



Research Repository

On the Smoothed Complexity of Combinatorial Local Search

Accepted for publication in Mathematics of Operations Research.

Research Repository link: <https://repository.essex.ac.uk/41246/>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the published version if you wish to cite this paper.

<https://pubsonline.informs.org/journal/moor>

On the Smoothed Complexity of Combinatorial Local Search

Yiannis Giannakopoulos*

Alexander Grosz[†]

Themistoklis Melissourgos[‡]

3 July, 2025

Abstract

We propose a unifying framework for smoothed analysis of combinatorial local optimization problems, and show how a diverse selection of problems within the complexity class PLS can be cast within this model. This abstraction allows us to identify key structural properties, and corresponding parameters, that determine the smoothed running time of local search dynamics. We formalize this via a black-box tool that provides concrete bounds on the expected maximum number of steps needed until local search reaches an exact local optimum. This bound is particularly strong, in the sense that it holds for any starting feasible solution, any choice of pivoting rule, and does not rely on the choice of specific noise distributions that are applied on the input, but it is parameterized by just a global upper bound ϕ on the probability density. The power of this tool can be demonstrated by instantiating it for various PLS-hard problems of interest to derive efficient smoothed running times (as a function of ϕ and the input size).

Most notably, we focus on the important local optimization problem of finding pure Nash equilibria in Congestion Games, that has not been studied before from a smoothed analysis perspective. Specifically, we propose novel smoothed analysis models for general and Network Congestion Games, under various representations, including explicit, step-function, and polynomial resource latencies. We study PLS-hard instances of these problems and show that their standard local search algorithms run in polynomial smoothed time.

Finally, we present further applications of our framework to a wide range of additional combinatorial problems, including local Max-Cut in weighted graphs, the Travelling Salesman problem (TSP) under the k -opt local heuristic, and finding pure equilibria in Network Coordination Games.

*University of Glasgow. Email: yiannis.giannakopoulos@glasgow.ac.uk

[†]Technical University of Munich. Email: alexander.grosz@tum.de. This author was supported by the Alexander von Humboldt Foundation with funds from the German Federal Ministry of Education and Research (BMBF).

[‡]University of Essex. Email: themistoklis.melissourgos@essex.ac.uk

Contents

1	Introduction	1
1.1	Our Results and Outline	2
2	Smoothed Combinatorial Local Optimization	3
2.1	Smoothed Combinatorial Local Search	6
3	Smoothed Analysis of Local Search	6
3.1	Limitations of our main theorem	9
4	Smoothed Analysis for Congestion Games	10
4.1	Nash Equilibria as Combinatorial Local Optimization Problems	11
4.2	Restrained Congestion Games	13
4.3	Network Congestion Games	16
5	Further Applications	19
5.1	TSP with the k -OPT Heuristic	19
5.2	Maximum Satisfiability	22
5.3	Maximum Cut	24
5.3.1	Equivalent Problems	26
5.4	Network Coordination Games	27
5.5	Weighted Set Problems	29
5.5.1	Weighted 3D-Matching	30
5.5.2	Exact Cover by 3-Sets and Set-Cover	31
5.5.3	Hitting Set	32
6	Conclusion and Open Problems	33
A	Omitted Proofs	34
A.1	Proof of Lemma 3.3	34
A.2	Proof of Lemma 3.4	34
B	PLS-Hardness of Restrained Congestion Games	35
C	PLS-Hardness of Compact Network Congestion Games	37
D	PLS-Hardness of LOCALMAX-k-CUT	41
E	PLS-Hardness of Hitting Set with Bounded Element Occurrence	42
F	Maximum Constraint Assignment	43

1 Introduction

Local search heuristics are some of the most prominent, and widely used in practice, algorithms for solving computationally hard, combinatorial problems [AL97; MKA07]. Their appeal stems not only from their simplicity and theoretical elegance, but also from the fact that, for many applications, they seem to perform remarkably well both in terms of their running time and the quality of the solutions they produce.

Theoreticians have long tried to rigorously study the performance of local search, but also explain its prevalence in practice. Johnson, Papadimitriou, and Yannakakis [JPY88] introduced the class PLS to capture the complexity of local optimization problems; since then, many important such problems have been shown to be PLS-complete, implying that they most likely cannot be solved (exactly) in polynomial time. This hardness applies not only to local search algorithms, but to arbitrary local optimization methods. For local search, in particular, PLS-hardness (under tight reductions) implies the provable existence of instances leading to exponentially slow convergence [Yan97]. Examples include the Travelling Salesman problem under the k -OPT heuristic [Kre89; CKT99] (TSP/ k -OPT in the following), Local Maximum Cut on weighted graphs [SY91] (LOCALMAXCUT), and the problem of finding pure Nash equilibria in Congestion Games [FPT04] (PNE-CONGESTION).

On the other hand, Orlin, Punnen, and Schulz [OPS04] designed a local-search-based polynomial-time scheme for efficiently computing *approximately* locally-optimal solutions for general combinatorial optimization problems (with linear objectives). Although this result provides concrete justification for the practical tractability of local optimization, there are still many important aspects that call for further investigation. First, if one requires exponential accuracy, their FPTAS still cannot provide polynomial running times. Secondly, approximate solutions do not always make sense for all local optimization problems; there are problems in PLS that are inherently exact, and they are not derived by simply considering the local version of some “master” global optimization problem. A notable example is PNE-CONGESTION [Rou16]. Finally, we would like to be able to argue about the more general family of “vanilla” local search, and to ideally get positive results that do not depend on additional details and specific choices of pivoting rules. Addressing these points is a key objective of the present paper.

Smoothed analysis was introduced by Spielman and Teng [ST04] as a more realistic alternative to traditional worst-case analysis, where now the adversarially selected input is submitted to small random shocks of its numerical parameters, *before* being presented to the algorithm; the running time is then measured *in expectation* with respect to these perturbations. Under this model, Spielman and Teng were able to show that Simplex, the archetypical method for solving linear programs, is guaranteed to terminate in polynomial time (under a shadow pivoting rule) – as opposed to its exponential complexity under worst-case analysis. This remarkable result established smoothed analysis as a canonical framework for studying the performance of algorithms beyond the worst-case (see [Rou21] for an overview of this field).

In particular, smoothed analysis has been applied successfully to important local search algorithms, providing thus a theoretical basis for the justification of their good performance in practice; these include, e.g., the k -OPT heuristic for TSP and the FLIP heuristic for LOCALMAXCUT. (A more detailed exposition of related work on this front is deferred to the following sections of this paper, where each of our local optimization problems of interest is explicitly studied; see, namely, Sections 4 and 5.) A common characteristic of this prior work, though, is that running-time analysis is usually tailored specifically to the local optimization problem at hand. This naturally creates the need for technically heavy derivations, from which it is generally not clear how to pin down the core properties of the underlying local-search structure that allow for the efficient smoothed complexity. Furthermore, as a result, it is often not easy to immediately generalize these results to capture interesting extensions, e.g., argue about the *asymmetric* version of TSP, or go from TSP/2-OPT to TSP/ k -OPT. Finally, this lack of sufficient abstraction is one of the reasons that smoothed analysis has not been yet considered at all for

prominent PLS-hard problems, including, e.g., PNE-CONGESTION. Dealing with this set of challenges is another driving force behind our paper.

1.1 Our Results and Outline

We start by proposing an abstract model for smoothed analysis of combinatorial local optimization (CLO) in [Section 2](#). Our family of CLO problems includes problems in PLS that have an arbitrary combinatorial neighbourhood structure and linear objective functions; essentially, our model generalizes that of [\[OPS04\]](#) beyond binary configurations. In [Section 2.1](#), we add our smoothness layer that introduces probabilistic noise (independently) to the cost parameters of the CLO problem. No further assumptions are made on the distributions of the perturbed costs; their densities are only parameterized by a global upper bound of ϕ . This is a standard model (employed, e.g., in [\[BV06; RV07; ERV16; ER17\]](#)) that makes positive smoothed-analysis results even stronger, and extends the seminal model of Gaussian perturbations from [\[ST04\]](#).

[Section 3](#) contains our key technical result for deriving upper bounds on the expected (under smoothness) number of local-search steps, until an *exact* local optimum is reached. Our black-box tool ([Theorem 3.2](#)) can be readily applied to an abstract CLO problem, once its underlying neighbourhood structure is appropriately captured; this is formalized through the notion of *separability* (see [Definition 3.1](#)), quantified by a collection of three parameters that are critical for the upper bound given by [Theorem 3.2](#). We note here that like many problem-specific state-of-the-art results, our bounds are robust against the specific choice of a starting point for the local search dynamics, as well as the pivoting rule utilized at every step to transition to an improving neighbour. In other words, our main black-box tool establishes bounds for the entire *family* of local search heuristics of a local optimization problem. At a technical level, our proof works by lower-bounding the probability that *all* steps of the local-search sequence improve the objective sufficiently, thereby applying a standard approach to a general, application independent abstract problem.

In [Section 4](#) we demonstrate the applicability of our general framework by instantiating it for PNE-CONGESTION, a prominent PLS-complete problem which has not been studied before from a smoothed analysis perspective. First, we propose smoothed analysis formulations for various representations of interest for the problem, namely explicit, step-function, or polynomial resource latencies.¹ Next, after formally establishing how PNE-CONGESTION is indeed a CLO problem in [Section 4.1](#), we identify a natural parameterization of the problem that we call *B-restrained games* ([Section 4.2](#)), where B is an upper bound on the number of resources that can be changed during a single-player deviation (see [Definition 4.1](#)). Interestingly enough, the case of constant B is still rich enough to encode the full PLS-hardness of PNE-CONGESTION ([Appendix B](#)), while at the same time it can be shown (see the proof of [Theorem 4.2](#)) to be appropriately separable in order to immediately provide polynomial smoothed running time bounds via our black-box tool developed in [Section 3](#). Similarly, in [Section 4.3](#) we also study a special class of *network* congestion games, which we term (A, B) -compact ([Definition 4.5](#)); we establish polynomial smoothed complexity for various families of instances, including the one where A is polynomial and B is constant ([Corollary 4.7](#)), which we pair with a complementing PLS-hardness proof ([Appendix C](#)).

In the remainder of our paper, we apply our high-level framework to various other local optimization problems of interest, by first formally establishing that they can be viewed as CLO problems and then identifying the proper separability structures that can be plugged into our black-box tool to provide good smoothed bounds for local search. For some of these problems we rederive existing bounds from the literature, in a much simpler way; for others we improve or extend existing results; and for others,

¹ Very recently, and after our paper had appeared online, Giannakopoulos [\[Gia24\]](#) built upon our model to establish the existence of a fully polynomial-time approximation scheme (FPTAS) for computing *approximate* pure Nash equilibria in smoothed congestion games.

we are the first to propose and perform smoothed analysis on them. An important underlying feature of all our proofs is the unifying and straightforward way in which our bounds are derived across the entire spectrum of local search problems studied in this paper.

In particular, in [Section 5.1](#) we study the smoothed running time of TSP/ k -OPT. We improve on this problem in two directions: we consider the most general version, namely *asymmetric* TSP [\[STV20\]](#), and without any metric constraints, while we study the k -OPT heuristic for all $k \geq 2$. In [Theorem 5.3](#), we show that the k -OPT heuristic has polynomial smoothed running time for every fixed k . To the best of our knowledge, our work is the first to extend the results of [\[ERV16\]](#) to the case beyond $k = 2$.

In [Section 5.2](#) we study MAXSAT/ k -FLIP [\[YI99; Sze11\]](#) in which we search for a truth assignment whose flipping of at most k variables cannot improve its sum of clause-weights. We show that, for any fixed k , when each variable appears constantly many times in the clauses, the smoothed running time is polynomial, while the problem remains PLS-hard [\[Kre89\]](#). To the best of our knowledge, prior to this work, only the case of $k = 1$ had been studied under smoothness in [\[Bra18\]](#), whose analysis is limited to the case of complete Boolean formulas; that is, given a fixed clause size, the formula must contain *all* possible clauses of that size.

Then, in [Section 5.3](#) we apply our black-box tool to LOCALMAXCUT, and show that when the input graph has logarithmic degree, the smoothed running time is polynomial. Our $\mathcal{O}(2^\Delta |V| |E|^2)$ running time, where Δ is the maximum degree of the input graph $G = (V, E)$, improves the bound given in [\[ET11\]](#), while at the same time extends it beyond Gaussian perturbations. Furthermore, for the generalization of that problem, namely LOCALMAX- k -CUT, for $k > 2$, we provide a bound that does *not* depend on k . This results in a smoothed polynomial running time for graphs with logarithmic degree. We note that Bibak, Carlson, and Chandrasekaran [\[BCC21\]](#) had shown smoothed super-polynomial running time for graphs with arbitrary degree.

Network Coordination Games are another problem on which we apply our unified smoothed analysis framework in [Section 5.4](#). Our results establish polynomial smoothed complexity for graphs with either (a) constant degree and arbitrary number of strategies per player, or (b) logarithmic degree and a constant number of strategies per player. Prior work [\[BKM20\]](#) provides *quasipolynomial* bounds under the assumption of a constant number of strategies.

A selection of weighted set problems has been put into the context of PLS by [\[DS10a\]](#) using canonical k -change neighbourhoods, where already for small constant values of k , PLS-hardness was established. In [Section 5.5](#), we define a smoothed problem formulation and prove polynomial smoothed complexity for Weighted 3D-Matching, Exact Cover by 3-Sets and Set-Cover. For the Hitting Set problem, we identify a property that allows us to parameterize the smoothed running time, and prove accompanying PLS-hardness for a smoothly efficient subclass of instances. Finally, we show that the related problem Maximum Constraint Assignment has polynomial smoothed complexity when its parameters are appropriately bounded.

2 Smoothed Combinatorial Local Optimization

In this section we formalize our model and fix the necessary notation.

We denote by \mathbb{N} and \mathbb{R} , the natural and real numbers, respectively. We also denote $\mathbb{N}^* := \mathbb{N} \setminus \{0\}$ and $[k] := \{1, 2, \dots, k\}$, $\llbracket k \rrbracket := \{0\} \cup [k]$ for $k \in \mathbb{N}$. We will use boldface notation for vectors, $\mathbf{s} = (s_1, s_2, \dots, s_n) \in \mathbb{R}^n$. For an index $i \in [n]$, we use \mathbf{s}_{-i} to denote the $(n - 1)$ -dimensional vector that results from an n -dimensional vector \mathbf{s} if we remove its i -component; in that way, we can express \mathbf{s} as (s_i, \mathbf{s}_{-i}) in order to easily denote deviations in the i -th component. More generally, for a set of indices $I \subseteq [n]$, \mathbf{s}_I denotes the $|I|$ -dimensional vector that we get if we keep only the components of \mathbf{s} whose indices are in I . For finite sets I, \mathcal{I} , we say that \mathcal{I} is a *cover* of I if $I \subseteq \bigcup_{I' \in \mathcal{I}} I'$. All logarithms appearing in our paper are of base 2.

Combinatorial local optimization (CLO) An instance of a *combinatorial local optimization (CLO)* problem is composed of:

- A set of feasible *configurations* $S \subseteq \llbracket M \rrbracket^\nu \times \{0, 1\}^{\bar{\nu}}$, where $M, \nu \in \mathbb{N}^*$, $\bar{\nu} \in \mathbb{N}$. A configuration s can be expressed as $s = (s^\bullet, s^\circ)$, where $s^\bullet = (s_1, \dots, s_\nu) \in \llbracket M \rrbracket^\nu$ is called its *cost part* and $s^\circ = (s_{\nu+1}, \dots, s_{\nu+\bar{\nu}}) \in \{0, 1\}^{\bar{\nu}}$ its *non-cost part*.
- A vector of *costs* $c = (c_1, c_2, \dots, c_\nu) \in [-1, 1]^\nu$. We call the c_i *cost coefficients*.²
- A *neighbourhood* function $N : S \longrightarrow 2^S$. For $s \in S$, any configuration $s' \in N(s)$ will be called a *neighbour* of s .

For the special case of $M = 1$ we will call our problem *binary*.

The *cost* of a configuration s (with respect to a fixed cost vector c) is given by

$$C(s) := c \cdot s^\bullet = \sum_{i=1}^{\nu} c_i s_i.$$

A configuration s is said to be a *local optimum* (*minimum*³) if there are no neighbours with better costs; formally,

$$C(s) \leq C(s') \quad \text{for all } s' \in N(s). \quad (1)$$

For every CLO problem we can define its *configuration* (or *neighbourhood*) *graph*, made up by all edges pointing from configurations to their neighbours. Formally, it is the directed graph $G = (V, E)$ with node set $V = S$ and edges $E = \{(s, s') \mid s \in S, s' \in N(s)\}$. Observe that the configuration graph does not depend on the costs c , but only on the combinatorial structure of the problem. Taking the cost vector c into consideration, we can now restrict the configuration graph edges to those that correspond to locally improving moves, i.e. take $E' = \{(s, s') \in E \mid C(s') < C(s)\}$. The resulting acyclic subgraph $G' = (V, E')$ is called the *transition graph*.

The transition graph provides an elegant and concise interpretation of local optimization: finding a local optimum of a combinatorial (local optimization) problem translates to finding a sink of its transition graph. This task is exactly the object of interest of our paper: we study (the running time of) algorithms that find local optima of such problems.

Complexity If the configuration set and the neighbourhoods $N(\cdot)$ were given explicitly in the input of a local optimization algorithm, then finding a locally optimal solution would have been a computationally trivial task: we could afford (in polynomial time) to exhaustively go over all nodes of the transition graph, until we find a sink. Notice that the existence of a sink is guaranteed, by the fact that S is finite. However, most problems of interest (including all the problems that we study here) have exponential-size configuration graphs, with respect to the dimension $\nu + \bar{\nu}$ of the problem. Therefore, S and N are usually instead described *implicitly*, via some succinct representation (that is polynomial on $\nu, \bar{\nu}$ and M). Then, the computational complexity of our algorithms is naturally measured as a function of the critical parameters $\nu, \bar{\nu}$ and M (and the bit representation of the costs c).

² Note that the restriction of costs to the range of $[-1, 1]$ has been made to facilitate the translation to smoothed CLO problems below (see Section 2.1). In fact, due to the linearity of our objective, scaling c by any positive constant does not change the structure of the problem, and we will implicitly use this fact when formulating other problems as CLO problems.

³ Here we choose to write the local optimality condition with respect to minimization problems. All the theory developed in this paper applies immediately to local maximization problems as well: just flip the inequality in (1) or, equivalently, negate the cost vector c .

Polynomial local optimization (PLS) In this paper we want to study problems contained in PLS, the “canonical” complexity class for local optimization problems, introduced by Johnson, Papadimitriou, and Yannakakis [JPY88]. Therefore, without further mention, from now on we will assume that our CLO problems further satisfy the following properties:

- An *initial* configuration $s_0 \in S$ can be computed in polynomial time (on the size of the input).
- There exists a polynomial-time algorithm that, given as input any configuration s , decides whether s is a local optimum and, if not, returns an improving neighbour $s' \in N(s)$ with $C(s') < C(s)$. Such an algorithm is called *pivoting rule*.

It is important to clarify that the polynomial time algorithms of the bullets above are formally *not* part of the description of the CLO problem, neither are they required to be specified in the input of a local optimization algorithm. The *existence* of these algorithms is merely a requirement for the *membership* of the problem in the class PLS (similarly to the existence of short certificates and an efficient verifier for membership in NP). As a result, PLS can be interpreted as the class of problems that correspond to looking for a sink in an (implicitly given, possibly exponentially large) directed acyclic graph⁴, where at least one node can be found efficiently and, for any given node, at least one neighbour (if such exists) can be found efficiently [DP11; FGMS20].

This interpretation naturally gives rise to the *standard local search* heuristic: start from an arbitrary initial configuration, and at every iteration perform an (arbitrary⁵) locally improving move until no such move exists any more. For a fixed pivoting rule, this corresponds to traversing a single path of the transition graph. Due to the definition of PLS membership above, the running time of such a process is thus determined by the *number* of these local search iterations. In the worst case, this amounts to bounding the length of the *longest path* in the transition graph.

It is important to emphasize that, regardless of being a very natural heuristic, standard local search is definitely not the only method for finding local optima: as a matter of fact, we know that there exist local optimization problems that are efficiently solvable via more involved, “centralized” methods (e.g., by using linear programming), but for which standard local search would provably require exponential time (see, e.g., [ARV08]): more precisely, there exist nodes in their transition graphs from which *all* paths have exponential length [SY91].

To assert the intractability of a problem, the common argumentation is via complexity theory techniques that prove the conditional inexistence of polynomial time algorithms. If a problem is PLS-hard, then unless $PLS = FP$, there is no algorithm that solves it in polynomial time under traditional worst-case analysis. However, since we are studying the performance of local search heuristics, we are in fact interested in specifically *their* running time. Schäffer and Yannakakis [SY91] introduced the notion of a stronger reduction among problems in PLS, named *tight PLS-reduction*, which preserves key structural properties of the initial instance to the target-problem’s instance. An important implication of an initial-problem P having a tight PLS-reduction to target-problem Q is the following: for any instance I of P , a path in the transition graph of the corresponding reduced instance J of Q induces a path of I of length no larger than its own. In particular, if for an instance I there exists a starting configuration from which all paths to solutions are exponentially long, this also applies to the reduced instance J . We show that all the problems we study have this property, either by referring to explicit bad instances for the problem at hand, or through (chains of) tight reductions, which we construct if they haven’t been established before.

For a more thorough treatment of the complexity of local optimization and the class PLS, the interested reader is referred to the excellent monograph of Yannakakis [Yan97].

⁴ This is the transition graph mentioned earlier.

⁵ Thus, standard local search is essentially a *family* of algorithms; different pivoting rules can give rise to different local search algorithms.

2.1 Smoothed Combinatorial Local Search

Under traditional, worst-case algorithmic analysis, the running time of an algorithm for a CLO problem would be evaluated against an adversarially selected cost vector $\mathbf{c} = (c_1, \dots, c_v)$. Instead, our goal here is to propose a systematic *smoothed analysis* [ST09; Rög08; Rou21] framework for local optimization. Therefore, for the remainder of this paper we assume that \mathbf{c} is not fixed, but drawn randomly from a product distribution. More specifically, each cost coefficient c_i is drawn independently from a continuous probability distribution with density $f_i : [-1, 1] \rightarrow [0, \phi]$. We note that $\phi \geq 1/2$ is necessary for a probability distribution. These distributions can be adversarially selected, but their realizations c_i are not; the running time is then computed *in expectation* with respect to the random cost vector \mathbf{c} . An efficient algorithm runs in time polynomial in the combinatorial-structure parameters v , \bar{v} and M of the problem, and in the smoothness parameter ϕ . We will sometimes refer to this model as *smoothed CLO*, if we want to give particular emphasis to the fact that we are performing a smoothed running time analysis (as opposed to worst-case analysis).

In that sense, smoothed analysis can be seen as interpolating between two extremes: an average-case analysis setting where all c_i 's are drawn i.i.d. from a uniform distribution over $[-1, 1]$, derived for $\phi = 1/2$; and traditional worst-case analysis that can be derived in the limit, via $\phi \rightarrow \infty$, as distributions f_i approximate adversarial, single-point-mass instances.

In this paper we focus on smoothed analysis for standard local search, and so our quantity of interest will be the expected number of improving moves (for any set of adversarially given input distributions) until a local optimum is reached; that is, the expected length of the longest path in the transition graph. This allows us to also avoid some delicate representation issues that are typical of smoothed analysis, and which have to do with how the realizations c_i of *continuous* distributions (which therefore produce irrational numbers almost surely) can be handled as inputs to a Turing-machine-based computational model. For a careful discussion about this topic we point to the papers of Beier and Vöcking [BV06] or Röglin and Vöcking [RV07]. In a nutshell, for our purposes it is safe to think of the polynomial-time improving-local-move oracle (from the PLS definition) as having access to real-number arithmetic.

3 Smoothed Analysis of Local Search

In this section we present our first main result, which is a black-box tool for upper-bounding the number of improving moves of standard local search, under smoothed analysis. To achieve this, we first highlight an appropriate underlying structure of CLO problems and identify key parameters that characterize it (see Definition 3.1). Then, our bounds for standard local search are directly expressed as a function of these parameters (see Theorem 3.2). As we will demonstrate in Sections 4.2, 4.3 and 5, for various CLO problems of interest these parameters are well-behaved enough to result in polynomial smoothed running times.

