

Approximating Bounded Occurrence Ordering CSPs*

VENKATESAN GURUSWAMI[†]YUAN ZHOU[‡]

Abstract

A theorem of Håstad shows that for every constraint satisfaction problem (CSP) over a fixed size domain, instances where each variable appears in at most $O(1)$ constraints admit a non-trivial approximation algorithm, in the sense that one can beat (by an additive constant) the approximation ratio achieved by the naive algorithm that simply picks a random assignment. We consider the analogous question for ordering CSPs, where the goal is to find a linear ordering of the variables to maximize the number of satisfied constraints, each of which stipulates some restriction on the local order of the involved variables. It was shown recently that without the bounded occurrence restriction, for *every* ordering CSP it is Unique Games-hard to beat the naive random ordering algorithm.

In this work, we prove that the CSP with monotone ordering constraints $x_{i_1} < x_{i_2} < \dots < x_{i_k}$ of arbitrary arity k can be approximated beyond the random ordering threshold $1/k!$ on bounded occurrence instances. We prove a similar result for all ordering CSPs, with arbitrary payoff functions, whose constraints have arity at most 3. Our method is based on working with a carefully defined Boolean CSP that serves as a proxy for the ordering CSP. One of the main technical ingredients is to establish that certain Fourier coefficients of this proxy constraint have substantial mass. These are then used to guarantee a good ordering via an algorithm that finds a good Boolean assignment to the variables of a low-degree bounded occurrence multilinear polynomial. Our algorithm for the latter task is similar to Håstad's earlier method but is based on a greedy choice that achieves a better performance guarantee.

1 Introduction

Constraint satisfaction. Constraint satisfaction problems (CSPs) are an important class of optimization problems. A CSP is specified by a finite set Π of relations, each of arity k , over a domain $\{0, 1, \dots, D-1\}$, where k, D are some fixed constants. An instance of such a CSP consists of a set of variables V and a collection of constraints (possibly with weights) each of which is a relation from Π applied to some k -tuple of variables from V . The goal is to find an assignment $\sigma : V \rightarrow D$ that maximizes the total weight of satisfied constraints. For example in the Max Cut problem, $k = D = 2$ and Π consists of the single relation $\text{CUT}(a, b) = \mathbf{1}(a \neq b)$. More generally, one can also

*This research was supported in part by NSF CCF 1115525 and US-Israel BSF grant 2008293. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

[†]Computer Science Department, Carnegie Mellon University, guruswami@cmu.edu

[‡]Computer Science Department, Carnegie Mellon University, yuanzhou@cs.cmu.edu

allow real-valued payoff functions $f : \{0, 1, \dots, D - 1\}^k \rightarrow \mathbb{R}^+$ in Π (instead of just $\{0, 1\}$ -valued functions), with the goal being to find an assignment maximizing the total payoff.

Most Max CSP problems are NP-hard, and there is by now a rich body of work on approximation algorithms and hardness of approximation results for CSPs. Algorithmically, semidefinite programming (SDP) has been the principal tool to obtain good approximation ratios. In fact, SDP is universal for CSPs in the sense that under the Unique Games conjecture a certain canonical SDP achieves the optimal approximation ratio [Rag08]. However, many CSPs, including Max 3SAT, Max 3LIN, Max NAE-4-SAT, etc., are *approximation resistant*, meaning that for any $\epsilon > 0$, even when given a $(1 - \epsilon)$ -satisfiable instance, it is hard to find an assignment that satisfies more than a fraction $r + \epsilon$ of the constraints, where r , the *random assignment threshold*, is the expected fraction of constraints satisfied by a random assignment [Hås01, AM09]. In other words, it is hard to improve upon the naive algorithm that simply picks a random assignment without even looking at the structure of the instance.

Let us call a CSP that is not approximation resistant as *non-trivially approximable*. In spite of a rich body of powerful algorithmic and hardness results, we are quite far from a complete classification of all CSPs into approximation resistant or non-trivially approximable. Several partial results are known; for example, the classification is known for Boolean predicates of arity 3. It is known that every binary CSP (i.e., whose constraints have arity 2), regardless of domains size (as long as it is fixed), is non-trivially approximable via a SDP-based algorithm [GW95, EG04, Hås08]. In a different vein, Håstad [Hås00] showed that for *every* Boolean CSP, when restricted to sparse instances where each variable participates in a bounded number B of constraints, one can beat the random assignment threshold (by an amount that is at least $\Omega(1/B)$). Trevisan showed that for Max 3SAT beating the random assignment threshold by more than $O(1/\sqrt{B})$ is NP-hard, so some degradation of the performance ratio with the bound B is necessary [Tre01].

Ordering CSPs. With this context, we now turn to ordering CSPs, which are the focus of this paper. The simplest ordering CSP is the well-known Maximum Acyclic Subgraph (MAS) problem, where we are given a directed graph and the goal is to order the vertices $V = \{x_1, \dots, x_n\}$ of the graph so that a maximum number of edges go forward in the ordering. This can be viewed as a “CSP” with variables V and constraints $x_i < x_j$ for each directed edge (x_i, x_j) in the graph; the difference from usual CSPs is that the variables are to be ordered, i.e., assigned values from $\{1, 2, \dots, n\}$, instead of being assigned values from a fixed domain (of size independent of n).

An ordering CSP of arity k is specified of a constraint $\Pi : S_k \rightarrow \{0, 1\}$ where S_k is the set of permutations of $\{1, 2, \dots, k\}$. An instance of such a CSP consists of a set of variables V and a collection of constraints which are (ordered) k -tuples. The constraint tuple $e = (x_{i_1}, x_{i_2}, \dots, x_{i_k})$ is satisfied by an ordering of V if the local ordering of the variables $x_{i_1}, x_{i_2}, \dots, x_{i_k}$, viewed as an element of S_k , belongs to the subset Π . The goal is to find an ordering that maximizes the number of satisfied constraint tuples. An example of an arity 3 ordering CSP is the *Betweenness* problem with constraints of the form x_{i_2} occurs between x_{i_1} and x_{i_3} (this corresponds to the subset $\Pi = \{123, 321\}$ of S_3). More generally, one can allow more than one kind of ordering constraint, or even a payoff function $\omega_e : S_k \rightarrow \mathbb{R}^+$ for each constraint tuple e . The goal in this case is to find an ordering \mathcal{O} that maximizes $\sum_e \omega_e(\mathcal{O}|_e)$ where $\mathcal{O}|_e$ is the relative ordering of vertices in e induced by \mathcal{O} .

For the problem to decide whether all the constraints of an ordering CSP can be satisfied or not

(i.e. the decision version), [GM06] showed a dichotomy theorem for 3-ary ordering CSPs. [BK10] proved a more generalised dichotomy theorem for a broader class of problems (the temporal CSPs) of all arities.

Now we turn to the optimization problem – to maximize the number of satisfied ordering constraints. Despite much algorithmic progress on CSPs, even for MAS there was no efficient algorithm known to beat the factor $1/2$ achieved by picking a random ordering. This was explained by the recent work [GMR08] which showed that such an algorithm does not exist under the Unique Games conjecture, or in other words, MAS is approximation resistant. This hardness result was generalized to all ordering CSPs of arity 3 [CGM09], and later to higher arities, showing that every ordering CSP is approximation resistant (under the UGC) [GHM⁺11]¹

In light of this pervasive hardness of approximating ordering CSPs, in this work we ask the natural question raised by Håstad’s algorithm for bounded occurrence CSPs [Hås00], namely whether bounded occurrence instances of ordering CSPs admit a non-trivial approximation. For the case of MAS, Berger and Shor [BS97] gave an efficient algorithm that given any directed graph of total degree D , finds an ordering in which at least a fraction $(1/2 + \Omega(1/\sqrt{D}))$ of the edges go forward. This shows that bounded occurrence MAS is non-trivially approximable. The algorithm is quite simple, though its analysis is subtle. The approach is to order the vertices randomly, and process vertices in this order. When a vertex is processed, if it has more incoming edges than outgoing edges (in the graph at that stage), all outgoing edges are removed, and otherwise all its incoming edges are removed. The graph remaining after all the vertices are processed is returned as the acyclic subgraph.

