Journal of Computer and System Sciences ••• (••••) •••–•••

Contents lists available at ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

Time-approximation trade-offs for inapproximable problems ☆

Édouard Bonnet [a], Michael Lampis [b,*], Vangelis Th. Paschos [b]

[a] Department of Computer Science, Middlesex University, London, United Kingdom
[b] Université Paris-Dauphine, PSL Research University, CNRS, UMR 7243, LAMSADE, 75016, Paris, France

ARTICLE INFO

Article history:
Received 14 September 2016
Received in revised form 20 August 2017
Accepted 18 September 2017
Available online xxxx

Keywords:
Approximation algorithms
Exponential algorithms
Sub-exponential algorithms
Hardness of approximation

ABSTRACT

In this paper we focus on problems inapproximable in polynomial time and explore how quickly their approximability improves as the allowed running time is gradually increased from polynomial to (sub-)exponential. We tackle a number of problems: For Min Independent Dominating Set, Max Induced Path, Forest and Tree, for any $r(n)$, a simple, known scheme gives an approximation ratio of r in time roughly $r^{n/r}$. We show that, if this running time could be significantly improved, the ETH would fail. For Max Minimal Vertex Cover we give a $\sqrt{r}$-approximation in time $2^{n/r}$. We match this with a similarly tight result. We also give a $\log r$-approximation for Min ATSP in time $2^{n/r}$ and an r-approximation for Max Grundy Coloring in time $r^{n/r}$. Finally, we investigate the approximability of Min Set Cover, when measuring the running time as a function of the number of sets in the input.

1. Introduction

One of the central questions in combinatorial optimization is how to deal efficiently with NP-hard problems, with approximation algorithms being one of the most widely accepted approaches. Unfortunately, for many optimization problems, even approximation has turned out to be hard to achieve in polynomial time. This has naturally led to a more recent turn towards super-polynomial and sub-exponential time approximation algorithms. The goal of this paper is to contribute to a systematization of this line of research, while adding new positive and negative results for some well-known optimization problems.

For many of the most paradigmatic NP-hard optimization problems the best polynomial-time approximation algorithm is known (under standard assumptions) to be the trivial algorithm. In the super-polynomial time domain, these problems exhibit two distinct types of behavior. On the one hand, APX-complete problems, such as MAX-3SAT, have often been shown to display a "sharp jump" in their approximability. In other words, the only way to obtain any improvement in the approximation ratios for such problems is to accept a fully exponential running time, unless the Exponential Time Hypothesis (ETH) is false [27].

A second, more interesting, type of behavior is displayed on the other hand by problems which are traditionally thought to be "very inapproximable", such as Clique. For such problems it is sometimes possible to improve upon the (poor) approximation ratios achievable in polynomial time with algorithms running only in *sub-exponential* time. In this paper, we

☆ A conference version of this paper was presented in STACS 2016.

* Corresponding author.
E-mail addresses: edouar.bonnet@lamsade.dauphine.fr (É. Bonnet), michail.lampis@dauphine.fr (M. Lampis), paschos@lamsade.dauphine.fr (V.Th. Paschos).

https://doi.org/10.1016/j.jcss.2017.09.009
0022-0000/

concentrate on such "hard" problems and begin to sketch out the spectrum of trade-offs between time and approximation that can be achieved for them.

On the algorithmic side, the goal of this paper is to design *time-approximation trade-off schemes*. By this, we mean an algorithm which, when given an instance of size n and an (arbitrary) approximation ratio $r > 1$ as a target, produces an r-approximate solution in time $T(n, r)$. The question we want to answer is what is the best function $T(n, r)$, for each particular value of r. Put more abstractly, we want to sketch out, as accurately as possible, the Pareto curve that describes the best possible relation between worst-case approximation ratio and running time for each particular problem. For several of the problems we examine the best known trade-off algorithm is some simple variation of brute-force search in appropriately sized sets. For some others, we present trade-off schemes with much better performance, using ideas from exponential-time and parameterized algorithms, as well as polynomial-time approximation.