We now introduce some terminology and notation that will be necessary for stating our main result in Theorem 3.2. Fix an instance of a CLO problem with cost coordinates $[v]$ and configurations S (see Section 2), and let $G = (S, E)$ be its neighbourhood graph. A *covering* $(\mathcal{E}, \mathcal{I})$ of this instance consists of a cover \mathcal{E} of the edges of its neighbourhood graph G , and a cover \mathcal{I} of its cost coordinates. That is, $\mathcal{E} \subseteq 2^E$ and $\mathcal{I} \subseteq 2^{[v]}$ such that $E = \bigcup_{E' \in \mathcal{E}} E'$ and $[v] = \bigcup_{I \in \mathcal{I}} I$. We call \mathcal{E} the *transition cover* and \mathcal{I} the *coordinate cover*, and they contain *transition clusters* and *coordinate clusters* respectively. Recall that, under our previous discussion (see Section 2), \mathcal{E} can be simply interpreted as covering all potential configuration transitions that can be made by standard local search, clustered appropriately in different groups $T \in \mathcal{E}$. For an arbitrary set of such transitions $T \subseteq E$, we also define its *core* to be the set of coordinates affected by any of the transitions:

$$\text{core}(T) := \{i \in [v] \mid s_i^\bullet \neq s_i'^\bullet \text{ for some } (s, s') \in T\},$$

and its *diversity* with respect to a given set of coordinates $I \subseteq [v]$ to be the number of different configuration changes when projected to I :

$$\delta_I(T) := |\text{range}_I(T)|, \quad \text{where } \text{range}_I(T) := \{s_I^\bullet - s_I'^\bullet \mid (s, s') \in T\}.$$

Definition 3.1 (Separable instances). An instance of a CLO problem will be called (λ, β, μ) -separable if it has a covering $(\mathcal{E}, \mathcal{I})$ with $|\mathcal{E}| \leq \lambda$, such that every transition cluster $T \in \mathcal{E}$ has:

- (a) a core that can be covered by using β many coordinate clusters from the cover \mathcal{I} ; formally, there exists an $\mathcal{I}_T \subseteq \mathcal{I}$ with $|\mathcal{I}_T| \leq \beta$ such that $\text{core}(T) \subseteq \bigcup_{I \in \mathcal{I}_T} I$, and
- (b) diversity at most μ with respect to all coordinate clusters; formally, $\max_{I \in \mathcal{I}} \delta_I(T) \leq \mu$.

Theorem 3.2. For every (λ, β, μ) -separable smoothed combinatorial local optimization instance, standard local search terminates after at most

$$3 \cdot \mu^\beta \lambda \cdot v^2 M \log(M+1) \cdot \phi$$

many steps (in expectation).

For the proof of [Theorem 3.2](#) we will need the following technical lemmas. Their proofs can be found in [Appendix A](#). We would like to highlight the fact that the structural quantities (λ, β, μ) appear as the expression $\mu^\beta \lambda$ in the statement of the theorem and are the deciding quantities for the running time complexity of the problem, as evidenced in the applications further below. As the choice of the covering is not unique (and often even allows for more than just one *natural* choice), this exact expression helps to explain the quantitative interaction between the properties of the covering. Finally, we believe it is interesting to note here that the *specific* proof of [Theorem 3.2](#) that we provide below, imposes a *tight* dependence on the μ^β parameter on the theorem's bound, due to the union-bound technique we deploy. In other words, in order to improve the dependence on β in the exponent (even asymptotically), a different proof technique would be required.

Lemma 3.3. Let $I \subseteq [v]$ be a set of cost coordinates and $\mathcal{J} \subseteq 2^{[v]}$ be a cover of I , i.e. $I \subseteq \bigcup_{J \in \mathcal{J}} J$. Then, for every set of transitions $T \subseteq E$,

$$\delta_I(T) \leq \prod_{J \in \mathcal{J}} \delta_J(T).$$

Lemma 3.4. Let $\phi > 0$ and let $X = (X_1, X_2, \dots, X_m)$ be a random real vector, where each component X_i is drawn independently from a continuous distribution with density $f_i : \mathbb{R} \rightarrow [0, \phi]$. Then, for every nonzero vector $\xi \in \mathbb{R}^m$ and every $\varepsilon \geq 0$,

$$\text{Prob}[0 \leq \xi \cdot X \leq \varepsilon] \leq \min\left(\frac{1}{\|\xi\|_\infty}, \frac{\sqrt{2}}{\|\xi\|_2}\right) \cdot \varepsilon \phi, \quad (2)$$

where $\|\cdot\|_2$ and $\|\cdot\|_\infty$ denote the Euclidean and maximum norms, respectively. For the special⁶ case where ξ is a (nonzero) integer vector we get

$$\text{Prob}[0 \leq \xi \cdot X \leq \varepsilon] \leq \varepsilon \phi. \quad (3)$$

⁶ In this paper we will be actually making use of [Lemma 3.4](#) only via its weaker bound (3), rather than the stronger form (2). This is sufficient for our purposes because, as it turns out, this has an asymptotically negligible effect on our bounds. However, we still choose to state (and prove) [Lemma 3.4](#) in its full generality, since we expect it to be of potential independent interest for future extensions, especially if one considers more involved structures, or non-integral configurations.

Proof of Theorem 3.2. Fix a (λ, β, μ) -separable smoothed CLO instance, with neighbourhood graph $G = (S, E)$, and let $(\mathcal{E}, \mathcal{I})$ be a covering satisfying the conditions of Definition 3.1.

We introduce the following notation:

$$\Delta(\mathbf{s}, \mathbf{s}') := \begin{cases} C(\mathbf{s}) - C(\mathbf{s}'), & \text{if this is positive,} \\ \infty, & \text{otherwise,} \end{cases}$$

for all $(\mathbf{s}, \mathbf{s}') \in E$, and

$$\Delta := \min_{(\mathbf{s}, \mathbf{s}') \in E} \Delta(\mathbf{s}, \mathbf{s}').$$

Notice that these are random variables, depending on the realizations of the cost vector \mathbf{c} .

Our first goal is to give a bound on the probability that there exists a local move that improves the cost only by (at most) $\varepsilon > 0$, as a function of this improvement bound ε . If this quantity is sufficiently small, then with high probability, standard local search will achieve improvements more than ε at *every* step, thus resulting in faster convergence. To upper bound this probability $\text{Prob}[\Delta \leq \varepsilon]$, we first use a union bound over the cover \mathcal{E} of all transitions $(\mathbf{s}, \mathbf{s}') \in E$ to get

$$\text{Prob}[\Delta \leq \varepsilon] = \text{Prob}\left[\bigcup_{(\mathbf{s}, \mathbf{s}') \in E} [\Delta(\mathbf{s}, \mathbf{s}') \leq \varepsilon]\right] \leq \sum_{T \in \mathcal{E}} \text{Prob}\left[\bigcup_{(\mathbf{s}, \mathbf{s}') \in T} [\Delta(\mathbf{s}, \mathbf{s}') \leq \varepsilon]\right]. \quad (4)$$

Next, for a fixed transition cluster $T \in \mathcal{E}$ we can express the inner union of events in (4) as

$$\begin{aligned} \bigcup_{(\mathbf{s}, \mathbf{s}') \in T} [\Delta(\mathbf{s}, \mathbf{s}') \leq \varepsilon] &= \bigcup_{(\mathbf{s}, \mathbf{s}') \in T} \left[0 < \mathbf{c}_{\text{core}(T)} \cdot (\mathbf{s}_{\text{core}(T)} - \mathbf{s}'_{\text{core}(T)}) \leq \varepsilon\right] \\ &= \bigcup_{\mathbf{x} \in \text{range}_{\text{core}(T)}(T)} [0 < \mathbf{c}_{\text{core}(T)} \cdot \mathbf{x} \leq \varepsilon]. \end{aligned} \quad (5)$$

By the separability assumption of our CLO instance (see Definition 3.1), for any cluster $T \in \mathcal{E}$ there exists a subset \mathcal{I}_T of \mathcal{I} with $|\mathcal{I}_T| \leq \beta$ that covers $\text{core}(T)$ such that, additionally, $\max_{J \in \mathcal{I}_T} \delta_J(T) \leq \max_{I \in \mathcal{I}} \delta_I(T) \leq \mu$. So, from Lemma 3.3 we can deduce that

$$|\text{range}_{\text{core}(T)}(T)| = \delta_{\text{core}(T)}(T) \leq \prod_{J \in \mathcal{I}_T} \delta_J(T) \leq \mu^\beta.$$

Furthermore, observe that each $\mathbf{x} \in \text{range}_{\text{core}(T)}(T)$ is a nonzero integral vector. Thus, applying (3) of Lemma 3.4 we can derive that, for a fixed $\mathbf{x} \in \text{range}_{\text{core}(T)}(T)$, we have

$$\text{Prob}[0 < \mathbf{c}_{\text{core}(T)} \cdot \mathbf{x} \leq \varepsilon] \leq \phi \varepsilon.$$

Therefore, using again a union bound, this time on event (5), we can see that

$$\text{Prob}\left[\bigcup_{(\mathbf{s}, \mathbf{s}') \in T} [\Delta(\mathbf{s}, \mathbf{s}') \leq \varepsilon]\right] \leq \sum_{\mathbf{x} \in \text{range}_{\text{core}(T)}(T)} \text{Prob}[0 < \mathbf{c}_{\text{core}(T)} \cdot \mathbf{x} \leq \varepsilon] \leq \mu^\beta \phi \varepsilon.$$

Plugging this into (4), and using again the separability of our instance, we can finally bound our desired probability

$$\text{Prob}[\Delta \leq \varepsilon] \leq \lambda \mu^\beta \varepsilon \phi. \quad (6)$$

Now we continue with the second part of the proof, in which we utilize the probability bound in (6) to derive a concrete bound on the expected number of iterations of standard local search. Let \mathcal{T} denote

the random variable of that maximum length among all paths in the transition graph of our instance. Then, our goal is to bound $\mathbb{E} [\mathcal{T}]$.

Recall that the number of different possible cost-part configurations is trivially upper-bounded by $(M+1)^v$. Also, at every step of standard local search, the configuration cost is *strictly* decreasing. Thus, $\mathcal{T} \leq (M+1)^v$. Furthermore, since the range of configuration costs is within $[-Mv, Mv]$, and the minimum cost improvement of every iteration is Δ , we also know that $\mathcal{T} \leq (2Mv/\Delta)$. Using these, we get

$$\begin{aligned}
\mathbb{E} [\mathcal{T}] &= \sum_{t=1}^{(M+1)^v} \text{Prob} [\mathcal{T} \geq t] \\
&\leq \sum_{t=1}^{(M+1)^v} \text{Prob} \left[\Delta \leq \frac{2Mv}{t} \right] \\
&\leq 2\lambda\mu^\beta \phi Mv \sum_{t=1}^{(M+1)^v} \frac{1}{t}, && \text{due to (6),} \\
&\leq 2\lambda\mu^\beta \phi Mv \frac{3}{2} \log(M+1)^v, && \text{since } 1 + \frac{1}{2} + \dots + \frac{1}{n} \leq \frac{3}{2} \log n \quad \forall n \geq 2, \\
&= 3\mu^\beta \lambda \phi v^2 M \log(M+1).
\end{aligned}$$

□

3.1 Limitations of our main theorem

The method we used in the proof of our main result, [Theorem 3.2](#), is based on bounding the improvement of a single step in the procedure. This approach has been successful in the derivation of a variety of prior work in smoothed analysis as we outline in [Section 4](#) and [Section 5](#), and it turns out that our generalized proof captures all the aforementioned works' results.

However, our model and proof do not capture some other recent results on smoothed analysis. We identified the following two different approaches in the literature:

1. For some problems, the results rely on bounds on the improvement of *sequences* of steps, rather than improvements of a single step. Notable examples are problems on graphs, and include [\[ABPW17; ER17; BCC21; CGVYZ20\]](#) on LOCALMAXCUT and [\[ERV16\]](#) on TSP/2-OPT. These results provide smoothed (quasi)polynomial bounds for the case of general graphs, whereas our method gives such bounds only for the case of graphs with at most (poly)logarithmic degree.
2. Other results, not only use the above sequence-improvement analysis, but also heavily rely on the rich random interaction of the model at hand. For example, the smoothed polynomial time results of [\[ABPW17; BCC21; CGVYZ20\]](#) on LOCALMAXCUT and [\[BKM20\]](#) on NETCOORDNASH rely on the fact that the graph is complete in the underlying problem.

To tackle these cases with a generalized methodology, it seems that significantly new techniques are required.

Left as open research directions for future work, the following approaches could potentially capture the aforementioned cases.

1. Instead of working on the base problem A , we could try to define a problem B for which a single step essentially realizes a sequence of polynomially many single steps in A . Thus, the additional properties could potentially carry over to the single step of problem B . A generalized notion of separability could now apply to the single step of B to achieve smoothed analysis results.

2. We conjecture that a modification of our framework, which would take advantage of the high amount of randomness introduced in every single step of the procedure, would be capable to establish improved smoothed complexity results for some problems, such as LOCALMAXCUT or NETCOORDNASH on graphs with n nodes and degree at least $n - \Theta(\log n)$. The modified framework would have to involve a complementary notion of separability that enables it to unify this proof approach.

4 Smoothed Analysis for Congestion Games

Congestion games are composed of finite nonempty sets of *players* $\mathcal{N} = [n]$ and *resources* \mathcal{R} . Each player $i \in \mathcal{N}$ has a *strategy set* $\Sigma_i \subseteq 2^{\mathcal{R}}$ and each resource $r \in \mathcal{R}$ has a *cost* (or *latency*) *function* $\kappa_r : [n] \rightarrow \mathbb{R}_{\geq 0}$. Each (pure) *strategy profile* (or *outcome*) $\sigma = (\sigma_1, \dots, \sigma_n) \in \Sigma := \Sigma_1 \times \dots \times \Sigma_n$ induces a *load* on each resource r , equal to the number of players that use it:

$$\ell_r(\sigma) := |\{i \in \mathcal{N} \mid r \in \sigma_i\}|.$$

Then, the *cost of player i* is the total cost she experiences from all resources that she is using:

$$C_i(\sigma) := \sum_{r \in \sigma_i} \kappa_r(\ell_r(\sigma)).$$

An (exact) pure *Nash equilibrium* (PNE) is an outcome σ^* from which no player can improve her cost by unilaterally deviating. Formally, for any player $i \in \mathcal{N}$ and any deviation $\sigma'_i \in \Sigma_i$:

$$C_i(\sigma^*) \leq C_i(\sigma'_i, \sigma_{-i}^*).$$

Thus, if a strategy profile σ is *not* a PNE, there has to exist a player i and a deviation $\sigma'_i \in \Sigma_i$ that reduces her cost, i.e.

$$C_i(\sigma'_i, \sigma_{-i}) < C_i(\sigma).$$

Such a strategy σ'_i is then called a *better-response* (of player i with respect to the profile σ).

By the seminal work of Rosenthal [Ros73] we know that function $\Phi : \Sigma \rightarrow \mathbb{R}_{\geq 0}$ defined by

$$\Phi(\sigma) = \sum_{r \in \mathcal{R}} \sum_{\ell=1}^{\ell_r(\sigma)} \kappa_r(\ell) \tag{7}$$

and commonly referred to as *Rosenthal's potential*, has the property that

$$C_i(\sigma'_i, \sigma_{-i}) - C_i(\sigma) = \Phi(\sigma'_i, \sigma_{-i}) - \Phi(\sigma) \quad \forall \sigma \in \Sigma \quad \forall i \in \mathcal{N} \quad \forall \sigma'_i \in \Sigma_i.$$

In other words, function Φ is an (exact) *potential* [MS96] of the corresponding congestion game. This implies that PNE of a congestion game correspond *exactly* to the set of *local* minimizers of its Rosenthal's potential (7), meaning that $\Phi(\sigma^*) \leq \Phi(\sigma'_i, \sigma_{-i}^*)$ for any player i and any deviation $\sigma'_i \in \Sigma_i$. This also immediately establishes the existence of PNE in *all* congestion games, since the potential function $\Phi(\sigma)$ can only take finitely many different values.

Depending on the type and representation of the cost functions, different classes of congestion games can arise. Below, we describe three prominent ones which we will focus on in this paper:

- *General*. The cost functions are given explicitly as a list of nonnegative values, one for each possible load on the resource; $(\kappa_r(1), \kappa_r(2), \dots, \kappa_r(n))$. Notice how in this model we do not impose a monotonicity constraint; this is deliberate, to maintain full generality. If one wants to focus on nondecreasing cost functions (as is many times the case in the literature), then the step-function representation (see below) can easily be used instead.

- *Polynomials (of degree d)*. The cost functions are polynomials of maximum degree $d \in \mathbb{N}$ with non-negative coefficients. More specifically, the cost functions are given implicitly by the coefficients $\{\alpha_{r,j}\}_{r \in \mathcal{R}, j \in \llbracket d_r \rrbracket} \subseteq \mathbb{R}_{\geq 0}$, where $d_r \leq d$, via

$$\kappa_r(\ell) = \sum_{j=0}^{d_r} \alpha_{r,j} \ell^j \quad \text{for all } \ell \in [n]. \quad (8)$$

- *Step functions (with d break-points)*. The cost functions are nondecreasing, piecewise constant, given by pairs of break-points and value-increases. More specifically, for each resource $r \in \mathcal{R}$ there is a list of break-points $\mathcal{B}_r \subseteq [n]$ and associated jumps $\{\alpha_{r,j}\}_{j \in \mathcal{B}_r}$. We denote the number of break-points of a resource r by $d_r := |\mathcal{B}_r|$, and we set $d := \max_{r \in \mathcal{R}} d_r$. Then, the cost functions are given via

$$\kappa_r(\ell) = \sum_{j \in \mathcal{B}_r \cap [\ell]} \alpha_{r,j} \quad \text{for all } \ell \in [n]. \quad (9)$$

To the best of our knowledge, no smoothed analysis model has been established so far for congestion games. In this paper, we propose and study the following perturbation semantics for the aforementioned classes. Note that as is commonly the case in smoothed analysis, the impact of randomness in the given input parameters depends crucially on the way they are used to compute the cost in the model. While the notion of *general congestion games* does typically capture all possible latency functions in the deterministic setting, the results for smoothed analysis do not directly translate between the models.⁷

- *General congestion games*. The costs $\kappa_r(\ell)$ are independently distributed according to densities $f_{r,\ell} : [0, 1] \rightarrow [0, \phi]$, for all $r \in \mathcal{R}$ and $\ell \in [n]$.
- *Polynomial games*. The coefficients $\alpha_{r,j}$ are independently distributed according to densities $f_{r,j} : [0, 1] \rightarrow [0, \phi]$, for all $r \in \mathcal{R}$ and $j \in \llbracket d_r \rrbracket$.
- *Step-function games*. The jump increases $\alpha_{r,j}$ are independently distributed according to densities $f_{r,j} : [0, 1] \rightarrow [0, \phi]$, for all $r \in \mathcal{R}$ and $j \in \mathcal{B}_r$. Notice, however, that the break-points \mathcal{B}_r themselves are not subjected to any noise, and they are assumed to be fixed (and adversarially selected).

4.1 Nash Equilibria as Combinatorial Local Optimization Problems

We now show how PNE-CONGESTION, the problem of finding a pure Nash equilibrium in congestion games, can actually be interpreted as a combinatorial local optimization problem (with respect to our definitions in [Section 2.1](#)), for any of the cost models described above. For all models, their randomness semantics translate directly to the respective randomness of cost coefficients in the smoothed CLO problem.

- *General congestion games*. By (7) we know that PNE correspond exactly to local minimizers of the potential function $\Phi(\sigma) = \sum_{r \in \mathcal{R}} \sum_{j=1}^{\ell_r(\sigma)} \kappa_r(j)$. Therefore, finding a PNE of a general congestion game can be viewed as a binary ($M = 1$) CLO problem, with cost dimension $v = |\mathcal{R}| n$ (the cost coordinates are given by $\mathcal{R} \times [n]$), where each strategy profile σ is mapped to a cost configuration $\mathbf{s}^\bullet = (s_{r,j})_{r \in \mathcal{R}, j \in [n]}$ given by the indicator functions:

$$s_{r,j} = \mathbb{1} [j \leq \ell_r(\sigma)]. \quad (10)$$

⁷ See e.g. the discussion in [Section 5.3](#) for LOCALMAXCUT, where the current state-of-the-art results for the complete graph are better than for arbitrary graphs.

The CLO cost coefficients are given by $c_{r,j} = \kappa_r(j)$.

Furthermore, we want to establish a one-to-one correspondence between configurations \mathbf{s} of the CLO problem and strategy profiles σ of our congestion games where, in particular, the neighbours of \mathbf{s} are exactly the configurations corresponding to single-player deviations $\{(\sigma'_i, \sigma_{-i})\}_{i \in \mathcal{N}, \sigma'_i \in \Sigma_i}$. In that way, better-responses of the congestion game would correspond exactly to local improvements in the CLO formulation. However, this cannot be achieved by using just the cost part defined above by (10), since in that case, different strategy profiles may end up being mapped to the same cost part configuration (when they induce the same resource loads). To overcome this technical pitfall, we also maintain a non-cost part \mathbf{s}° , which keeps track of the actual strategies of the players: this can be easily achieved, with only an additional polynomial size burden.⁸

Finally, notice that the neighbourhoods of the CLO problem we created can be explicitly listed and efficiently searched for a better (smaller cost) value: they have a maximum size of $n \cdot \max_{i \in \mathcal{N}} |\Sigma_i|$, which is polynomial in the description of the original congestion game.

- *Polynomial games.* Using (8), Rosenthal's potential (7) can now be written as

$$\Phi(\sigma) = \sum_{r \in \mathcal{R}} \sum_{\ell=1}^{\ell_r(\sigma)} \sum_{j=0}^{d_r} \alpha_{r,j} \ell^j = \sum_{r \in \mathcal{R}} \sum_{j=0}^{d_r} \alpha_{r,j} \sum_{\ell=1}^{\ell_r(\sigma)} \ell^j = \sum_{r \in \mathcal{R}} \sum_{j=0}^{d_r} \alpha_{r,j} \mathfrak{S}_j(\ell_r(\sigma)),$$

where

$$\mathfrak{S}_j(\ell) := \sum_{k=1}^{\ell} k^j \leq \ell^{j+1} \leq \ell^{d+1} \leq n^{d+1}$$

for any $\ell \in \mathbb{N}$. This induces a CLO problem with parameters

$$v = \sum_{r \in \mathcal{R}} (d_r + 1) \leq |\mathcal{R}| (d + 1) \quad \text{and} \quad M = n^{d+1},$$

with each strategy profile σ corresponding to a cost configuration $\mathbf{s}^\bullet = (s_{r,j})_{r \in \mathcal{R}, j \in [d_r]}$ given by:

$$s_{r,j} = \mathfrak{S}_j(\ell_r(\sigma)).$$

The costs coefficients are given by $c_{r,j} = \alpha_{r,j}$. Notice that again all neighbourhoods are efficiently searchable since they have a polynomial maximum size of $n \cdot \max_{i \in \mathcal{N}} |\Sigma_i|$.

- *Step-function games.* Using (9), Rosenthal's potential (7) can now be written as

$$\begin{aligned} \Phi(\sigma) &= \sum_{r \in \mathcal{R}} \sum_{\ell=1}^{\ell_r(\sigma)} \sum_{j \in \mathcal{B}_r \cap [\ell]} \alpha_{r,j} = \sum_{r \in \mathcal{R}} \sum_{j \in \mathcal{B}_r \cap [\ell_r(\sigma)]} (\ell_r(\sigma) - j + 1) \alpha_{r,j} \\ &= \sum_{r \in \mathcal{R}} \sum_{j \in \mathcal{B}_r} \max(0, \ell_r(\sigma) - j + 1) \alpha_{r,j}. \end{aligned}$$

This induces a CLO problem with dimension

$$v = \sum_{r \in \mathcal{R}} |\mathcal{B}_r| = \sum_{r \in \mathcal{R}} d_r \leq |\mathcal{R}| d \quad \text{and} \quad M = n,$$

⁸ E.g. we can choose $\bar{v} = |\mathcal{N}|$ and let $s_i^\circ \in [|\Sigma_i|]$ be the index of strategy σ_i deployed by player i in profile σ . To simplify our exposition, in the remaining congestion game classes studied below, we will avoid explicitly discussing these non-cost parts; they are identical to the general congestion game model.

with each strategy profile σ corresponding to configuration $\mathbf{s}^\bullet = (s_{r,j})_{r \in \mathcal{R}, j \in [d_r]}$ given by

$$s_{r,j} = \max(0, \ell_r(\sigma) - \mathcal{B}_r(j) + 1),$$

where we use $\mathcal{B}_r(j)$ to denote the j -th break-point⁹ of resource r . The costs are given by $c_{r,j} = \alpha_{r, \mathcal{B}_r(j)}$. Again, it is straightforward to see that all neighbourhoods are efficiently searchable.