Evidently, this algorithm is tailored to the MAS problem, and heavily exploits its underlying graph-theoretic structure. It therefore does not seem amenable for extensions to give non-trivial approximations to other ordering CSPs.

Our results. In this work, we prove that important special cases of ordering CSPs do admit non-trivial approximation on bounded occurrence instances. In particular, we prove this for the following classes of ordering CSPs:

1. The monotone ordering k -CSP for arbitrary k with constraints of the form $x_{i_1} < x_{i_2} < \dots < x_{i_k}$ (i.e., the CSP defined by the constraint subset $\{123 \dots k\} \subseteq S_k$ consisting of the identity permutation). This can be viewed as the arity k generalization of the MAS problem. (Note that we allow multiple constraint tuples on the same *set* of k variables, just as one would allow 2-cycles in a MAS instance given by a directed graph.)
2. All ordering CSPs of arity 3, even allowing for arbitrary payoff functions as constraints.

Our proofs show that these ordering CSPs admit an ordering into “4 slots” that beats the random ordering threshold. We remark that CSP instances which are satisfiable for orderings into n slots but do not admit good “ c slot” solutions for any fixed constant c are the basis of the Unique Games hardness results for ordering CSPs [GMR08, GHM⁺11]. Our results show that for arity 3 CSPs and monotone ordering k -ary CSPs such gap instances cannot be bounded occurrence.

¹This does not rule out non-trivial approximations for *satisfiable* instances. Of course for satisfiable instances of MAS, which correspond to DAGs, topological sorting satisfies all the constraints. For Betweenness, a factor $1/2$ approximation for satisfiable instances is known [CS98, Mak09].

Our methods. As mentioned above, the combinatorial approach of the Berger-Shor algorithm for MAS on degree-bounded graphs seems to have no obvious analog for more complicated ordering constraints. We prove our results by applying (after some adaptations) Håstad’s algorithm [Hås00] to certain “proxy” Boolean CSPs that correspond to solutions to the ordering CSP that map the variables into a domain of size 4. The idea is to only consider orderings where the range is a small constant (like [4]) instead of $[n]$. This idea was also used in recent hardness results on ordering [GMR08, CGM09, GHM⁺11]. But the fact that one can afford to restrict the range even for algorithm design (in the case of some CSPs) is a surprise.

For the case of monotone ordering constraints (of arbitrary arity k), we prove that for this proxy payoff function on the Boolean hypercube, a specific portion of the Fourier spectrum carries non-negligible mass. This is the technical core of our argument. Once we establish this, the task becomes finding a Boolean assignment to the variables of a bounded-occurrence low-degree multilinear polynomial (namely the sum of the Fourier representations of all the constraints) that evaluates to a real number that is non-negligibly larger than the constant term (which is the random assignment threshold). We present a greedy algorithm for this latter task which is similar to Håstad’s algorithm [Hås00], but yields somewhat better quantitative bounds.

Our result on general ordering 3-CSPs faces an additional complication since it can happen that the concerned part of Fourier spectrum is in fact zero for certain kinds of constraints. We identify all the cases when this troublesome phenomenon occurs, proving that in such cases the pay-off function can be expressed as a linear combination of arity 2 pay-off functions (accordingly, we call these cases as “binary representable” pay-off functions). If the binary representable portion of the pay-offs is bounded away from 1, then the remaining pay-offs (which we call “truly 3-ary”) contribute a substantial amount to the Fourier spectrum. Fortunately, the binary representable portion of pay-offs can be handled by our argument for monotone ordering constraints (specialized to arity two). So in the case when they comprise most of the constraints, we prove that their contribution to the Fourier spectrum is significant and cannot be canceled by the contribution from the truly 3-ary pay-offs.

1.1 Outline for the rest of the paper

In Section 2, we formally define the ordering CSPs with bounded occurrence, and the proxy problems (the t -ordering version). We also introduce the notation and analytic tools we will need in the remainder of the paper. In Section 3, we present an algorithm which is a variant of Håstad’s algorithm in [Hås00], and is used to solve the proxy problems. In Section 4 and Section 5, we prove the two main theorems (Theorems 1 and 2) of the paper.

2 Preliminaries

2.1 Ordering CSPs, bounded occurrence ordering CSPs

An *ordering* over vertex set V is an *injective* mapping $\mathcal{O} : V \rightarrow \mathbb{Z}^+$. An instance of k -ary monotone ordering problem $G = (V, E, \omega)$ consists of vertex set V , set E of k -tuples of distinct vertices, and

weight function $\omega : E \rightarrow \mathbb{R}^+$. The weight satisfied by ordering \mathcal{O} is

$$\text{Val}^{\mathcal{O}}(G) \stackrel{\text{def}}{=} \sum_{e=(v_{i_1}, v_{i_2}, \dots, v_{i_k}) \in E} \omega(e) \cdot \mathbf{1}_{\mathcal{O}(v_{i_1}) < \mathcal{O}(v_{i_2}) < \dots < \mathcal{O}(v_{i_k})}.$$

We also denote the value of the optimal solution by

$$\text{Val}(G) \stackrel{\text{def}}{=} \max_{\text{injective } \mathcal{O}: V \rightarrow \mathbb{Z}} \{\text{Val}^{\mathcal{O}}(G)\}.$$

We can extend the definition of the monotone ordering problem to ordering CSPs $\mathcal{I} = (V, E, \Omega)$ with general pay-off functions, where V and E are similarly defined. For each k -tuple $e = (v_1, v_2, \dots, v_k) \in E$, a general pay-off function $\omega_e \in \Omega$, mapping from all $k!$ possible orderings among $\mathcal{O}(v_1), \mathcal{O}(v_2), \dots, \mathcal{O}(v_k)$ to $\mathbb{R}^{\geq 0}$, is introduced. That is, for an ordering \mathcal{O} , its pay-off $\omega_e(\mathcal{O})$ for constraint tuple e only depends on $\mathcal{O}|_e$, the relative ordering on vertices of e induced by \mathcal{O} . The overall pay-off achieved by an ordering \mathcal{O} is defined as $\text{Val}^{\mathcal{O}}(\mathcal{I}) \stackrel{\text{def}}{=} \sum_{e \in E} \omega_e(\mathcal{O})$. The optimal pay-off for the instance is then given by

$$\text{Val}(\mathcal{I}) \stackrel{\text{def}}{=} \max_{\text{injective } \mathcal{O}: V \rightarrow \mathbb{Z}} \{\text{Val}^{\mathcal{O}}(\mathcal{I})\}.$$

An ordering CSP problem $\mathcal{I} = (V, E, \Omega)$ (or a monotone ordering problem $G = (V, E, \omega)$) is called *B-occurrence bounded* if each vertex $v \in V$ occurs in at most B tuples in E .

2.2 The t -ordering version of ordering CSPs

We start this section with several definitions. Two orderings \mathcal{O} and \mathcal{O}' are *essentially the same* if $\forall u, v \in V, \mathcal{O}(u) < \mathcal{O}(v) \Leftrightarrow \mathcal{O}'(u) < \mathcal{O}'(v)$, otherwise we call them *essentially different*. For a positive integer m , denote $[m] = \{1, 2, \dots, m\}$. For integer $t > 0$, a t -ordering on V is a mapping $\mathcal{O}_t : V \rightarrow [t]$, not necessarily injective. An ordering \mathcal{O} is *consistent* with a t -ordering \mathcal{O}_t , denoted by $\mathcal{O} \sim \mathcal{O}_t$, when $\forall u, v \in V, \mathcal{O}_t(u) < \mathcal{O}_t(v) \Rightarrow \mathcal{O}(u) < \mathcal{O}(v)$.

