Journal of Computer and System Sciences 93 (2018) 30–40

Contents lists available at ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

Parameterized approximation via fidelity preserving transformations ☆

Michael R. Fellows [a], Ariel Kulik [b], Frances Rosamond [a], Hadas Shachnai [b,*,1]

[a] Department of Informatics, University of Bergen, Norway
[b] Department of Computer Science, Technion, Haifa 3200003, Israel

ARTICLE INFO

Article history:
Received 10 December 2013
Received in revised form 10 October 2017
Accepted 26 October 2017
Available online 13 November 2017

Keywords:
Fidelity preserving transformation
Fixed parameter tractability
Kernelization
Parameterized complexity
Approximation algorithms

ABSTRACT

We motivate and describe a new parameterized approximation paradigm which studies the interaction between approximation ratio and running time for any parametrization of a given optimization problem. As a key tool, we introduce the concept of an *α-shrinking transformation*, for $\alpha \geq 1$. Applying such transformation to a parameterized problem instance decreases the parameter value, while preserving the approximation ratio of α (or *α-fidelity*). Moving even beyond the approximation ratio, we call for a new type of *approximative kernelization race*. Our α-shrinking transformations can be used to obtain *approximative* kernels which are smaller than the best known for a given problem. The smaller "α-fidelity" kernels allow us to obtain an exact solution for the *reduced instance* more efficiently, while obtaining an approximate solution for the original instance. We show that such fidelity preserving transformations exist for several fundamental problems, including *Vertex Cover*, *d-Hitting Set*, *Connected Vertex Cover* and *Steiner Tree*.

1. Introduction

Given the common belief that most NP-hard problems cannot be solved, or even well-approximated, in polynomial time, it is natural for us to turn to a generalization of polynomial time, *fixed-parameter tractability*, to develop a paradigm of *parameterized approximation*.

Parameterized complexity approaches hard computational problems through a multivariate analysis of the running time. Instead of expressing the running time as a function of the input size n only, the running time is expressed as a function of n and k, where k is a well-defined parameter of the input instance (that can be an aggregate of several measurements). We say that a problem (with a particular parameter k) is *fixed-parameter tractable (FPT)* if it can be solved in time $f(k) \cdot \text{poly}(n)$, where f is an arbitrary function depending only on k. Thus we relax polynomial time by committing the exponential explosion to the parameter k. For further background on parameterized complexity we refer the reader to the textbooks [31,44,22], and the surveys in [23,25].

☆ A preliminary version of this paper appeared in Proc. of the 39th International Colloquium on Automata, Languages and Programming (ICALP), Warwick, July 2012.

* Corresponding author.
E-mail addresses: michael.fellows@uib.no (M.R. Fellows), ariel.kulik@gmail.com (A. Kulik), frances.rosamond@uib.no (F. Rosamond), hadas@cs.technion.ac.il (H. Shachnai).

[1] Work partially supported by the Loewengart Research Fund, and by the Ministry of Trade and Industry MAGNET program through the NEGEV Consortium (www.negev-initiative.org).

https://doi.org/10.1016/j.jcss.2017.11.001

Table 1
Approximations via α-fidelity shrinking: four examples.

Problem	Parameter	Kernel size	Running time	Best FPT algorithm
Vertex cover	solution size	$2(2-\alpha)k$	$O^*(1.273^{(2-\alpha)k})$	$O^*(1.273^k)$ [14]
Connected vertex cover	solution size	No $k^{O(1)}$	$O^*(2^{k(2-\alpha)})$	$O^*(2^k)$ [17]
3-Hitting set	solution size	$\frac{5(3-\alpha)^2}{4}k^2+\frac{3-\alpha}{2}k$	$O^*(2.076^{k(3-\alpha)/2})$	$O^*(2.076^k)$ [48]
Steiner tree	size of terminal set	No $k^{O(1)}$	$O^*(2^{(3-\alpha)k/2})$	$O^*(2^k)$ [6]