4.2 Restrained Congestion Games

Although congestion games are guaranteed to have (at least one) PNE, the computational problem of actually finding one is considered hard; as a matter of fact, the problem PNE-CONGESTION is one of the most prominent PLS-complete problems. Our goal in this section is to investigate whether this computational barrier can be bypassed, under the lens of the more optimistic complexity model of smoothed analysis.

To achieve this, we establish an upper bound (Theorem 4.2) on the expected number of better-responses that need to be performed until a PNE is found in a congestion game, as a function of a critical structural parameter of its action space that we identify (see Definition 4.1). Then, we can deduce that for congestion games in which this parameter is appropriately bounded, the smoothed running time becomes tractable (Corollary 4.3). At the same time, we show how such a restriction does not make the problem trivially tractable, by proving that PNE-CONGESTION remains PLS-complete even for this subclass of games (Theorem 4.4).

Definition 4.1 (Restrained Congestion Games). A congestion game will be called B -restrained, where $B \in \mathbb{N}$, if the maximum number of resources changed by any single-player deviation is at most B . Formally,

$$\max_{i \in \mathcal{N}} \max_{\sigma, \sigma' \in \Sigma_i} |\sigma \Delta \sigma'| \leq B,$$

where Δ denotes the standard symmetric difference operator (recall that σ, σ' are subsets of resources).

Theorem 4.2. Consider a B -restrained n -player congestion game, under any of the smoothed-analysis models described in Section 4 (namely general, polynomial, or step-function latencies), with maximum density parameter ϕ . Then, performing any better-response dynamics, starting from an arbitrary strategy profile, converge to an (exact) PNE of the congestion game in an expected number of iterations that is bounded by

- $O(n^{B+3} k^2 m^2 \phi)$ for general games,
- $O(n^{B+d+2} \log(n) (d+1)^3 k^2 m^2 \phi)$ for polynomial games with degree at most d , and
- $O((d+1)^{B+2} n^2 \log(n) k^2 m^2 \phi)$ for step-function games with at most d break-points,

where $m = |\mathcal{R}|$ is the number of resources and $k = \max_{i \in \mathcal{N}} |\Sigma_i|$ is the maximum strategy set size.

Proof. In Section 4.1 we already showed how congestion games can be interpreted as CLO problems. In particular, we established a one-to-one correspondence between better-responses of the players to local improvements of the CLO cost objective (which corresponds to the value of Rosenthal's potential). Using this interpretation, we can now make use of our main black-box tool from Section 3: bounding the expected number of local search steps will induce the same bound in the expected iterations of better-response dynamics in the original congestion game. Therefore, the gist of our proof is to construct coverings $(\mathcal{E}, \mathcal{I})$ so that the induced CLO problem can be shown to be (λ, β, μ) -separable (see Definition 3.1) for parameters with appropriately small magnitude.

⁹ That is, if $\mathcal{B}_r = \{b_1, b_2, \dots, b_{d_r}\} \subseteq [n]$, then $\mathcal{B}_r(j) = b_j$, for every $j \in [d_r]$.

We start by establishing some properties that will be shared across all three different cost models. Fix a congestion game and its corresponding CLO problem, as described in [Section 4.1](#).¹⁰ For convenience, we denote by $k_i \coloneqq |\Sigma_i|$ the size of the strategy set of a player i , and let $k \coloneqq \max_{i \in \mathcal{N}} k_i$. Also, in all models we have a similar index structure in the cost part; the cost-part is given by $(s_{r,j})_{r \in \mathcal{R}, j \in J_r}$, where J_r depends on the cost function model ($J_r = [n]$ for general, $J_r = \llbracket d_r \rrbracket$ for polynomial, and $J_r = [d_r]$ for step-function costs).

The transition cover \mathcal{E} is constructed from clusters that collect all edges in the neighbourhood graph (of the CLO problem) that correspond to a fixed deviation of a player, regardless of the configuration of the remaining players. Formally, we let

$$\mathcal{E} \coloneqq \{E(i, \sigma_i, \sigma'_i) \mid i \in \mathcal{N}, \sigma_i, \sigma'_i \in \Sigma_i\}, \quad \text{where} \\ E(i, \sigma_i, \sigma'_i) \coloneqq \{(s(\sigma_i, \sigma_{-i}), s(\sigma'_i, \sigma_{-i})) \in E \mid \sigma_{-i} \in \Sigma_{-i}\},$$

and $s(\sigma)$ is used to denote the CLO configuration corresponding to strategy profile σ in the congestion game. Now we immediately get the bound

$$|\mathcal{E}| \leq nk(k-1),$$

which will be used as the value for our separability parameter λ (see [Definition 3.1](#)).

Next, for the coordinate cover \mathcal{I} , we cluster the indices with respect to each resource, i.e. we choose

$$\mathcal{I} \coloneqq \{I_r \mid r \in \mathcal{R}\}, \quad \text{where} \quad I_r = \{(r, j) \mid j \in J_r\}.$$

In congestion games, a deviation $\sigma_i \rightarrow \sigma'_i$ only affects the resources $r \in \sigma_i \Delta \sigma'_i$. Their loads are changed to increase by 1 for $r \in \sigma'_i \setminus \sigma_i$ and decrease by 1 for $r \in \sigma_i \setminus \sigma'_i$; the load of all other resources does not change. Our choice of the cover \mathcal{I} , therefore, will allow us to settle β for all models due to the B -restrain assumption on the size of $\sigma_i \Delta \sigma'_i$. In more detail, recall that the cost parts of a configuration depend only on the loads of the resources, thus all components associated with resources $r \notin \sigma_i \Delta \sigma'_i$ (specifically, the components with coordinates I_r) remain unchanged during the transition, since the load of r does not change either. The fact that the size of those sets $\sigma_i \Delta \sigma'_i$ is universally bounded, by assumption, will let us use $\beta = B$ as a separability parameter (see [Definition 3.1](#)).

The remaining parameter μ depends on the structure of the configurations in the cost part. Again, we emphasize that for all cost models in [Section 4.1](#), the sub-configuration $s_{I_r}(\sigma)$, which comprises all components of $s(\sigma)$ that correspond to a resource r , depends only on the load $\ell_r(\sigma)$ of resource r (under strategy profile σ). We can therefore represent it as a function $\mathbf{h}_r : \llbracket n \rrbracket \rightarrow \llbracket M \rrbracket^{J_r}$, i.e., $s_{I_r}(\sigma) = \mathbf{h}_r(\ell_r(\sigma))$. For ease of notation we write $\mathbf{s}_r = \mathbf{s}_{I_r}$ in the following.

To discuss $\text{range}_{I_r}(E(i, \sigma_i, \sigma'_i))$, we need to consider the configuration changes given by

$$\mathbf{s}_r(\sigma) - \mathbf{s}_r(\sigma') = \begin{cases} \mathbf{h}_r(\ell_r(\sigma)) - \mathbf{h}_r(\ell_r(\sigma')) = \mathbf{h}_r(\ell_r(\sigma)) - \mathbf{h}_r(\ell_r(\sigma) + 1), & r \in \sigma'_i \setminus \sigma_i, \\ \mathbf{h}_r(\ell_r(\sigma)) - \mathbf{h}_r(\ell_r(\sigma')) = \mathbf{h}_r(\ell_r(\sigma)) - \mathbf{h}_r(\ell_r(\sigma) - 1), & r \in \sigma_i \setminus \sigma'_i. \end{cases}$$

The only variable in this expression is therefore the initial load $\ell_r(\sigma)$. In either case, there are n possible initial loads for each resource¹¹ and therefore also at most n difference-vectors within $\text{range}_{I_r}(E(i, \sigma_i, \sigma'_i))$; thus $\mu \leq n$. In the following, we will discuss the actual configuration difference structure for each model and whether we can improve μ over this basic bound.

¹⁰ We will use standard notation for the various components and parameters of the game and the CLO problem, as introduced above in [Section 4](#).

¹¹ Note that although there are $n+1$ different loads $0, \dots, n$, a load-increasing resource cannot already have load n and a decreasing one cannot have 0.

- *General.* The CLO representation follows a binary model with $M = 1$ and $v = mn$: each resource r corresponds to components $s_{r,j}$, $j = 1, \dots, n$ (the indices from I_r), with a value of $s_{r,j} = 1$ for $j \leq \ell_r(\sigma)$ and $s_{r,j} = 0$ otherwise. The function $\mathbf{h}_r : \llbracket n \rrbracket \rightarrow \{0, 1\}^n$ is thus given by

$$(\mathbf{h}_r(\ell))_j = \mathbb{1}[j \leq \ell], \quad \text{for all } j \in [n].$$

In particular, the vectors $\delta \in \text{range}_{I_r}(E(i, \sigma_i, \sigma'_i))$ are a result of moving the rightmost entry with value 1 within the vector \mathbf{s}_r to the next larger or smaller load in the configuration, i.e. they are given by: $\delta_{\ell_r(\sigma)+1} = -1$ for $r \in \sigma'_i \setminus \sigma_i$; $\delta_{\ell_r(\sigma)} = +1$ for $r \in \sigma_i \setminus \sigma'_i$; and zeroes elsewhere. Since there are n many such vectors for every resource r , we cannot improve over $\mu = n$.

In this case, thus, [Theorem 3.2](#) yields an expected running time of at most

$$3 \cdot n^B n k (k-1) \cdot (mn)^2 \cdot \phi = O\left(n^{B+3} k^2 m^2 \phi\right).$$

- *Polynomial games.* For this model, each (cost-part) configuration component $s_{r,j}$ is given by accumulated monomials $\mathfrak{S}_j(\ell_r(\sigma))$, for degrees $j = 0, 1, \dots, d_r$. Thus, $v = m(d+1)$, and also $M = n^{d+1}$, in order to capture all possible values of these functions. Therefore, we now get $\mathbf{h}_r : \llbracket n \rrbracket \rightarrow \llbracket n^{d+1} \rrbracket^{d_r+1}$ with

$$\mathbf{h}_r(\ell) = (\mathfrak{S}_0(\ell), \dots, \mathfrak{S}_{d_r}(\ell))$$

for the configuration component of I_r . Again, we cannot do better than the basic bound, so we use $\mu = n$.

Similarly to the previous case for general latency functions, using [Theorem 3.2](#) we can bound the expected number of better-response iterations by

$$3 \cdot n^B n k (k-1) \cdot (m(d+1))^2 n^{d+1} \log(n^{d+1}) \cdot \phi = O\left(n^{B+d+2} \log(n) (d+1)^3 k^2 m^2 \phi\right).$$

- *Step-function games.* The configuration mapping is now given by $s_{r,j} = \max(0, \ell_r(\sigma) - \mathcal{B}_r(j) + 1)$ and $M = n$, $v = m \cdot d$. Therefore, the (cost-part) configuration components of a resource r are represented by the function $\mathbf{h}_r : \llbracket n \rrbracket \rightarrow \llbracket n \rrbracket^{d_r}$ given by

$$\mathbf{h}_r(\ell) = (\max(0, \ell - \mathcal{B}_r(1) + 1), \dots, \max(0, \ell - \mathcal{B}_r(d_r) + 1)).$$

In this case, we can do even better than the basic $\mu = n$ bound. We investigate the structure of the differences with respect to each coordinate $j \in J_r$, for an increase in the load of a resource r , i.e. for $r \in \sigma'_i \setminus \sigma_i$ (the decreasing case follows analogously):

$$\begin{aligned} (\mathbf{h}_r(\ell) - \mathbf{h}_r(\ell+1))_j &= \max(0, \ell - \mathcal{B}_r(j) + 1) - \max(0, (\ell+1) - \mathcal{B}_r(j) + 1) \\ &= \begin{cases} 0 - 0 = 0, & \text{if } \ell < \mathcal{B}_r(j) - 1, \\ -((\ell+1) - \mathcal{B}_r(j) + 1 - 0) = -1, & \text{if } \ell = \mathcal{B}_r(j) - 1, \\ (\ell - \mathcal{B}_r(j) + 1) - ((\ell+1) - \mathcal{B}_r(j) + 1) = -1, & \text{if } \ell > \mathcal{B}_r(j) - 1. \end{cases} \end{aligned}$$

Because the jump points $\mathcal{B}_r(j)$ are ordered increasingly with respect to j , the resulting vectors $\mathbf{h}_r(\ell) - \mathbf{h}_r(\ell+1)$ are of the form $(-1, \dots, -1, 0, \dots, 0)$, including the zero vector. Therefore, with respect to the coordinate cluster I_r , there are at most $d_r + 1$ possible distinct vectors for the (cost-part) configuration differences, and we choose $\mu = \max_r d_r + 1 = d + 1$. By [Theorem 3.2](#) we can thus bound the expected number of iterations by

$$3 \cdot (d+1)^B n k (k-1) \cdot (m \cdot d)^2 n \log(n) \cdot \phi = O\left((d+1)^{B+2} n^2 \log(n) k^2 m^2 \phi\right).$$

□

An immediate corollary of [Theorem 4.2](#) is that, when congestion games are sufficiently restrained, PNE can be found efficiently via better-response dynamics. In what follows, by a *constantly-* and *polylogarithmically-*restrained congestion game we mean a B -restrained game with $B \in O(1)$ and $B \in O(\log^c N)$ for some constant $c > 0$, respectively, where N is the size of the input.

Corollary 4.3. *Better-response dynamics terminate in polynomial smoothed time for the class of constantly-restrained congestion games, and in quasipolynomial smoothed time for polylogarithmically-restrained games, under any cost model. Under the step-function cost model, in $O(\log N)$ -restrained congestion games with a constant number of steps d , better-response dynamics terminate in polynomial smoothed time.*

We now show that the class of constantly-restrained congestion games, for which [Corollary 4.3](#) provides efficient smooth running time, constitutes a computationally meaningful restriction of arbitrary congestion games, since they can still encode the PLS-completeness of the original problem. The hardness is a straightforward adaptation of that in [\[Rou16\]](#), and so the proof of [Theorem 4.4](#) is deferred to [Appendix B](#) for completeness. It makes use of the fact that LOCALMAXCUT is PLS-complete even for constant-degree graphs.

Theorem 4.4. *The problem of computing a PNE of a constantly-restrained congestion game is PLS-complete, for all the input models described in [Section 4](#) (namely general, polynomial, or step-function cost representations).*

In addition to the conditional intractability that PLS-hardness implies for these families of congestion games, we show an unconditional lower bound on the worst-case running time of the standard local search algorithm of the problem. We do this by using the notion of a *tight PLS-reduction*, as discussed in [Section 2](#). Since we reduce from the problem LOCALMAXCUT- d for $d \geq 5$ (defined in [Appendix B](#)), which admits a configuration starting from which the standard local search algorithm needs exponentially many iterations (see discussion after proof of [Theorem 5.10](#)), our tight PLS-reduction implies that standard local search of our families of congestion games – under *any* pivoting rule – takes exponential time in the worst case.

4.3 Network Congestion Games

An interesting, and very well studied, variation on the vanilla representation model for congestion games (which we presented at the start of [Section 4](#) above) is that of *network* congestion games. In such games, the strategy sets Σ_i of the players are not given explicitly in the input, but implicitly via an underlying directed graph G whose edges constitute the resources of the game. More precisely, for each player i we are given an origin o_i and a destination d_i node of G . Then, Σ_i is defined (implicitly) as the set of all (simple) $o_i \rightarrow d_i$ paths in G . Importantly, this means that now players may have exponentially many strategies available to them.

A critical implication is that Σ_i cannot be searched *exhaustively* for better-responses. However, a better-response can still be found efficiently: keeping all other players fixed, a minimum-cost strategy of player i is a shortest path on graph G with edge costs equal to the cost $c_r(\ell_r(\sigma_{-i}) + 1)$ of an edge/resource r when used by player i . This means that actually a polynomial-time *best-response* oracle does exist. This immediately places network congestion games in the complexity class PLS (since neighbourhoods can be searched efficiently for a local cost improvement; see our discussion in [Section 2](#)) and thus it constitutes a valid CLO problem (via a similar interpretation as we did for general congestion games in [Section 4.1](#)). To emphasize this, we will refer to these “canonical” best-response dynamics of network congestion games as *shortest-path* dynamics.

Finding PNE of network congestion games remains a PLS-complete problem [\[ARV08\]](#). Given the prominence of these games, both in the theoretical and applied literature, in this section we want to

identify conditions under which network congestion games inherit the desirable properties of their general counterparts that allow them to be tractable under smoothed analysis. More precisely, can our positive result from [Theorem 4.2](#) be applied to network congestion games in a straightforward way? And which structural parameters are now relevant for the running time bound? At the same time, the network congestion game instances that allow for efficient smoothed solutions should still be interesting enough to remain PLS-hard under traditional worst-case analysis. We introduce the following family of network congestion games that are defined by two parameters.

Definition 4.5 (Compact Network Congestion Games). For A, B positive integers, a network congestion game is called (A, B) -compact if (a) each player has at most A different best-response strategies, and (b) all such strategies are paths of length at most B . Formally, there exist strategy sets $\Sigma_i^* \subseteq \Sigma_i$, such that:

- (a) $|\Sigma_i^*| \leq A$ for all $i \in \mathcal{N}$ and $\arg\min_{\sigma_i \in \Sigma_i} C_i(\sigma_i, \sigma_{-i}) \subseteq \Sigma_i^*$ for all $i \in \mathcal{N}, \sigma_{-i} \in \Sigma_{-i}$, and
- (b) $|\sigma_i| \leq B$ for all $i \in \mathcal{N}, \sigma_i \in \Sigma_i^*$.

Property (b) above will serve the purpose of imposing the restraint condition (see [Definition 4.1](#)) needed to deploy [Theorem 4.2](#). Property (a) will help us constrain the exponentiality of the strategy space of network games, in order to be able to handle them using tools designed for general games. Intuitively, this property can be related to classical vehicle routing settings in which players can a priori exclude unreasonable detours or paths that involve a road with a construction site with large delay. Both properties are illuminated within the proof of our following positive result for network congestion games.

Theorem 4.6. *Consider an (A, B) -compact n -player network congestion game, under any of the smoothed-analysis models described in [Section 4](#) (namely general, polynomial, or step-function latencies), with maximum density parameter ϕ . Then, performing shortest-path dynamics, starting from an arbitrary strategy profile, converges to a PNE of the game in an expected number of iterations that is polynomial in $\phi, (d + 1)^B, A$, and the description of the game, where the parameter d depends on the cost function representation. In particular: for general latencies, $(d + 1)$ can be replaced by n ; for polynomial latencies, d is the maximum degree; and for step-functions, d is the maximum number of break-points.*

Proof. Let us count iterations, i.e., steps of the dynamics, starting from step 0 when the dynamics are in the initial arbitrary strategy profile. We will say that the dynamics are in step t if t improving moves have taken place since step 0.

We introduce a random variable S_i for the step of the dynamics in which player i best-responded (that is, chose a shortest path) for the first time. Observe that once player i best-responds, by definition, she plays a strategy in Σ_i^* and from then on, whatever strategy she plays will be in that set. Without loss of generality, let us rearrange the players according to the order with which each best-responds for the first time, meaning that player i best-responds only after all players j with $j < i$ did so. Focusing on a player $i + 1, i < n$, that has not best-responded yet in a given step of the dynamics, notice that neither have all players $j > i + 1$.

Let us now define the subgame G_i in which only players in $[i]$ participate and where each edge has its load artificially increased according to the fixed strategies of the other players. The initial strategy profile of each subgame G_i is the profile of the original game's dynamics after step S_i , which is in $\prod_{j \in [i]} \Sigma_j^*$, by definition. G_i allows us to exactly reflect the original game's dynamics between time steps S_i and $S_{i+1} - 1$, while all strategies of the participating players are in $\prod_{j \in [i]} \Sigma_j^*$.

Denote by T_i the random variable of the number of iterations until a PNE in G_i is reached. As soon as this is the case, the next time step in the original game must involve a new player, so $S_{i+1} \leq S_i + T_i + 1$. Note that not all T_i steps of G_i are necessarily played in the original game's dynamics, as player $i + 1$ can deviate from their original action at time step S_{i+1} before an equilibrium is reached in G_i .

Now we can express the number of better-response steps in the original network congestion game as $T = S_n + T_n$, and by the observation above, we can recurse to bound this variable by

$$T = S_n + T_n \leq S_{n-1} + T_{n-1} + 1 + T_n \leq \dots \leq (n-1) + \sum_{i=1}^n T_i.$$

Now notice that in subgame G_i of an (A, B) -compact network congestion game, any better-response dynamics (such as best-response dynamics) are equivalent to the same dynamics on a $2B$ -restrained¹² i -player congestion game with maximum strategy set size A (see [Section 4.2](#)). Therefore, we get that

$$\mathbb{E}[T] \leq (n-1) + \sum_{i=1}^n \mathbb{E}[T_i],$$

where the expectation $\mathbb{E}[T_i]$ is bounded polynomially due to [Theorem 4.2](#) for the various latency representations. \square

An immediate consequence of [Theorem 4.6](#) is that (analogously to [Corollary 4.3](#) for general congestion games) in (A, B) -compact network games with sufficiently small parameters A, B , a PNE can be found efficiently under smoothness.

Corollary 4.7. *Let N be the size of the input, and A be a polynomial in N . Shortest-path dynamics on (A, B) -compact network congestion games within the general, polynomial, and step-function cost model terminate in polynomial smoothed time when $B \in O(1)$ and in quasipolynomial smoothed time when $B \in O(\log^c N)$ for some constant $c > 0$. Under the step-function cost model, when A is a polynomial in N , $B \in O(\log N)$, and the number of steps d is constant, shortest-path dynamics terminate in polynomial smoothed time.*

The following hardness result establishes that such games, even for $A, B \in O(1)$, are PLS-hard. The proof is deferred to [Appendix C](#). It is based on the reduction constructed by Ackermann, Röglin, and Vöcking [[ARV08](#)], with special care taken in order to incorporate constant-length paths that can be established by making use of the fact that LOCALMAXCUT is PLS-complete even for constant degree graphs.

Theorem 4.8. *The problem of computing a PNE of an (A, B) -compact network congestion game is PLS-complete, even for $A, B \in O(1)$, for all the input models described in [Section 4](#) (namely general, polynomial, or step-function latencies).*

Similarly to our PLS-hardness reduction of [Theorem 4.4](#), the above theorem's PLS-reduction is tight, in the sense of [[SY91](#)] (see [Section 2](#)). The chain of tight PLS-reductions that leads to Network Congestion Games starts from LOCALMAXCUT with maximum degree 5, and includes Congestion Games. As we discuss after the proof of [Theorem 5.10](#), in such LOCALMAXCUT instances there is a starting configuration from which all improvement sequences of standard local search have exponential length in the worst case. Therefore, shortest-path dynamics on our family of network congestion games – under *any* pivoting rule – need exponentially many iterations in the worst case. In contrast, our [Theorem 4.6](#) and [Corollary 4.7](#) show that under smoothness, even for significantly wider instance families that include these problematic cases, shortest-path dynamics terminate after polynomially many steps in expectation.

¹² When each strategy consists of at most B edges, their pairwise symmetric differences have cardinality at most $2B$.

5 Further Applications

For the remainder of the paper, we will apply our main tool from [Section 3](#) to directly get bounds on the expected number of iteration steps for local search algorithms, for various well-known local optimization problems. It is important to emphasize that we will be looking into PLS-complete problems, that is, problems that are the hardest ones in PLS: unless $\text{PLS} = \text{FP}$, there is no (worst-case) polynomial time algorithm to find a local optimum of these problems. Furthermore, due to their reductions being tight PLS-reductions (see [\[SY91\]](#) for the definition), their individual standard local search algorithms are unconditionally inefficient: in each of these problems there exists an initial configuration, such that if the standard local search algorithm starts from there, every sequence of improving moves until termination is exponentially long.