The monotone ordering problem G can be naturally extended to its t -ordering version, which is a regular CSP problem over domain $[t]$ defined as follows. For each constraint $e = (v_1, v_2, \dots, v_k) \in E$, we introduce a pay-off function

$$\pi_e(\mathcal{O}_t) \stackrel{\text{def}}{=} \mathbf{E}_{\mathcal{O} \sim \mathcal{O}_t} [\mathbf{1}_{\mathcal{O}(v_1) < \mathcal{O}(v_2) < \dots < \mathcal{O}(v_k)}],$$

where the expectation is uniformly taken over all the essentially different orderings \mathcal{O} that are consistent with \mathcal{O}_t . (In this paper, when \mathcal{O} becomes a random variable for total ordering without further explanation, it is always uniformly taken over all essentially different orderings (that satisfy certain criteria).) Note that although π_e receives an n -dimensional vector as parameter in the equation above, its value depends only on the k values to v_1, v_2, \dots, v_k . Then the k -ary CSP problem (t -ordering version of G) is to find the t -ordering \mathcal{O}_t to maximize the objective function

$$\text{Val}_t^{\mathcal{O}_t}(G) \stackrel{\text{def}}{=} \sum_{e \in E} \omega(e) \cdot \pi_e(\mathcal{O}_t).$$

We denote the value of the optimal solution by

$$\text{Val}_t(G) \stackrel{\text{def}}{=} \max_{\mathcal{O}_t \in [t]^n} \{\text{Val}_t^{\mathcal{O}_t}(G)\}.$$

We can also extend the ordering CSP problem \mathcal{I} with general pay-off functions to its t -ordering version. For each constraint $e \in E$, the pay-off function in the t -ordering version is defined as $\pi_e(\mathcal{O}_t) \stackrel{\text{def}}{=} \mathbf{E}_{\mathcal{O} \sim \mathcal{O}_t}[\omega_e(\mathcal{O})]$. The pay-off achieved by a particular t -ordering \mathcal{O}_t is given by $\text{Val}_t^{\mathcal{O}_t}(\mathcal{I}) \stackrel{\text{def}}{=} \sum_{e \in E} \pi_e(\mathcal{O}_t)$, and the value of the optimal t -ordering solution is $\text{Val}_t(\mathcal{I}) \stackrel{\text{def}}{=} \max_{\mathcal{O}_t \in [t]^n} \{\text{Val}_t^{\mathcal{O}_t}(\mathcal{I})\}$.

Our approach to getting a good solution for (occurrence bounded) ordering CSPs is based on the following fact.

Fact 1. For all positive integers t , $\text{Val}(\mathcal{I}) \geq \text{Val}_t(\mathcal{I})$.

Note that for $t = 1$, $\text{Val}_1(\mathcal{I})$ equals the expected pay-off of a random ordering. Since the monotone ordering problem is a special case of ordering CSP with general pay-off functions, Fact 1 is also true for the monotone ordering problem. By fact 1, it is enough to find a good solution for t -ordering version of \mathcal{I} (or G) to show that $\text{Val}(\mathcal{I})$ (or $\text{Val}(G)$) is large.

2.3 Fourier transform of Boolean functions

For every $f : \{-1, 1\}^d \rightarrow \mathbb{R}$, we write the Fourier expansion of f as

$$f(x) = \sum_{S \subseteq [d]} \hat{f}(S) \chi_S(x),$$

where $\hat{f}(S)$ is the Fourier coefficient of f on S , and $\chi_S(x) = \prod_{i \in S} x_i$.

The Fourier coefficients can be computed by the inverse Fourier transform, i.e., for every $S \subseteq [d]$,

$$\hat{f}(S) = \mathbf{E}_{x \in \{-1, 1\}^d} [f(x) \chi_S(x)].$$

3 Finding good assignments for bounded occurrence polynomials

Let f be a polynomial in n variables x_1, x_2, \dots, x_n containing only multilinear terms of degree at most k with coefficients $\hat{f}(S)$. In other words, let $f(x) = \sum_{S \subseteq [n], |S| \leq k} \hat{f}(S) \chi_S(x)$. We say that f is D -occurrence bounded if for each coordinate $i \in [n]$, we have $|\{S \ni i : \hat{f}(S) \neq 0\}| \leq D$. We also define

$$|f| \stackrel{\text{def}}{=} \sum_{\emptyset \neq S \subseteq [n]} |\hat{f}(S)|.$$

Then, the following proposition shows us how to find a good assignment for f .

Proposition 1. Given a D -occurrence bounded polynomial f of degree at most k , it is possible, in $\text{poly}(n, 2^k)$ time, to find $x \in \{-1, 1\}^n$ such that

$$f(x) \geq \hat{f}(\emptyset) + |f|/(2kD).$$

Proof. We use the following algorithm to construct x .

Algorithm. As long as $|f| > 0$, the algorithm finds a non-empty set T that maximizes $|\hat{f}(T)|$, and let $\gamma = \hat{f}(T)$. We want to make sure we get $|\gamma|$ for credit while not losing too much other terms in $|f|$.

Note that for all $\emptyset \neq U \subsetneq T$, we have $\mathbf{E}_{z \in \{-1,1\}^T: \chi_T(z) = \text{sgn}(\hat{f}(T))} [\chi_U(z)] = 0$, and therefore

$$\mathbf{E}_{z \in \{-1,1\}^T: \chi_T(z) = \text{sgn}(\hat{f}(T))} \left[\sum_{U \subseteq T} \hat{f}(U) \chi_U(z) \right] = \hat{f}(\emptyset) + |\hat{f}(T)| = \hat{f}(\emptyset) + |\gamma|.$$

We can enumerate all the $z \in \{-1,1\}^T$ such that $\chi_T(z|_T) = \text{sgn}(\hat{f}(T))$ to find a particular z^* , with $\sum_{U \subseteq T} \hat{f}(U) \chi_U(z^*) \geq \hat{f}(\emptyset) + |\gamma|$. We fix $x|_T = z^*$. For the rest of the coordinates, let $g : \{-1,1\}^{[n] \setminus T} \rightarrow \mathbb{R}$ be defined as,

$$g(y) \stackrel{\text{def}}{=} f(y, z^*), \forall y \in \{-1,1\}^{[n] \setminus T}.$$

We note that g is also a D -occurrence bounded polynomial f of degree at most k , and by fixing all the variables in T , we have

$$\hat{g}(\emptyset) = \sum_{U \subseteq T} \hat{f}(U) \chi_U(z^*) \geq \hat{f}(\emptyset) + |\gamma|.$$

On the other hand, observing that $|T| \leq k$ and $|\gamma|$ is an upper bound of all $|\hat{f}(S)|$ with $S \neq \emptyset$, we have

$$\begin{aligned} |g| &= \sum_{\emptyset \neq S \subseteq [n] \setminus T} |\hat{g}(S)| = \sum_{\emptyset \neq S \subseteq [n] \setminus T} \left| \sum_{U \subseteq T} \hat{f}(S \cup U) \chi_U(z^*) \right| \\ &\geq \sum_{\emptyset \neq S \subseteq [n] \setminus T} |\hat{f}(S)| - \sum_{\emptyset \neq S \subseteq [n] \setminus T} \sum_{\emptyset \neq U \subseteq T} |\hat{f}(S \cup U)| \\ &\geq |f| - 2 \sum_{S: S \cap T \neq \emptyset} |\hat{f}(S)| \\ &\geq |f| - 2 \sum_{i \in T} \sum_{S \ni i} |\hat{f}(S)| \geq |f| - 2|T|D|\gamma| \geq |f| - 2kD|\gamma|. \end{aligned}$$

Then we can use the two inequalities above to establish $\hat{g}(\emptyset) + |g|/(2kD) \geq \hat{f}(\emptyset) + |f|/(2kD)$.

By recursively applying this algorithm on g , we can eventually fix all the coordinates in x , and get a constant function whose value is at least $\hat{f}(\emptyset) + |f|/(2kD)$. \square

Remark 1. The algorithm is similar to Håstad's algorithm in [Hås00] but we make a greedy choice of the term $\chi_T(x)$ to satisfy (the one with the largest coefficient $|\hat{f}(T)|$) at each stage. Our analysis of the loss in $|g|$ is more direct and leads to a better quantitative bound, avoiding the loss of a "scale" factor (which divides all non-zero coefficients of the polynomial) in the advantage over $\hat{f}(\emptyset)$.

