

# Complexity of Local Search for CSPs Parameterized by Constraint Difference

Aditya Anand 

University of Michigan, Ann Arbor, MI, USA

Tommaso D’Orsi 

Bocconi University, Milan, Italy

Euiwoong Lee 

University of Michigan, Ann Arbor, MI, USA

Sijin Peng 

Massachusetts Institute of Technology,  
Cambridge, MA, USA

Vincent Cohen-Addad 

Google Research, NYC, USA

Anupam Gupta 

New York University, NY, USA

Debmalya Panigrahi 

Duke University, Durham, NC, USA

---

## Abstract

In this paper, we study the parameterized complexity of local search, whose goal is to find a good nearby solution from the given current solution. Formally, given an optimization problem where the goal is to find the largest *feasible* subset  $S$  of a universe  $U$ , the new input consists of a *current solution*  $P$  (not necessarily feasible) as well as an ordinary input for the problem. Given the existence of a feasible solution  $S^*$ , the goal is to find a feasible solution as good as  $S^*$  in parameterized time  $f(k) \cdot n^{O(1)}$ , where  $k$  denotes the *distance*  $|P \Delta S^*|$ . This model generalizes numerous classical parameterized optimization problems whose parameter  $k$  is the minimum number of elements removed from  $U$  to make it feasible, which corresponds to the case  $P = U$ .

We apply this model to widely studied Constraint Satisfaction Problems (CSPs), where  $U$  is the set of constraints, and a subset  $U'$  of constraints is feasible if there is an assignment to the variables satisfying all constraints in  $U'$ . We give a complete characterization of the parameterized complexity of all boolean-alphabet *symmetric* CSPs, where the predicate’s acceptance depends on the number of true literals.

**2012 ACM Subject Classification** Theory of computation  $\rightarrow$  Parameterized complexity and exact algorithms

**Keywords and phrases** Constraint Satisfaction Problems, Parameterized Local Search, Optimization

**Digital Object Identifier** 10.4230/LIPIcs.IPEC.2025.26

**Related Version** *Full Version*: <https://arxiv.org/abs/2512.03275>

**Funding** *Anupam Gupta*: Supported in part by NSF awards CCF-2224718 and CCF-2422926.

*Euiwoong Lee*: Supported in part by NSF award CCF-2236669 and by Google, Inc.

*Debmalya Panigrahi*: Supported in part by NSF awards CCF-1955703 and CCF-2329230.

## 1 Introduction

The classical theory of NP-hardness sets limits on the polynomial-time tractability of a vast array of algorithmic problems. This raises the question: *what additional information about the problem instance can one use to help overcome complexity barriers?* A very successful line of research that aims to address this broad question is that of *fixed parameter tractable* (FPT) algorithms. The idea is to identify a parameter that has a small value in *typical* instances, and design an algorithm whose running time is exponential in the value of the parameter but polynomial in the overall size of the instance.

This philosophy has been widely applied to the following general class of optimization problems: let  $U$  be a universe and let  $\mathcal{S} \subseteq 2^U$  be an (often implicitly given) collection of *feasible* subsets of  $U$ , and the goal is to find  $S \in \mathcal{S}$  that maximizes  $|S|$ , or equivalently to find



© Aditya Anand, Vincent Cohen-Addad, Tommaso D’Orsi, Anupam Gupta, Euiwoong Lee, Debmalya Panigrahi, and Sijin Peng;

licensed under Creative Commons License CC-BY 4.0

20th International Symposium on Parameterized and Exact Computation (IPEC 2025).

Editors: Akanksha Agrawal and Erik Jan van Leeuwen; Article No. 26; pp. 26:1–26:17



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

the smallest  $T$  that puts  $U \setminus T \in \mathcal{S}$ . For instance, given a graph  $G = (V, E)$ , if  $U = E$  and  $F \in \mathcal{S}$  if  $(V, F)$  is 2-colorable, the problem corresponds to the famous MAXCUT/MINUNCUT pair. Or given  $G = (V, E)$ , if  $U = V$  and  $\mathcal{S}$  denotes the collection of all independent sets, then we have INDEPENDENT SET/VERTEX COVER. In this class, there are many situations where typical (or at least interesting) instances have the optimal value for the minimization version much smaller than  $|U|$ . (I.e., the universe  $U$  is already close to feasible.) Consequently, numerous results in the FPT literature study a problem from the above class - we let the parameter  $k$  be the optimal value for the minimization version, and give an algorithm of running time  $f(k)n^{O(1)}$  (better than the naive running time of  $n^{O(k)}$ ).

Can we extend the applicability of this framework? We can interpret the algorithms in the above paragraph as one-step *local search*, where we are given a current solution (not necessarily feasible) *close* to some feasible solution  $S^*$  with the optimal value, and the algorithm finds a feasible solution as good as  $S^*$ . This motivates us to ask: *if the algorithm is provided any current solution  $P \subseteq U$  (as a solution for the max version) and there is a feasible solution  $S^* \subseteq U$  close to  $P$ , can the algorithm efficiently recover a solution as good as  $S^*$ ?* Formally, we define the parameter  $k$  as the distance  $|P \Delta S^*|$ , and our goal is to design an algorithm whose runtime is arbitrary in  $k$  but polynomial in the size of the problem instance.<sup>1</sup>

We apply this model to Constraint Satisfaction Problems (CSPs), one of the most important classes of optimization problems in the theory of computation. Given a predicate  $R \subseteq \{0, 1\}^r$ , MAXCSP( $R$ ) is the computational problem whose input consists of the set of variables  $V$  and the set of constraints  $\mathcal{C}$ . Given an assignment  $\alpha : V \rightarrow \{0, 1\}$ , each constraint  $C \in \mathcal{C}$  is satisfied when the values of the  $r$  specified literals (variables or their negations) belong to  $R$ . (See Section 2 for formal definitions.) The goal is to find an assignment that satisfies as many constraints as possible. Generalizing the MAXCUT example above, we apply MAXCSP( $R$ ) to our model where  $U = \mathcal{C}$  and a subset of constraints  $\mathcal{C}' \subseteq \mathcal{C}$  is feasible if there exists an assignment  $\alpha$  that satisfies all of  $\mathcal{C}'$  (i.e.,  $\mathcal{C}'$  is *satisfiable*); so the current solution  $\mathcal{P}$  is a set of constraints and we assume that it is close to some set of satisfiable constraints. It yields the following computational task.

**Problem:** IMPROVEMAXCSP( $R$ )

**Input:** A CSP( $R$ ) instance  $I(V, \mathcal{C})$ , an integer  $k \in \mathbb{N}$ , and  $\mathcal{P} \subseteq \mathcal{C}$ .

**Parameter:**  $k$

**Promise:** There exists an assignment  $\alpha_0$  with the property  $|\mathcal{C}_I(\alpha_0) \Delta \mathcal{P}| \leq k$ , where  $\mathcal{C}_I(\alpha_0)$  is the set of clauses  $\alpha_0$  satisfies, and  $\Delta$  represents symmetric difference.

**Output:** A *good* assignment  $\alpha$ , where the word *good* simply means  $|\mathcal{C}_I(\alpha)| \geq |\mathcal{C}_I(\alpha_0)|$ . (We do not require  $\alpha$  to satisfy  $|\mathcal{C}_I(\alpha) \Delta \mathcal{P}| \leq k$ .)

Note that, given an instance  $(I(V, \mathcal{C}), k, \mathcal{P})$ , the problem is essentially equivalent to find  $\alpha$  such that  $|\mathcal{C}_I(\alpha)|$  is at least  $\max_{\beta: |\mathcal{C}_I(\beta) \Delta \mathcal{P}| \leq k} |\mathcal{C}_I(\beta)|$ ; in words, find a feasible solution as good as any feasible  $k$ -neighbor of  $\mathcal{P}$ . (Formally, we still will keep the fixed promised assignment  $\alpha_0$  as part of the instance and define a good assignment based on  $\alpha_0$ .)