All the proofs follow a similar structure: for each optimization problem, we first show that it can be formulated as a CLO problem, establishing one-to-one correspondence between feasible solutions and configurations, respecting the intended neighbourhood structure and the parameters to be perturbed under smoothness; then, we apply [Theorem 3.2](#) to the smoothed CLO problem, via an appropriate covering, to bound the expected running time. Coming up with the right choice for a covering is the technical part of our proofs.

Many of the following problems have a parameterized neighbourhood (k -OPT, k -FLIP, etc.). Their neighbourhood size usually increases exponentially in this parameter and therefore it is often assumed as being constant: this implies efficiently searchable neighbourhoods and, importantly, membership in PLS. In principle, however, it might be of interest to also consider instance-dependent neighbourhood sizes, either in terms of superpolynomial local improvements or for special types of instances where, e.g., neighbourhoods for logarithmically large values of k might actually be efficiently searchable for the specific problem at hand. For that reason, we make a point for all our running-time bounds to *not* hide any dependencies on these parameters in the $O(\cdot)$ -notation. In that way, even when the bound depends on k in the exponent, one can readily deduce a bound on the expected number of local-search iterations, beyond constant k ; in particular, for polylogarithmic k we derive quasipolynomial bounds in the input size.

A technical remark is in order here: in the following, we will not explicitly differentiate between maximization and minimization problems, as [Theorem 3.2](#) can be applied analogously in either case. In order to appeal to intuition though, for maximization problems we will look at the configuration difference $s'_I - s_I$, instead of $s_I - s'_I$, when discussing the cover properties (see [Section 3](#)), as this corresponds to the expression giving a positive improvement for a transition (s, s') .

5.1 TSP with the k -OPT Heuristic

The Travelling Salesman problem (TSP) is one of the best studied problems in combinatorial optimization [\[Coo12; ABCC06\]](#). TSP is the problem of finding a Hamiltonian cycle – also called *tour* in this setting – with minimum sum of edge weights in a weighted undirected graph. This problem is NP-hard as a global optimization problem, which is a standard result in classical complexity theory [\[GJ90\]](#). One of the most commonly used methods to approximate the solution is the k -OPT heuristic, where the neighbourhood is induced by replacing at most k edges of the current tour on n vertices with another set of edges, resulting in a new tour. Despite the worst-case results on its approximation ratio, such as the tight $\Theta(n^{1/k})$ ratio in Metric TSP for any constant k due to [\[Zho20\]](#), more promising probabilistic observations exist and this heuristic seems to perform well in practice, see e.g. [\[GP07; JM97; EM09\]](#).

Krentel [\[Kre89\]](#) settled that the local optimization problem of TSP/ k -OPT is PLS-complete, for a sufficiently large constant k , as well as the existence of instances that require exponentially many iterations of local improvements. Klauck [\[Kla96\]](#) improves this result to a tight PLS-reduction to TSP/ k -OPT. Englert, Röglin, and Vöcking [\[ERV14\]](#) constructed two-dimensional Euclidean TSP instances in which 2-OPT needs exponentially many iterations, under the L_p metric for any $p > 0$. Furthermore,

Chandra, Karloff, and Tovey [CKT99] provide instances where k -OPT takes exponentially long to converge, for every $k > 2$. Note that these explicit constructions provide *a* path of exponential length, but not *all* paths starting at the same configuration have this property.

As for the smoothed complexity of TSP, [ERV14] established polynomial smoothed running time for the 2-OPT heuristic for TSP in the input model where the graph nodes are given as points in \mathbb{R}^d , drawn from independent distributions over $[0, 1]^d$ with density bounded by ϕ , and their pair-wise distances follow the L_p metric. This result has later been improved in [MV13; MR23] for the special case of smoothness by Gaussian noise.

Furthermore, in [ERV16] a different model of smoothness was discussed for 2-OPT in the general TSP (on undirected graphs), where the cost of each edge is drawn from independent distributions, and a polynomial bound on the smoothed running time was established. Our analysis extends the polynomial smoothed complexity of the k -OPT heuristic for TSP from $k = 2$ to arbitrary $k \geq 3$. Using distributions that are independent of the (undirected) edges reflects the application better, where we can think of the location of each city as fixed, but the travelling times between them can vary due to traffic, which is uncorrelated between different cities. Our analysis covers the case of TSP on a directed graph, known as *Asymmetric TSP* (ATSP) [KP80; GP07; MKA07; STV20], which has not been studied before in the context of smoothed analysis. Under this asymmetric model, the travelling times in different directions, between the same two cities, can differ arbitrarily.

When discussed in the context of a local search heuristic for the global optimization problem, the k -OPT heuristic also includes every 2-OPT improvement, and is therefore often considered a stronger method (which comes at the price of higher per-step complexity). In terms of the running time of local search, however, the methods are not as easily comparable: every solution of the k -OPT heuristic also provides a solution for 2-OPT on the same instance, but it might take more steps to actually terminate the local search procedure for k -OPT due to the larger neighbourhood; a solution for k -OPT not only needs to be at least as good as every possible exchange of 2 edges, but is compared to every k -OPT move and local search may terminate only a long time after the first local optimum sufficient for the 2-OPT heuristic has been encountered.

Definition 5.1 (TSP/ k -OPT). By TSP/ k -OPT we denote the following local optimization problem. Let $G = (V, E)$ be a weighted undirected graph. Find a locally minimal tour C on G with respect to the k -OPT neighbourhood. The cost of a tour is given by the sum of its edge weights. The k -OPT neighbourhood of C is the set of all tours that can be generated from C by replacing some set of $r \in \{2, 3, \dots, k\}$ edges from it with a different set of r edges in the graph. A solution is a tour that does not have a neighbour with smaller cost.

Analogously, we define the problem ATSP/ k -OPT on directed graphs. We will use the notion of *edges* to refer to both, edges in an undirected graph and arcs in a directed graph.

Definition 5.2 (ATSP/ k -OPT). Given a weighted directed graph $G = (V, E)$, ATSP/ k -OPT is the problem of finding a locally minimal directed tour (a directed Hamiltonian cycle) under the k -OPT neighbourhood. The cost of a directed tour is given by the sum of its edge weights. The neighbours of a given tour are all directed tours that can be generated by replacing at most k edges from it by other edges.

Note that to find a neighbour of C using the k -OPT heuristic we are removing any set of $2 \leq r \leq k$ edges of a tour (for which there are $\binom{|V|}{r}$ possible choices of edges for each r , as C consists of $|V|$ many edges). Let us fix a set of r edges from the current tour to be removed, thus splitting the tour into r disjoint paths (or isolated vertices, which we consider as degenerate paths). Then, for each such path p , its endpoints can only be reconnected via an edge to an endpoint of the other $r - 1$ paths to form a new tour; in other words, at most $2(r - 1) \leq 2r$ vertices need to be considered for each of the $2r$ endpoints. Any full reconnection of the set of endpoint vertices can be thought of as a permutation of this set, so there are at most $(2r)!$ possible new tours. In total, for each fixed value of r we thus get

at most $\binom{|V|}{r} \cdot (2r)!$ neighbouring tours, resulting in at most $(k-1) \cdot \binom{|V|}{k} \cdot (2k)! \in O(|V|^k (2k+1)!)$ many neighbours for any given tour. Therefore, when k is constant, the k -OPT-neighbourhood of any configuration is efficiently searchable. The same bound holds for k -OPT of ATSP, where fewer of the reconnections outlined above are feasible. In particular, for ATSP a valid heuristic is only given for $k \geq 3$, as the unique new (undirected) tour given by any application of the 2-OPT procedure does not form a valid directed cycle.

Similarly to [ERV16], we assume that an initial tour of the graph is given or can be queried by an oracle, because even the problem of deciding whether a Hamiltonian cycle exists in a graph is NP-complete.

For the smoothed version of TSP/ k -OPT and ATSP/ k -OPT, the edge weights w_e are drawn from independent distributions $f_e : [0, 1] \rightarrow [0, \phi]$ each. Our analysis only makes use of the concept of a subset of edges that forms a cycle and thus applies to both, undirected and directed edges. Note that TSP can be reduced to ATSP in standard complexity by choosing symmetric edge weights, which introduces some overhead by doubling the representation size, as every edge is represented by a pair of opposing edges. However, under smoothed analysis, these two representations of TSP differ substantially, because both the opposing edges' weights appear in the input of the problem, they experience independent noise in the ATSP problem, contrary to the single edge's weight that appears in the input of the original TSP problem.

Theorem 5.3. *The problem of computing a local optimum of TSP or ATSP with the k -OPT heuristic (for constant k), on a graph with m edges, whose weights $w_e, e \in E$ are drawn independently from distributions with densities $f_e : [0, 1] \rightarrow [0, \phi]$, has polynomial smoothed complexity, i.e. the expected number of iterations of the local search algorithm can be bounded by $O\left(4^{k^2} m^{k+2} \phi\right)$.*

Proof. We define a binary ($M = 1$) CLO problem. Each component of the cost-part corresponds to an edge in the graph, i.e. $v = m$ and the cost part is $\mathbf{s}^\bullet = (s_e)_{e \in E}$. A tour C is mapped to a CLO configuration $\mathbf{s}(C)$ by indicating which edges are being used in it:

$$s_e(C) = \mathbb{1} [e \in C].$$

The cost coefficients are given by $c_e = w_e$. By this model, the one-to-one correspondence between the tours and the configurations of the CLO problem is immediate, as either is a representation of a subset of edges and their costs coincide. The neighbourhood relation for the CLO problem is derived directly from the neighbouring tours in the TSP/ k -OPT instance. Because the CLO cost coefficients are equal to the edge weights, the smoothed costs of the TSP instance also translate directly to the smoothed CLO model in Theorem 3.2. We establish the remaining properties of the theorem (a covering $(\mathcal{E}, \mathcal{I})$ and its parameters $\lambda = m^k, \mu \in \{1, 2\}, \beta = 2k^2$) in the following.

For the coordinate cover, we can simply use the singleton coordinate clusters: $\mathcal{I} = \{\{e\} \mid e \in E\}$. To cover the neighbourhood graph, we define transition clusters $E(e_1, \dots, e_k)$ for every choice of $e_i \in E$. Each cluster captures the set of transitions where the given edges e_i are removed from the tour and replaced by (equally many) other edges to form a new tour. Note that we allow for some edges to coincide, $e_i = e_j$, to account for transitions where fewer than k edges are replaced. Therefore, the number of clusters is $\lambda = m^k$.

Next, we investigate for every transition cluster $E(e_1, \dots, e_k)$ which edges can possibly replace the removed ones to form a new tour. This is equivalent to determining the core of each transition cluster and allows us to settle β . Note that this is *not* the same as the number of neighbours of a given tour. When a subset of k edges is removed from a tour, there remains a set of k paths. Note that some of these paths might be degenerate in the sense that they only contain a single vertex, which occurs when two adjacent e_i, e_j are removed from the tour. We call these vertices *isolated*. Denote by k_1 the number of isolated vertices and by k_2 the number of proper paths that remain, thus $k = k_1 + k_2$. In order to

form a new tour, each of the endpoints and isolated vertices needs to be reconnected. Clearly, we need to reconnect the vertices within this set, as every other vertex in the graph still has 2 incident edges. Additionally, we also do not reconnect them to their previous neighbour (or neighbours in the case of an isolated vertex). For every end point of a non-degenerate path, the edge that leads to the other end point of the same path also cannot be used, as this would create a short cycle. In total, we thus have $k_1 + 2k_2$ vertices that need to be reconnected and each of them has at most $k_1 + 2k_2 - 3$ possible edges to reconnect them properly. Note that due to symmetry, every edge we count appears twice in the analysis, so we can bound the number of possible edges that can be added to form a new tour by

$$\frac{1}{2}(k_1 + 2k_2)(k_1 + 2k_2 - 3) \leq 2k^2 - k,$$

and together with the k edges that are being removed, at most $\beta = 2k^2$ different edges can be affected by any transition cluster. This argument, and therefore also the bound, still holds when a set of $r < k$ edges are replaced from a tour.

Regarding μ , we need to bound the number of possible values of $s_e - s'_e$ for any transition (s, s') within a given transition cluster $E(e_1, \dots, e_k)$. For edges $e \in \{e_1, \dots, e_k\}$ that are removed, we observe $s_e - s'_e = 1$, whereas for any of the previously discussed edges e that can possibly be added to form a new tour, we have $s_e - s'_e \in \{0, -1\}$, depending on whether it is actually used in the new tour represented by the configuration s' or not. Therefore we conclude with $\mu = 2$ and apply [Theorem 3.2](#) with the derived parameters to get a bound on the expected number of iterations for TSP/ k -OPT of

$$3 \cdot 2^{2k^2} m^k \cdot m^2 \cdot \phi = 3 \cdot 4^{k^2} m^{k+2} \phi.$$

□

Remark. A slightly more refined analysis allows to improve the constant for TSP/2-OPT: When removing two edges from a given tour, there exists only one other possibility to construct a new tour, as from the three neighbouring vertices available for reconnection, one is the previous neighbour and another is the other end of the same path. Therefore, we can choose $\mu = 1$ instead, as we know precisely which edges are being added to the tour, and get a bound on the expected number of iterations of $3m^{k+2}\phi$ instead.

5.2 Maximum Satisfiability

MAXSAT is the optimization problem of finding an assignment of a CNF formula that maximizes the weight of satisfied clauses. This problem is NP-hard and therefore heuristics are often used to approximate solutions in practice. The k -FLIP local search algorithm is one of the most commonly used heuristics, both by itself or as part of a more general heuristic, see e.g. [\[YI99; Sze11\]](#) for theoretical discussions on it. Using this heuristic, we can define the local optimization problem MAXSAT/ k -FLIP:

Definition 5.4. MAXSAT/ k -FLIP is the problem of finding a locally maximal truth assignment of a CNF formula with clauses $C = \{C_1, \dots, C_m\}$ over binary variables $X = \{x_1, \dots, x_n\}$. There are weights w_1, \dots, w_m associated with the clauses and the weight of an assignment is given by the sum of weights of the satisfied clauses. The neighbours of a truth assignment are given by the k -FLIP neighbourhood, which are all assignments generated by changing the truth value of up to k variables.

When k is a constant, this neighbourhood has polynomial size and is therefore efficiently searchable. MAXSAT/1-FLIP has been shown to be PLS-complete in a series of works [\[Kre90; Kre89; Kla96\]](#), even for the case where the size of the clauses (i.e. the number of literals) and the number of occurrences of any single variable across all clauses are constant. These hardness results were proven via tight PLS-reductions (see [Section 2](#)), starting from the maximization or its equivalent minimization problem Circuit

Flip [JPY88], for which there exists a configuration with only exponentially long sequences to solutions. This implies directly that the MAXSAT/1-FLIP instances we study have also such a configuration whose improvement sequences are all exponentially long under the standard local search algorithm.

In [Bra18], a smoothed version of MAXSAT/1-FLIP has been analysed, by adapting the method for LOCALMAXCUT from [ABPW17], claiming a bound on the expected running time of $O(n^{5\kappa+5})$, where κ is a bound on the size of the clauses. Note, however, that similarly to how the original method relies on the complete graph for the result, this result applies only to the complete formula, in which every possible clause of given fixed size is present.

Our contribution shows polynomial smoothed time for MAXSAT/1-FLIP with bounded variable occurrence, but without any other restrictions on the formula.

Theorem 5.5. *The expected number of iterations of smoothed MAXSAT/1-FLIP, where each clause weight w_i is drawn from independent distributions with densities $f_i : [0, 1] \rightarrow [0, \phi]$, is bounded by $O(2^B n m^2 \phi)$, where B is the maximum number of different clauses that any variable occurs in.*

Proof of Theorem 5.5. MAXSAT/1-FLIP can be stated as a binary ($M = 1$) CLO problem. We identify the cost part with the clauses ($v = m$) by defining a variable s_i for every clause C_i with cost coefficient $c_i = w_i$. Denote by $a : \mathcal{X} \rightarrow \{0, 1\}$ a truth assignment and by $s(a)$ its induced CLO configuration, defined as follows: In the cost part, we assign $s_i(a) = 1$ if the associated clause C_i is satisfied, and $s_i(a) = 0$ otherwise. Thus, the cost of the CLO configuration $s(a)$ is equal to the weight of the assignment a in the MAXSAT/1-FLIP problem. Because this does not necessarily represent all the information of the underlying truth assignment, we use the non-cost part of $s(a)$ to represent a explicitly as a binary string of length $\bar{v} = n$. The CLO neighbourhood can now be derived from the assignments in the 1-FLIP neighbourhood and be searched efficiently. For the smoothed problems, the random distributions of the weights in the smoothed MAXSAT/1-FLIP instance are mapped identically to the distributions of the costs in the smoothed CLO problem. In the following, we apply Theorem 3.2 to achieve a bound on the smoothed running time.

To do so, we first specify the singleton coordinate cover $\mathcal{I} = \{\{i\}, i = 1, \dots, m\}$ of the indices of the cost part. The transition cover \mathcal{E} comprises clusters $E(x_j, b)$ for each variable x_j and $b \in \{0, 1\}$. The cluster $E(x_j, b)$ collects all transitions where the truth assignment of variable x_j flips from $b - 1$ to value b , regardless of the assignment of the remaining variables. Formally, we can cover the edges of the neighbourhood graph using $\lambda = 2n$ many transition clusters by

$$\mathcal{E} = \{E(x_j, b) \mid j \in [n], b \in \{0, 1\}\},$$

where $E(x_j, b) = \{(s(1 - b, a_{-j}), s(b, a_{-j})) \mid a : \mathcal{X} \rightarrow \{0, 1\} \text{ arbitrary}\},$

and the notation (b, a_{-j}) refers to the truth assignment that maps variable x_j to value b and all other variables according to assignment a . A change in the assignment of variable x_j can only affect the satisfaction of clauses that contain the literal x_j or \bar{x}_j . Therefore, we know that $\text{core}(E(x_j, b))$ comprises precisely the indices of those clauses. By assumption, there are at most B such clauses for every x_j , and we have established $\beta = B$.

Finally, we specify $\mu = 2$ by inspecting, for fixed (x_j, b) , the entries s_i, s'_i of a transition $(s, s') \in E(x_j, b)$: For $b = 1$, when clause C_i contains the literal x_j , it may either become satisfied by the changed assignment (i.e. $s_i = 0$ and $s'_i = 1$ after the transition) or remain satisfied ($s_i = s'_i = 1$). For a clause that contains the negative literal \bar{x}_j , either it becomes unsatisfied ($s_i = 1$ and $s'_i = 0$) or remains satisfied ($s_i = s'_i = 1$). Analogously, the same observation holds with exchanged roles in the literals for $b = 0$. In any case, the difference $s'_i - s_i$ can only take two possible values for every transition in $E(x_j, b)$ and $\mu = 2$ is justified.

Thus, by application of Theorem 3.2, the expected number of steps is bounded by $3 \cdot 2^B 2n \cdot m^2 \cdot \phi$. \square

Our proof method can easily be adapted to account for more general k -FLIP neighbourhoods as well, where at most k variables can change their value simultaneously. As the k -FLIP neighbourhood for $k > 1$ also includes every 1-FLIP move, every local optimum to the former is also optimal for the latter; thus MAXSAT/ k -FLIP is still PLS-hard.

Theorem 5.6. *The expected number of iterations of smoothed MAXSAT/ k -FLIP, where each clause weight w_i is drawn from independent distributions with densities $f_i : [0, 1] \rightarrow [0, \phi]$ is bounded by $O(3^{kB} n^k m^2 \phi)$, where B is the maximum number of clauses that a variable occurs in.*

Proof. The proof is analogous to [Theorem 5.5](#) and we only state the necessary changes to deal with the k -FLIP neighbourhood instead of 1-FLIP:

For the transition cover, we now use clusters $E(x_{j_1}, \dots, x_{j_k})$, $x_{j_\ell} \in \mathcal{X}$. Each cluster collects all transitions where all given variables change their truth assignment, regardless of their initial assignment or the assignment of the remaining variables. We do not require all x_{j_ℓ} to be different, so as to account for neighbouring assignments where fewer than k variables flip. Therefore, we can now bound the number of transition clusters by $\lambda = n^k$. Every variable occurs in at most B clauses, therefore the transitions in $E(x_{j_1}, \dots, x_{j_k})$ can affect at most $\beta = kB$ cost components.

Finally, because the CLO problem is binary, we know that within every component s_i the transition difference $s'_i - s_i$ can take at most the three values $\{-1, 0, 1\}$, which gives $\mu = 3$. Applying [Theorem 3.2](#) yields a bound for the expected number of steps of $3 \cdot 3^{kB} n^k \cdot m^2 \cdot \phi$. \square

5.3 Maximum Cut

LOCALMAXCUT is among the problems in PLS that has been most extensively studied in the context of smoothed analysis. The result that our method generalizes is the polynomial smoothed complexity for graphs of up to logarithmic degree due to [\[ET11\]](#), who studied a smoothness model with Gaussian noise. By application of our black-box tool, we are able to extend this result to the smoothness model of arbitrary distributions with density bounded by ϕ and also prove it directly for the expected running time instead of the weaker *probably smoothed complexity* they use. Further polynomial results under smoothness exist for LOCALMAXCUT on complete graphs, both for a version of “metric” LOCALMAXCUT in [\[ER15\]](#) and for the noise model with general edge weights in [\[ABPW17; BCC21\]](#). For general graphs on the other hand, quasipolynomial smoothed complexity was shown in [\[ER17\]](#) and further improved by [\[CGVYZ20\]](#), as well as parametrized by the arboricity of the graph in [\[Sch24\]](#).

It is important to note that none of the aforementioned graph models is stronger than any other, and that is due to a delicate difference between worst-case and smoothed analysis in the equivalence between input representations. For example, in worst-case analysis, having a general graph is equivalent to having the complete graph as input to LOCALMAXCUT, in the sense that the complexity of the problem remains the same (up to a constant factor): any general graph can be turned into a complete graph with zero weights for the non-existing edges. On the contrary, in smoothed analysis this is not the case, since an instance in which an edge does not exist cannot be straightforwardly turned into an – in terms of smoothed complexity – equivalent instance by setting that edge’s weight to zero: as soon as the zero weight appears in the input it is supposed to undergo smoothness, something that we wish to disallow for an essentially non-existing edge.

Our method allows us to easily extend the analysis to LOCALMAX- k -CUT, a generalization of LOCALMAXCUT that considers partitions of the vertex set into k blocks instead of only 2. The FLIP neighbourhood for LOCALMAX- k -CUT is similarly given by a single vertex changing from one of the cut-sets to another. LOCALMAX- k -CUT under this neighbourhood has been analysed by [\[BCC21\]](#), who established polynomial smoothed complexity for the complete graph and a bound of $O(\phi n^{2(2k-1)k \log(kn)+3+\eta})$, for every $\eta > 0$, in the general case. Our method approaches this problem from a different perspective through the parameterization on the degree of the graph and thereby improves on this result for graphs of up to logarithmic degree with a bound that does not depend on k at all.

We first analyse the case of $k = 2$, namely, the classical LOCALMAXCUT problem.

Definition 5.7. LOCALMAXCUT is the problem of finding a locally maximal cut of a weighted, undirected graph $G = (V, E)$, $n = |V|$, $m = |E|$ with weights w_e for $e \in E$. A cut is a partition (V_1, V_2) of V and the neighbours of a cut are given by the FLIP neighbourhood, in which any single vertex can move from one partition block to the other. The weight of a cut is given by the sum of all edge weights across the cut, i.e.

$$W(V_1, V_2) = \sum_{e \in E(V_1, V_2)} w_e, \quad \text{where} \quad E(V_1, V_2) = \{uv \in E \mid u \in V_1, v \in V_2\}$$

is the cut-set of the cut (V_1, V_2) and we call the edges $e \in E(V_1, V_2)$ cut-edges.

LOCALMAXCUT is PLS-complete, as shown in [SY91]; this is still the case for graphs whose degree is bounded as low as five [ET11]. Independently of this reduction, [MT10] have shown that even when the degree is bounded by four, there exist instances and initial cuts such that every local search sequence is exponentially long.