Remark 2. In the performance guarantee of the algorithm, $\hat{f}(\emptyset)$ corresponds to the value of random assignments in the later sections, while $|f|/(2kD)$ corresponds the advantage we get over random assignments. Because of the $1/D$ factor, our algorithm gives weaker guarantee than Berger-Shor gives, but our algorithm extends to permutation CSPs of larger arities.

4 Bounded occurrence monotone ordering problem

Our main result in this section is the following.

Theorem 1. *For any constant $k > 1$, given a B -occurrence bounded k -ary monotone ordering problem $G = (V, E, \omega)$, there is a poly-time randomized algorithm to find a solution satisfying at least $\text{Val}(G)(1/k! + \Omega_k(1/B))$ weight (in expectation).*

To prove the above theorem, we will show the following lemma.

Lemma 1. *For any constant $k > 1$, given a B -occurrence bounded k -ary monotone ordering problem $G = (V, E, \omega)$ with total weight W . Then it is possible, in polynomial time, to find a 4-ordering solution \mathcal{O}_4 with $\text{Val}(G)(1/k! + \Omega_k(1/B))$ weight.*

Note that given Lemma 1, the randomized algorithm that samples ordering $\mathcal{O} \sim \mathcal{O}_4$ fulfills the task promised in the theorem.

Lemma 1 also implies the following.

Corollary 1. *For any B -occurrence bounded k -ary monotone ordering problem G , we have $\text{Val}_4(G) \geq \text{Val}(G)(1/k! + \Omega_k(1/B))$.*

Proof of Lemma 1. We begin the proof with the analysis of the pay-off function $\pi_e : [4]^{\{v_1, v_2, \dots, v_k\}} \rightarrow \mathbb{R}$ for some $e = (v_1, v_2, \dots, v_k) \in E$. We can also view π_e as a real-valued function defined on Boolean cube $\{-1, 1\}^{2k}$, so that

$$\begin{aligned} \pi_e(x_1, x_2, \dots, x_{2k}) \\ = \pi_e\left((1 - x_1) + \frac{(1 - x_2)}{2} + 1, \dots, (1 - x_{2k-1}) + \frac{(1 - x_{2k})}{2} + 1\right). \end{aligned}$$

If we let $\Gamma(e)$ be the set of all $k!$ permutations of e , then

$$\begin{aligned} & \sum_{e' \in \Gamma(e)} \mathbf{E}_{\mathcal{O}_4 \in [4]^k} [\pi_{e'}(\mathcal{O}_4)] \\ = & \sum_{e' = (v_{i_1}, v_{i_2}, \dots, v_{i_k}) \in \Gamma(e)} \mathbf{E}_{\mathcal{O}_4 \in [4]^k} \left[\mathbf{E}_{\mathcal{O} \sim \mathcal{O}_4} [\mathbf{1}_{\mathcal{O}(v_{i_1}) < \mathcal{O}(v_{i_2}) < \dots < \mathcal{O}(v_{i_k})}] \right] \\ = & \mathbf{E}_{\mathcal{O}_4 \in [4]^k} \left[\mathbf{E}_{\mathcal{O} \sim \mathcal{O}_4} \left[\sum_{e' = (v_{i_1}, v_{i_2}, \dots, v_{i_k}) \in \Gamma(e)} \mathbf{1}_{\mathcal{O}(v_{i_1}) < \mathcal{O}(v_{i_2}) < \dots < \mathcal{O}(v_{i_k})} \right] \right] = 1. \end{aligned}$$

Since $\mathbf{E}_{\mathcal{O}_4 \in [4]^k} [\pi_{e'}(\mathcal{O}_4)]$ is the same for all $e' \in \Gamma(e)$, we know that $\mathbf{E}_{\mathcal{O}_4 \in [4]^k} [\pi_e(\mathcal{O}_4)] = 1/k!$. Hence we have the following fact.

Fact 2. $\widehat{\pi}_e(\emptyset) = \mathbf{E}_{x \in \{-1, 1\}^{2k}} [\pi_e(x)] = \mathbf{E}_{\mathcal{O}_4 \in [4]^k} [\pi_e(\mathcal{O}_4)] = \frac{1}{k!}$.

By Fact 2, if we apply the algorithm in Proposition 1, to the objective function $f(x) = \sum_{e \in E} \omega(e) \pi_e(x)$ of the 4-ordering version, we are guaranteed to have a solution that is no worse than the random

threshold $(1/k!)$. Then, we only need to identify some non-negligible weights on the rest of the Fourier spectrum of f .

Let $S_{\text{odd}} = \{2i - 1 | i \in [k]\}$, and $S_{\text{odd}}^+ = S_{\text{odd}} \cup \{2k\}$. We make the following claim which we will prove in Appendix A.

Claim 1. $\widehat{\pi}_e(S_{\text{odd}}^+) = \frac{-2^{1-k} + 2^{2-2k}}{k!}$.

The above claim makes sure there is indeed non-negligible mass on non-empty-set Fourier coefficients for each constraint. Then we prove that, when summing up these constraints, either of the following two cases happens.

- Some weights shown in Claim 1 are not canceled by others, and finally appears in the non-empty-set Fourier coefficients for the final objective function f .
- Some weights are canceled by others, but in this case, the guarantee by $\widehat{f}(\emptyset)$ itself beats $1/k!$ in terms of approximation ratio.

We define

$$\widehat{\|\pi_e\|} = \sum_{S \subseteq [2k]: \forall i \in [k], S \cap \{2i-1, 2i\} \neq \emptyset} |\widehat{\pi}_e(S)|.$$

Now Claim 1 implies $\widehat{\|\pi_e\|} = \Omega_k(1)$ for all $k \geq 2$. Let $\Gamma \subseteq E$ be a set of constraints sharing the same $\Gamma(e)$, and let us define, by abusing notation slightly

$$\omega(\Gamma) = \sum_{e \in \Gamma} \omega(e), \quad \omega_{\max}(\Gamma) = \max_{e \in \Gamma} \{\omega(e)\}, \quad \text{and} \quad \pi_\Gamma(x) = \sum_{e \in \Gamma} \omega(e) \pi_e(x).$$

We treat $\pi_\Gamma : \{-1, 1\}^{2k} \rightarrow \mathbb{R}$ as a real-valued function defined on a Boolean cube.

The idea of defining $\widehat{\|\pi_e\|}$ and Γ is as follows. The Fourier mass identified in Claim 1 could be canceled within Γ , but once the mass goes into $\widehat{\|\pi_\Gamma\|}$, it cannot be canceled by $\widehat{\|\pi_{\Gamma'}\|}$ for a different Γ' , and will finally go into $|f|$. Then the following lemma shows that either $\widehat{\|\pi_\Gamma\|}$, or $\widehat{\pi_\Gamma}(\emptyset)$ alone, beats $\omega_{\max}(\Gamma)/k!$, where $\omega_{\max}(\Gamma)$ is an upperbound on the optimal solution's performance on the constraints in Γ .

Lemma 2. For all $\alpha, 0 < \alpha < 1$, we have $\widehat{\pi_\Gamma}(\emptyset) + \alpha \widehat{\|\pi_\Gamma\|} \geq \omega_{\max}(\Gamma) \left(\frac{1}{k!} + \alpha \cdot \Omega_k(1) \right)$.