We also observe that if the decision version of CSP( $R$ ) (i.e., finding an assignment that satisfies every constraint) can be solved in polynomial time, an easy  $|\mathcal{C}|^k \cdot \text{poly}(|V|, |\mathcal{C}|)$ -time algorithm for IMPROVEMAXCSP( $R$ ) follows by exhaustively guessing  $(\mathcal{C}_I(\alpha_0) \Delta \mathcal{P})$  and

<sup>1</sup> Actually, since  $|S \Delta S^*| = |(U \setminus S) \Delta (U \setminus S^*)|$ , this parameter remains the same for both maximization and minimization versions.

solving the decision version with constraints  $\mathcal{C}_I(\alpha_0)$ ; otherwise it is NP-hard even when  $k = 0$  (just letting  $\mathcal{P} = \mathcal{C}$ ), which concludes that  $\text{IMPROVEMAXCSP}(R)$  is in  $\mathbf{XP}$  if and only if the decision version is in  $\mathbf{P}$ . Also, if  $\text{MAXCSP}(R)$  is NP-hard, there is no  $\text{poly}(|V|, |\mathcal{C}|, k)$ -time algorithm for  $\text{IMPROVEMAXCSP}(R)$ .

In this paper, we provide a complete characterization of the parameterized complexity of  $\text{IMPROVEMAXCSP}(R)$ , where  $R$  is restricted to a *symmetric* predicate, whose acceptance only depends on the number of true literals. This still includes many well-known CSPs including  $r\text{LIN}$  (which generalizes  $\text{MAXCUT}$ ),  $r\text{SAT}$ ,  $r\text{AND}$ ,  $r\text{NAE}$ , and  $r\text{AE}$ , where AE and NAE abbreviate All-Equal and Not-All-Equal respectively. Namely, the main contribution of this paper is the following theorem:

► **Theorem 1** (Main Theorem (Informal)). *For any symmetric predicate  $R$ ,  $\text{IMPROVEMAXCSP}(R)$  is either FPT or  $W[1]$ -hard.*

It turns out that among nontrivial predicates,  $r\text{AND}$  for any  $r \geq 1$  and  $2\text{AE}$  (generalizing  $\text{MAXCUT}$ ) are the only ones in FPT. We design parameterized algorithms for them extending the previous algorithms for  $\text{MINCSP}$ . While many predicates inherit their hardness result from the hardness of  $\text{MINCSP}$ , we also prove hardness for the problems whose  $\text{MINCSP}$  is tractable, namely  $r\text{AE}$  for  $r \geq 3$  and  $\leq 1$ -out-of- $r\text{SAT}$  for  $r \geq 2$ . See Section 2.3 for our formal characterization.

In the context of local search, a natural special case of  $\text{IMPROVEMAXCSP}(R)$  is when the input  $\mathcal{P}$  is restricted to be  $\mathcal{C}_I(\beta)$  for some assignment  $\beta : V \rightarrow \{0, 1\}$ . Theorem 1 addresses this special case as well, in the sense that our  $W[1]$ -hard reductions indeed have  $\mathcal{P} = \mathcal{C}_I(\beta)$  for some assignment  $\beta$ .

## Related Work and Discussion

The parameterized complexity of local search, whose goal is to improve the current solution to a strictly better solution nearby, was indeed studied in the parameterized complexity literature. Fellows, Fomin, Lokshtanov, Rosamond, Saurabh, and Villanger [8] study various problems like  $r\text{-CENTER}$ ,  $\text{VERTEX COVER}$ ,  $\text{ODD CYCLE TRANSVERSAL}$ ,  $\text{MAX-CUT}$ , and  $\text{MAX-BISECTION}$  on special classes of sparse graphs like planar, minor-free, and bounded-degree graphs, and Szeider [25] considers a similar framework for  $\text{SAT}$  and  $\text{MAXSAT}$ .

This (strict) local search model can be relaxed to another model called *permissive* local search [19]. Here we are given a solution  $S$ , and the goal is to find a solution  $S'$  which has lower cost (which may differ from  $S$  in more than  $k$  vertices), or correctly conclude that every solution which differs from  $S$  in at most  $k$  vertices has higher cost. This model has been studied for a few problems - including a variant of stable marriage [19],  $\text{VERTEX COVER}$  [9] and also in the context of  $\text{MIN ONES CSPs}$  [14], extending Marx’s complete classification of the parameterized complexity of CSPs based on if there is a solution with exactly  $k$  ones [17].

One crucial difference between our model and all previous models on CSPs is how the distance between the two solutions is defined. The previous papers define it as the number of variables with different values in two solutions, but our definition concerns the set of constraints whose satisfaction changes between two solutions. In fact, recall that our definition of a solution is an arbitrary set  $\mathcal{P}$  of constraints not necessarily feasible (i.e., corresponding to a variable assignment), which allows  $\mathcal{P} = \mathcal{C}$  and captures algorithms parameterized by the number of unsatisfied constraints. In that sense, our model is slightly broader than traditional local search where one maintains a feasible solution and tries to strictly improve in every step.

Fixed parameter tractability of CSPs parameterized by the number of unsatisfied constraints has been an important and active research area [2, 11, 12, 13], where a complete classification for Boolean relations was obtained very recently. Their definition of CSPs is strictly more general than ours, where each constraint language contains multiple predicates, even with different arities. It is an exciting research direction to extend our framework for every CSP.

Our negative results can be interpreted as demonstrating the importance of *near-perfect instances* for classical FPT algorithms. For instance, we show  $\text{IMPROVEMAXCSP}(2\text{SAT})$  is  $W[1]$ -hard while  $\text{MINCSP}(2\text{SAT})$  is known to admit an FPT algorithm [23, 22, 6, 21], so the fact that the optimal solution indeed satisfies all but  $k$  constraints affects the parameterized complexity.

There are several natural research directions based on this paper. We believe that the following directions will be interesting to study further:

1. The most immediate one is to extend our characterization for all CSPs. One notable set of boolean relations whose complexity is still open in our model is  $\{x = y, x \wedge \neg y\}$ .
2. A natural variant of our model is when the solution is represented as an assignment to the variables, which is guaranteed to have a Hamming distance at most  $k$  to a strictly better solution. For  $\text{MIN ONES CSP}$ , where the goal is to find a satisfying assignment with the minimum number of true variables, this variant has been studied previously [14].
3. Other than CSPs, one framework that captures numerous results in the FPT literature is covering (resp. packing) problems, especially on graphs, where we want to select the smallest (resp. largest) subset  $S \subseteq U$  subject to some covering (resp. packing) constraints. Local search naturally extends to this setting (i.e., we are given  $S' \subseteq U$  with  $|S' \Delta S| \leq k$ ), and it will be interesting to see if classical FPT results on graph covering/packing problems extend to this model, including  $\text{MULTICUT}$  [3, 18],  $\text{MIN BISECTION}$  [5],  $k$ - $\text{CUT}$  [10],  $\text{DISJOINT PATHS}$  [24].
4. In the approximation algorithms literature, there are numerous studies about CSPs on *structured instances*, most notably dense and expanding instances (see [1] and references therein). It would be interesting to see whether these structures, if suitably parameterized, help the problems become more tractable in our model.

## 2 Our Results

### 2.1 Constraint Satisfaction Problems

In this section we recall basic definitions of constraint satisfaction problems.

For integer  $r \in \mathbb{N}^+$ , a *boolean relation*  $R$  is a subset of  $\mathbb{F}_2^r$ . The integer  $r$  is called the *arity* of  $R$ . For a boolean relation  $R$  of arity  $r$  and  $b \in \mathbb{F}_2^r$ , define the boolean relation  $R \oplus b := \{\alpha \oplus b \mid \alpha \in R\}$ , where  $\oplus$  is the addition operator over  $\mathbb{F}_2^r$ . When considering  $R$  as a clause in a CSP instance,  $\oplus$  operator can be seen as variable negations. For a relation  $R$  with index  $[r]$  and  $S \subseteq [r]$ , define the projection relation  $\pi_S(R) = \{\alpha_S \mid \alpha \in R\}$ , where  $\alpha_S$  means extracting the value on coordinates in  $S$ .