Theorem 5.8. *The expected number of iterations of smoothed LOCALMAXCUT with weights w_e drawn from independent distributions $f_e : [-1, 1] \rightarrow [0, \phi]$ is bounded by $O(2^{\Delta(G)} nm^2 \phi)$, where $\Delta(G)$ denotes the maximum degree of the graph.*

Proof. To formulate this problem as a CLO problem, we define a binary ($M = 1$) variable s_e in the cost-part for every edge $e \in E$ in the graph of the LOCALMAXCUT instance, thus $v = m$. For the CLO configuration $s(V_1, V_2)$ induced by a cut (V_1, V_2) , we set $s_e = 1$ if and only if the edge e is a cut-edge and 0 otherwise. Even though the vertex partition can be efficiently recovered from this representation (up to the arbitrary assignment of connected components to either part of the partition), we disregard the discussion and represent the partition explicitly in the non-cost part instead, using a binary string of length $\bar{v} = n$ that indicates which vertices belong to V_1 . The edge weights are translated directly to the cost coefficients by $c_e = w_e$ (which also makes smoothed LOCALMAXCUT correspond directly to the smoothed CLO problem) and the neighbourhood relation is derived from the original problem as well to be in one-to-one correspondence. This neighbourhood is efficiently searchable and the cost of the CLO configuration induced by a cut is equal to the weight of the cut in LOCALMAXCUT.

Now, we apply Theorem 3.2 on the smoothed CLO problem defined above. We derive the remaining properties in the following. The covering uses a singleton coordinate cover $\mathcal{I} = \{\{e\}, e \in E\}$. For the transition cover, \mathcal{E} comprises clusters $E(v)$ for every vertex $v \in V$, capturing all transitions where the vertex v moves from one side to the other, and thus $\lambda = n$. When vertex v moves, all edges that are not incident to v either stay cut-edges or non-cut-edges, so the only components possibly affected are those of edges to neighbours of v . Therefore, we have $\beta = \Delta(G)$. Moreover, all edges that are incident to v do in fact change their configuration: Every edge to a neighbour of v in the same partition block as v before the transition is a cut-edge afterwards and also the other way around. Alas, we can choose $\mu = 2$, as the difference $s'_e - s_e$ for every affected edge e is always ± 1 . Thus, Theorem 3.2 provides a bound on the expected number of iterations of $3 \cdot 2^{\Delta(G)} n \cdot m^2 \cdot \phi$. \square

The result also applies with little modification to LOCALMAX- k -CUT.

Definition 5.9. LOCALMAX- k -CUT is the problem of finding a locally maximal cut of a weighted, undirected graph $G = (V, E)$, $n = |V|$, $m = |E|$ with weights w_e for $e \in E$. A cut is a partition (V_1, \dots, V_k) of V and the neighbours of a cut are given by the FLIP neighbourhood, in which a single vertex can move from one partition block to another. The weight of a cut is given by the sum of all edge weights across the k -cut, i.e.

$$W(V_1, \dots, V_k) = \sum_{e \in E'} w_e, \quad \text{where} \quad E' = \bigcup_{i < j \leq k} E(V_i, V_j),$$

and we call the edges $e \in E'$ cut-edges.

This generalization only has a minor impact on our bound, where the exponential part of the bound increases from $2^{\Delta(G)}$ to $3^{\Delta(G)}$. Note that compared to the result by [BCC21], the parameter k does not appear in our bound.

Theorem 5.10. *The expected number of iterations of smoothed LOCALMAX- k -CUT with densities bounded by ϕ is bounded by $O(3^{\Delta(G)}nm^2\phi)$, where $\Delta(G)$ denotes the maximum degree of the graph.*

Proof. The proof is analogous to the proof of Theorem 5.8, except for the changes outlined in the following. The non-cost part now needs to encode which of the k partition blocks every vertex belongs to. It is straightforward to do this using a binary vector of polynomial size. The transition clusters remain structurally the same as $E(v), v \in V$, and each of these sets covers all transitions of a vertex between any of the partition blocks.

The argument, that only the incident edges are affected still holds, but for $k > 2$ not all neighbouring edges of a moving vertex must change between being a cut-edge or not. Therefore, we can bound $\mu = 3$ by observing that in this binary CLO problem, the difference $s'_e - s_e$ for any cost component can only take a value from $\{-1, 0, 1\}$. Thus we conclude with an application of Theorem 3.2 to achieve a bound on the expected number of iterations of $3 \cdot 3^{\Delta(G)}n \cdot m^2 \cdot \phi$. \square

The family of instances for which we prove smoothed efficiency are PLS-hard. In particular, in Appendix D we prove that LOCALMAX- k -CUT is PLS-hard for any $k \geq 2$ and even when the input graph has maximum degree $k+3$. LOCALMAXCUT is its special case $k = 2$, which was already known by [ET11] to be PLS-hard. This means that, unless PLS = FP, there is no worst-case polynomial time algorithm to solve LOCALMAX- k -CUT for any $k \geq 2$.

Apart from the aforementioned conditional intractability, there are results on the unconditional worst-case performance of the standard local search algorithms for LOCALMAXCUT. Monien and Tscheuschner [MT10] constructed instances for which there exists a configuration (i.e., a cut) such that, if the standard local search algorithm starts from there, it will need exponential time, regardless of what path of improvements it chooses.

Theorem 5.11 ([MT10], Theorem 5). *For every $n \in \mathbb{N}^*$, there is a graph on $O(n)$ vertices with maximum degree four, and an initial cut s_0 , such that every sequence of improving FLIP moves starting from s_0 and ending in a local optimum contains a vertex that is moved 2^{n+1} times.*

In contrast, our Theorem 5.8 shows that, under smoothness, the expected number of iterations of standard local search is polynomial even for graphs with logarithmically large degree. Furthermore, in Theorem 5.10 we prove that this still holds for the general problem LOCALMAX- k -CUT with any $k \geq 2$.

5.3.1 Equivalent Problems

The following problems are straightforward equivalent reformulations of LOCALMAX- k -CUT, which allows us to derive analogous smoothness results from Theorem 5.8 and Theorem 5.10 by just restating the problem as its LOCALMAX- k -CUT counterpart instead of applying the CLO results to them directly.

Party Affiliation Games In these games, as outlined by e.g. [BCK10], players correspond to vertices in the graph $G = (V, E_f \cup E_e)$ and choose one of two sides of a cut (one of two *parties*). Each edge in the weighted game graph is considered either a “friendly” or an “enemy” edge, denoted by $e \in E_f$ or $e \in E_e$. The weights w_e are non-negative. The utility of a player is given as the sum of edge weights of friendly edges within the same party and of enemy edges across the two parties. If we denote the parties by V_1, V_2 and use the cut notation from above, the game admits a potential function for a strategy profile s as

$$\Phi(s) = \sum_{e \in E_f \setminus E_f(V_1, V_2)} w_e + \sum_{e \in E_e(V_1, V_2)} w_e.$$

This puts the problem of finding a PNE in this game into PLS by using best-response dynamics. By subtracting $\sum_{e \in E_f} w_e$ from this function, we achieve another potential

$$\Phi'(s) = - \sum_{e \in E_f(V_1, V_2)} w_e + \sum_{e \in E_e(V_1, V_2)} w_e.$$

It is immediate that by changing the signs of the weights of all edges in E_f , the resulting problem can be considered an instance of LOCALMAXCUT. Therefore, we state the following result.

Corollary 5.12. *The expected number of steps in a best-response dynamic of a smoothed party affiliation game with weights w_e drawn from independent distributions $f_e : [0, 1] \rightarrow [0, \phi]$ is bounded by $O(2^{\Delta(G)} nm^2 \phi)$, where $\Delta(G)$ denotes the maximum degree of the graph.*

Note that the assumption of non-negative weights and the separation of the edge set into friendly and enemy edges simply corresponds to the sign of the weights in the corresponding LOCALMAXCUT instance. The PLS-hardness for this problem is usually attributed to [SY91] due to the equivalence with LOCALMAXCUT. Following the PLS-hardness discussion for LOCALMAXCUT above, we deduce that even for graphs with constantly bounded degree, party affiliation games are PLS-hard.

Symmetric Additively Separable Hedonic Games Again, we consider the players as vertices in a weighted graph $G = (V, E)$. In these games, players partition themselves into sets V_1, \dots, V_n , $n = |V|$, and achieve utility according to the sum of weights of edges within their respective party. If we denote by $E(V_i)$ the set of edges of the subgraph induced by V_i , this game has a potential function for a strategy profile s given by

$$\Phi(s) = \sum_{i \in [n]} \sum_{e \in E(V_i)} w_e.$$

Therefore, the problem NASHSTABLE (see [GS19]) of finding a PNE (also called a Nash-stable outcome in this setting due to the existence of alternative stability properties) for these games is in PLS using best-response dynamics. By subtracting $\sum_{e \in E} w_e$ from this function and denoting $E' = \bigcup_{i < j \leq n} E(V_i, V_j)$, we achieve another potential given by

$$\Phi'(s) = - \sum_{e \in E'} w_e.$$

By changing the sign of all weights, this is precisely the potential function of the LOCALMAX- k -CUT problem on the same graph. In other words, working on the same graph, asking to find a PNE in these games is equivalent to finding a solution to LOCALMAX- k -CUT for $k = n$. Thus, we state the following bound.

Corollary 5.13. *The expected number of steps in a best-response dynamic of a smoothed symmetric additively separable hedonic game with weights w_e drawn from independent distributions $f_e : [-1, 1] \rightarrow [0, \phi]$ is bounded by $O(3^{\Delta(G)} nm^2 \phi)$, where $\Delta(G)$ denotes the maximum degree of the graph.*

While the discussion in Appendix D can be used to confirm the PLS-hardness result for this problem first given by Gairing and Savani [GS19], neither suffices to immediately deduce hardness results for bounded degree graphs. To our knowledge, the complexity in the bounded degree case is open.

5.4 Network Coordination Games

Network Coordination Games are a special class of potential games (see [MS96]), where players want to coordinate with each other locally with respect to some graph structure. The game is given by an undirected graph $G = (V, E)$, $n = |V|$, $m = |E|$, where V represents the set of players. Each player $v \in V$

chooses one of k actions and plays a separate matrix coordination game with each of their neighbours simultaneously, i.e. the same action is chosen in each of these games. For every edge $uv \in E$ of the graph, there is a $k \times k$ payoff matrix given that we denote by A_{uv} , according to which the players receive their partial payoffs of the game on this edge. A player's total payoff is then given by the sum of payoffs over all the games on their incident edges; for a strategy profile σ , player u 's payoff is $\text{payoff}_u(\sigma) = \sum_{v \in \Gamma(u)} A_{uv}(\sigma_u, \sigma_v)$, where $\Gamma(u)$ is the set of neighbours of u in the graph G . Note that we specify each matrix only once for each edge, so $A_{vu} = A_{uv}^T$.

The potential structure of the game implies that, just like in congestion games, better-response dynamics converge to a PNE. The dynamics are derived from the PNE definition; a single player can deviate unilaterally to increase their payoff. The local search neighbourhood is thus given by all possible single player deviations and a solution is a local optimum of the potential with respect to this neighbourhood, or equivalently a PNE of the game. The potential function is given by

$$\Phi(\sigma) = \sum_{uv \in E} A_{uv}(\sigma_u, \sigma_v) = \frac{1}{2} \sum_{v \in V} \text{payoff}_v(\sigma)$$

where σ denotes a strategy profile and $A_{uv}(i, j)$ denotes the entry at coordinates (i, j) of the $k \times k$ payoff matrix in the game between u and v . The problem NETCOORDNASH of finding a PNE in a Network Coordination Games is PLS-complete by a reduction from LOCALMAXCUT due to [CD11].

In Boodaghians, Kulkarni, and Mehta [BKM20], the authors conducted smoothed analysis on NETCOORDNASH. Their smoothness model draws each matrix entry for each of the games from independent distributions $f_{uv,i,j} : [-1, 1] \rightarrow [0, \phi]$, where $uv \in E$ denotes an edge in the graph and i, j refer to the entry of the matrix A_{uv} at coordinates (i, j) . The authors have established that, in general graphs, the expected running time of a better-response algorithm under this smoothness model is bounded by $\phi \cdot (nk)^{O(k \log(nk))}$.

We investigate this problem by introducing the degree of the graph as a new parameter, which allows us to establish polynomial smoothed complexity for graphs with constantly bounded degree. This approach allows us to achieve a bound that only depends polynomially on k and the graph size, and the only property that appears in its exponent is the degree of the graph. For instances with either (a) constant degree and arbitrary k , or (b) logarithmic degree and constant k , we provide polynomial smoothed time bounds, which improves upon the bounds induced by [BKM20] for those instances.

Note that the PLS-hardness reduction for Network Coordination Games from LOCALMAXCUT in [CD11] preserves the graph structure of the LOCALMAXCUT instance in the game graph. As LOCALMAXCUT is PLS-complete even for graphs of maximum degree 5 (see Section 5.3), our result establishes another PLS-complete problem class with polynomial smoothed complexity: NETCOORDNASH on graphs with bounded degree. We remark here that the PLS-hardness reduction in the aforementioned paper is tight, in the sense of [SY91] (see Section 2). Since the standard local search of LOCALMAXCUT with maximum degree 5 has a starting configuration whose all improvement paths to solutions are exponentially long (see our discussion after the proof of Theorem 5.10), the same holds for the standard local search of the NETCOORDNASH instances we study here. Our smoothed analysis results show that, contrary to the worst-case performance, these algorithms terminate after polynomially many steps in expectation, even for logarithmically large degree of the game graph.

Theorem 5.14. *The expected running time of better-response dynamics for smoothed Network Coordination Games on a graph $G = (V, E)$ with costs drawn from independent distributions $f_{uv,i,j} : [-1, 1] \rightarrow [0, \phi]$ for every $uv \in E, i, j \in [k]$ is bounded by $O(k^{4\Delta(G)+6}nm^2\phi)$, where $\Delta(G)$ is the degree of the game graph G .*

Proof. We define a binary CLO problem ($M = 1$) to represent NETCOORDNASH. The structure of the cost part corresponds to the different games played (the edge set E of the game graph) and their matrix entries: for every edge $uv \in E$ and every entry of its payoff matrix, $i, j \in [k]$ we define a cost component

$s_{uv,i,j}$ with its corresponding cost coefficient $c_{uv,i,j} = A_{uv}(i, j)$. The size of the cost part therefore is $v = mk^2$.

From each strategy profile σ of the game we can induce a CLO configuration $s(\sigma)$. This configuration follows an indicator structure within the components for a fixed edge uv : we set $s_{uv,i,j}(\sigma) = 1$ if and only if $\sigma_u = i \wedge \sigma_v = j$ and $s_{uv,i,j}(\sigma) = 0$ otherwise. Thus, within every set of the k^2 components of a fixed edge uv , there is exactly one entry of 1 for every valid CLO configuration. The CLO cost of these induced CLO configurations $s(\sigma)$ is equal to the value of the potential $\Phi(\sigma)$ of NETCOORDNASH. From this representation, we can recover the strategy profile of the game in a straightforward way and therefore also compute all neighbouring configurations according to the better-response dynamics efficiently.

Next, we establish all properties for an application of [Theorem 3.2](#). The assumed smoothness model on the matrix entries translates directly to our smoothed CLO model, as every cost coefficient $c_{uv,i,j}$ is equal to a single random variable $A_{uv}(i, j)$ in smoothed NETCOORDNASH. In the following, we specify the covering $(\mathcal{E}, \mathcal{I})$ and its parameters $\lambda = nk(k-1)$, $\mu = k^4$, $\beta = \Delta(G)$. Let

$$\mathcal{I} = \{I_{uv}, uv \in E\}, \quad \text{where} \quad I_{uv} = \{(uv, i, j) \mid i, j \in [k]\}$$

is the coordinate cluster of all components associated with edge uv , i.e. the game between u and v . As outlined above, every CLO configuration has exactly one entry of 1 within the coordinates of any I_{uv} . Therefore, when considering a pair (s, s') of configurations, the expression $s'_{I_{uv}} - s_{I_{uv}}$ can take at most $\mu = k^4$ different values, as both $s_{I_{uv}}, s'_{I_{uv}}$ are one of k^2 configurations each.

Next, we cover the neighbourhood graph by $\mathcal{E} = \{E(u, a, a') \mid u \in V, a \neq a' \in [k]\}$. The transition cluster $E(u, a, a')$ contains all transitions where player u deviates from action a to action a' . There are $\lambda = nk(k-1)$ different transition clusters in this cover. Finally, we establish β : Every deviation of a player u only affects the outcome of the games they participate in. These games are specified by the graph G , therefore each player participates in at most $\Delta(G)$ games. In the CLO problem, this corresponds to a change in the components related to the coordinate clusters $I_{uv}, v \in \Gamma(u)$, so every game a player participates in corresponds to a single coordinate cluster. This gives $\beta = \Delta(G)$, and we conclude the proof by applying [Theorem 3.2](#) to bound the expected number of iterations by

$$3 \cdot (k^4)^{\Delta(G)} nk(k-1) \cdot (mk^2)^2 \cdot \phi \leq 3k^{4\Delta(G)+6} nm^2 \phi.$$

□

5.5 Weighted Set Problems

In this section we discuss a selection of standard weighted set problems. The global decision versions of these problems are NP-complete even for the unweighted case, dating back to [\[GJ90\]](#). The discussion of the local optimization versions of the problems in our context is inspired by [\[DS10a\]](#), who classified all of them as PLS-hard by tight PLS-reductions from Maximum Constraint Assignment MCA- (p, q, r) , a problem shown to be PLS-hard in [\[DM13\]](#) for fixed parameters $(p, q, r) \in \{(3, 2, 3), (2, 3, 6), (6, 2, 2)\}$. This section's focus is on several weighted set problems and not on MCA- (p, q, r) . For our smoothed analysis results on the latter problem itself, see [Appendix F](#).

It is important to highlight the fact that all the problems in this section have been shown in [\[DS10a\]](#) to be PLS-hard via tight PLS-reductions (see [Section 2](#)). The starting problem in the chain of reductions is Circuit Flip [\[JPY88\]](#). Since the standard local search algorithm for Circuit Flip has exponentially long sequences in the worst case (see [\[Yan97, Lemma 12\]](#)), by our argumentation right after [Theorem 4.4](#), the standard local search algorithms for all the weighted set problems we study here have a configuration from which all improving sequences to solutions have exponential length. In contrast, we prove that under smoothed analysis, these algorithms terminate after polynomially many iterations in expectation.

5.5.1 Weighted 3D-Matching

An instance of $W3DM-(p, q)$ is given by an integer n and a weight function $w : [n]^3 \rightarrow \mathbb{R}$ on triples. A 3-dimensional matching is a subset $S \subseteq [n]^3$, $|S| = n$ of triples $T = (T_1, T_2, T_3)$ that covers all values in every component; that is, for all $i \in [n]$ and $k \in \{1, 2, 3\}$ there exists a $T \in S$ such that $T_k = i$. For the sake of intuition, we call the first component the *boys*, the second the *girls* and the third one *homes* and think of the problem as finding an allocation that puts a boy and a girl together into their respective home for each of the triples. The weight of a matching S is given by the sum of weights of the triples in use, i.e. $W(S) = \sum_{T \in S} w(T)$. A solution of $W3DM-(p, q)$ is a matching that is locally maximal in weight regarding the (p, q) -neighbourhood, defined as follows.

The neighbourhood of a matching S is given by all matchings created by replacing up to p triples by other (exactly p many) triples. Additionally, we require that the total number of boys or girls relocated into new homes is at most q . Note that any triple that is being replaced moves the boy or the girl to a new home, but not necessarily both; the parameter q can be thought of as restricting how often both of them may be relocated. When p is considered a constant, this neighbourhood can be efficiently searched. Furthermore, the second restriction q does not appear in our bound, but we have stated the problem as it appears in [DS10a], who established PLS-completeness for all $p \geq 6, q \geq 12$.

Theorem 5.15. *The expected number of iterations of local search for smoothed $W3DM-(p, q)$, where the weights $w(T)$ for each triple $T \in [n]^3$ are drawn from independent distributions with density $f_T : [0, 1] \rightarrow [0, \phi]$, is bounded by $O(n^{6p+6}\phi)$.*

Proof. We formulate this problem as a CLO problem while maintaining the weight structure of $W3DM-(p, q)$ in the CLO cost coefficients and then apply Theorem 3.2. Consider a binary ($M = 1$) CLO problem where the cost part corresponds to the set of all triples, i.e. we choose $v = n^3$ and the cost variables are given by $s_T, T \in [n]^3$. The cost coefficients are set to the weights of the corresponding triples, $c_T = w(T)$. The CLO configuration $s(S)$ induced by a matching S has $s_T(S) = 1$ whenever $T \in S$ and $s_T(S) = 0$ otherwise. Thus, the weight of a matching corresponds to the cost of the induced CLO configuration. The neighbourhood is derived directly and efficiently from the original instance.

To apply Theorem 3.2, we establish a covering $(\mathcal{E}, \mathcal{I})$ and the parameters $\lambda = (n^3)^p, \beta = p, \mu = 1$ in the following. We start with the singleton coordinate cover, $\mathcal{I} = \{\{T\}, T \in [n]^3\}$. The neighbourhood graph is covered by transition clusters $E(T_1, \dots, T_p, U_1, \dots, U_p), T_i, U_i \in [n]^3$. Each of these clusters contains all valid transitions where the triples T_i are replaced by U_i in the matching, regardless of the remaining triples in the matching. There are at most $\lambda = (n^3)^{2p}$ such transition clusters when accounting every possible choice of the $2p$ triples. Note that some triples $T_i = U_j$ may coincide to account for fewer than k tuple replacements in the solution, and that each permutation within the triples T_i or U_j refers to the same set of transitions.

Now, we investigate the structure of this covering to determine the remaining parameters. Within the transition cluster $E(T_1, \dots, T_p, U_1, \dots, U_p)$, the only cost components that can change their assignment are the ones corresponding to a triple T_i or U_j . Because we add and remove at most p triples each, we therefore have $\beta = 2p$. More importantly, we know that for any choice of $(s, s') \in E(T_1, \dots, T_p, U_1, \dots, U_p)$ for a fixed transition cluster, the difference vector $s' - s$ is uniquely determined: Outside of the selected tuples, there is no change in the configuration. When $T_i \neq U_j$ or $U_i \neq T_j$ for all $j = 1, \dots, k$ (i.e. T_i is actually removed from and U_i is actually being added to the matching), we know that $s_{T_i} = 1$ before the transition and $s'_{T_i} = 0$ afterwards, as well as $s_{U_i} = 0$ and $s'_{U_i} = 1$, so the difference in these components has only one possible value. Otherwise, when some $T_i = U_j$ coincide and thus T_i stays in the matching, we know that $s_{T_i} = s'_{T_i}$ and this component difference is again fixed as 0. Thus, we have settled $\mu = 1$.

Applying Theorem 3.2 gives a bound on the expected number of steps of $3 \cdot (n^3)^{2p} \cdot (n^3)^2 \cdot \phi$. \square

5.5.2 Exact Cover by 3-Sets and Set-Cover

The similar structure of the following cover problems allows for shared analysis. Both use the k -differ neighbourhood, where the neighbours may differ in up to k selected sets for the cover. For constant k , this is an efficiently searchable neighbourhood. First, let us define them separately:

- Exact Cover by 3-Sets (X3C- k) is the problem of finding a maximum weight exact cover of a finite set \mathcal{B} with $|\mathcal{B}| = 3q$. An exact cover in this problem is a subset $S \subseteq C = \{C_1, \dots, C_n\}$, $|S| = q$ of sets $C_i \subseteq \mathcal{B}$ with $|C_i| = 3$, such that $\bigcup_{C_i \in S} C_i = \mathcal{B}$. Thus, the sets in any S must be disjoint and S is a partition of \mathcal{B} into sets of size 3. With every set C_i there is given a weight w_i and we want to maximize the total weight $w(S) = \sum_{i: C_i \in S} w_i$.

The neighbours of an exact cover S are given by adding or removing at most k sets in total to/from S so as to generate another exact cover. Because we know that $|S| = q$ is fixed, this means we replace up to $k/2$ sets from the solution by the same number of other sets from C .

