accl rej as a fan of the k -bit response.

in ϕ_i
there are $\leq 2^k$ clauses. They are all satisfied
iff Verifier accepts.

"Yes" instance of VC \Rightarrow all clauses satisfied.

"No" instance of VC \Rightarrow every assignment fails to satisfy
 $\leq 1/2$ of the ϕ_i .

ϕ_i not satisfied $\Rightarrow \geq 1$ unsatisfied clause.

unsat clauses $\geq \underbrace{1/2(2^k)}_{\epsilon}$ fraction