Proof. First, by Fact 2, we know that $\widehat{\pi_\Gamma}(\emptyset) = \sum_{e \in \Gamma} \widehat{\pi}_e(\emptyset) = \omega(\Gamma) \cdot \frac{1}{k!}$. Let $e^* \in \Gamma$ be the constraint

with the most weight. If $\omega(e^*) = \omega_{\max}(\Gamma) \geq 2/3 \cdot \omega(\Gamma)$, we have

$$\begin{aligned}
\widehat{\pi}_\Gamma(\emptyset) + \alpha \|\widehat{\pi}_\Gamma\| &= \widehat{\pi}_\Gamma(\emptyset) + \alpha \left\| \sum_{e \in \Gamma} \omega(e) \pi_e \right\| \\
&\geq \widehat{\pi}_\Gamma(\emptyset) + \alpha \left(\|\omega(e^*) \pi_{e^*}\| - \sum_{e \in \Gamma \setminus \{e^*\}} \|\omega(e) \pi_e\| \right) \\
&= \omega(\Gamma) \cdot \frac{1}{k!} + \alpha \left(\omega(e^*) - \sum_{e \in \Gamma \setminus \{e^*\}} \omega(e) \right) \|\pi_{e^*}\| \\
&\geq \omega_{\max}(\Gamma) \left(\frac{1}{k!} + \frac{\alpha}{2} \|\pi_{e^*}\| \right) = \omega_{\max}(\Gamma) \left(\frac{1}{k!} + \alpha \cdot \Omega_k(1) \right).
\end{aligned}$$

where the last step follows from Claim 1. On the other hand, when $\omega(e^*) = \omega_{\max}(\Gamma) < 2/3 \cdot \omega(\Gamma)$,

$$\begin{aligned}
\widehat{\pi}_\Gamma(\emptyset) + \alpha \|\widehat{\pi}_\Gamma\| &\geq \widehat{\pi}_\Gamma(\emptyset) = \omega(\Gamma) \cdot \frac{1}{k!} \\
&> \omega_{\max}(\Gamma) \left(\frac{1}{k!} + \frac{1}{2} \cdot \frac{1}{k!} \right) = \omega_{\max}(\Gamma) \left(\frac{1}{k!} + \Omega_k(1) \right). \quad \square
\end{aligned}$$

Given a k -ary monotone ordering problem $G = (V, E, \omega)$, we partition $E = \Gamma_1 \cup \Gamma_2 \cup \dots \cup \Gamma_m$ into m disjoint groups, so that constraints e_j in each group Γ_i share a distinct $\Gamma(e)$ value. Then we write the objective function of its 4-ordering version as

$$f(x) = \sum_{e \in E} \omega(e) \pi_e(x) = \sum_{i=1}^m \sum_{e \in \Gamma_i} \pi_e(x) = \sum_{i=1}^m \pi_{\Gamma_i}(x),$$

where $f : \{-1, 1\}^{2n} \rightarrow \mathbb{R}$ is defined on Boolean cube. For each $1 \leq i \leq m$, let $\{v_{i_1}, v_{i_2}, \dots, v_{i_k}\}$ be the k vertices participating in Γ_i , then we note that for each $S \in \{2i_t - 1, 2i_t : t \in [k]\}$ that intersects with $\{2i_t - 1, 2i_t\}$ for each $t \in k$, we have $\hat{f}(S) = \widehat{\pi}_{\Gamma_i}(S)$, since all other constraints will have 0 as its Fourier coefficient over S . Then, for $\alpha \in (0, 1)$, we have

$$\hat{f}(\emptyset) + \alpha \cdot |f| \geq \sum_{i=1}^m \left(\widehat{\pi}_{\Gamma_i}(\emptyset) + \alpha \cdot \|\widehat{\pi}_{\Gamma_i}\| \right) \geq \sum_{i=1}^m \omega_{\max}(\Gamma_i) \left(\frac{1}{k!} + \alpha \cdot \Omega_k(1) \right), \quad (1)$$

where the last inequality is because of Lemma 2.

For each $\Gamma_i (1 \leq i \leq m)$, a total ordering \mathcal{O} will satisfy at most $\omega_{\max}(\Gamma_i)$ weight of constraints. This give an upper bound of the optimal solution

$$\text{Val}(G) \leq \sum_{i=1}^m \omega_{\max}(\Gamma_i). \quad (2)$$

Fix a coordinate $i \in [2k]$, each constraint π_e contributes at most 2^{k-1} non-zero Fourier coefficients containing i . Since $G = (V, E, \omega)$ is B -occurrence bounded, there are at most $B2^{k-1}$ non-zero Fourier coefficients of f containing i , therefore f is $B2^{k-1}$ -occurrence bounded.

Applying Proposition 1 to f , the polynomial time algorithm gets a vector $x \in \{-1, 1\}^{2n}$, which corresponds to a 4-ordering \mathcal{O}_4 (recall that every two consecutive bits in x encode a value in $[4]$),

such that

$$\begin{aligned} \text{Val}_4^{\mathcal{O}_4}(G) = f(x) &\geq \hat{f}(\emptyset) + \frac{1}{k2^k B} |f| \\ &\geq \sum_{i=1}^m \omega_{\max}(\Gamma_i) \left(\frac{1}{k!} + \Omega_k(1/B) \right) \geq \text{Val}(G) \left(\frac{1}{k!} + \Omega_k(1/B) \right), \end{aligned}$$

where the last two inequalities use (1) and (2) separately. \square

5 Bounded occurrence 3-ary ordering CSP with general pay-off functions

For a ordering CSP problem $\mathcal{I} = (V, E, \Omega)$ with general pay-off functions, we define

$$\text{Rand}(\mathcal{I}) \stackrel{\text{def}}{=} \mathbf{E}_{\text{injective } \mathcal{O}: V \rightarrow \mathbb{Z}} [\text{Val}^{\mathcal{O}}(\mathcal{I})],$$

as the performance of random ordering. Then we prove our main result for 3-ary ordering CSPs:

Theorem 2. *Given a B -occurrence bounded 3-ary ordering CSP problem $\mathcal{I} = (V, E, \Omega)$ with general pay-off functions, there is a poly-time randomized algorithm, to find a solution satisfying at least $\text{Rand}(\mathcal{I}) + (\text{Val}(\mathcal{I}) - \text{Rand}(\mathcal{I})) \cdot \Omega(1/B)$ weight (in expectation).*

To prove Theorem 2, it is enough to prove the following lemma.

Lemma 3. *Given a B -occurrence bounded 3-ary ordering CSP problem $\mathcal{I} = (V, E, \Omega)$ with general pay-off functions, it is possible, in polynomial time, to find a 4-ordering solution \mathcal{O}_4 with $\text{Rand}(\mathcal{I}) + (\text{Val}(\mathcal{I}) - \text{Rand}(\mathcal{I})) \cdot \Omega(1/B)$ weight.*

We call a set E of constraints *simple set* if there are no two constraints $e_1, e_2 \in E$ with $\Gamma(e_1) = \Gamma(e_2)$. We can assume in the proof that E is simple, or we can combine the two constraints sharing the same $\Gamma(\cdot)$ into a new constraint (the new pay-off function is just an addition of two old pay-off functions, perhaps with some permutations), and this does not increase the occurrence bound B .

Ideal proof sketch of Lemma 3. Similarly as we did with monotone ordering problems, our ideal goal is to first show that for each constraint $e \in E$, the 4-ordering pay-off function π_e ensures that $\|\hat{\pi}_e\|$ is proportional to the maximum possible value of ω_e , specifically to argue that there exists some constant $c > 0$, such that $\|\hat{\pi}_e\| \geq c(\max_{\mathcal{O}}\{\omega_e(\mathcal{O})\} - \mathbf{E}_{\mathcal{O}}[\omega_e(\mathcal{O})])$. Then, because this part of the Fourier spectrum cannot be canceled with coefficients of other constraints, they will appear in the objective function $f(x) = \sum_{e \in E} \pi_e(x)$. This will give a good lower bound on $|f|$, as follows:

$$\begin{aligned} |f| &\geq \sum_{e \in E} \|\hat{\pi}_e\| \geq \sum_{e \in E} c(\max_{\mathcal{O}}\{\omega_e(\mathcal{O})\} - \mathbf{E}_{\mathcal{O}}[\omega_e(\mathcal{O})]) \\ &= c \left(\sum_{e \in E} \max_{\mathcal{O}}\{\omega_e(\mathcal{O})\} - \mathbf{E}_{\mathcal{O}} \left[\sum_e \omega_e(\mathcal{O}) \right] \right) \geq c(\text{Val}(\mathcal{I}) - \text{Rand}(\mathcal{I})). \end{aligned}$$

At this point, we can use Proposition 1 to get a non-trivial gain over random solution.