X3C- k is PLS-complete due to [DS10a] for all $k \geq 12$.

- Set-Cover (SC- k) is the problem of finding a minimum weight cover from a collection of sets $C = \{C_1, \dots, C_n\}$, $C_i \subseteq \mathcal{B}$, i.e. a subset $S \subseteq C$ such that $\bigcup_{C_i \in S} C_i = \mathcal{B}$. Each set C_i has some weight w_i . The cost of a cover S is given by $c(S) = \sum_{i: C_i \in S} w_i$.

The neighbourhood of a cover is given by changing for up to k sets whether they are being used. This means that for any $k_1 + k_2 \leq k$ there may be k_1 sets added to and k_2 sets removed from S in order to construct a new cover of \mathcal{B} .

SC- k is PLS-complete due to [DS10a] for all $k \geq 2$.

Theorem 5.16. *The expected number of iterations of local search for X3C- k and SC- k , when the weights w_i for each set $C_i \in C$ are drawn from independent distributions with density $f_i : [0, 1] \rightarrow [0, \phi]$, is bounded by $\mathcal{O}(3^k n^{k+2} \phi)$.*

Proof. We formulate these problems as CLO problems, while directly translating the random distributions of the weights to the distributions of smoothed CLO costs and apply Theorem 3.2.

Consider a binary ($M = 1$) CLO problem where the cost part corresponds to the elements of the collection C , i.e. $v = n$ and every set $C_i \in C$ has a corresponding cost component s_i . The cost coefficients are set to the weights of the respective sets, $c_i = w_i$. The CLO configuration $s(S)$ induced by a cover S for X3C- k and SC- k has $s_i(S) = 1$ whenever $C_i \in S$ and $s_i(S) = 0$ otherwise. Thus the weight of any cover is equal to the cost of its induced CLO configuration. The CLO neighbourhood is derived directly and efficiently from the original instance of the problem.

To apply Theorem 3.2, we establish a covering $(\mathcal{E}, \mathcal{I})$ and the parameters $\lambda = n^k$, $\beta = k$, $\mu = 3$ in the following. We start with the singleton coordinate cover, $\mathcal{I} = \{\{i\}, i = 1, \dots, n\}$. The transition cover \mathcal{E} comprises clusters $E(D_1, \dots, D_k)$, $D_i \in C$, each indicating which D_i are changed, regardless of its initial status; this means that if $D_i \in S$, it is being removed, and otherwise it is being added to the cover. There are at most $\lambda = n^k$ such clusters by accounting for every possible choice of D_i in k components. Note that the same D_i may be chosen multiple times to account for changes of fewer than k sets in the cover and that each permutation within the D_i refers to the same transition cluster. Clearly, with each of these clusters we affect only the cost configuration of the sets D_i that we change. Because we change at most k sets in the transition, we have $\beta = k$.

Finally, as these are binary problems, we know that $s'_i - s_i \in \{-1, 0, 1\}$ and therefore $\mu = 3$. By Theorem 3.2 this gives a bound on the expected number of steps of $3 \cdot 3^k n^k \cdot n^2 \cdot \phi$. \square

5.5.3 Hitting Set

An instance of Hitting Set (HS- k) is given by a collection of sets $C = \{C_1, \dots, C_n\}$, $C_i \subseteq \mathcal{B}$ where \mathcal{B} is a finite ground set, weights w_i for each $C_i \in C$ and some integer m . A solution is a locally maximal subset $S \subseteq \mathcal{B}$ that is bounded in size $|S| \leq m$. The weight of a subset S is given by $w(S) = \sum_{i: S \cap C_i \neq \emptyset} w_i$, i.e. the sum of all sets in C that are *hit* by the subset S , which means that at least a single element of C_i appears in S .

The local search neighbourhood of a subset S is given by changing for at most k elements from \mathcal{B} whether they occur in S or not. This means for any $k_1 + k_2 \leq k$ we can add k_1 new elements to and remove k_2 elements from S to create one of its neighbours.

HS- k is PLS-complete due to [DS10a] for all $k \geq 1$ (note that the proof can only be found in the full version of the paper [DS10b]). In fact, the reduction allows for a slight adaptation that gives PLS-completeness of HS- k for every $k \geq 1$ even when the maximum number of occurrences of an element within C is bounded. In our proof deferred to Appendix E we reduce from MAXSAT/1-FLIP with bounded variable occurrence (every variable appears in at most \tilde{B} many clauses) instead of the standard MAXSAT/1-FLIP; the additional property translates to the fact that the number of sets that contain a given element is bounded by $\tilde{B} + 1$, and therefore also by a constant. MAXSAT/1-FLIP with bounded variable occurrence was shown to be PLS-complete in [Kre89] and [Kla96] via tight PLS-reductions.

The starting problem of these reduction paths needs exponentially many steps to find a solution in the worst case, under its standard local search algorithm. Therefore, the HS-1 instances we study here inherit this property, meaning that the standard local search algorithm has exponentially long improvement sequences in the worst case (see the relevant discussion in Section 2). However, we prove that under smoothness, even in significantly wider families of instances, standard local search terminates after polynomially many iterations in expectation.

Theorem 5.17. *The expected number of iterations of local search for HS- k , when the weights w_i for each set $C_i \in C$ are drawn from independent distributions with density $f_i : [0, 1] \rightarrow [0, \phi]$, is bounded by $O\left(3^{kB} |\mathcal{B}|^k n^2 \phi\right)$. The parameter B denotes the maximum number of occurrences of an element within C , i.e., $B = \max_{b \in \mathcal{B}} |\{C_i \in C \mid b \in C_i\}|$.*

Proof. We formulate this problem as a CLO problem apply Theorem 3.2 to its smoothed counterpart:

Consider a binary ($M = 1$) CLO problem, where the cost part corresponds to the elements of the collection C , i.e. $v = n$ and every set C_i corresponds to a cost component s_i . A subset S induces a CLO configuration $\mathbf{s}(S)$ with $s_i(S) = 1$ if and only if $\exists b \in S : b \in C_i$, $s_i = 0$ otherwise. The cost coefficients are set to $c_i = w_i$ and thus, the cost of $\mathbf{s}(S)$ corresponds to the weight of S in the HS- k instance. The subset S itself can be maintained in the non-cost configuration as an indicator vector over the ground set \mathcal{B} , i.e. $\bar{v} = |\mathcal{B}|$. Also, the neighbourhood of the CLO problem is efficiently derived from the neighbourhood of the HS- k instance and the smoothed costs of the HS- k instance correspond to the derived smoothed CLO model because the definition of the CLO cost coefficients.

To apply Theorem 3.2, we need to establish a covering $(\mathcal{E}, \mathcal{I})$ and the parameters $\lambda = |\mathcal{B}|^k$, $\beta = kB$, $\mu = 3$ in the following. We start with the singleton coordinate cover, $\mathcal{I} = \{\{i\}, i = 1, \dots, n\}$. The neighbourhood graph is covered by transition clusters $E(b_1, \dots, b_k)$, $b_j \in \mathcal{B}$, each indicating for which elements in \mathcal{B} the subset changes, regardless of the initial subset. Every such cluster therefore contains all valid transitions from an initial subset S , where each b_j is removed from the set if S already contains b_j , or added to the set if it doesn't. There are at most $\lambda = |\mathcal{B}|^k$ different transition clusters, accounting for each choice of b_j in k components. Note that we do allow the same $b \in \mathcal{B}$ to appear multiple times in order to represent changes of fewer than k elements. Each permutation of the b_j gives the same transition cluster.

Lastly, we settle the properties β and μ for a fixed transition cluster $E(b_1, \dots, b_k)$, $b_j \in \mathcal{B}$. When changing any element b in the subset, this can only affect the cost components s_i of the C_i , where $b \in C_i$;

the weight contribution of all other sets is independent of b . Because, by assumption, each b occurs in at most B different sets C_i , we know that in total there are at most $\beta = kB$ components affected by any change in k different elements. Finally, within the configuration change in a transition for any of the singleton coordinate clusters, we simply use that the CLO problem is binary, and therefore $s'_i - s_i \in \{-1, 0, 1\}$ in any case, directly giving $\mu = 3$. Applying [Theorem 3.2](#) with these parameters yields a bound on the expected numbers of iterations of $3 \cdot 3^{kB} |\mathcal{B}|^k \cdot n^2 \cdot \phi$. \square

6 Conclusion and Open Problems

In this work, we have introduced the class of CLO problems, which captures many relevant problems in PLS and allows us to define a corresponding smoothed version of each problem in a unifying way. This framework allowed us to establish a black-box tool to give a bound on the smoothed running time for such problems by merely supplying some structural parameters of the problem, which can often be found naturally. We have made use of this newly established framework by applying it to a multitude of existing problems. Thereby, we rederive some existing results on the smoothed complexity of these problems. Furthermore, the framework allows us to apply the same structural ideas to problems that have not been discussed in the field of smoothed complexity before. For all these problems, we complement the positive results of smoothed running time with negative worst-case results. In particular, we show that they are PLS-hard, and additionally, that the local-search algorithms used for the positive results actually admit exponential time instances.

A variety of more recent results in the field of smoothed analysis apply a technique that is not captured by our framework: instead of considering a single step in the improvement dynamics, the path through the neighbourhood graph is structurally clustered into sequences, for which the *total* improvement is bounded by a problem-specific rank-based argument. A natural question that arises is whether our framework can be generalized to capture this technique in a streamlined fashion as well, in order to unify even more of the landscape of smoothed analysis or establish better bounds than we can currently provide. This open question is related to [Point 1](#) of our discussion about the limitations of our key technical tool ([Theorem 3.2](#)), in [Section 3.1](#).

There have been some observations on smoothed models, where considering a fully parameterized instance enabled improving the bounds on the smoothed running times of the standard local search algorithm compared to the general instances. An example of this phenomenon is LOCALMAXCUT for complete graphs for which smoothed polynomial time has been established in [\[ABPW17\]](#), whereas for general graphs there is no better than quasi-polynomial smoothed time known. The former kind of graph input is fully parameterized (all possible edges are considered to undertake smoothed perturbation), as opposed to the latter. We conjecture that there is a unifying framework of a similar flavour to ours which would work for these fully parameterized instances of PLS-complete problems. This is related to [Point 2](#) of our discussion in [Section 3.1](#).

In this paper, we show that the problem of computing (pure, exact) equilibria in congestion games admits a smoothed polynomial running time algorithm for special classes of general and network congestion games (see [Theorem 4.2](#) and [Theorem 4.8](#), respectively). However, outside of these restricted instances of the problem, the complexity remains unresolved. An efficient algorithm for the computation of *approximate* equilibria (FPTAS) under smoothed analysis, was recently presented in [\[Gia24\]](#).

To the best of our knowledge, the published literature on smoothed analysis of problems in PLS has been limited to establishing upper bounds on the smoothed running time of the standard local search algorithm. The question that arises here is whether we can show that a problem remains PLS-hard even if its input is perturbed (within some range) according to the smoothed model. From the positive side, an interesting question is to characterize the class of PLS-complete problems whose standard local search algorithm terminates in smoothed polynomial time.

Acknowledgements

Y. Giannakopoulos is grateful to Diogo Poças for many useful discussions and inspiration during the early stages of this project. The authors thank Rahul Savani for pointing towards the inclusion of hedonic games to this paper.

A Omitted Proofs

A.1 Proof of Lemma 3.3

First, we show that the statement holds when \mathcal{J} is a partition instead of a cover, i.e. we assume the sets in \mathcal{J} are pairwise disjoint and $J \subseteq I$ for every $J \in \mathcal{J}$. Let us fix a subset of transitions $T \subseteq E$. Denote by $\llbracket M \rrbracket_{\pm} = \{-M, \dots, -1, 0, 1, \dots, M\}$ the set of possible values for the transition difference $s_i - s'_i$ in every component. The set $\text{range}_I(T)$ thus contains vectors from $\llbracket M \rrbracket_{\pm}^I$ and $\text{range}_J(T) \subseteq \llbracket M \rrbracket_{\pm}^J, J \in \mathcal{J}$ analogously. We identify each $\mathbf{x} \in \llbracket M \rrbracket_{\pm}^I$ with the family $(\mathbf{x}_J)_{J \in \mathcal{J}}$, as every index $i \in I$ is contained in exactly one of the $J \in \mathcal{J}$, and observe that $(\mathbf{x}_J)_{J \in \mathcal{J}} \in \prod_{J \in \mathcal{J}} \llbracket M \rrbracket_{\pm}^J$.

We prove that $\text{range}_I(T) \subseteq \prod_{J \in \mathcal{J}} \text{range}_J(T)$ under this identification. For the sake of contradiction, suppose that there exists $\mathbf{x} \in \text{range}_I(T)$ with $\mathbf{x} \notin \prod_{J \in \mathcal{J}} \text{range}_J(T)$. Therefore there exists $J \in \mathcal{J}$ such that $\mathbf{x}_J \notin \text{range}_J(T)$. But by definition, $\text{range}_J(T)$ contains \mathbf{x}_J due to the same transition $(\mathbf{s}, \mathbf{s}')$ that generated \mathbf{x} itself, a contradiction.

Now, we show that for $A \subseteq B$ we have $\delta_A(T) \leq \delta_B(T)$. Every $\mathbf{x}, \tilde{\mathbf{x}} \in \text{range}_A(T), \mathbf{x} \neq \tilde{\mathbf{x}}$ are coming from transitions $(\mathbf{s}, \mathbf{s}'), (\tilde{\mathbf{s}}, \tilde{\mathbf{s}}')$. These transitions also provide $\mathbf{x}', \tilde{\mathbf{x}}' \in \text{range}_B(T)$. Since $A \subseteq B$ and for the projections of \mathbf{x}' and $\tilde{\mathbf{x}}'$ on A we have $\mathbf{x}'_A = \mathbf{x}, \tilde{\mathbf{x}}'_A = \tilde{\mathbf{x}}$, it holds that also $\mathbf{x}' \neq \tilde{\mathbf{x}}'$. Thus, all elements in $\text{range}_A(T)$ have distinct representatives in $\text{range}_B(T)$ and therefore $\delta_A(T) \leq \delta_B(T)$.

Combining the observations above proves the main statement: Every cover \mathcal{J} of I can be modified to be a partition \mathcal{J}_p of I by removing every index $i \in I$ that appears more than once from all but a single set and remove every $i \notin I$ entirely. Therefore, every set $A \in \mathcal{J}_p$ has a corresponding superset $B \in \mathcal{J}$, and we conclude

$$\delta_I(T) \leq \prod_{A \in \mathcal{J}_p} \delta_A(T) \leq \prod_{B \in \mathcal{J}} \delta_B(T).$$

A.2 Proof of Lemma 3.4

For an arbitrary (absolutely) continuous real random variable X with density function $f : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ we will denote $M(X) := \text{ess sup}_{x \in \mathbb{R}} f(x)$.¹³ For the random variables in the statement of our Lemma we have $M(X_i) \leq \phi$ for all $i \in [m]$. Let also $Y_i = \xi_i X_i$, for all $i \in [m]$, and

$$Y = \sum_{i \in I^*} Y_i \quad \text{where } I^* = \{i \in [m] \mid \xi_i \neq 0\}.$$

Notice that $I^* \neq \emptyset$, since $\xi = (\xi_1, \dots, \xi_m)$ is nonzero.

Then, it holds that $M(Y_i) = \frac{1}{\xi_i} M(X_i)$ for all $i \in I^*$. Furthermore, since Y_i 's are independent, by Bobkov and Chistyakov [BC14, Remark, p. 105] we know that the following inequality applies:

$$\frac{1}{M^2(Y)} \geq \frac{1}{2} \sum_{i \in I^*} \frac{1}{M^2(Y_i)}.$$

¹³ Here ess sup denotes the *essential supremum* operator (see, e.g., Halmos [Hal74, Sec. 21].); that is, roughly speaking, $M(X)$ is upper-bounding function f *almost everywhere* in its domain, excluding possibly a zero-measure set of points. Formally, $\text{ess sup}_{x \in \mathbb{R}} f(x) := \inf \{c \in \mathbb{R} \mid \mu(\{x \in \mathbb{R} \mid f(x) > c\}) = 0\}$ where μ is the standard Lebesgue measure on \mathbb{R} .

From this we get:

$$M(Y) \leq \sqrt{2} \frac{1}{\sqrt{\sum_{i \in I^*} \frac{1}{M^2(Y_i)}}} = \sqrt{2} \frac{1}{\sqrt{\sum_{i \in I^*} \frac{\xi_i^2}{M^2(X_i)}}} \leq \sqrt{2} \phi \frac{1}{\sqrt{\sum_{i \in I^*} \xi_i^2}} = \sqrt{2} \phi \frac{1}{\|\xi\|_2}.$$

Finally, using the above inequality we get that, for any $\varepsilon \geq 0$,

$$\text{Prob}[0 \leq \xi \cdot X \leq \varepsilon] = \text{Prob}[0 \leq Y \leq \varepsilon] \leq \int_0^\varepsilon M(Y) \leq \varepsilon \cdot \sqrt{2} \phi \frac{1}{\|\xi\|_2}.$$

We will now also derive an alternative bound. Let $i^* \in \arg\max_{i \in I^*} |\xi_i|$, i.e., $|\xi_{i^*}| = \|\xi\|_\infty$. Since $I^* \neq \emptyset$, we know that $\xi_{i^*} \neq 0$. Then, we get:

$$\begin{aligned} \text{Prob}[0 \leq \xi \cdot X \leq \varepsilon] &= \text{Prob}\left[-\sum_{i \neq i^*} \xi_i X_i \leq \xi_{i^*} X_{i^*} \leq \varepsilon - \sum_{i \neq i^*} \xi_i X_i\right] \\ &= \begin{cases} \text{Prob}\left[X_{i^*} \in \left(Z, Z + \frac{\varepsilon}{|\xi_{i^*}|}\right)\right], & \text{if } \xi_{i^*} > 0, \\ \text{Prob}\left[X_{i^*} \in \left(Z - \frac{\varepsilon}{|\xi_{i^*}|}, Z\right)\right], & \text{if } \xi_{i^*} < 0, \end{cases} \end{aligned} \quad (11)$$

where

$$Z = -\frac{1}{\xi_{i^*}} \sum_{i \neq i^*} \xi_i X_i.$$

Since random variables X_{i^*} and Z are independent, both probabilities in (11) can be uniformly bounded, for each realization of Z , by $\frac{\varepsilon}{|\xi_{i^*}|} \phi$.

Merging the two bounds derived above, we finally get the desired

$$\text{Prob}[0 \leq \xi \cdot X \leq \varepsilon] \leq \min\left(\frac{1}{\|\xi\|_\infty}, \frac{\sqrt{2}}{\|\xi\|_2}\right) \varepsilon \phi$$

For nonzero *integral* vectors ξ , we have $\min\left(\frac{1}{\|\xi\|_\infty}, \frac{\sqrt{2}}{\|\xi\|_2}\right) \leq 1$, since $\|\xi\|_\infty \geq 1$.

B PLS-Hardness of Restrained Congestion Games

Here we prove [Theorem 4.4](#), by showing that the problem of finding a PNE in a congestion game is PLS-complete even for constantly-restrained games and for the simplest versions of general, polynomial, and step-function costs. Particularly, in [Corollary B.2](#) and [Theorem B.3](#) we show that even 10-restrained congestion games are PLS-hard under all considered cost function models.

We start by presenting the main reduction from the problem LOCALMAXCUT- d to congestion games with at most two strategies per player and access to at most d resources per strategy. LOCALMAXCUT- d is the restricted version of LOCALMAXCUT (see [Definition 5.7](#)) where the input graph has maximum degree d . Since LOCALMAXCUT-5 is PLS-complete ([\[ET11\]](#)), our reduction shows that congestion games with two strategies per player and access to at most five resources is already PLS-hard. The reduction we present is identical to the one that appears in the textbook of Roughgarden [[Rou16](#), Theorem 19.4] which reduces LOCALMAXCUT to finding a PNE in congestion games. For completeness, we present the reduction here while keeping track of the important parameters of the resulting instance: the number of strategies per player and the number of resources per strategy. Note that within our reduction, the strategies and the local search neighbourhood of the reduced instance correspond one-to-one to the

cuts and neighbourhood of the LOCALMAXCUT- d instance. Therefore, the reduction is tight in the sense of [SY91].

Given a LOCALMAXCUT- d instance on the graph $G = (V, E)$ we will create a congestion game with player set V , where each player corresponds to a vertex of the graph. For each edge $e \in E$ we create two resources r_e, \bar{r}_e . Let the (graph) neighbourhood of a vertex $i \in V$ be denoted by $N(i) := \{j \in V \mid ij \in E\}$. Let us also denote the edges incident to i by $E(i) := \{ij \in E \mid j \in N(i)\}$. Each player $i \in V$ has two strategies, namely, $S_1(i) := \{r_e\}_{e \in E(i)}$ and $S_2(i) := \{\bar{r}_e\}_{e \in E(i)}$. For $e \in E$, the cost for using a resource is

$$\kappa_{r_e}(1) = \kappa_{\bar{r}_e}(1) = 0, \quad \text{and} \quad \kappa_{r_e}(\ell) = \kappa_{\bar{r}_e}(\ell) = w_e, \quad \text{for } \ell \geq 2. \quad (12)$$

Notice that, by construction, at most two players can use the same resource.

We claim that any PNE of the congestion game corresponds to a solution to the initial LOCALMAXCUT- d instance, that is, a cut (V_1, V_2) in the following way: For any player i that plays strategy $S_1(i)$, the corresponding node $i \in V$ is placed in V_1 , otherwise she is placed in V_2 . Consider a PNE of the congestion game, and for the sake of contradiction, suppose there exists a vertex $i \in V_1$ which, if moved to V_2 , yields a greater cut weight. Let $E_1(i) := \{ij \in E \mid j \in V_1\}$ and $E_2(i) := \{ij \in E \mid j \in V_2\}$ be the edges between i and her neighbours in V_1 and V_2 , respectively. Then $\sum_{e \in E_1(i)} w_e - \sum_{e \in E_2(i)} w_e > 0$. Now consider the difference in experienced cost of player i when deviating from $S_1(i)$ to $S_2(i)$. This would be

$$\begin{aligned} \sum_{e \in E_1(i)} \kappa_{\bar{r}_e}(1) + \sum_{e \in E_2(i)} \kappa_{\bar{r}_e}(2) - \sum_{e \in E_1(i)} \kappa_{r_e}(2) - \sum_{e \in E_2(i)} \kappa_{r_e}(1) &= \sum_{e \in E_2(i)} \kappa_{\bar{r}_e}(2) - \sum_{e \in E_1(i)} \kappa_{r_e}(2) \\ &= \sum_{e \in E_2(i)} w_e - \sum_{e \in E_1(i)} w_e \\ &< 0. \end{aligned}$$

This means that player i would decrease her cost if she unilaterally deviated to $S_2(i)$, therefore, the original state was not a PNE, a contradiction. Similarly, we get a contradiction by assuming that there is a vertex $i \in V_2$ which, if moved to V_1 , yields a greater cut weight. Hence, any PNE of the congestion game corresponds to a solution of LOCALMAXCUT- d . In this reduction, the constructed congestion game has step-function costs (with only a single step at $\ell = 2$), implying immediately that it holds for general cost functions (see (12)). Since for any $d \geq 5$ LOCALMAXCUT- d is PLS-complete ([ET11]), the above reduction yields the following.

Theorem B.1. *Finding a PNE, in congestion games with step function or general cost representation, is PLS-complete even when every player has 2 strategies, each consisting of at most 5 resources.*

Recall now that, by construction, $|E(i)| \leq d$ for all $i \in V$. Therefore, $|S_1(i)| = |S_2(i)| \leq d$, which implies that the symmetric difference between any pair of strategies of a player in such congestion games is at most $2d$ (see Definition 4.1). Therefore, we get the following.

Corollary B.2. *Finding a PNE in 10-restrained congestion games, with step-function or general cost representation, is PLS-complete.*

We will now show how they also hold for affine cost functions with nonnegative coefficients, that is, the simplest class of polynomial functions, excluding constant functions. For $e \in E$, let the cost for using the corresponding resource be

$$\kappa_{r_e}(\ell) = \kappa_{\bar{r}_e}(\ell) = \ell \cdot w_e, \quad \text{for } \ell \geq 1, \quad (13)$$

and recall that, by construction, at most 2 players can be using the same resource.