A boolean relation  $R$  of arity  $r$  is *symmetric* if for all  $(x_1, \dots, x_r) \in \mathbb{F}_2^r$  and any permutation  $p: [r] \rightarrow [r]$ ,  $(x_1, \dots, x_r) \in R \iff (x_{p(1)}, \dots, x_{p(r)}) \in R$ . Equivalently,  $R$  is symmetric if there exists  $S \subseteq \{0, 1, \dots, r\}$  such that  $(x_1, \dots, x_r) \in R \iff \sum_{i=1}^r x_i \in S$  (where the addition is performed in  $\mathbb{Z}$ ). We define  $\text{SymRel}(r, S)$  to be such  $R$ .

A *boolean constraint language* is a set of boolean relations. Two constraint languages are *equivalent* if they contain exactly the same set of relations. A *clause*  $C$  over a boolean constraint language  $\Gamma$  is a pair  $(V_C, R_C)$ , where  $R_C \in \Gamma$  and  $V_C = (v_{C,1}, \dots, v_{C,r})$  is an

$r$ -tuple of variables where  $r$  is the arity of  $R_C$ . We refer to  $V_C$  as the *scope* of the clause  $C$ . An assignment  $\alpha : V_C \rightarrow \{0, 1\}$  *satisfies* clause  $(V_C, R_C)$  if  $(\alpha(v_{C,1}), \dots, \alpha(v_{C,r})) \in R_C$ , and  $\alpha$  *violates* the clause otherwise.

In this paper, we only consider a restricted set of boolean constraint languages, where the language  $\Gamma$  includes all possible negation patterns  $R \oplus b$  for some symmetric predicate  $R$  of arity  $r$  and  $b \in \mathbb{F}_2^r$ . We use the notation  $\text{SymLang-N}(R) = \{R \oplus b \mid b \in \mathbb{F}_2^r\}$  for such  $\Gamma$ , where  $\text{SymLang-N}$  stands for “symmetric language with negation”. We further define  $\text{SymLang-N}(r, S)$  to be  $\text{SymLang-N}(\text{SymRel}(r, S))$ .

For a boolean constraint language  $\Gamma$ , a  $\text{CSP}(\Gamma)$  instance  $I$  is a pair  $(V_I, \mathcal{C}_I)$ , where  $V_I$  is the set of variables, and  $\mathcal{C}_I$  is the set of clauses over the boolean constraint language  $\Gamma$ , all of which have scope inside  $V_I$ . We say a variable  $v$  is incident to a clause  $C$  if  $v$  belongs to the scope of  $C$ .

For an assignment  $\alpha : V_I \rightarrow \{0, 1\}$ , define  $\mathcal{C}_I(\alpha)$  to be the set of clauses  $\alpha$  satisfies in  $I$ . We further define  $\text{cost}_I(\alpha) = |\mathcal{C} \setminus \mathcal{C}_I(\alpha)|$ , the number of unsatisfied clauses, to be the *cost* of the assignment. We define the *size* of an instance  $I$  as  $\max(|V_I|, |\mathcal{C}_I|)$  and denote it by  $n_I$ . We ignore the subscripts when the instance is clear in the context.

Now we give formal definitions for the problems we consider. We first define  $\text{MINCSP}(\Gamma)$ , and then define our central problem  $\text{IMPROVEMAXCSP}(\Gamma)$ . In this paper, our main goal is to characterize the parameterized complexity of  $\text{IMPROVEMAXCSP}(\Gamma)$  when  $\Gamma = \text{SymLang-N}(r, S)$  for some  $r \in \mathbb{N}^+$  and  $S \subseteq \{0, 1, \dots, r\}$ .

**Problem:**  $\text{MINCSP}(\Gamma)$

**Input:** A  $\text{CSP}(\Gamma)$  instance  $I(V, \mathcal{C})$ , an integer  $k \in \mathbb{N}$ .

**Parameter:**  $k$ .

**Output:** An assignment  $\alpha$  with  $\text{cost}(\alpha) \leq k$  if it exists, otherwise report that such an assignment does not exist.

**Problem:**  $\text{IMPROVEMAXCSP}(\Gamma)$

**Input:** A  $\text{CSP}(\Gamma)$  instance  $I(V, \mathcal{C})$ , an integer  $k \in \mathbb{N}$ , and  $\mathcal{P} \subseteq \mathcal{C}$

**Parameter:**  $k$

**Promise:** There exists an assignment  $\alpha_0 : V_I \rightarrow \{0, 1\}$  such that  $|\mathcal{C}_I(\alpha_0) \Delta \mathcal{P}| \leq k$ .

**Output:** A *good* assignment  $\alpha$ , where the word *good* simply means  $|\mathcal{C}_I(\alpha)| \geq |\mathcal{C}_I(\alpha_0)|$ . (We do not require the output  $\alpha$  to satisfy  $|\mathcal{C}_I(\alpha) \Delta \mathcal{P}| \leq k$ .)

From the definition,  $\text{MINCSP}(\Gamma)$  is a special case of  $\text{IMPROVEMAXCSP}(\Gamma)$  by setting  $\mathcal{P}$  to be the constraint set  $\mathcal{C}$ .

A boolean constraint language  $\Gamma$  is *trivial* if  $\Gamma = \emptyset$  or  $\Gamma = \{\{0, 1\}^r\}$  for some  $r \in \mathbb{N}^+$  and *nontrivial* otherwise. When  $\Gamma$  is trivial,  $\text{MINCSP}(\Gamma)$  and  $\text{IMPROVEMAXCSP}(\Gamma)$  can be solved by trivial algorithms. So in the following we only consider nontrivial boolean constraint languages.

## 2.2 Abbreviations for Special CSPs

For simplicity, we abbreviate boolean constraint languages mentioned in the paper. Here we introduce the following simple but useful claim in identifying equivalence between boolean constraint languages.

▷ **Claim 2.**  $\text{SymLang-N}(r, S)$  and  $\text{SymLang-N}(r, \{r - x \mid x \in S\})$  are equivalent boolean constraint languages.

*Proof.* From  $\text{SymRel}(r, S) + (1, 1, \dots, 1) = \text{SymRel}(r, \{r - x \mid x \in S\})$ , where  $(1, 1, \dots, 1)$  is the element in  $\mathbb{F}_2^r$  that has value 1 in each coordinate, we have  $\text{SymRel}(r, S) + b = \text{SymRel}(r, \{r - x \mid x \in S\}) + (b + (1, 1, \dots, 1))$  for any  $b \in \mathbb{F}_2^r$ . ◁

We will use the following abbreviations in this paper.

- $r\text{AND} = \text{SymLang-N}(r, \{0\}) = \text{SymLang-N}(r, \{r\})$ .
- $r\text{SAT} = \text{SymLang-N}(r, \{1, 2, \dots, r\}) = \text{SymLang-N}(r, \{0, 1, \dots, r - 1\})$ .
- $r\text{AE}$  (All Equal)  $= \text{SymLang-N}(r, \{0, r\})$ .
- $\leq 1\text{-out-of-}r\text{SAT} = \text{SymLang-N}(r, \{0, 1\}) = \text{SymLang-N}(r, \{r - 1, r\})$ .

We also use abbreviations for the corresponding MINCSP and IMPROVEMAXCSP problems. For example,  $\text{IMPROVEMAXCSP}(r\text{AND})$  corresponds to the problem  $\text{IMPROVEMAXCSP}(\text{SymLang-N}(r, \{0\}))$ .

### 2.3 Our Characterization

Having defined our framework, we first provide the formal version of Theorem 1:

► **Theorem 3 (Main Theorem (Formal)).** *For any  $r \in \mathbb{N}^+$  and  $S \subseteq \{0, 1, \dots, r\}$ , the problem  $\text{IMPROVEMAXCSP}(\text{SymLang-N}(r, S))$  is either FPT or W[1]-hard.*

As  $\text{MINCSP}(\Gamma)$  is a special case of  $\text{IMPROVEMAXCSP}(\Gamma)$ , the W[1]-hardness of the former implies that of the latter. The recent characterization of parameterized tractability of  $\text{MINCSP}(\Gamma)$  [13] implies the following theorem.