But unfortunately, sometimes $\hat{\|\pi_e\|}$ can be 0 even when there is a large gap between $\max_{\mathcal{O}}\{\omega_e(\mathcal{O})\}$ and $\mathbf{E}_{\mathcal{O}}[\omega_e(\mathcal{O})]$.

Fact 3. Let $e = (v_i, v_j, v_k)$, the following pay-off functions are such kind of examples (for the statement above).

- $\omega_e(\mathcal{O}) = \mathbf{1}_{\mathcal{O}(v_i) < \mathcal{O}(v_j)}$, and $\omega_e(\mathcal{O}) = \mathbf{1}_{\mathcal{O}(v_i) > \mathcal{O}(v_j)}$;
- $\omega_e(\mathcal{O}) = \mathbf{1}_{\mathcal{O}(v_j) < \mathcal{O}(v_k)}$, and $\omega_e(\mathcal{O}) = \mathbf{1}_{\mathcal{O}(v_k) > \mathcal{O}(v_j)}$;
- $\omega_e(\mathcal{O}) = \mathbf{1}_{\mathcal{O}(v_k) < \mathcal{O}(v_i)}$, and $\omega_e(\mathcal{O}) = \mathbf{1}_{\mathcal{O}(v_i) > \mathcal{O}(v_k)}$;
- $\omega_e(\mathcal{O}) = \mathbf{1}_{\mathcal{O}(v_i) < \mathcal{O}(v_j) < \mathcal{O}(v_k)} + \mathbf{1}_{\mathcal{O}(v_j) < \mathcal{O}(v_k) < \mathcal{O}(v_i)} + \mathbf{1}_{\mathcal{O}(v_k) < \mathcal{O}(v_i) < \mathcal{O}(v_j)}$, and
 $\omega_e(\mathcal{O}) = \mathbf{1}_{\mathcal{O}(v_k) < \mathcal{O}(v_j) < \mathcal{O}(v_i)} + \mathbf{1}_{\mathcal{O}(v_i) < \mathcal{O}(v_k) < \mathcal{O}(v_j)} + \mathbf{1}_{\mathcal{O}(v_j) < \mathcal{O}(v_i) < \mathcal{O}(v_k)}$.

It is easy to see that first three pairs of pay-off functions have $\hat{\|\pi_e\|} = 0$ since they only depend on two out of three coordinates. For the last pair, note we can rewrite them as summation of pay-off functions that depend on no more than two coordinates: $\omega_e(\mathcal{O}) = \mathbf{1}_{\mathcal{O}(v_i) < \mathcal{O}(v_j)} + \mathbf{1}_{\mathcal{O}(v_j) < \mathcal{O}(v_k)} + \mathbf{1}_{\mathcal{O}(v_k) < \mathcal{O}(v_i)} - 1$ and $\omega_e(\mathcal{O}) = \mathbf{1}_{\mathcal{O}(v_k) < \mathcal{O}(v_j)} + \mathbf{1}_{\mathcal{O}(v_j) < \mathcal{O}(v_i)} + \mathbf{1}_{\mathcal{O}(v_i) < \mathcal{O}(v_k)} - 1$, respectively. We call the pay-off functions shown in Fact 3 as *binary representable* pay-off functions.

Proof of Lemma 3. In view of the above, we claim the following crucial lemma which shows that the above binary-representable pay-offs are the only obstacles to our plan. The proof of the lemma is deferred to Appendix B.

Lemma 4. *There exists a constant $c > 0$, such that for every 3-ary pay-off function ω_e on some triple $e = (v_i, v_j, v_k)$ the following holds. If for every binary representable pay-off function ω , there exists some ordering \mathcal{O} of v_i, v_j, v_k such that $\omega(\mathcal{O}) > 0$ and $\omega_e(\mathcal{O}) = 0$ (i.e., ω_e does not “contain” any binary representable pay-off functions), then*

$$\hat{\|\pi_e\|} \geq c \cdot \max_{\mathcal{O}}\{\omega_e(\mathcal{O})\}.$$

Using Lemma 4, it is not hard to see that we can rewrite the set E of constraints (together with their pay-off functions) as $E_3 \cup E_2$ together with a constant $b \in \mathbb{R}$, such that

$$\sum_{e \in E} \omega_e = \sum_{e_3 \in E_3} \omega_{e_3} + \sum_{e_2 \in E_2} \omega_{e_2} + b,$$

where E_3 is a simple set of 3-ary pay-off functions none of which contain binary representable pay-off functions, and E_2 is a simple set of (weighted) monotone 2-ordering constraints.

The constraints in E_2 can be handled by virtue of Claim 1, since they just correspond to monotone k -ordering constraints for $k = 2$.

Corollary 2 (of Claim 1). *For every $e_2 \in E_2$, $\hat{\|\pi_{e_2}\|} \geq 1/2 \cdot \max_{\mathcal{O}}\{\omega_{e_2}(\mathcal{O})\}$.*

We emphasize that pay-off functions are always non-negative, and therefore

Fact 4. $\text{Rand}(\mathcal{I}) = \mathbf{E}_{\mathcal{O}} \left[\sum_{e \in E} \omega_e(\mathcal{O}) \right] = \mathbf{E}_{\mathcal{O}} \left[\sum_{e_3 \in E_3} \omega_{e_3}(\mathcal{O}) + \sum_{e_2 \in E_2} \omega_{e_2}(\mathcal{O}) + b \right] \geq b.$

Remember our 4-ary objective function $f : [4]^n \rightarrow \mathbb{R}$ is

$$f(\mathcal{O}_4) = \mathbf{E}_{\mathcal{O} \sim \mathcal{O}_4} \left[\sum_{e \in E} \omega_e(\mathcal{O}) \right] = \sum_{e_3 \in E_3} \pi_{e_3}(\mathcal{O}_4) + \sum_{e_2 \in E_2} \pi_{e_2}(\mathcal{O}_4) + b,$$

by a similar argument as Fact 2, we know that, viewing f as a function defined on Boolean cube, we have

Fact 5. $\hat{f}(\emptyset) = \mathbf{E}_x[f(x)] = \mathbf{E}_{\mathcal{O}} \left[\sum_{e \in E} \omega_e(\mathcal{O}) \right] = \text{Rand}(\mathcal{I}).$

Since f is a $O(B)$ -occurrence bounded polynomial, in order to apply Proposition 1 to finish this proof, we only need to prove that

$$\left| \sum_{e_3 \in E_3} \pi_{e_3} + \sum_{e_2 \in E_2} \pi_{e_2} \right| \geq \left(\sum_{e_3 \in E_3} \max_{\mathcal{O}} \{\omega_{e_3}(\mathcal{O})\} + \sum_{e_2 \in E_2} \max_{\mathcal{O}} \{\omega_{e_2}(\mathcal{O})\} \right) \cdot \Omega(1), \quad (3)$$

as this would imply that

$$|f| = \left| \sum_{e_3 \in E_3} \pi_{e_3} + \sum_{e_2 \in E_2} \pi_{e_2} \right| \geq (\text{Val}(\mathcal{I}) - b) \cdot \Omega(1) \geq (\text{Val}(\mathcal{I}) - \text{Rand}(\mathcal{I})) \cdot \Omega(1),$$

where the last step used Fact 4.

To prove (3), we first establish the following upper bound.

Fact 6. For each $e_3 \in E_3$, we have

$$\begin{aligned} |\pi_{e_3}| &\leq \sum_{S \subseteq [6]} |\widehat{\pi}_{e_3}(S)| \leq \sqrt{2^6 \left(\sum_{S \subseteq [6]} \widehat{\pi}_{e_3}(S)^2 \right)} \\ &= 8 \sqrt{\mathbf{E}_{x \in \{-1,1\}^6} [\pi_{e_3}(x)^2]} \leq 8 \max_{x \in \{-1,1\}^6} \{\pi_{e_3}(x)\} \leq 8 \max_{\mathcal{O}} \{\omega_{e_3}(\mathcal{O})\}. \end{aligned}$$

We discuss the following two cases to establish (3).