Similarly to the previous reduction, we show that any PNE of this congestion game corresponds to a solution to the initial LOCALMAXCUT- d instance. In particular, for every player i using $S_1(i)$ (resp.

$S_2(i)$) we place its corresponding vertex i in V_1 (resp. V_2). Consider a PNE of the congestion game, and suppose there exists a vertex $i \in V_1$ which, if moved to V_2 , yields a greater cut weight. Therefore, $\sum_{e \in E_1(i)} w_e - \sum_{e \in E_2(i)} w_e > 0$. The difference in experienced cost of player i when deviating from $S_1(i)$ to $S_2(i)$ is

$$\begin{aligned} & \sum_{e \in E_1(i)} \kappa_{\bar{r}_e}(1) + \sum_{e \in E_2(i)} \kappa_{\bar{r}_e}(2) - \sum_{e \in E_1(i)} \kappa_{r_e}(2) - \sum_{e \in E_2(i)} \kappa_{r_e}(1) = \\ &= \sum_{e \in E_1(i)} w_e + \sum_{e \in E_2(i)} 2w_e - \sum_{e \in E_1(i)} 2w_e - \sum_{e \in E_2(i)} w_e \\ &= \sum_{e \in E_2(i)} w_e - \sum_{e \in E_1(i)} w_e \\ &< 0. \end{aligned}$$

So, player i would decrease her cost if she unilaterally deviated to $S_2(i)$, contradicting the assumption that the original state was a PNE. Similarly, we are led to a contradiction by assuming that there is a node $i \in V_2$ which, if moved to V_1 , yields a greater cut weight. Therefore, any PNE of the congestion game with linear cost function corresponds to a solution of LOCALMAXCUT- d , and again, since for any $d \geq 5$ LOCALMAXCUT- d is PLS-complete ([ET11]) we get the following.

Theorem B.3. *Congestion games with affine cost functions are PLS-complete even when every player has 2 strategies, each consisting of at most 5 resources, and each resource being available to at most 2 players. In particular, such games are 10-restrained.*

C PLS-Hardness of Compact Network Congestion Games

Here we prove that computing a PNE in network congestion games is PLS-hard even for games with the simplest versions of general, polynomial, and step cost functions. Our reduction is similar to that of [ARV08], however, our reduction route is simpler: we start from LOCALMAXCUT- d (see introduction of Appendix B) and we reduce to network congestion games via general congestion games. We present it for completeness while remarking how the parameters of interest, namely, the maximum number of best-responses per player, A , and the maximum path-length per player, B (see Definition 4.5), are affected by the corresponding parameters of the initial problem.

We will take the reduction of Appendix B a step further and end up with a network congestion game with at most 2 cost levels per edge-cost; the number of cost levels for a given edge cost function is defined as the number $d + 1$, where d is the number of break-points when the cost function is in its step function representation (see Section 4). The reduction in the proof of Theorem B.1 had as a starting point the PLS-complete problem of computing a local max cut in a LOCALMAXCUT- d instance, so we will be essentially reducing this problem to computing PNEs in our family of network congestion games. Since the latter is also in PLS ([FPT04]), our reduction implies its PLS-completeness. Furthermore, the reduction is tight, since \mathcal{R} (the set of reasonable feasible solutions in tight reductions according to Definition 3.2 in [SY91]) is our set $\prod_{i \in \mathcal{N}} \Sigma_i^*$ of strategy profiles that consist of best-responses.

Let us continue the reduction from the proof of Theorem B.1. In this section we will be calling I the congestion game instance constructed in the aforementioned proof. Let the set of players be $[n]$, and the two strategies of player $i \in [n]$ be $S_1(i)$ and $S_2(i)$, where $|S_1(i)| = |S_2(i)| \leq d$. We now construct the network congestion game instance I' . To avoid confusion with the terms “vertices” and “edges” of the local max cut problems, we use instead “nodes” and “arcs”, respectively, for network congestion games. We therefore also attain the vertex set and edge set notation V and E of the former problem, and denote by N and A the corresponding sets of the latter problem. The set of players in I' is identical to that of I . Let us make a pair of nodes for each player i , an origin node o_i and a destination node d_i . For any pair $i, j \in [n]$ with $i < j$, if edge $ij \in E$, we create a *supernode* U_{ij} as shown in Figure 2.




Figure 1: The graph of a LOCALMAXCUT-3 instance.




Figure 2: Supernode U_{ij} . It admits at most two incoming arcs and at most two outgoing arcs. The arc between the two nodes inside it has cost function $\kappa_{r_{ij}}$ (see (12)). Supernode \bar{U}_{ij} is defined analogously.




Figure 3: Half of the network congestion game instance I' corresponding to the LOCALMAXCUT-3 instance of Figure 1; the other half is symmetric and consists of supernodes \bar{U}_{ij} . Supernodes U_{ij} have a circular shape. The origin nodes have a triangular shape, while the destination nodes have a square shape. The black/dashed arcs are light, while the green/solid ones are heavy.

We will now define the arcs of I' , and let us remark that all the arcs are directed. We have three kinds of arcs: the *regular* ones inside each supernode U_{ij} which have cost function $\kappa_{r_{ij}}$ as defined in (12)¹⁴; the *light* ones that have cost 0 for any number of players using them; and the *heavy* ones between supernodes $U_{ij}, U_{i'j'}$ which have cost $(i' - i + j - 1) \cdot W$, for any number of players using them, where $W > \sum_{ij \in E} w_{ij}$. Consider now all existing supernodes U_{ij} for a fixed i , and let R_i be the increasingly ordered tuple of the elements in $\{j \mid ij \in E \wedge i < j\}$. Note that $i \in [n - 1]$. We create a line graph, called *row i* , by connecting each supernode U_{ij} with a light arc to $U_{ij'}$, where j, j' are two consecutive elements of tuple R_i . Now consider all existing supernodes U_{ij} for a fixed j , and let C_j be the increasingly ordered tuple of the elements in $\{i \mid ij \in E \wedge i < j\}$. Note also that $j \in [n] \setminus \{1\}$. We connect each supernode U_{ij} with a heavy arc to $U_{i'j}$, where i, i' are two consecutive elements of tuple C_j .

For each $i \in [n]$, let $U_{k\ell}$ be a supernode such that either $k = i$ or $\ell = i$, and among such existent supernodes consider $U_{k^-\ell^-}$, where $k^- + \ell^- \leq k + \ell$. Similarly, consider $U_{k^+\ell^+}$, where $k^+ + \ell^+ \geq k + \ell$. We connect via light arcs o_i to $U_{k^-\ell^-}$, and $U_{k^+\ell^+}$ to d_i . Also, if both C_i and R_i are non-empty, let k^*, ℓ^* be the largest element in C_i and the smallest element in R_i , respectively. We connect U_{k^*i} to $U_{i\ell^*}$ via a heavy arc.

We have described half of the construction, where we translated the resources $r_e, e \in E$ of the intermediate congestion game instance I into paths of I' . We have to implement an entirely symmetric construction where we also translate resources \bar{r}_e into paths by adding an identically connected set of supernodes \bar{U}_{ij} for $ij \in E$, which connect to the existing set of o_i 's and d_i 's. As an example of the reduction, we present in Figure 1 a LOCALMAXCUT-3 instance, and in Figure 3 the respective network congestion game instance I' . For presentation simplicity, we only show half of the network since the one consisting of supernodes \bar{U}_{ij} is symmetric.

The strategy set of each player $i \in [n]$ consists of all possible $o_i - d_i$ paths. For fixed i , let $P(i)$ be the path consisting of all existing $U_{k\ell}$'s for which either $k = i$ or $\ell = i$, and similarly, let $\bar{P}(i)$ be defined analogously using $\bar{U}_{(k,\ell)}$'s. We will call $Q_1(i)$ and $Q_2(i)$ the strategies $o_i - P(i) - d_i$ and $o_i - \bar{P}(i) - d_i$, respectively.

We now show how to translate a solution of I' directly to a solution of the initial LOCALMAXCUT- d instance in polynomial time. For any player i that plays $Q_1(i)$ (resp. $Q_2(i)$) in I' , we place vertex i in V_1 (resp. V_2), and observe that due to Claim C.2 only $Q_1(i)$ or $Q_2(i)$ can be played in a PNE. The following auxiliary definition and claims will be needed in the correctness proof of our reduction.

Definition C.1 (Strictly dominant set). For a given player i with strategy set Σ_i , the set $D_i \subseteq \Sigma_i$ is strictly dominant if every strategy in D_i strictly dominates all strategies outside D_i ; formally, for every $\sigma_i^* \in D_i$, $\sigma_i' \in \Sigma_i \setminus D_i$, we have $C_i(\sigma_i^*, \sigma_{-i}) < (\sigma_i', \sigma_{-i})$, for all strategies σ_{-i} of the other players.

Observe that, for any fixed player i , the strictly dominant set D_i is unique, but it can be empty.

Claim C.2. For any player $i \in [n]$, $Q(i) := \{Q_1(i), Q_2(i)\}$ is strictly dominant.

Proof. For some player i , if C_i is non-empty let k be its smallest item, otherwise $k = 0$. The cost experienced by i when playing $Q_1(i)$ is at most $\sum_{ij \in E} w_{ij} + (i - k + i - 1) \cdot W < (2i - k) \cdot W$, where in the left-hand side of the inequality the leftmost term is due to the regular arcs, and the rightmost term is due to the heavy arcs. The exact same upper bound holds if she played $Q_2(i)$, by definition of the costs on the corresponding arcs.

If i played any strategy other than $Q_1(i)$ or $Q_2(i)$, then she would be using a path that reaches row i , otherwise she would not have access to d_i . Therefore, her path passes through heavy arcs starting from row $k \leq i - 1$ and reaching row i . Note that if $k = i$ then the path is $Q_1(i)$ or $Q_2(i)$, and if $k > i$ then i

¹⁴ In Claim C.3 we will show that at most two players can simultaneously use a regular arc. Similarly to U_{ij} , in each supernode \bar{U}_{ij} there exists a regular arc with cost function $\kappa_{\bar{r}_{ij}}$.

does not have access to d_i . Also, at least one of the heavy arcs must be between two supernodes that are on columns ℓ, ℓ' , respectively, where $\ell' \geq \ell \geq i + 1$, otherwise the path is $Q_1(i)$ or $Q_2(i)$. Therefore, her cost would be at least $0 + (i - k + \ell - 1) \cdot W \geq (2i - k) \cdot W$, where again, in the left-hand side of the inequality the leftmost term is due to the regular arcs, and the rightmost term is due to the heavy arcs.

We conclude that for any player i , each of the strategies in $Q(i)$ strictly dominates all strategies outside $Q(i)$. \square

Claim C.3. In any PNE of the network congestion game, at most two players can be using the same regular arc.

Proof. Suppose we are in a PNE. Since no strictly dominated strategy can be played by any player in a PNE, we know that any player i can only be playing either $Q_1(i)$ or $Q_2(i)$ by Claim C.2. For any regular arc consider its supernode U_{ij} . By definition of the aforementioned strategies, this supernode can only be used by players i and j . \square

Now we will show the correctness of the reduction. First, note that the construction described can be performed in polynomial time. It remains to prove that any PNE of the network congestion game corresponds to a solution of the LOCALMAXCUT- d instance. Consider a PNE in the constructed game and observe that due to Claim C.2, no player i can be playing some strategy other than $Q_1(i)$ or $Q_2(i)$. For the sake of contradiction, suppose that there is a node $i \in V_1$ which yields greater cut value if moved to V_2 . Then $\sum_{e \in E_1(i)} w_e - \sum_{e \in E_2(i)} w_e > 0$.

Recall that there exists a supernode U_{mt} only if $mt \in E$. If C_i is non-empty let k be its smallest item, otherwise $k = 2i$. Let us compute the difference in experienced cost of player i when she deviates from $Q_1(i)$ to $Q_2(i)$ (the only other strategy that can yield lower cost). This is

$$\begin{aligned} \sum_{e \in E_1(i)} \kappa_{\bar{r}_e}(1) + \sum_{e \in E_2(i)} \kappa_{\bar{r}_e}(2) + (2i - k) \cdot W - \sum_{e \in E_1(i)} \kappa_{r_e}(2) - \sum_{e \in E_2(i)} \kappa_{r_e}(1) - (2i - k) \cdot W = \\ = \sum_{e \in E_2(i)} \kappa_{\bar{r}_e}(2) - \sum_{e \in E_1(i)} \kappa_{r_e}(2) \\ = \sum_{e \in E_2(i)} w_e - \sum_{e \in E_1(i)} w_e \\ < 0. \end{aligned}$$

This means that the initial strategy profile was not a PNE, a contradiction. Assuming there is a vertex in V_2 in the LOCALMAXCUT- d instance that could improve the cut value by moving to V_1 , via a similar reasoning we can arrive to a contradiction. Notice that by construction of $Q_1(i)$ and $Q_2(i)$, each player i passes through at most d supernodes in a PNE. Note also that since every best-response of a player is necessarily a strategy in the strictly dominant set D , we have that for parameter B of Definition 4.5 it holds $B \leq |D| = 2$. In this reduction, the constructed network congestion game has step cost functions, implying immediately that it holds for general cost functions (see (12)). Finally, since LOCALMAXCUT- d is PLS-complete for any $d \geq 5$ and network congestion games are in PLS ([FPT04]), we get the following theorem.

Theorem C.4. Network congestion games, with step function or general cost representation, are PLS-complete even when every player has 2 strategies that are in the strictly dominant set, each having maximum length 11, and each of the network's edges has at most 2 cost levels. In particular, such games are (2, 11)-compact.

We will now show how the same result also holds for affine cost functions with nonnegative coefficients, that is, the simplest class of polynomial functions, excluding constant functions. Similarly to the proof of Theorem B.3, let us set the cost functions of regular arcs to be (13). Consider a PNE of

the Network Congestion Game, and recall that no player i can be playing some strategy other than $Q_1(i)$ or $Q_2(i)$ due to [Claim C.2](#). Suppose now that there is a node $i \in V_1$ which yields greater cut value if moved to V_2 . Therefore, $\sum_{e \in E_1(i)} w_e - \sum_{e \in E_2(i)} w_e > 0$.

Again, if C_i is non-empty let k be its smallest item, otherwise $k = 2i$. The difference in experienced cost of player i when she deviates from $Q_1(i)$ to $Q_2(i)$ is

$$\begin{aligned} \sum_{e \in E_1(i)} \kappa_{\bar{r}_e}(1) + \sum_{e \in E_2(i)} \kappa_{\bar{r}_e}(2) + (2i - k) \cdot W - \sum_{e \in E_1(i)} \kappa_{r_e}(2) - \sum_{e \in E_2(i)} \kappa_{r_e}(1) - (2i - k) \cdot W = \\ = \sum_{e \in E_1(i)} w_e + \sum_{e \in E_2(i)} 2w_e - \sum_{e \in E_1(i)} 2w_e - \sum_{e \in E_2(i)} w_e \\ = \sum_{e \in E_2(i)} w_e - \sum_{e \in E_1(i)} w_e \\ < 0. \end{aligned}$$

Therefore, this contradicts the assumption that the initial strategy profile was a PNE. Similarly, by assuming there is a vertex in V_2 in the LOCALMAXCUT- d instance that could improve the cut value by moving to V_1 , we get a contradiction. Hence, we get the following theorem.

Theorem C.5. *Network Congestion Games with affine cost functions are PLS-complete, even when every player has 2 strategies that are in the strictly dominant set, each having maximum length 11. In particular, such games are $(2, 11)$ -compact.*

D PLS-Hardness of LOCALMAX- k -CUT

Here we present a PLS-hardness reduction from LOCALMAXCUT- d to LOCALMAX- k -CUT for any $k \geq 2$ under the FLIP neighbourhood. LOCALMAXCUT- d is the restricted version of LOCALMAXCUT where the input graph has maximum degree $d \in \mathbb{N}^*$.

Theorem D.1. *LOCALMAX- k -CUT is PLS-hard for any $k \geq 2$, even when the maximum degree of the input graph is $k + 3$.*

Proof. Consider an instance of LOCALMAXCUT- d on graph $G' = (V, E)$ and weights $w_e, e \in E$, where $|V| = n \in \mathbb{N}^*$, and without loss of generality, let $V = \{v_1, \dots, v_n\}$. For each vertex $v_i \in V$ we introduce k special vertices v_i^1, \dots, v_i^k . For each fixed $i \in [n]$, we introduce a special edge $v_i^j v_i^\ell$ for every $j, \ell \in [k]$ with $j \neq \ell$, where $w_{v_i^j v_i^\ell} := W^* > \sum_{e \in E} w_e$. Additionally, for each fixed $i \in [n]$ we introduce a special edge $v_i v_i^j$ for every $j \in \{3, 4, \dots, k\}$, where $w_{v_i v_i^j} := W^* > \sum_{e \in E} w_e$. This instance creation clearly takes polynomial time. Suppose the cut (V_1, \dots, V_k) is a solution to the constructed LOCALMAX- k -CUT instance, where each V_j is a *partition block*. We translate this solution to a solution of LOCALMAXCUT- d by considering as one partition block $V_1 \cap V$ and the other block $V_2 \cap V$.

It remains to show that the latter is indeed a solution to the initial LOCALMAXCUT- d instance. First, observe that no two special vertices v_i^j, v_i^ℓ are in the same partition block. If that were the case, then there would exist a vertex-part V_j that does not contain any special vertex from the set $\{v_i^1, \dots, v_i^k\}$, and moving v_i^j to V_j increases the weight of the cut by at least $W^* - \sum_{e \in E} w_e > 0$, contradicting the local optimality of the cut. Therefore, without loss of generality (by renaming the vertex-parts), for any fixed $j \in [k]$, let V_j contain v_i^j for all $i \in [n]$. Now we claim that no vertex $v_i \in V$ is contained in any V_j for $j \in \{3, 4, \dots, k\}$. If that were the case, then moving v_i to V_1 (or V_2) would increase the weight of the cut by at least $W^* - \sum_{e \in E} w_e > 0$, contradicting the local optimality of the cut.

To prove that this instance's solution induces a solution to the initial LOCALMAXCUT- d instance, we examine the vertices $v_i \in V$. According to the above, all v_i 's are in $V_1 \cup V_2$. Let us now consider some

$v_i \in (V_1 \cup V_2) \cap V$ of the LOCALMAX- k -CUT solution, and suppose that in the induced LOCALMAXCUT- d cut, by moving v_i from its current cut-set to the other cut-set, resulted in a cut with greater weight. Then this would also be the case in the LOCALMAX- k -CUT solution, since there are no special edges between v_i and any vertex of $V_1 \cup V_2$. This contradicts its local optimality, and completes the correctness of the reduction.

Notice now that the LOCALMAX- k -CUT instance has maximum degree $k + d - 2$: the maximum degree of any vertex v_i is $d + (k - 2)$, while that of any vertex v_i^j is $(k - 1) + 1$. Since LOCALMAXCUT- d is PLS-hard for any $d \geq 5$ ([ET11]), we get the statement of our theorem. \square

E PLS-Hardness of Hitting Set with Bounded Element Occurrence

Here we present an adaptation of the proof of PLS-hardness due to [DS10b]. Instead of reducing from MAXSAT/1-FLIP (or CNFSat as it is called in the aforementioned work), we show a tight PLS-reduction from MAXSAT/1-FLIP with bounded variable occurrence, which is also PLS-complete due to [Kre89] and [Kla96] via tight PLS-reductions.

Theorem E.1. *HS-1 is PLS-hard, even when the number of sets in which any single element occurs is bounded by a constant.*

Proof. We reduce from MAXSAT/1-FLIP with bounded variable occurrence, i.e. we are given a CNF formula with clauses $\tilde{C} = \{\tilde{C}_1, \dots, \tilde{C}_{\tilde{m}}\}$ over binary variables $\tilde{X} = \{x_1, \dots, x_{\tilde{n}}\}$ and weights \tilde{w}_i for every \tilde{C}_i . Additionally, every variable only occurs in at most \tilde{B} clauses.

Given such an instance of MAXSAT/1-FLIP, we define an instance of HS-1 and show that every locally optimal subset corresponds to a locally optimal assignment of the MAXSAT/1-FLIP instance. The ground set is given by a positive and negative element for every variable; $\mathcal{B} = \bigcup_{x \in \tilde{X}} \{x, \bar{x}\}$. For each of the variables we also define a set $C_x = \{x, \bar{x}\}$, $x \in \tilde{X}$ with large weight W exceeding the sum of all clause weights \tilde{w}_i . We restrict the subset size to the number of variables, $m = \tilde{m}$, and the subset corresponding to a truth assignment is given by choosing the respective positive or negative element for every variable according to the assignment. Finally, for every clause \tilde{C}_i of the MAXSAT/1-FLIP instance, we define a set C_i for HS-1 that contains exactly the variables of the clause (x for positive and \bar{x} for negative literals) and its weight corresponds to the weight of the clause, $w_i = \tilde{w}_i$. The auxiliary sets C_x imply that every locally optimal subset S contains exactly one of x, \bar{x} for every variable: Otherwise, because of our choice of m , it must hold that for some variable both x, \bar{x} do not occur in S , while for another one both y, \bar{y} do occur. Because the sets C_x have large weight, it is a local improvement to replace y by x , thus making both C_x and C_y contribute their weight to the total weight of the subset.

Now, every locally optimal subset S corresponds to a truth assignment and their weights also coincide up to the constant weight $m \cdot W$ of the auxiliary sets: A clause \tilde{C}_i is satisfied as soon as any of its literals is true, and analogously the set C_i contributes its weight to the solution as soon as any of its elements are contained in S . The neighbourhood graph contains all changes of an element x to its negative copy \bar{x} or vice-versa. However, when starting from a subset that corresponds to a truth assignment, a change to a neighbouring subset that does not correspond to one is never an improving move, which concludes the reduction as being tight.

Up to this point, this is exactly the reduction from [DS10b]. When the original instance of MAXSAT/1-FLIP has bounded variable occurrence, it follows immediately that the HS-1 instance has bounded element occurrence $B = \max_{b \in \mathcal{B}} |\{C \in \mathcal{C} \mid b \in C\}|$: The maximum number of element occurrence in the HS-1 instance is $B \leq \tilde{B} + 1$, as apart from the sets C_i that correspond to the clauses, we have only added the auxiliary set $\{x, \bar{x}\}$ for every variable to the problem. \square

F Maximum Constraint Assignment

This section deals with the problem of Maximum Constraint Assignment (MCA- (p, q, r)) as defined in [DM13]. MCA- (p, q, r) is a general optimization problem over integer variables. A heavily restricted variant of this problem has undergone smoothed analysis by the name of Binary Max-CSP in [CGVYZ20]. Compared to (even non-binary) Max-CSP, MCA- (p, q, r) covers many possible weights for the different assignments of a single constraint instead of just a single one. Our focus will be on translating the parameters of MCA- (p, q, r) directly to parameters in the smoothed running time.

Definition F.1 (MCA- (p, q, r)). By MCA- (p, q, r) we denote the following local optimization problem. Let $C = \{C_1, \dots, C_m\}$ be a set of constraints over variables $\mathcal{X} = \{x_1, \dots, x_n\}$, i.e. each constraint C_i is associated with a set of variables and written symbolically as $C_i(x_{i_1}, \dots, x_{i_{p_i}})$. When assigning integers from $[r]$ to the variables, these constraints are evaluated as functions $C_i : [r]^{p_i} \rightarrow [0, 1]$ with $p_i \leq p$. The feasible solutions are the integer assignments of the variables, i.e. functions $a : X \rightarrow [r]$. The sum of all constraints evaluated for the assignment is the weight of the solution:

$$w(a) = \sum_{i=1}^m C_i(a(x_{i_1}), \dots, a(x_{i_{p_i}})).$$

Every variable x_i appears in at most q constraints from C . The local search neighbourhood is given by changing the mapping of a single variable of a given assignment to a different value. This is an efficiently searchable neighbourhood.

As shown in [DM13], MCA- $(3, 2, 3)$ and MCA- $(2, 3, 6)$ are PLS-complete. Therefore, we will treat the parameters (p, q, r) as constant in the following, leaving of interest only the instance size parameters m and n . This also justifies the assumption that the every function C_i can be given as a list of values for all r^{p_i} possible input assignments.