► **Theorem 4.** *For nontrivial  $\Gamma = \text{SymLang-N}(r, S)$  for some  $r \in \mathbb{N}^+$  and  $S \subseteq \{0, 1, \dots, r\}$  that is not equivalent to  $r\text{AND}$ ,  $r\text{AE}$  or  $\leq 1\text{-out-of-}r\text{SAT}$ ,  $\text{IMPROVEMAXCSP}(\Gamma)$  is W[1]-hard.*

It remains to study the remaining problems.

► **Remark 5.** We say that an algorithm  $A$  solves  $\text{IMPROVEMAXCSP}(\Gamma)$ , if for any instance that satisfies the promise,  $A$  outputs a good assignment. However, the run time is measured on all possible instances. In other words,  $A$  runs in  $f(k)n^{O(1)}$  time if it terminates and outputs some assignment in  $f(k)n^{O(1)}$  time no matter whether the input instance  $(I, k, \mathcal{P})$  satisfies the promise. Clearly, only the run time on instances satisfying the promise is critical, as we can always set a run time limit to the algorithm to rule out instances that require too much time to indicate their violation of promise.

For  $\text{IMPROVEMAXCSP}(r\text{AND})$ , a color-coding combined with the LP rounding for  $\text{DENSEST SUBHYPERGRAPH}$  yields the following result proved in Section 3.

► **Theorem 6.** *There exists an algorithm that solves  $\text{IMPROVEMAXCSP}(r\text{AND})$  in  $2^{O(k \log k)}n^{O(1)}$  time for all  $r \geq 1$ .*

For  $r\text{AE}$ , its tractability depends on  $r$ . When  $r = 2$ , we use the celebrated *unbreakability* framework [10, 4, 5, 16] to design a parameterized algorithm in Section 4.

► **Theorem 7.** *There exists an algorithm that solves  $\text{IMPROVEMAXCSP}(2\text{AE})$  in time  $2^{O(k^2)}n^{O(1)}$ .*

However, for  $r\text{AE}$  with  $r \geq 3$ , we show nontrivial reductions from Paired Minimum *st*-Cut to show the following hardness.

► **Theorem 8.** *IMPROVEMAXCSP( $r$ AE) is  $W[1]$ -hard for all  $r \geq 3$ .*

Finally, for  $\leq 1$ -out-of- $r$ SAT, we prove that even the first nontrivial case  $r = 2$ , which is equivalent to 2SAT, is already  $W[1]$ -hard. Even though  $\text{MINCSP}(2\text{SAT})$  is in FPT, this result follows from the (slightly informal) fact that *Vertex Cover and Independent Set are equivalent in our model*.

► **Theorem 9.** *IMPROVEMAXCSP( $\leq 1$ -out-of- $r$ SAT) is  $W[1]$ -hard for any  $r \geq 2$ .*

The hardness proofs can be found in the full version. It is easy to see that the other theorems imply Theorem 3 in a straightforward way.

### 3 FPT Algorithm for $r$ AND

The main goal of this section is to prove Theorem 6. During the course of the algorithm, we create instances involving AND constraints of different arities. Hence, we will show a slightly stronger result:

► **Theorem 10.** *There is an algorithm that solves  $\text{IMPROVEMAXCSP}(\cup_{r' \leq r} r'\text{AND})$  in  $2^{O(k \log k)} n^{O(1)}$  time for all  $r \geq 1$ .*

Since  $\cup_{r' \leq r} r'\text{AND} \supseteq r\text{AND}$ , Theorem 10 immediately shows that for any  $r \geq 1$ , the problem  $\text{IMPROVEMAXCSP}(r\text{AND})$  is FPT.

#### 3.1 Preliminaries

In this section we recall basic definitions about hypergraphs. A *hypergraph*  $H(V, E)$  is a generalization of graph where an edge can contain multiple vertices. For  $e \in E$  and  $v \in V$ , say  $v$  is incident to  $e$  and  $e$  is incident to  $v$  if  $e$  contains  $v$ . For a hypergraph  $H$  and  $V_0 \subseteq V$ , we denote  $H(V_0)$  as the *induced subhypergraph* of  $V_0$  in  $H$ , and denote  $E(H(V_0))$  to be the edge set of the induced subhypergraph.

In the main algorithm in this section (provided in Theorem 14), we need to solve the following problem:

**Problem:** MAXIMUM INDUCED SUBHYPERGRAPH WITH VERTEX WEIGHTS.

**Input:** A hypergraph  $H(V, E)$  and a weight function  $w : V \rightarrow \mathbb{Z}$ .

**Output:** A vertex subset  $V_0 \subseteq V$  that maximizes  $|E(H(V_0))| - \sum_{v \in V_0} w(v)$ .

The problem is closely related to DENSEST SUBGRAPH (See [15] for a survey) and has a polynomial-time algorithm based on similar algorithms for DENSEST SUBGRAPH. One algorithm is to observe that  $|E(H(V_0))| - \sum_{v \in V_0} w(v)$  is supermodular and run a submodular minimization algorithm.<sup>2</sup> The below is an LP-based algorithm specific to our setting.

► **Lemma 11.** *There exists a polynomial-time exact algorithm that solves MAXIMUM INDUCED SUBHYPERGRAPH WITH VERTEX WEIGHTS.*

<sup>2</sup> We thank an anonymous reviewer for this observation.

**Proof.** Consider the following LP:

$$\begin{array}{ll}
 \text{maximize} & \sum_{e \in E} x_e - \sum_{v \in V} w(v)y_v \\
 \text{subject to} & x_e \leq y_v \quad \forall e \in E, v \in V \text{ s.t. } e \text{ is incident to } v \\
 & x_e \in [0, 1] \quad \forall e \in E \\
 & y_v \in [0, 1] \quad \forall v \in V
 \end{array}$$

Since we use  $x_e \leq y_v$  to mimic the constraint that “one should select a vertex before selecting incident edges”, if the LP is integral (i.e.  $x_e, y_v \in \{0, 1\}$ ) then the set of vertices  $v$  with  $y_v = 1$  in the optimal solution of the LP corresponds to the correct output to the MAXIMUM INDUCED SUBHYPERGRAPH WITH VERTEX WEIGHTS problem.

The algorithm solves the LP, providing optimal fractional solution  $(x_e^*, y_v^*)$ , and rounds the solution to an integer solution in the following way: Arbitrarily select  $p \in (0, 1]$  and set  $x_e^p = [x_e^* \geq p]$ ,  $y_v^p = [y_v^* \geq p]$ , where  $[P]$  denotes the Iverson bracket and equals to 1 if  $P$  is true otherwise 0 for some statement  $P$ . The solution is clearly an integral solution to the LP.

Now we show that any such  $p$  produces an integral solution that has the same value as that of  $(x_e^*, y_v^*)$ , denoted as  $f^*$ . Define  $f(p)$  to be the objective function value of the solution  $(x_e^p, y_v^p)$ . Clearly we have  $f(p) \leq f^*$ . We further have

$$\begin{aligned}
 \int_0^1 f(p) dp &= \int_0^1 \left( \sum_{e \in E} [x_e^* \geq p] - \sum_{v \in V} w(v)[y_v^* \geq p] \right) dp \\
 &= \sum_{e \in E} x_e^* - \sum_{v \in V} w(v)y_v^* = f^*
 \end{aligned} \tag{1}$$

So the set  $\{x \in (0, 1] \mid f(x) < f^*\}$  has zero measure. Further from the fact that  $f(p)$  is a step function of finite number of “steps” and each “step” has nonzero length, such set must be empty, and  $f(p) = f^*$ .  $\blacktriangleleft$

Our FPT algorithms in Section 3 and Section 4 will use the technique of *color coding* and its derandomization given below.