When $\sum_{e_3 \in E_3} \max_{\mathcal{O}} \{\omega_{e_3}(\mathcal{O})\} \geq \frac{1}{32} \cdot \sum_{e_2 \in E_2} \max_{\mathcal{O}} \{\omega_{e_2}(\mathcal{O})\}$, by Lemma 4, we have

$$\begin{aligned} \left| \sum_{e_3 \in E_3} \pi_{e_3} + \sum_{e_2 \in E_2} \pi_{e_2} \right| &\geq \sum_{e_3 \in E_3} \|\widehat{\pi}_{e_3}\| \geq c \sum_{e_3 \in E_3} \max_{\mathcal{O}} \{\omega_{e_3}(\mathcal{O})\} \\ &\geq \frac{c}{33} \cdot \left(\sum_{e_3 \in E_3} \max_{\mathcal{O}} \{\omega_{e_3}(\mathcal{O})\} + \sum_{e_2 \in E_2} \max_{\mathcal{O}} \{\omega_{e_2}(\mathcal{O})\} \right). \end{aligned}$$

On the other hand, when $\sum_{e_3 \in E_3} \max_{\mathcal{O}} \{\omega_{e_3}(\mathcal{O})\} < 1/32 \cdot \sum_{e_2 \in E_2} \max_{\mathcal{O}} \{\omega_{e_2}(\mathcal{O})\}$, by Corollary 2 and Fact 6,

$$\begin{aligned}
& \left| \sum_{e_3 \in E_3} \pi_{e_3} + \sum_{e_2 \in E_2} \pi_{e_2} \right| \geq \sum_{e_2 \in E_2} \hat{\|\pi_{e_2}\|} - \sum_{e_3 \in E_3} |\pi_{e_3}| \\
& \geq \frac{1}{2} \sum_{e_2 \in E_2} \max_{\mathcal{O}} \{\omega_{e_2}(\mathcal{O})\} - 8 \sum_{e_3 \in E_3} \max_{\mathcal{O}} \{\omega_{e_3}(\mathcal{O})\} \geq \left(\frac{1}{2} - \frac{1}{4}\right) \sum_{e_2 \in E_2} \max_{\mathcal{O}} \{\omega_{e_2}(\mathcal{O})\} \\
& \geq \frac{1}{4} \cdot \frac{32}{33} \cdot \left(\sum_{e_3 \in E_3} \max_{\mathcal{O}} \{\omega_{e_3}(\mathcal{O})\} + \sum_{e_2 \in E_2} \max_{\mathcal{O}} \{\omega_{e_2}(\mathcal{O})\} \right). \quad \square \text{ (Lemma 3)}
\end{aligned}$$

6 Concluding remarks

In this paper, we investigated the problem whether there are non-trivial approximation algorithms for bounded occurrence ordering CSPs. By reducing the problem to a CSP over a fixed size domain, and applying a variant of Håstad's algorithm [Hås00], we give a positive answer for monotone ordering problems, and 3-ary ordering CSPs. The obvious open question left by our work is whether we can extend the technique presented in this paper to ordering CSPs with arbitrary arity. Given our approach, the following natural question arises in this vein: given maximum occurrence B , and arity k , does there exist a constant $t = t(B, k)$ so that it is enough to solve the t -ordering version to get a non-trivial approximate solution for the original ordering CSP? At first glance, one might believe that the answer to this question is "no", as t being independent of n seems too strong a restriction. But as Lemma 1 showed, even under a stronger restriction that $t = 4$ (which is independent of k as well), the answer is still "yes" for monotone ordering problems. In view of this special case, we believe that there is a generalization of Lemma 1 and Lemma 3 to general bounded occurrence ordering CSPs, and leave the resolution of this as an open problem.

References

- [AM09] Per Austrin and Elchanan Mossel. Approximation resistant predicates from pairwise independence. *Computational Complexity*, 18(2):249–271, 2009. Preliminary version in CCC'08. 2
- [BK10] Manuel Bodirsky and Jan Kára. The complexity of temporal constraint satisfaction problems. *J. ACM*, 57:9:1–9:41, February 2010. 3
- [BS97] Bonnie Berger and Peter W. Shor. Tight bounds for the Maximum Acyclic Subgraph problem. *J. Algorithms*, 25(1):1–18, 1997. 3
- [CGM09] Moses Charikar, Venkatesan Guruswami, and Rajsekar Manokaran. Every permutation CSP of arity 3 is approximation resistant. In *Proceedings of the 24th IEEE Conference on Computational Complexity*, pages 62–73, July 2009. 3, 4
- [CS98] Benny Chor and Madhu Sudan. A geometric approach to betweenness. *SIAM J. Discrete Math.*, 11(4):511–523, 1998. 3

- [EG04] Lars Engebretsen and Venkatesan Guruswami. Is constraint satisfaction over two variables always easy? *Random Structures and Algorithms*, 25(2):150–178, 2004. 2
- [GHM⁺11] Venkatesan Guruswami, Johan Håstad, Rajsekar Manokaran, Prasad Raghavendra, and Moses Charikar. Beating the random ordering is hard: Every ordering CSP is approximation resistant. *SIAM J. Comput.*, 40(3):878–914, 2011. 3, 4
- [GM06] Walter Guttmann and Markus Maucher. Variations on an ordering theme with constraints. In Gonzalo Navarro, Leopoldo E. Bertossi, and Yoshiharu Kohayakawa, editors, *IFIP TCS*, volume 209 of *IFIP*, pages 77–90. Springer, 2006. 3
- [GMR08] Venkatesan Guruswami, Rajsekar Manokaran, and Prasad Raghavendra. Beating the random ordering hard: Inapproximability of maximum acyclic subgraph. In *Proceedings of the 49th IEEE Symposium on Foundations of Computer Science*, pages 573–582, 2008. 3, 4
- [GW95] Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, 1995. 2
- [Hås00] Johan Håstad. On bounded occurrence constraint satisfaction. *Inf. Process. Lett.*, 74(1-2):1–6, 2000. 2, 3, 4, 7, 14
- [Hås01] Johan Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, 2001. 2
- [Hås08] Johan Håstad. Every 2-CSP allows nontrivial approximation. *Computational Complexity*, 17(4):549–566, 2008. 2
- [Mak09] Yury Makarychev. Simple linear time approximation algorithm for betweenness. *Microsoft Research Technical Report MSR-TR-2009-74*, 2009. 3
- [Rag08] Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the 40th ACM Symposium on Theory of Computing*, pages 245–254, 2008. 2
- [Tre01] Luca Trevisan. Non-approximability results for optimization problems on bounded degree instances. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pages 453–461, 2001. 2

A Proof of Claim 1

In this section, we prove that for monotone k -ary constraint π_e ,

Claim 1 (restated). $\hat{\pi}_e(S_{\text{odd}}^+) = \frac{-2^{1-k} + 2^{2-2k}}{k!}$.

Proof. Since the coordinates of π_e has the same order as they are shown in e , we can write the value of π_e explicitly as

$$\pi_e(x) = \begin{cases} \frac{1}{n_1!n_2!n_3!n_4!} & \text{when } x = (1, 1)^{n_1} \circ (1, -1)^{n_2} \circ (-1, 1)^{n_3} \circ (-1, -1)^{n_4} \\ 0 & \text{otherwise} \end{cases}.$$

Therefore

$$\begin{aligned} \widehat{\pi}_e(S_{\text{odd}}^+) &= \mathbf{E}_{x \in \{-1, 1\}^{2k}} [\pi_e(x) \chi_{S_{\text{odd}}^+}(x)] \\ &= \frac{1}{4^k} \sum_{n_1+n_2+n_3+n_4=k} \frac{(-1)^{n_3+n_4}}{n_1!n_2!n_3!n_4!} \cdot (-1)^{\mathbf{1}_{n_4>0} \vee (n_3=n_4=0 \wedge n_2>0)} \\ &= \frac{1}{4^k} \left(- \sum_{n_1+n_2+n_3+n_4=k, n_4>0} \frac{(-1)^{n_3+n_4}}{n_1!n_2!n_3!n_4!} \right. \\ &\quad \left. + \sum_{n_1+n_2+n_3=k, n_3>0} \frac{(-1)^{n_3}}{n_1!n_2!n_3!} - \sum_{n_1+n_2=k, n_2>0} \frac{1}{n_1!n_2!} + \frac{1}{k!} \right). \end{aligned}$$

This is just the k -th coefficient of the polynomial

$$\frac{1}{4^k} \left(-e^{2x} e^{-x} (e^{-x} - 1) + e^{2x} (e^{-x} - 1) - e^x (e^x - 1) + e^x \right) = \frac{-2e^{2x} + 4e^x - 1}{4^k}.$$

Thus, for $k > 0$, we have $\widehat{\pi}_e(S_{\text{odd}}^+) = \frac{-2^{1-k} + 2^{2-2k}}{k!}$. □

B Proof of Lemma 4

For reader's convenience, we restate the lemma as follows.