Are the trade-off schemes we present optimal? A naive way to answer this question could be to look at an extreme, already solved case: set r to a value that makes the running time polynomial and observe that the approximation ratios of our algorithms generally match (or come close to) the best-known polynomial-time approximation ratios. However, this observation does not alone imply satisfactorily the optimality of a trade-off scheme: it leaves open the possibility that much better performance can be achieved when r is restricted to a different range of values. Thus, the second, perhaps more interesting, direction of this paper is to provide lower bound results (almost) matching several of our algorithms *for any point in the trade-off curve*. For a number of problems, these results show that the known schemes are (essentially) the best possible algorithms, everywhere in the domain between polynomial and exponential running time. We stress that we obtain these much stronger *sub-exponential inapproximability* results relying only on standard, appropriately applied, PCP machinery, as well as the ETH.

1.1. Previous work

Moderately exponential and sub-exponential time approximation algorithms are relatively new topics, but most of the standard graph problems have already been considered in the trade-off setting of this paper. For Max Independent Set and Min Coloring an r-approximation in time $2^{O(n/r)}$ was given by Bourgeois et al. [6,4]. For Min Set Cover, a $\log r$-approximation in time $c^{n/r}$ and an r-approximation in time $c^{m/r}$, where n and m are the number of elements and sets respectively, were given by Cygan, Kowalik and Wykurz [10,5]. For Min Independent Dominating Set an r-approximation in $c^{n \log r/r}$ is given in [3]. An algorithm with similar performance is given for Bandwidth in [11] and for Capacitated Dominating Set in [12]. In all the results above, c denotes some appropriate constant.

On the hardness side, the direct inspiration of this paper is the recent work of Chalermsook, Laekhanukit and Nanongkai [7] where the following was proved.

Theorem 1. *[7] For all fixed $\varepsilon > 0$, for all sufficiently large $r = O(n^{1/2-\varepsilon})$, if there exists an r-approximation for* Max Independent Set *running in $2^{n^{1-\varepsilon}/r^{1+\varepsilon}}$ then there exists a randomized sub-exponential algorithm for* 3-SAT.

Theorem 1 essentially shows that the very simple approximation scheme of [6] is probably "optimal", up to an arbitrarily small constant in the second exponent, for a large range of values of r (not just for polynomial time). The hardness results we present in this paper follow the same spirit and in fact also rely on the technique of appropriately combining PCP machinery with the ETH, as was done in [7]. To the best of our knowledge, Max Independent Set and Max Induced Matching (for which similar results are given in [7]) are the only problems for which the trade-off curve has been so accurately bounded. The only other problem for which the optimality of a trade-off scheme has been investigated is Min Set Cover. For this problem the work of Moshkovitz [26] and Dinur and Steurer [13] showed that there is a constant $c > 0$ such that $\log r$-approximating Min Set Cover requires time $2^{(n/r)^c}$ (under ETH). It is not yet known if this constant c can be brought arbitrarily close to 1.

Let us note that trading execution time for solution's precision can also be investigated under complexity models other than the classical one, for instance under the parameterized complexity model. We mention here [15] where a formal framework for parameterized approximation starts to be drawn and some problems like Min Vertex Cover, or Steiner Tree are studied and the recent paper by [8] where the constant inapproximability of the parameterized Min Dominating Set is shown. Finally, let us mention the recent notion of α-approximate kernels [23] which can be seen as another way to deal with parameterized approximation.

1.2. Summary of results

In this paper we want to give upper and lower bound results for trade-off schemes that match as well as the algorithm of [6] and Theorem 1 do for Max Independent Set; we achieve this for the following problems.

Max Independent Set: Given a graph $G = (V, E)$, Max Independent Set consists of finding a set $S \subseteq V$ of maximum size such that for any $(u, v) \in S \times S$, $(u, v) \notin E$.

Min Set Cover: Given a ground set C of cardinality n and a system $\mathcal{S} = \{S_1, \ldots, S_m\} \subset 2^C$, Min Set Cover consists of determining a minimum size subsystem $\mathcal{S}'$ such that $\cup_{S \in \mathcal{S}'} S = C$.

Please cite this article in press as: É. Bonnet et al., Time-approximation trade-offs for inapproximable problems, J. Comput. Syst. Sci. (2018), https://doi.org/10.1016/j.jcss.2017.09.009