► **Theorem 12** ([20]). *Given a set  $U$  of size  $n$  and integers  $0 \leq a, b \leq n$ , one can in time  $2^{O(\min(a,b) \log(a+b))} \cdot n \log n$  construct a family  $\mathcal{F}$  of at most  $2^{O(\min(a,b) \log(a+b))} \log n$  colorings  $U \rightarrow \{0, 1\}$  such that the following holds: for any sets  $A, B \subseteq U$ ,  $A \cap B = \emptyset$ ,  $|A| \leq a$ ,  $|B| \leq b$ , there exists a coloring  $\chi \in \mathcal{F}$  with  $\chi(a) = 1$  for all  $a \in A$  and  $\chi(b) = 0$  for all  $b \in B$ .*

### 3.2 Algorithm with Variable Solution

Given an instance  $(I, k, \mathcal{P})$ , we say that  $\mathcal{P}$  is *satisfiable* if some assignment satisfies all clauses in  $\mathcal{P}$ . We will give an algorithm that runs in  $2^{O(k \log k)} n^{O(1)}$  time and solves IMPROVEMAXCSP( $\cup_{r' \leq r} r'$ AND) when the instance satisfies an additional promise that  $\mathcal{P}$  is satisfiable.

When this additional promise is satisfied, it is easy to find an assignment satisfying all the clauses of  $\mathcal{P}$  in polynomial time, so in the following we assume that the instance  $(I, k, \mathcal{P})$  is equipped with such an assignment  $\alpha$ . In the full version, we show how to transform any instance into instances where this additional promise is true.

Equipped with this additional assignment, we then restructure the instance in the following way: Define  $k' := k + |\mathcal{P} \Delta C_I(\alpha)|$  and  $\mathcal{P}' := C_I(\alpha)$ , and replace  $(I, k, \mathcal{P}, \alpha)$  with  $(I, k', \mathcal{P}', \alpha)$ , to keep  $\mathcal{P}$  to be maximal in accordance with  $\alpha$ . In other words, we reset  $\mathcal{P}$  to be the set of clauses satisfied by  $\alpha$ . If the instance satisfies the promise (that some good assignment is

close to the  $\mathcal{P}$ ), the good assignment satisfies at most  $|\mathcal{P}| + k$  clauses, which implies that  $|\mathcal{P} \Delta C_I(\alpha)| \leq k$ , and  $k' \leq 2k$ . If  $k' > 2k$ , we return an arbitrary assignment (to handle cases where the promise is not met).

Now we show that, there exists an good assignment close to  $\mathcal{P}$  whose Hamming distance to  $\alpha$  is small.

► **Lemma 13.** *For a  $\text{IMPROVEMAXCSP}(\cup_{r' \leq r} r' \text{AND})$  instance  $(I, k, \mathcal{P}, \alpha)$ , there exists a good assignment  $\beta : V \rightarrow \{0, 1\}$  with  $|C_I(\beta) \Delta \mathcal{P}| \leq k$  and  $|\{v \in V \mid \alpha(v) \neq \beta(v)\}| \leq rk$ .*

**Proof.** Choose a good  $\beta$  with  $|C_I(\beta) \Delta \mathcal{P}| \leq k$ , and if there are multiple ones, choose the one with the smallest hamming distance to  $\alpha$ .

For any  $v$  with  $\alpha(v) \neq \beta(v)$ , consider assignment  $\beta'$  that differs from  $\beta$  only on  $v$ .  $\beta'$  has smaller Hamming distance to  $\alpha$  than  $\beta$ , so either  $\beta'$  is not good, or  $\beta'$  is good but does not satisfy the promise. Both cases require that some clause  $C \in \mathcal{C}$  is satisfied in  $\beta$  but violated in  $\beta'$ . Since all clauses are AND clauses and  $\beta'(v) = \alpha(v)$ , such clause  $C$  is violated in  $\alpha$ . Therefore, all variables contributing to the hamming distance are incident to some clauses in  $C_I(\beta) \Delta \mathcal{P}$ , and the statement follows from the fact that  $|C_I(\beta) \Delta \mathcal{P}| \leq k$  and that the arity of all relations in  $\cup_{r' \leq r} r' \text{AND}$  is no more than  $r$ . ◀

Now we use the (derandomized) color coding technique along with the algorithm of Lemma 11 for MAXIMUM INDUCED SUBHYPERGRAPH WITH VERTEX WEIGHTS to solve our problem  $\text{IMPROVEMAXCSP}(\cup_{r' \leq r} r' \text{AND})$ .

► **Theorem 14.** *There is an algorithm solving  $\text{IMPROVEMAXCSP}(\cup_{r' \leq r} r' \text{AND})$  in  $2^{O(k \log k)} n^{O(1)}$  time when the input instance satisfies an additional promise that  $\mathcal{P}$  is satisfiable.*

**Proof.** First use Theorem 12 with  $a = b = rk$  to obtain a family of at most  $2^{O(k \log k)} \log n$  colorings with binary labels for variables. The algorithm then produces an assignment for each coloring and reports the best among them. In the following we fix a particular coloring.

Let  $L_1$  be the set of variables with label 1. Construct a binary relation  $\sim$  on  $L_1$ , where  $u \sim v$  if  $u = v$  or some clause  $C \in \mathcal{P}$  is incident to both  $u$  and  $v$ . Since  $\sim$  is reflexive and symmetric, its transitive closure is an equivalence relation. Denote  $L_1 / \sim$  as the quotient set of the equivalence relation which includes all equivalence classes of the transitive closure. The intuition is that, the algorithm will produce the answer from  $\alpha$  by flipping several variables with label 1, and variables in one equivalence class are considered a group and would be flipped or kept at the same time.

We then construct a hypergraph  $H(V_H, E_H)$  with weight function  $w : V_H \rightarrow \mathbb{Z}$  as follows:

- For each equivalence class in  $L_1 / \sim$ , introduce a vertex  $v$  in  $V_H$ , whose weight  $w(v)$  equals to the number of clauses in  $\mathcal{P}$  incident to any variable in the equivalence class. Intuitively, the weight measures that how many clauses we will lose if we flip this equivalence class.
- For each clause  $C \in \mathcal{C}_I \setminus \mathcal{P}$ , if  $C$  can be satisfied by flipping  $L_1$  from  $\alpha$ , introduce a hyperedge incident to all equivalence classes containing incident vertices of  $C$ . So the clause will be satisfied if we flip all equivalence classes incident to this hyperedge.

Finally, the algorithm uses the algorithm for MAXIMUM INDUCED SUBHYPERGRAPH WITH VERTEX WEIGHTS proposed in Lemma 11 to find  $V_0 \subseteq V_H$  that maximizes  $|E(H(V_0))| - \sum_{v \in V_0} w(v)$ , and returns the assignment produced by flipping in  $\alpha$  all variables belonging to equivalence classes in  $V_0$ .

The algorithm runs in  $2^{O(k \log k)} n^{O(1)}$  time since the whole process can be done in polynomial time except trying all  $2^{O(k \log k)} n^{O(1)}$  colorings produced by Theorem 12. Now we show that the algorithm produces a good assignment for instances satisfying the promise

that some good assignment is close to the  $\mathcal{P}$ . Since we finally choose the best assignment among all colorings, it is sufficient to show that there exists some coloring that produces a good assignment.

Fix any good assignment  $\beta$  with  $|C_I(\beta)\Delta\mathcal{P}| \leq k$  and  $|\{v \in V \mid \alpha(v) \neq \beta(v)\}| \leq rk$ , where  $\alpha$  is an assignment such that  $\mathcal{P} = C_I(\alpha_0)$ . According to Lemma 13 such  $\beta$  exists. Define  $N(C_I(\beta)\Delta\mathcal{P})$  to be all variables incident to any clause in  $C_I(\beta)\Delta\mathcal{P}$ . Define  $V_0 = \{v \in N(C_I(\beta)\Delta\mathcal{P}) \mid \alpha(v) = \beta(v)\}$  and  $V_1 = N(C_I(\beta)\Delta\mathcal{P}) \setminus V_0$ . Note that both  $|V_0|, |V_1| \leq rk$ .