Extensive research since the beginning of the 70's has led to results exhibiting limits to the approximability of NP-hard problems. Comprehensive surveys of works on classical approximation algorithms can be found, e.g., in [34,47,49]. Formally, given a maximization (minimization) problem Π, we say that $\mathcal{A}$ is an r-approximation algorithm for some $r \geq 1$, if for any instance I of Π $\mathcal{A}$ yields a solution that satisfies $OPT(I)/\mathcal{A}(I) \leq r$ ($\mathcal{A}(I)/OPT(I) \leq r$), where $OPT(I)$ is the value of an optimal solution for I. Thus, for instance, *Maximum Independent Set* on a graph $G = (V, E)$, with $|V| = n$, is inapproximable within a ratio better than $n^{1-\epsilon}$, for some $\epsilon > 0$, unless P = NP [50]. Assuming the Unique Games Conjecture (UGC), the *Vertex Cover* problem cannot be approximated in polynomial time to within any constant factor better than 2, and the best constant-factor approximation for *d-Hitting Set* is d [37]. These results lead to the question that is at the heart of our study.

> "Given an optimization problem, Π, that is hard to approximate within factor ρ, for some $\rho > 1$: can we devise a family of α-approximation algorithms, A_α, such that A_ρ is polynomial, A_1 has the running time of the best FPT algorithm for Π, and A_α defines a continuous tradeoff between approximation ratios and running times?"

We show below that our parameterized approximation algorithms have approximation ratios better than the best possible polynomial-time approximation algorithms (under the common assumption that $P \neq NP$, and assuming that UGC holds). Our technique enables us to obtain any ratio $\alpha \in [1, \rho(\Pi)]$ for a given problem Π, where $\rho(\Pi)$ is the best known polynomial-time approximation ratio for the problem, and α is the approximation ratio achieved, depending on the desired running-time of the algorithm. In developing a general paradigm for parameterized approximation, we combine tools used in approximation algorithms with the framework of parameterized complexity. We move now to an overview of our results, after which will follow an in-depth presentation of α-shrinking transformations.

1.1. Our results

In this paper we describe a new parameterized approximation paradigm which relates parameterized complexity and polynomial-time approximation. While many earlier studies refer to parametrization by solution size, or, more generally, by the value of the objective function, our approximation approach can be applied for *any parametrization* of a given problem. We demonstrate our techniques with several fundamental problems, including *Vertex Cover*, *d-Hitting Set*, *Connected Vertex Cover*, and *Steiner Tree*.

We summarize our results in Table 1. For each of the studied problems, we specify the kernel size obtained by our algorithms (when applicable), as well as the running time of the algorithm as function of the approximation ratio, $\alpha \geq 1$, and the best known running time of an exact FPT algorithm for the problem.[2]

One of the most important practical techniques in parameterized complexity is *kernelization*. Here one takes a problem instance specified by $(x, k) \in \Sigma^* \times \mathbb{N}$ and produces, typically in polynomial time, a small instance of the problem: (x', k') such that (x, k) is a 'yes' instance iff (x', k') is a 'yes' instance; moreover, $|x'| \leq g(k)$ and usually $k' \leq k$. This technique is widely used in practice, as it usually relies on a number of easily implementable reduction rules.

There are two types of *races* in parameterized complexity research: the race for the smallest possible function $f(k)$ in the running time of an exact algorithm, and the race for the smallest possible function $g(k)$ to bound the size of a kernel. These races are well-established, and the current leader boards are exhibited on the FPT community Wiki [45]. Our parameterized approximation paradigm gives rise to a new kind of race, *approximative kernelization*.

As a key tool in our study, we introduce (in Section 2) the concept of *α-shrinking transformation*, for $\alpha \geq 1$. We shall see that applying such transformation to a parameterized problem instance decreases the parameter value, while preserving α-fidelity in the approximation ratio. We show that α-shrinking transformations can be used also as a tool for *approximative kernelization*, to obtain kernels which are smaller than the best known for a given problem. Thus, we define the notion of *α-fidelity kernel*, for $\alpha \geq 1$, where the special case of $\alpha = 1$ is a standard kernel. Such smaller α-fidelity kernels will allow us to solve exactly problem instances more efficiently, while obtaining an α-approximate solution for the original instance.

[2] The notation O^* hides factors polynomial in the input size.