Lemma 4 (restated). *There exists a constant $c > 0$, such that for every 3-ary pay-off function ω_e on some triple $e = (v_i, v_j, v_k)$ the following holds. If for every binary representable pay-off function ω , there exists some ordering \mathcal{O} of v_i, v_j, v_k such that $\omega(\mathcal{O}) > 0$ and $\omega_e(\mathcal{O}) = 0$ (i.e., ω_e does not "contain" any binary representable pay-off functions), then*

$$\|\widehat{\pi}_e\| \geq c \cdot \max_{\mathcal{O}} \{\omega_e(\mathcal{O})\}.$$

Proof. W.l.o.g. suppose that the 3-ary constraint $e = (v_1, v_2, v_3)$, and its pay-off function

$$\begin{aligned} \omega_e(\mathcal{O}) &= a_1 \cdot \mathbf{1}_{\mathcal{O}(v_1) < \mathcal{O}(v_2) < \mathcal{O}(v_3)} + a_2 \cdot \mathbf{1}_{\mathcal{O}(v_1) < \mathcal{O}(v_3) < \mathcal{O}(v_2)} + a_3 \cdot \mathbf{1}_{\mathcal{O}(v_2) < \mathcal{O}(v_1) < \mathcal{O}(v_3)} \\ &\quad + a_4 \cdot \mathbf{1}_{\mathcal{O}(v_2) < \mathcal{O}(v_3) < \mathcal{O}(v_1)} + a_5 \cdot \mathbf{1}_{\mathcal{O}(v_3) < \mathcal{O}(v_1) < \mathcal{O}(v_2)} + a_6 \cdot \mathbf{1}_{\mathcal{O}(v_3) < \mathcal{O}(v_2) < \mathcal{O}(v_1)}. \end{aligned}$$

We can check by definition that

$$\begin{aligned} \widehat{\pi}_e(\{1, 2, 4, 6\}) &= 12/(2^6 \cdot 3!) \cdot (-a_1 - a_2 + 2a_3 - a_4 + 2a_5 - a_6), \\ \widehat{\pi}_e(\{2, 3, 4, 6\}) &= 12/(2^6 \cdot 3!) \cdot (2a_1 - a_2 - a_3 - a_4 - a_5 + 2a_6), \\ \widehat{\pi}_e(\{2, 4, 5, 6\}) &= 12/(2^6 \cdot 3!) \cdot (-a_1 + 2a_2 - a_3 + 2a_4 - a_5 - a_6). \end{aligned}$$

Let $\mathbf{a} = (a_1, a_2, a_3, a_4, a_5, a_6) \in (\mathbb{R}^{\geq 0})^6$, and $\mathbf{u}_1 = 1/2\sqrt{3} \cdot (-1, -1, 2, -1, 2, -1)$, $\mathbf{u}_2 = 1/2\sqrt{3} \cdot (2, -1, -1, -1, -1, 2)$, $\mathbf{u}_3 = 1/2\sqrt{3} \cdot (-1, 2, -1, 2, -1, -1)$ be unit vectors. Let $\text{span}(\mathbf{v})$ be the linear space spanned by vector \mathbf{v} . For any linear space $D \subseteq \mathbb{R}^6$, denote by D^\perp the orthogonal complement of D (in \mathbb{R}^6). Use $\text{dist}(\mathbf{v}, D)$ denote the Euclidean distance from vector \mathbf{v} to space D . By the above identities

$$\begin{aligned} \|\hat{\pi}_e\| &\geq |\hat{\pi}_e(\{1, 2, 4, 6\})| + |\hat{\pi}_e(\{2, 3, 4, 6\})| + |\hat{\pi}_e(\{2, 4, 5, 6\})| \\ &= 1/32 \cdot (\text{dist}(\mathbf{a}, \text{span}(\mathbf{u}_1)^\perp) + \text{dist}(\mathbf{a}, \text{span}(\mathbf{u}_2)^\perp) + \text{dist}(\mathbf{a}, \text{span}(\mathbf{u}_3)^\perp)). \end{aligned} \quad (4)$$

Since ω_e does not include binary representable pay-off functions listed in Fact 3, we know that \mathbf{a} belongs to the set A defined by

$$\begin{aligned} A = \{ \mathbf{x} = (x_1, x_2, x_3, x_4, x_5, x_6) \in (\mathbb{R}^{\geq 0})^6 : x_1x_2x_5 = x_3x_4x_6 = 0, \\ x_1x_2x_3 = x_4x_5x_6 = 0, x_1x_3x_4 = x_2x_5x_6 = 0, x_1x_4x_5 = x_2x_3x_6 = 0 \}. \end{aligned}$$

One can check that

$$\begin{aligned} A = & \{ (x_1, x_2, 0, x_4, 0, x_6) : x_1, x_2, x_4, x_6 \geq 0 \} \\ & \cup \{ (x_1, 0, x_3, 0, x_5, x_6) : x_1, x_3, x_5, x_6 \geq 0 \} \\ & \cup \{ (0, x_2, x_3, x_4, x_5, 0) : x_2, x_3, x_4, x_5 \geq 0 \}. \end{aligned} \quad (5)$$

We now prove the following fact which will help us lower bound (4).

Lemma 5. *Given a linear space $K \subseteq \mathbb{R}^d$ (we only use the case $d = 6$ in the proof of Lemma 4), and $T \subseteq \mathbb{R}^d$, a intersection of finite closed half-spaces (which pass the origin). If $T \cap K = \{\mathbf{0}\}$, then there exists $c_0 > 0$, such that $\forall \mathbf{x} \in T$, $\text{dist}(\mathbf{x}, K) \geq c_0 \cdot \|\mathbf{x}\|_2$.*

Proof. By the property of T , we know that 1) T is closed, and 2) $\forall \mathbf{x} \neq \mathbf{0}, \lambda > 0, \mathbf{x} \in T \Leftrightarrow \lambda \mathbf{x} \in T$. By the second property, we only need to prove $\forall \mathbf{x} \in T, \|\mathbf{x}\|_2 = 1, \text{dist}(\mathbf{x}, K) \geq c_0$. Then by the first property, $f(\mathbf{x}) = \text{dist}(\mathbf{x}, K)$, being a continuous mapping defined on closed set $T \cap \mathbb{S}^{d-1}$, has its image set $I = \{\text{dist}(\mathbf{x}, K) : \mathbf{x} \in T \cap \mathbb{S}^{d-1}\}$ closed (where \mathbb{S}^{d-1} is the unit sphere in \mathbb{R}^d). Since $K \cap (T \cap \mathbb{S}^{d-1}) = \emptyset$, we can set $c_0 = \inf I > 0$. \square

We can view each of the three components of A in (5) as intersection of closed half-spaces (passing through the origin). Then we check that the i -th ($i = 1, 2, 3$) component only intersects $\text{span}(\mathbf{u}_i)^\perp$ at $\mathbf{0}$. Therefore by Lemma 5 we conclude that there exists $c > 0$, such that for all $\mathbf{a} \in A$,

$$\|\hat{\pi}_e\| \geq \sum_{i=1,2,3} \text{dist}(\mathbf{a}, \text{span}(\mathbf{u}_i)^\perp) \geq c \cdot \|\mathbf{a}\|_2 \geq c \cdot \max_i |a_i| = c \cdot \max\{\omega_e(\mathcal{O})\}. \quad \square$$