According to Theorem 12, there exists a coloring assigning  $V_0$  to 0 and  $V_1$  to 1. We consider the behavior of the algorithm when trying this coloring.

We first claim that, there exists a set of equivalence classes in  $L_1 / \sim$  whose union is exactly  $V_1$ . This is equivalent to the statement that there is no  $u \in V_1$  and  $v \in L_1 \setminus V_1$  such that  $u \sim v$ . Suppose  $u \sim v$  and  $u \in V_1$ , then some  $C \in \mathcal{P}$  is incident to both of them. Since  $C$  is incident to  $u \in V_1$  which satisfies  $\alpha(u) \neq \beta(u)$ ,  $C$  is violated by  $\beta$ , so  $C \in C_I(\beta)\Delta\mathcal{P}$ . Then  $v \in N(C_I(\beta)\Delta\mathcal{P})$ , which finishes the proof since it implies  $v \notin L_1 \setminus V_1$ .

The previous statement shows that, there exists some  $V^* \subseteq V_H$ , such that by flipping from  $\alpha$  all variables in equivalence classes corresponding to  $V^*$ , one can produce  $\beta$ . We further show the following two claims:

- For any  $V \subseteq V_H$ , suppose  $\alpha'$  is produced by flipping from  $\alpha$  all variables in every equivalence class in  $V$ , then  $cost_I(\alpha) - cost_I(\alpha') \leq |E(H(V))| - \sum_{v \in V} w(v)$ . Notice that the latter is the objective of the MAXIMUM INDUCED SUBHYPERGRAPH WITH VERTEX WEIGHTS problem. In other words, the cost of  $\alpha'$  would not be underestimated. To see this, consider clauses in  $C_I(\alpha)\Delta C_I(\alpha')$ . Clauses satisfied in  $\alpha$  but unsatisfied in  $\alpha'$  are those belonging to  $\mathcal{P}$  and incident to  $V$ , counted exactly by  $\sum_{v \in V} w(v)$ . Clauses satisfied in  $\alpha'$  but unsatisfied in  $\alpha$  are partially counted by  $|E(H(V))|$ , since it does not include satisfied clauses in  $\alpha'$  violated by the assignment produced from flipping  $L_1$  in  $\alpha$ . However, all hyperedges in  $E(H(V))$  correspond to clauses satisfied in  $\alpha'$  but unsatisfied in  $\alpha$ .
- $cost_I(\alpha) - cost_I(\beta) = |E(H(V^*))| - \sum_{v \in V^*} w(v)$ , i.e. the cost of  $\beta$  is correctly estimated by the MAXIMUM INDUCED SUBHYPERGRAPH WITH VERTEX WEIGHTS objective. We only need to show that  $E(H(V^*))$  contains all clauses satisfied in  $\beta$  but unsatisfied in  $\alpha$ . Since the coloring colors  $V_0$  with 0 and  $V_1$  with 1, all clauses satisfied in  $\beta$  but unsatisfied in  $\alpha$  have label-1 incident variables inside  $V_1 \subseteq L_1$ , so the clauses will introduce hyperedges in  $H$  whose scope is inside  $V^*$ , so they are all counted by  $|E(H(V^*))|$ .

The previous two claims imply that,  $V^*$  is one of the optimal solutions to the MAXIMUM INDUCED SUBHYPERGRAPH WITH VERTEX WEIGHTS problem for the constructed hypergraph  $5H$ , and any optimal solution of the MAXIMUM INDUCED SUBHYPERGRAPH WITH VERTEX WEIGHTS problem corresponds to a good assignment to the CSP instance, which finishes the proof.  $\blacktriangleleft$

## 4 FPT Algorithms for 2AE

In this section we prove Theorem 7.

### 4.1 Preliminaries

We recall useful definitions about graphs. Let  $G = (V, E)$  be a multigraph, where parallel edges and loops are allowed. For two disjoint vertex sets  $A, B \subseteq V$ , we denote by  $E(A, B)$  the set of edges between  $A$  and  $B$ . When  $(A, B)$  is a partition, then  $E(A, B)$  is the set

of edges in the corresponding cut. For a vertex subset  $V' \subseteq V$ , we denote by  $G(V')$  the subgraph of  $G$  induced by  $V'$ . We write  $E(G(V'))$  for the set of edges in  $G(V')$ . For a set  $E' \subseteq E$ , we denote by  $G \setminus E'$  the multigraph  $(V, E \setminus E')$ . For  $S \subseteq V$ , we let  $N(S) := \{v \in V \setminus S : \exists u \in S \text{ such that } (u, v) \in E\}$  be the open neighborhood of  $S$ . We say two cuts  $(A, B)$  and  $(X, Y)$  are  $k$ -close if  $|E(A, B) \Delta E(X, Y)| \leq k$ . When the graph is clear in the context, we define  $n = \max(|V|, |E|)$  to be the size of the graph.

Notice that we can see clauses as edges connecting variables, assigning boolean values to variables as partitioning the variable set into two sets, and the equality and non-equality constraints correspond to uncut and cut constraints respectively. This motivates us to introduce the following graph problem. Given a graph  $G = (V, E)$ , an edge type function  $t : E \rightarrow \{0, 1\}$  and a partition  $(A, B)$  of  $V$ , define  $C(A, B) = (t^{-1}(1) \cap E(A, B)) \cup (t^{-1}(0) \cap (E \setminus E(A, B)))$  to be the set of edges satisfied by  $(A, B)$ , where edges of type 1 need to be separated, while edges of type 0 should be inside some vertex set.

**Problem:** IMPROVEMAXCUT-UNCUT

**Input:** A connected multigraph  $G(V, E)$ , an integer  $k$ ,  $\mathcal{P} \subseteq E$ , and an edge type function  $t : E \rightarrow \{0, 1\}$ .

**Parameter:**  $k$

**Promise:** There exists some  $(A, B)$  that maximizes  $|C(A, B)|$  and satisfies  $|C(A, B) \Delta \mathcal{P}| \leq k$ .

**Output:** Any partition  $(A, B)$  of vertices that maximizes  $|C(A, B)|$ . (Notice that we do not require  $|C(A, B) \Delta \mathcal{P}| \leq k$  for the output.)

Clearly, if the given graph for IMPROVEMAXCUT-UNCUT is not connected, we can solve each connected component separately, so connectivity of  $G$  does not lose any generality for the problem. The following claim is straightforward.

▷ **Claim 15.** IMPROVEMAXCSP(2AE) and IMPROVEMAXCUT-UNCUT are equivalent.

In the following, we present the algorithm for IMPROVEMAXCUT-UNCUT.

## 4.2 From Edge Solution to Vertex Solution

We first apply a pre-processing algorithm to the given IMPROVEMAXCUT-UNCUT instance, intending to make it in accordance with some partition of the vertex set.

