TD: FPRAS and Inapproximability

Algorithmique et Programmation

Nicolas Schabanel
nicolas.schabanel@cnrs.fr

Pierre Senellart
pierre.senellart@ens.fr

December 8, 2016

The goal of this exercise session is to study two important aspects of approximation algorithms: first, fully polynomial randomized approximation schemes, and second gap-preserving reductions as a tool to obtain inapproximability results.

1 Counting the solutions to a DNF formula

In this exercise, we want to approximate the number of solutions to a Boolean formula in DNF (disjunctive normal form): Let $f = C_1 \vee \cdots \vee C_m$ be a DNF Boolean formula (an `OR` of `AND`-clauses as opposed to conjunctive normal form formulas, used in SAT, which considers an `AND` of `OR`-clauses) on n variables $x_1, \ldots, x_n$. Each clause C_j is of the form $C_j = l_1 \wedge \cdots \wedge l_{k_j}$, where each l_i is a litteral (i.e., a variable or its negation) and k_j denotes the number of literals in C_j. We assume that every clause is *satisfiable and non redundant* (i.e., contains each variable at most once, negated or not). The goal is to compute $\#f$, the number of truth assignments of the variables (x_i) that satisfy f.

We want to evaluate $\#f$ by sampling randomly the 2^n possible truth assignments of variables (x_i).

1a. We first consider uniform Monte-Carlo sampling: Draw a truth assignment τ uniformly at random in $\{0, 1\}^n$, and set $X = 2^n$ if τ satisfies f and $X = 0$ otherwise.

   Show that $\mathbb{E}[X] = \#f$. How many random bits does every draw of X use?

1b. However, X fails to estimate $\#f$ in polynomial time. Indeed, even if $\#f > 0$, the probability for X to be non-zero can be exponentially small. Thus, a polynomial number of draws of X is not enough to estimate $\#f$ within any constant factor. Indeed, assume *for the next two questions only* that n is even and that f consists of the single clause: $f = x_1 \wedge x_2 \wedge \cdots \wedge x_{n/2}$.

   1bα) What is the value of $\#f$? What is the probability for X to be non-zero for this f?

   1bβ) Let $X_1, \ldots, X_k$ be the values of k independent draws of X for this f. What is the probability that at least one of the draws is non-zero? Show that if k is polynomial in n, this probability is a $o(1)$, i.e., the probability goes to 0 when n goes to ∞.

1c. Let us now go back to the general case. Since uniform sampling does not work, we decide to use a biased variable which *samples only satisfying assignments*. Let $\mathcal{T}_j$ denote the set of truth assignments that satisfy clause C_j (which has k_j litterals). Remark that $\#f = \#(\bigcup_j \mathcal{T}_j)$, our goal is thus to count the number of elements in the union of the $\mathcal{T}_j$.

1cα) What is the size of $\mathcal{T}_j$?

1cβ) Let c_τ denote the number of clauses satisfied by truth assignment τ. Express $M := \sum_\tau c_\tau$ in terms of the $\#\mathcal{T}_j$'s.

1cγ) Design a (polynomial-time, in n and m) randomized algorithm that draws a truth assignment τ with probability c_τ/M, i.e. with probability proportional to c_τ. Explain in details your random sampling procedure. How many random bits does your algorithm use (as a $O(\cdot)$ of n and m)?

Hint: Remark that you can first draw a set $\mathcal{T}_j$ and then choose τ in $\mathcal{T}_j$.

1cδ) Consider the random variable Y defined as follows: draw a truth assignment τ with probability c_τ/M, and set $Y = M/c_\tau$.

Show that $\mathbb{E}[Y] = \#f$.

1cε) Let $\sigma^2(Z) = \mathbb{E}[(Z - \mathbb{E}[Z])^2]$ denote the *variance* of a random variable Z.

Show that $\sigma^2(Y) \leqslant ((m-1)\mathbb{E}[Y])^2$, where m is the number of clauses in f.

Hint: Show that Y belongs to interval $[M/m, M]$.

1cζ) Let Z_1 and Z_2 be two independent random non-negative variables and $Z = Z_1 + Z_2$. Show that $\mathbb{E}[Z_1 Z_2] = \mathbb{E}[Z_1]\mathbb{E}[Z_2]$ and that $\sigma^2(Z) = \sigma^2(Z_1) + \sigma^2(Z_2)$.

1cη) Recall Chebyshev's inequality which claims that for all random variable Z and all $a \in \mathbb{R}_+$,

$$\Pr\{|Z - \mathbb{E}[Z]| \geqslant a\} \leqslant \frac{\sigma^2(Z)}{a^2}.$$

Let $Y_1, \ldots, Y_k$ be the values of k independent draws of Y, and set $Z = \frac{Y_1 + \cdots + Y_k}{k}$. Show that for all $\varepsilon > 0$, and all $k \geqslant 4(m-1)^2/\varepsilon^2$,

$$\Pr\{|Z - \#f| < \varepsilon \cdot \#f\} \geqslant 3/4.$$

1cθ) Write down a polynomial time (in n, m and $1/\varepsilon$) randomized algorithm which outputs, for all $\varepsilon > 0$, a value θ such that $(1-\varepsilon)\,\#f \leqslant \theta \leqslant (1+\varepsilon)\,\#f$ with constant probability $> \frac{1}{2}$ (independent of n, m and ε).

1d. Can we use this algorithm to solve SAT? Justify your answer. Can we obtain from it a PTAS for Max-SAT? Justify your answer.