Theorem F.2. *Given a smoothed MCA- (p, q, r) problem, where the cost values $C_i(v_i)$ for every integer assignment $v_i \in [r]^{p_i}$ are drawn from independent distributions $f_{i,v_i} : [0, 1] \rightarrow [0, \phi]$, the expected number of iterations of local search is bounded by $O(r^{2p(q+1)+2}nm^2\phi)$.*

Proof. We first define a binary ($M = 1$) CLO problem, where the cost part corresponds to all constraints and their possible assignments, i.e. for each $i \in [m]$ and every $v_i \in [r]^{p_i}$ we have a binary cost variable s_{i,v_i} with associated cost coefficient $c_{i,v_i} = C_i(v_i)$. The model works on an indicator basis, i.e. for every assignment a of MCA- (p, q, r) , the cost variables of the induced CLO configuration $s(a)$ take value $s_{i,v_i} = 1$ only for $v_i = (a(x_{i_j}))_{j=1}^{p_i}$, and $s_{i,v} = 0$ for every other integer vector $v \neq v_i$, for all $i \in [m]$. Intuitively, the single entry of 1 within every set of coordinates for fixed i indicates which of the possible integer assignments is currently chosen, and thus this value contributes its weight to the total weight of the configuration. Therefore, we have in total $\nu = \sum_{i=1}^m r^{p_i}$ many cost components. Now it holds that for every integer assignment a we have

$$c(a) = \sum_{i=1}^m C_i(a(x_{i_1}), \dots, a(x_{i_{p_i}})) = \sum_{i=1}^m \sum_{v_i \in [r]^{p_i}} c_{i,v_i} \mathbb{1}[v_i = (a(x_{i_1}), \dots, a(x_{i_{p_i}}))] = C(s(a))$$

and the cost functions coincide. The neighbourhood relation for the CLO problem is derived directly from MCA- (p, q, r) as well and is also efficiently searchable.

We now translate the cost distributions of smoothed MCA- (p, q, r) to the distributions of the induced smoothed CLO problem and apply Theorem 3.2 to it. We specify the covering $(\mathcal{E}, \mathcal{I})$ and the properties λ, β and μ in the following. To capture the indicator structure defined above, we define a coordinate cluster $I_i = \{(i, v_i) \mid v_i \in [r]^{p_i}\}$ for all cost components associated with constraint C_i . There are in

total r^{p_i} different indicator assignments for the configuration of \mathbf{s}_{I_i} . In the difference $\mathbf{s}'_{I_i} - \mathbf{s}_{I_i}$, both \mathbf{s}_{I_i} and \mathbf{s}'_{I_i} are one of these assignments, and we therefore have at most $(r^{p_i})^2$ distinct differences by taking all possible combinations. Thus, we choose $\mu = r^{2p}$.

Now, we specify the transition cover \mathcal{E} . We denote by $E(j, v, v')$ the transition cluster that captures all transitions where the assignment of variable x_j changes from integer v to v' . Formally, this is

$$E(j, v, v') = \{(\mathbf{s}(v, a_{-j}), \mathbf{s}(v', a_{-j})) \mid a_{-j} : \mathcal{X} \setminus \{x_j\} \longrightarrow [r]\}$$

where a_{-j} denotes an arbitrary assignment of the remaining variables to integers and (v, a_{-j}) denotes the assignment that maps $\mathcal{X} \setminus \{x_j\}$ as defined by a_{-j} and x_j to v . Now, $\mathcal{E} = \{E(j, v, v') \mid j \in [n], v, v' \in [r], v \neq v'\}$ and therefore $\lambda = nr(r-1)$.

Finally, in order to specify β , we investigate which coordinate clusters are affected by the transitions in a fixed cluster $E(j, v, v')$: In such a transition, the configuration of \mathbf{s}_{I_i} changes if and only if x_j occurs in constraint C_i . By assumption, every variable appears in at most q different constraints. This gives $\beta = q$ and concludes the proof with a bound on the expected number of steps of

$$3 \cdot (r^{2p})^q nr(r-1) \cdot \left(\sum_{i=1}^m r^{p_i} \right)^2 \cdot \phi \leq 3(r^{2p})^q nr^2 \phi m^2 r^{2p} = 3r^{2p(q+1)+2} nm^2 \phi.$$

□

Note that a variant of $\text{MCA-}(p, q, r)$ can be defined which bridges the gap between its vast description of the constraints and the constraints in Max-CSP: instead of separate weights for every possible assignment of the variables in a constraint in $\text{MCA-}(p, q, r)$ or just a set of *satisfying* assignments with a single weight, regardless of the specific assignment, we can use several (disjoint) sets of assignments with a given weight each for every constraint. For the smoothed problem, each of these weights – and therefore the contribution of several assignments in a given constraint – is perturbed as a single value. By doing so, many problems can be directly expressed by an instance of $\text{MCA-}(p, q, r)$, which gives another unified proof method for smoothed complexity. This model requires only slight adaptations from our analysis and also gives polynomial smoothed complexity under the variable occurrence constraints derived from $\text{MCA-}(p, q, r)$.

References

- [ABCC06] David L. Applegate, Robert E. Bixby, Vašek Chvátal, and William J. Cook. *The Traveling Salesman Problem: A Computational Study*. Princeton University Press, 2006. URL: <http://www.jstor.org/stable/j.ctt7s8xg> (visited on 11/09/2022) (cit. on p. 19).
- [ABPW17] Omer Angel, Sébastien Bubeck, Yuval Peres, and Fan Wei. “Local Max-Cut in Smoothed Polynomial Time”. In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*. ACM, 2017, pp. 429–437. DOI: [10.1145/3055399.3055402](https://doi.org/10.1145/3055399.3055402) (cit. on pp. 9, 23, 24, 33).
- [AL97] Emile Aarts and Jan Karel Lenstra, eds. *Local Search in Combinatorial Optimization*. John Wiley, 1997. DOI: [10.2307/j.ctv346t9c](https://doi.org/10.2307/j.ctv346t9c) (cit. on p. 1).
- [ARV08] Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. “On the Impact of Combinatorial Structure on Congestion Games”. In: *Journal of the ACM* 55.6 (2008), pp. 1–22. DOI: [10.1145/1455248.1455249](https://doi.org/10.1145/1455248.1455249) (cit. on pp. 5, 16, 18, 37).
- [BC14] S. G. Bobkov and G. P. Chistyakov. “Bounds on the Maximum of the Density for Sums of Independent Random Variables”. In: *Journal of Mathematical Sciences* 199.2 (2014), pp. 100–106. DOI: [10.1007/s10958-014-1836-9](https://doi.org/10.1007/s10958-014-1836-9) (cit. on p. 34).
- [BCC21] Ali Bibak, Charles Carlson, and Karthekeyan Chandrasekaran. “Improving the Smoothed Complexity of FLIP for Max Cut Problems”. In: *ACM Transactions on Algorithms* 17.3 (2021), 19:1–19:38. DOI: [10.1145/3454125](https://doi.org/10.1145/3454125) (cit. on pp. 3, 9, 24, 26).
- [BCK10] Anand Bhalgat, Tanmoy Chakraborty, and Sanjeev Khanna. “Approximating Pure Nash Equilibrium in Cut, Party Affiliation, and Satisfiability Games”. In: *Proceedings of the 11th ACM Conference on Electronic Commerce (EC)*. 2010, pp. 73–82. DOI: [10.1145/1807342.1807353](https://doi.org/10.1145/1807342.1807353) (cit. on p. 26).
- [BKM20] Shant Boodaghians, Rucha Kulkarni, and Ruta Mehta. “Smoothed Efficient Algorithms and Reductions for Network Coordination Games”. In: *11th Innovations in Theoretical Computer Science Conference (ITCS)*. Vol. 151. LIPIcs. 2020, 73:1–73:15. DOI: [10.4230/LIPIcs.ITCS.2020.73](https://doi.org/10.4230/LIPIcs.ITCS.2020.73) (cit. on pp. 3, 9, 28).
- [Bra18] Alexandre Brandts-Longtin. “Smoothed Analysis of Max-k-Sat”. MA thesis. Montreal: McGill University, 2018. URL: <https://escholarship.mcgill.ca/concern/theses/f4752k14b> (cit. on pp. 3, 23).
- [BV06] Rene Beier and Berthold Vöcking. “Typical Properties of Winners and Losers in Discrete Optimization”. In: *SIAM Journal on Computing* 35.4 (2006), pp. 855–881. DOI: [10.1137/s0097539705447268](https://doi.org/10.1137/s0097539705447268) (cit. on pp. 2, 6).
- [CD11] Yang Cai and Constantinos Daskalakis. “On Minmax Theorems for Multiplayer Games”. In: *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. USA, 2011, pp. 217–234. DOI: [10.5555/2133036.2133056](https://doi.org/10.5555/2133036.2133056) (cit. on p. 28).
- [CGVYZ20] Xi Chen, Chenghao Guo, Emmanouil-Vasileios Vlatakis-Gkaragkounis, Mihalis Yannakakis, and Xinzhi Zhang. “Smoothed Complexity of Local Max-Cut and Binary Max-CSP”. In: *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*. ACM, 2020, pp. 1052–1065. DOI: [10.1145/3357713.3384325](https://doi.org/10.1145/3357713.3384325) (cit. on pp. 9, 24, 43).
- [CKT99] Barun Chandra, Howard Karloff, and Craig Tovey. “New Results on the Old k -opt Algorithm for the Traveling Salesman Problem”. In: *SIAM Journal on Computing* 28.6 (1999), pp. 1998–2029. DOI: [10.1137/s0097539793251244](https://doi.org/10.1137/s0097539793251244) (cit. on pp. 1, 20).

- [Coo12] William J. Cook. *In Pursuit of the Traveling Salesman: Mathematics at the Limits of Computation*. Princeton University Press, 2012. URL: <http://www.jstor.org/stable/j.ctt7t8kc> (visited on 11/09/2022) (cit. on p. 19).
- [DM13] Dominic Dumrauf and Burkhard Monien. “On the PLS-complexity of maximum constraint assignment”. In: *Theoretical Computer Science* 469 (2013), pp. 24–52. DOI: [10.1016/j.tcs.2012.10.044](https://doi.org/10.1016/j.tcs.2012.10.044) (cit. on pp. 29, 43).
- [DP11] Constantinos Daskalakis and Christos Papadimitriou. “Continuous Local Search”. In: *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2011. DOI: [10.1137/1.9781611973082.62](https://doi.org/10.1137/1.9781611973082.62) (cit. on p. 5).
- [DS10a] Dominic Dumrauf and Tim Süß. “On the Complexity of Local Search for Weighted Standard Set Problems”. In: *Programs, Proofs, Processes*. Ed. by Fernando Ferreira, Benedikt Löwe, Elvira Mayordomo, and Luís Mendes Gomes. 2010, pp. 132–140. DOI: [10.1007/978-3-642-13962-8_15](https://doi.org/10.1007/978-3-642-13962-8_15) (cit. on pp. 3, 29–32).
- [DS10b] Dominic Dumrauf and Tim Süß. “On the Complexity of Local Search for Weighted Standard Set Problems”. In: *CoRR* abs/1004.0871 (2010). arXiv: [1004.0871](https://arxiv.org/abs/1004.0871) (cit. on pp. 32, 42).
- [EM09] Christian Engels and Bodo Manthey. “Average-case approximation ratio of the 2-opt algorithm for the TSP”. In: *Operations Research Letters* 37.2 (2009), pp. 83–84. DOI: [10.1016/j.orl.2008.12.002](https://doi.org/10.1016/j.orl.2008.12.002) (cit. on p. 19).
- [ER15] Michael Etscheid and Heiko Röglin. “Smoothed Analysis of the Squared Euclidean Maximum-Cut Problem”. In: *Algorithms - ESA 2015*. Springer Berlin Heidelberg, 2015, pp. 509–520. DOI: [10.1007/978-3-662-48350-3_43](https://doi.org/10.1007/978-3-662-48350-3_43) (cit. on p. 24).
- [ER17] Michael Etscheid and Heiko Röglin. “Smoothed Analysis of Local Search for the Maximum-Cut Problem”. In: *ACM Transactions on Algorithms* 13.2 (2017), 25:1–25:12. DOI: [10.1145/3011870](https://doi.org/10.1145/3011870) (cit. on pp. 2, 9, 24).
- [ERV14] Matthias Englert, Heiko Röglin, and Berthold Vöcking. “Worst Case and Probabilistic Analysis of the 2-Opt Algorithm for the TSP”. In: *Algorithmica* 68.1 (2014), pp. 190–264. DOI: [10.1007/s00453-013-9801-4](https://doi.org/10.1007/s00453-013-9801-4) (cit. on pp. 19, 20).
- [ERV16] Matthias Englert, Heiko Röglin, and Berthold Vöcking. “Smoothed Analysis of the 2-Opt Algorithm for the General TSP”. In: *ACM Transactions on Algorithms* 13.1 (2016), pp. 1–15. DOI: [10.1145/2972953](https://doi.org/10.1145/2972953) (cit. on pp. 2, 3, 9, 20, 21).
- [ET11] Robert Elsässer and Tobias Tscheuschner. “Settling the Complexity of Local Max-Cut (Almost) Completely”. In: *International Colloquium on Automata, Languages, and Programming (ICALP)*. Ed. by Luca Aceto, Monika Henzinger, and Jiří Sgall. 2011, pp. 171–182. DOI: [10.1007/978-3-642-22006-7_15](https://doi.org/10.1007/978-3-642-22006-7_15) (cit. on pp. 3, 24–26, 35–37, 42).
- [FGMS20] John Fearnley, Spencer Gordon, Ruta Mehta, and Rahul Savani. “Unique end of potential line”. In: *Journal of Computer and System Sciences* 114 (2020), pp. 1–35. DOI: [10.1016/j.jcss.2020.05.007](https://doi.org/10.1016/j.jcss.2020.05.007) (cit. on p. 5).
- [FPT04] Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. “The Complexity of Pure Nash Equilibria”. In: *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC)*. 2004, pp. 604–612. DOI: [10.1145/1007352.1007445](https://doi.org/10.1145/1007352.1007445) (cit. on pp. 1, 37, 40).
- [Gia24] Yiannis Giannakopoulos. “A Smoothed FPTAS for Equilibria in Congestion Games”. In: *Proceedings of the 25th ACM Conference on Economics and Computation (EC)*. 2024, pp. 401–413. DOI: <https://doi.org/10.1145/3670865.3673615> (cit. on pp. 2, 33).

- [GJ90] Michael R. Garey and David S. Johnson. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., 1990. DOI: [10.5555/574848](https://doi.org/10.5555/574848) (cit. on pp. 19, 29).
- [GP07] Gregory Gutin and Abraham P. Punnen, eds. *The Traveling Salesman Problem and Its Variations*. Springer US, 2007. DOI: [10.1007/b101971](https://doi.org/10.1007/b101971) (cit. on pp. 19, 20).
- [GS19] Martin Gairing and Rahul Savani. “Computing Stable Outcomes in Symmetric Additively Separable Hedonic Games”. In: *Mathematics of Operations Research* 44.3 (2019), pp. 1101–1121. DOI: [10.1287/moor.2018.0960](https://doi.org/10.1287/moor.2018.0960) (cit. on p. 27).
- [Hal74] Paul R. Halmos. *Measure Theory*. Vol. 18. Graduate Texts in Mathematics. Springer, 1974 (cit. on p. 34).
- [JM97] David S. Johnson and Lyle A. McGeoch. “The traveling salesman problem: a case study”. In: *Local Search in Combinatorial Optimization*. Ed. by Emile Aarts and Jan Karel Lenstra. John Wiley, 1997. Chap. 8. DOI: [10.2307/j.ctv346t9c](https://doi.org/10.2307/j.ctv346t9c) (cit. on p. 19).
- [JPY88] David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. “How easy is local search?” In: *Journal of Computer and System Sciences* 37.1 (1988), pp. 79–100. DOI: [10.1016/0022-0000\(88\)90046-3](https://doi.org/10.1016/0022-0000(88)90046-3) (cit. on pp. 1, 5, 23, 29).
- [Kla96] Hartmut Klauck. “On the Hardness of Global and Local Approximation”. In: *Proceedings of the 5th Scandinavian Workshop on Algorithm Theory (SWAT)*. 1996, pp. 88–99. DOI: [10.1007/3-540-61422-2_123](https://doi.org/10.1007/3-540-61422-2_123) (cit. on pp. 19, 22, 32, 42).
- [KP80] Paris-C. Kanellakis and Christos H. Papadimitriou. “Local Search for the Asymmetric Traveling Salesman Problem”. In: *Operations Research* 28.5 (1980), pp. 1086–1099. DOI: [10.1287/opre.28.5.1086](https://doi.org/10.1287/opre.28.5.1086) (cit. on p. 20).
- [Kre89] Mark W. Krentel. “Structure in Locally Optimal Solutions”. In: *Proceedings of the 30th Annual Symposium on Foundations of Computer Science (FOCS)*. 1989, pp. 216–221. DOI: [10.1109/SFCS.1989.63481](https://doi.org/10.1109/SFCS.1989.63481) (cit. on pp. 1, 3, 19, 22, 32, 42).
- [Kre90] Mark W. Krentel. “On Finding and Verifying Locally Optimal Solutions”. In: *SIAM Journal on Computing* 19.4 (1990), pp. 742–749. DOI: [10.1137/0219052](https://doi.org/10.1137/0219052) (cit. on p. 22).
- [MKA07] Wil Michiels, Jan Korst, and Emile Aarts. *Theoretical Aspects of Local Search*. Springer, 2007. DOI: [10.1007/978-3-540-35854-1](https://doi.org/10.1007/978-3-540-35854-1) (cit. on pp. 1, 20).
- [MR23] Bodo Manthey and Jesse van Rhijn. “Improved Smoothed Analysis of 2-Opt for the Euclidean TSP”. In: *34th International Symposium on Algorithms and Computation (ISAAC 2023)*. Ed. by Satoru Iwata and Naonori Kakimura. Vol. 283. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023, 52:1–52:16. DOI: [10.4230/LIPIcs.ISAAC.2023.52](https://doi.org/10.4230/LIPIcs.ISAAC.2023.52) (cit. on p. 20).
- [MS96] Dov Monderer and Lloyd S. Shapley. “Potential games”. In: *Games and Economic Behavior* 14.1 (1996), pp. 124–143. DOI: [10.1006/game.1996.0044](https://doi.org/10.1006/game.1996.0044) (cit. on pp. 10, 27).
- [MT10] Burkhard Monien and Tobias Tscheuschner. “On the Power of Nodes of Degree Four in the Local Max-Cut Problem”. In: *Proceedings of the 7th International Conference in Algorithms and Complexity (CIAC)*. 2010, pp. 264–275. DOI: [10.1007/978-3-642-13073-1_24](https://doi.org/10.1007/978-3-642-13073-1_24) (cit. on pp. 25, 26).
- [MV13] Bodo Manthey and Rianne Veenstra. “Smoothed Analysis of the 2-Opt Heuristic for the TSP: Polynomial Bounds for Gaussian Noise”. In: *Algorithms and Computation*. Springer Berlin Heidelberg, 2013, pp. 579–589. DOI: [10.1007/978-3-642-45030-3_54](https://doi.org/10.1007/978-3-642-45030-3_54) (cit. on p. 20).

- [OPS04] James B. Orlin, Abraham P. Punnen, and Andreas S. Schulz. “Approximate Local Search in Combinatorial Optimization”. In: *SIAM Journal on Computing* 33.5 (2004), pp. 1201–1214. DOI: [10.1137/S0097539703431007](https://doi.org/10.1137/S0097539703431007) (cit. on pp. 1, 2).
- [Rög08] Heiko Röglin. “The Complexity of Nash Equilibria, Local Optima, and Pareto-Optimal Solutions”. PhD thesis. RWTH Aachen University, 2008. URL: <https://publications.rwth-aachen.de/record/50046> (cit. on p. 6).
- [Ros73] Robert W. Rosenthal. “A Class of Games Possessing Pure-Strategy Nash Equilibria”. In: *International Journal of Game Theory* 2.1 (1973), pp. 65–67. DOI: [10.1007/BF01737559](https://doi.org/10.1007/BF01737559) (cit. on p. 10).
- [Rou16] Tim Roughgarden. “Pure Nash Equilibria and PLS-Completeness”. In: *Twenty Lectures on Algorithmic Game Theory*. Cambridge University Press, 2016, pp. 261–278. DOI: [10.1017/CB09781316779309.020](https://doi.org/10.1017/CB09781316779309.020) (cit. on pp. 1, 16, 35).
- [Rou21] Tim Roughgarden, ed. *Beyond the Worst-Case Analysis of Algorithms*. Cambridge University Press, 2021. DOI: [10.1017/9781108637435](https://doi.org/10.1017/9781108637435) (cit. on pp. 1, 6).
- [RV07] Heiko Röglin and Berthold Vöcking. “Smoothed analysis of integer programming”. In: *Mathematical Programming* 110.1 (2007), pp. 21–56. DOI: [10.1007/s10107-006-0055-7](https://doi.org/10.1007/s10107-006-0055-7) (cit. on pp. 2, 6).
- [Sch24] Gregory Schwartzman. “Local Max-Cut on Sparse Graphs”. In: *32nd Annual European Symposium on Algorithms (ESA 2024)*. Vol. 308. Leibniz International Proceedings in Informatics (LIPIcs). 2024, 98:1–98:6. DOI: [10.4230/LIPIcs.ESA.2024.98](https://doi.org/10.4230/LIPIcs.ESA.2024.98) (cit. on p. 24).
- [ST04] Daniel A. Spielman and Shang-Hua Teng. “Smoothed analysis of algorithms”. In: *Journal of the ACM* 51.3 (2004), pp. 385–463. DOI: [10.1145/990308.990310](https://doi.org/10.1145/990308.990310) (cit. on pp. 1, 2).
- [ST09] Daniel A. Spielman and Shang-Hua Teng. “Smoothed Analysis: An Attempt to Explain the Behavior of Algorithms in Practice”. In: *Commun. ACM* 52.10 (2009), pp. 76–84. DOI: [10.1145/1562764.1562785](https://doi.org/10.1145/1562764.1562785) (cit. on p. 6).
- [STV20] Ola Svensson, Jakub Tarnawski, and László A. Végh. “A Constant-Factor Approximation Algorithm for the Asymmetric Traveling Salesman Problem”. In: *Journal of the ACM* 67.6 (2020). DOI: [10.1145/3424306](https://doi.org/10.1145/3424306) (cit. on pp. 3, 20).
- [SY91] Alejandro A. Schäffer and Mihalis Yannakakis. “Simple Local Search Problems That are Hard to Solve”. In: *SIAM J. Comput.* 20.1 (1991), pp. 56–87. DOI: [10.1137/0220004](https://doi.org/10.1137/0220004) (cit. on pp. 1, 5, 18, 19, 25, 27, 28, 36, 37).
- [Sze11] Stefan Szeider. “The parameterized complexity of k -flip local search for SAT and MAX SAT”. In: *Discrete Optimization* 8.1 (2011), pp. 139–145. DOI: [10.1016/j.disopt.2010.07.003](https://doi.org/10.1016/j.disopt.2010.07.003) (cit. on pp. 3, 22).
- [Yan97] Mihalis Yannakakis. “Computational Complexity”. In: *Local Search in Combinatorial Optimization*. Ed. by Emile Aarts and Jan Karel Lenstra. John Wiley, 1997. Chap. 2. DOI: [10.2307/j.ctv346t9c](https://doi.org/10.2307/j.ctv346t9c) (cit. on pp. 1, 5, 29).
- [YI99] M. Yagiura and T. Ibaraki. “Analyses on the 2 and 3-Flip Neighborhoods for the MAX SAT”. In: *Journal of Combinatorial Optimization* 3.1 (1999), pp. 95–114. DOI: [10.1023/a:1009873324187](https://doi.org/10.1023/a:1009873324187) (cit. on pp. 3, 22).
- [Zho20] Xianghui Zhong. “On the Approximation Ratio of the k -Opt and Lin-Kernighan Algorithm for Metric and Graph TSP”. In: *28th Annual European Symposium on Algorithms (ESA 2020)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. DOI: [10.4230/LIPICS.ESA.2020.83](https://doi.org/10.4230/LIPICS.ESA.2020.83) (cit. on p. 19).