► **Lemma 16.** *There exists an algorithm that given a IMPROVEMAXCUT-UNCUT instance, runs in  $4^k n^{O(1)}$  time and finds a bipartition  $(A, B)$  of the vertex set  $V$  such that there exists a good partition  $(A^*, B^*)$  which satisfies  $|C(A^*, B^*) \Delta \mathcal{P}| \leq k$  and  $|C(A^*, B^*) \Delta C(A, B)| \leq 3k$ .*

**Proof.** Fix any  $(A^*, B^*)$  satisfying  $|C(A^*, B^*) \Delta \mathcal{P}| \leq k$ . We first find the largest subset  $\mathcal{P}_1 \subseteq \mathcal{P}$  that can be simultaneously satisfied by some bipartition using a  $4^k n^{O(1)}$ -time algorithm for MINCSP(2AE) [7]: Since  $C(A^*, B^*) \cap \mathcal{P}$  is one possible solution, this means that  $|\mathcal{P}_1| \geq |C(A^*, B^*) \cap \mathcal{P}| \geq \max(\mathcal{P}, |C(A^*, B^*)|) - k$ . Let  $(A, B)$  be any bipartition satisfying  $\mathcal{P}_1$ . Consider  $C(A, B)$ : It contains all edges in  $\mathcal{P}_1$  and perhaps more. However, since the size cannot be larger than  $|C(A^*, B^*)|$ , there are at most  $k$  edges in  $C(A, B) \setminus \mathcal{P}_1$ . Now we have

$$|C(A, B) \Delta C(A^*, B^*)| \leq |\mathcal{P} \Delta C(A^*, B^*)| + |\mathcal{P} \setminus \mathcal{P}_1| + |C(A, B) \setminus \mathcal{P}_1| \leq 3k.$$

This completes the proof. ◀

From Lemma 16, we can set  $\mathcal{P}$  to be such  $C(A, B)$  and triple  $k$  to change the original IMPROVEMAXCUT-UNCUT instance to an equivalent IMPROVEMAXCUT-UNCUT instance, whose  $\mathcal{P}$  is given by some  $C(A, B)$  for a vertex partition  $(A, B)$ .

### 4.3 Algorithms for Instances with Vertex Solution

In the rest of this section, we prove the following theorem:

► **Theorem 17.** *There exists an algorithm that, given a IMPROVEMAXCUT-UNCUT instance  $(G(V, E), k, \mathcal{P}, t)$  where  $\mathcal{P} = C(A, B)$  for some partition  $(A, B)$  of  $V$ , finds a partition  $(X, Y)$  that maximizes  $|C(X, Y)|$  in time  $2^{O(k^2)}n^{O(1)}$ .*

We first show how this theorem implies the main theorem.

**Proof of Theorem 7.** According to Claim 15, we only need to find such algorithm for IMPROVEMAXCUT-UNCUT. For a IMPROVEMAXCUT-UNCUT instance, we first apply the algorithm stated in Lemma 16 to produce an instance equivalent to the input whose  $\mathcal{P}$  is given by some  $C(A, B)$  for a vertex partition. Then we apply the algorithm stated in Theorem 17 to solve the new instance. The runtime is  $4^k n^{O(1)} + 2^{O(k^2)} n^{O(1)} = 2^{O(k^2)} n^{O(1)}$ . ◀

Recall that we are given a vertex partition  $(A, B)$  as the solution, while there exists a good partition  $(X, Y)$  with  $|C(A, B) \Delta C(X, Y)| \leq k$ . Let  $AX := A \cap X$  and define  $AY, BX, BY$  similarly.

Our high-level approach follows the framework first introduced by Kawarabayashi and Thorup [10] for Minimum  $k$ -cut and its refinement by Chitnis et al. [4]. Let us define the following *terminal version* of the problem. We will fix  $k$  to be the parameter of the initial instance, which *remains invariant throughout the recursive calls to the terminal version*.

**Problem:** IMPROVEMAXCUT-UNCUT TERMINAL VERSION

**Input:** A connected multigraph  $G(V, E)$ , an integer  $k' \leq k$ , a  $\mathcal{P} \subseteq E$ , an edge type function  $t : E \rightarrow \{0, 1\}$ , a set of terminals  $T \subseteq V$  with  $|T| \leq 2k$  (Notice that it is  $k$  rather than  $k'$ ) and a *marked edge set*  $M \subseteq E$ .

**Parameter:**  $k'$

**Promise:**

- $\mathcal{P} = C(A, B)$  for some partition  $(A, B)$ ;
- $M$  is a matching allowing parallel edges: Any two edges in  $M$  are either completely disjoint or parallel copies of each other.
- There exists a good cut  $(X, Y)$  with  $|C(A, B) \Delta C(X, Y)| \leq k'$  and  $M \subseteq E(X, Y) \cap E(A, B)$ .

**Requirement:** For *each* function  $f : T \rightarrow \{X, Y\}$  and  $0 \leq k'' \leq k$ , output a cut  $(X_{f, k''}, Y_{f, k''})$  satisfying all the following or output nothing:

1. consistent with  $f$ :  $f(v) = X$  iff  $v \in X_{f, k''}$  for all  $v \in T$ ,
2. consistent with  $M$ :  $M \subseteq E(A, B) \cap E(X_{f, k''}, Y_{f, k''})$ ,
3.  $k''$ -close to  $(A, B)$ :  $|C(A, B) \Delta C(X_{f, k''}, Y_{f, k''})| \leq k''$ ,

For any  $(X, Y)$  which is a good cut that satisfies  $|C(A, B) \Delta C(X, Y)| \leq k''$  and  $M \subseteq E(X, Y) \cap E(A, B)$  for some  $k'' \leq k'$ , the output  $(X_{f^*, k''}, Y_{f^*, k''})$ , where  $f^*$  is consistent to  $(X, Y)$ , should additionally be a good cut (not necessarily equal to  $(X, Y)$ ).

Clearly, IMPROVEMAXCUT-UNCUT is a special case with  $M = T = \emptyset$  after doing the pre-processing in Lemma 16. In the following we introduce the algorithm solving IMPROVEMAXCUT-UNCUT TERMINAL VERSION.

Let  $q := k^2 2^{2k+2}$ . Before introducing the algorithm, we require one additional definition.

► **Definition 18** ( $(k, q)$ -cut). Let  $G = (V, E)$  be a multigraph and let  $M \subseteq E$  be the set of marked edges which is a matching allowing parallel edges. A cut  $L \subseteq V$  is called  $(k, q)$ -balanced (more simply a  $(k, q)$ -cut) if the following conditions are met.

- $|E(L, V \setminus L)| \leq k$ .
- Both  $G(L)$  and  $G(V \setminus L)$  are connected (using both edges inside and outside  $M$ ).
- Both  $G(L)$  and  $G(V \setminus L)$  contains at least  $q$  non-marked edges (edges outside  $M$ ).

Our algorithm follows a simple win-win strategy. If there is no  $(k, q)$ -cut, a simple color-coding algorithm will solve the terminal problem. If there is a  $(k, q)$ -cut, by recursively solving a smaller subproblem, one can reduce the number of unmarked edges. Combining this with the algorithm to compute a  $(k, q)$ -cut, one can solve the terminal version in FPT time.

Now we present each of the three parts. We start by directly solving the case where there is no  $(k, q)$ -cut, using the (derandomized) coloring coding technique.

▷ **Claim 19.** One can solve IMPROVEMAXCUT-UNCUT TERMINAL VERSION in  $2^{O(k \log q)} n^{O(1)}$  time when there is no  $(k, q)$ -cut in the instance.

Proof. Given an IMPROVEMAXCUT-UNCUT TERMINAL VERSION instance  $I = (G, k', \mathcal{P} = C(A, B), t, T, M)$ , the algorithm considers every  $f : T \rightarrow \{X, Y\}$  and  $0 \leq k'' \leq k'$  and does the following. Assume that  $f$  and  $k''$  is consistent with some good cut  $(X, Y)$ , because apart from running time, the requirement only concerns for the correct  $f$  and  $k''$ .

Note that the set of edges  $S = C(A, B) \Delta C(X, Y) = E(A, B) \Delta E(X, Y) = E(AX, BX) \cup E(AY, BY) \cup E(AX, AY) \cup E(BX, BY)$  forms a cut between  $L := AX \cup BY$  and  $R := AY \cup BX$ , no matter how  $t$  assigns type to edges. Since  $G$  is connected and  $|S| \leq k''$ ,  $G \setminus S$  has at most  $k'' + 1$  connected components.