2 Gap preserving reductions: Vertex-Cover(d) and Steiner-Tree

We admit the following hardness of approximation result for Max-3SAT(k), the restriction of Max-3SAT to instances where each variable occurs at most in k clauses:

Theorem 1. *There exist an absolute constant $\varepsilon_0 > 0$ and a polynomial-time gap-introducing reduction from* SAT *to* Max-3SAT(29) *that transforms a Boolean formula φ to a Boolean formula ψ such that:*

- *if φ is satisfiable, then $\mathrm{OPT}_{\mathrm{SAT}}(\psi) = m$, and*
- *if φ is not satisfiable, then $\mathrm{OPT}_{\mathrm{SAT}}(\psi) < (1 - \varepsilon_0)m$,*

where m is the number of clauses in ψ.

2a. Let Vertex-Cover(d) be the restriction of unweighted Vertex Cover to instances in which each vertex has degree at most d. We will prove the following result:
Theorem 2. *There exists a polynomial-time gap preserving reduction from* Max-3SAT(29) *to* Vertex-Cover(30) *that transforms a Boolean formula ψ to a graph $G = (V, E)$ such that:*

- *if $\mathrm{OPT}_{\mathrm{SAT}}(\psi) = m$, then $\mathrm{OPT}_{\mathrm{VC}}(G) \leqslant \frac{2}{3}|V|$, and*
- *if $\mathrm{OPT}_{\mathrm{SAT}}(\psi) < (1 - \varepsilon_0)m$, then $\mathrm{OPT}_{\mathrm{VC}}(G) > (1 + \varepsilon_{\mathrm{VC}})\frac{2}{3}|V|$,*

where m is the number of clauses in ψ, ε_0 is as in Theorem 1, and $\varepsilon_{\mathrm{VC}} = \varepsilon_0/2$.

Assume w.l.o.g. that each clause of ψ has exactly 3 literals (this can be easily accomplished by repeating the literals within a clause, if necessary). We will use the following standard transformation: G has 3 vertices for each clause of ψ; each of these vertices is labeled with one literal of the clause (thus, $|V| = 3m$); G has two types of edges:

- for each clause, G has 3 edges connecting its 3 vertices, and
- for each $u, v \in V$, if the literals labeling u and v are negations of each other, then (u, v) is an edge in G.

Each vertex of G has two edges of the first type and at most 28 edges of the second type. Hence, G has degree at most 30. As an illustration, consider the formula $(x_1 \vee \overline{x}_2 \vee x_3) \wedge (\overline{x}_1 \vee x_2 \vee x_3)$. The graph produced by the reduction is given below.

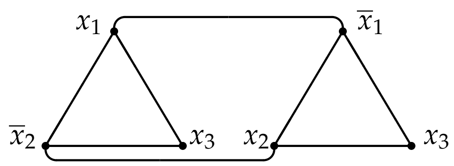

2aα) Show that the size of a maximum independent set in G is precisely $\mathrm{OPT}_{SAT}(\psi)$.

2aβ) Show that, in any graph, the complement of a maximum independent set is a minimum vertex cover. Conclude.

2b. The Steiner-Tree problem is as follows: given two disjoint sets of vertices R and S (*required* and *Steiner*) and a *cost* metric c on $R \sqcup S$, find a tree whose nodes $V = R \sqcup S'$ are all required vertices of R and a subset $S' \subseteq S$ of Steiner vertices, such that the sum of edge cost is minimum.

We now show the following result:

Theorem 3. *There exists a polynomial-time gap preserving reduction from* Vertex-Cover(30) *to the* Steiner-Tree *problem that transforms an instance* $G = (V, E)$ *of* Vertex-Cover(30) *to an instance* $H = (R, S, c)$ *of* Steiner-Tree *such that:*

- *if* $\mathrm{OPT}_{\mathrm{VC}}(G) \leqslant \frac{2}{3}|V|$, *then* $\mathrm{OPT}_{\mathrm{ST}}(H) \leqslant |R| + \frac{2}{3}|S| - 1$, *and*
- *if* $\mathrm{OPT}_{\mathrm{VC}}(G) > (1 + \varepsilon_{\mathrm{VC}})\frac{2}{3}|V|$, *then* $\mathrm{OPT}_{\mathrm{ST}}(H) > (1 + \varepsilon_{\mathrm{ST}})(|R| + \frac{2}{3}|S| - 1)$,

where $\varepsilon_{\mathrm{VC}}$ *is as in Theorem 2 and* $\varepsilon_{\mathrm{ST}} = 2\varepsilon_{\mathrm{VC}}/47$.

The instance $H = (R, S, c)$ will be such that G has a vertex cover of size c iff H has a Steiner tree of cost $|R| + c - 1$. H will have a required vertex r_e corresponding to each edge $e \in E$ and a Steiner vertex s_v corresponding to each vertex $v \in V$. The edge costs are as follows. An edge between a pair of Steiner vertices is of cost 1, and an edge between a pair of required vertices is of cost 2. An edge $\{r_e, s_v\}$ is of cost 1 if edge e is incident at vertex v in G, and it is of cost 2 otherwise.

2bα) Prove that the cost function is indeed a metric.

2bβ) Show that if G has a vertex cover of size c then H has a Steiner tree of cost $|R| + c - 1$.

2bγ) Assume that H has a Steiner tree T of cost $|R| + c - 1$ and that the edges of G are not all disjoint. Explain how to transform T into a Steiner tree of the same cost but using only edges of weight 1. Conclude that G has a vertex cover of size c.

2bδ) Conclude.