Assume towards a contradiction that there are two connected components  $L', R'$  of  $G \setminus S$  such that  $L' \subseteq L$  and  $R' \subseteq R$  and both  $G(L')$  and  $G(R')$  have at least  $q$  non-marked edges. Then the minimum cut separating  $L'$  and  $R'$  in  $G$  will cut at most  $k'' \leq k$  edges where the two resulting parts are connected and contain  $L'$  and  $R'$  respectively, which contradicts the fact that there is no  $(k, q)$ -cut. <sup>3</sup>

Therefore, at least one of  $L$  and  $R$  (say  $R$ ) has the property that every connected component of  $G \setminus S$  contained in  $R$  has at most  $q$  non-marked edges. Since marked edges form a matching, the number of vertices in such a component can be at most  $2q + 2$ , so we have  $|R| = O(qk)$ .

Let  $\{R_i\}_{i=1}^\ell$  be the connected components of  $G(R)$ . Define  $a_i^+ = |C(A \Delta R_i, B \Delta R_i) \setminus C(A, B)|$  and  $a_i^- = |C(A, B) \setminus C(A \Delta R_i, B \Delta R_i)|$ . Notice that  $|C(X, Y)| = |C(A, B)| + \sum_{i=1}^\ell (a_i^+ - a_i^-)$  and  $|C(X, Y) \Delta C(A, B)| = \sum_{i=1}^\ell (a_i^+ + a_i^-) \leq k'$ . Guessing the correct values of  $\ell$  and  $a_i^+, a_i^-$  takes at most  $k^{O(k)}$  time.

Let  $N = N(R)$  be the open neighbor of  $R$  in  $G$ . Note that  $|N| \leq |E(R, V \setminus R)| = |C(A, B) \Delta C(X, Y)| \leq k''$ . Use Theorem 12 with  $a = |R|, b = k$  to have a family of at most  $2^{O(k \log q)} \log n$  colorings that color each vertex using  $\{0, 1\}$ . There exists a coloring  $\chi$  that colors  $R$  with 1 and colors  $N$  with 0. For each  $i \in [\ell]$ , find all candidates  $R'_i \subseteq V$  that is *identical* to  $R_i$  with respect to the guessed information about  $R_i$ . Formally,  $R'_i$  is a connected component of  $G$  after deleting all vertices of color 0, and in line with the guessing  $a_i^+$  and  $a_i^-$ .

For the correct coloring  $\chi$ , such  $R'_i$  exists since  $R_i$  satisfies the condition. Also, finding all such candidates  $R'_i$  is easy in polynomial time since one only needs to scan connected components after deleting color-0 vertices.

<sup>3</sup> We use the fact that any minimum  $st$ -cut in a connected graph have both sides connected.

## 26:14 Complexity of Local Search for CSPs Parameterized by Constraint Difference

For any  $R'_1, \dots, R'_\ell$  that are disjoint, since  $R'_i$  and  $N(R'_j)$  are disjoint for any  $i, j \in [\ell]$ , the set of edges newly cut/uncut by updating  $A \leftarrow A \Delta R'_i$  and  $B \leftarrow B \Delta R'_i$  are disjoint for all  $i$ . Therefore, letting  $R' := \cup_i R'_i$  and letting  $A' \leftarrow A \Delta R'$  and  $B' \leftarrow B \Delta R'$  the resulting partition has  $|C(A', B')| = |C(A, B)| + \sum_i (a_i^+ - a_i^-) = |C(X, Y)|$  and  $|C(A', B') \Delta C(A, B)| = \sum_i (a_i^+ + a_i^-) \leq k''$ , so the solution is good and  $k''$ -close to  $(A, B)$ .

We finally need the solution  $(A', B')$  to be consistent with both  $f$  and  $M$ . For any vertex  $v \in T$ , if  $v$  has color 0, then  $v \in A$  if and only if  $v \in X$  (recall  $R = AY \cup BX$ ), otherwise by including  $v$  in  $R$ , one can change the belonging of  $v$ . So the consistency of  $f$  requires certain components to be included into or excluded from  $R'_i$ . For  $(u, v) \in M$ , if  $u$  and  $v$  has the same color then  $M \in E(A, B) \cap E(X, Y)$  if and only if  $M \in E(A, B)$  which is always satisfied, otherwise  $M \in E(A, B) \cap E(X, Y)$  if and only if the label-1 vertex is not in  $R$ , so consistency of  $M$  requires certain components to be excluded from  $R'_i$ .

The consistency constraints are all in the form of “removing a component from the candidates” or “impose certain components to be selected by  $R'_i$ ”, which is easy to find out and handle in polynomial time.

We finally calculate the running time of the algorithm. The algorithm first tries all possible  $(f, k'')$  in  $2^{O(k)}$  time, tries all colorings in  $2^{O(k \log q)} n^{O(1)}$  time, and finally guesses the correct values of  $\ell$  and  $a_i^+, a_i^-$  in  $2^{O(k \log k)}$  time. All other processes are in polynomial time. So the total running time is  $2^{O(k \log q)} n^{O(1)}$ .  $\triangleleft$

Next we show how to compute a  $(k, q)$ -cut if it exists.

$\triangleright$  **Claim 20.** There exists an algorithm that runs in time  $2^{O(k \log q)} n^{O(1)}$  and returns a  $(k, q)$ -balanced cut for a connected graph if it exists.

*Proof.* Suppose that  $L \subseteq V$  is a  $(k, q)$ -cut and  $R = V \setminus L$ . We first show that there exists  $E_L \subseteq E(L)$  that induces a connected subgraph and has at most  $O(q)$  edges but at least  $q$  non-marked edges. Consider a procedure where we start from  $V_L = \{v\}$  for an arbitrary vertex  $v \in L$ , and iteratively add one arbitrary vertex  $u \in L \cap N(V_L)$  to  $V_L$  until  $E(G(V_L))$  contains at least  $q$  non-marked edges. Since  $E(G(V_L))$  is connected and increased by at least one while marked edges form a matching,  $|V_L| \leq 2q + 2$ . Then we choose  $E_L \subseteq E(G(V_L))$  that contains at least  $q$  non-marked edges while  $(V_L, E_L)$  is connected and  $|E_L| \leq O(q)$ . Define  $V_R$  and  $E_R$  similarly.

Use Theorem 12 with  $a = O(q)$  and  $b = k$  to try at most  $2^{O(k \log q)} \log n$  colorings of edges with 2 colors. One coloring  $\chi$  colors  $E(L, R)$  with 0 and  $E_L \cup E_R$  with 1. After deleting all edges of color 0, there are two connected components  $V'_L$  and  $V'_R$  such that  $E_L \subseteq E(G(V'_L))$  and  $E_R \subseteq E(G(V'_R))$ . Note that  $V'_L \neq V'_R$  because  $L$  and  $R$  are separated by removing edges of color 0.

Then, trying every pair  $(U', V')$  of connected components after deleting 0-colored edges and computing the minimum cut separating  $U'$  and  $V'$  in the original graph will yield a cut of size at most  $k$  where each side is connected has at least  $q$  non-marked edges. We use the fact that in a connected graph, any minimum  $s$ - $t$  cut has both of its sides connected.  $\triangleleft$

Finally, we prove that the existence of a  $(k, q)$ -cut, combined with a recursion, allows us to make progress by reducing the number of unmarked edges. Let  $T(m)$  be the running time of our algorithm to solve the terminal version with  $m$  unmarked edges.

$\triangleright$  **Claim 21.** Suppose there exists a  $(k, q)$ -cut, in time  $T(\ell) + n^{O(1)}$ , one can reduce the current instance to another IMPROVEMAXCUT-UNCUT TERMINAL VERSION instance with at least  $\ell - q/2$  fewer unmarked edges for some  $\ell \in [q, m - q]$ .

Proof. Let  $(L, R)$  be a  $(k, q)$ -cut. Without loss of generality, suppose that  $L$  has at most  $k$  terminals from  $T$ , and let  $T_L \subseteq L$  be the set of vertices that have an edge to  $R$ . Letting  $T' = (T \cap L) \cup T_L \subseteq L$  be the new set of terminals for  $L$ , we know that  $|T'| \leq 2k$ . Let  $\ell$  be the number of unmarked edges in  $L$  and  $m$  be the total number of unmarked edges.

Recursively solve the terminal version with input  $(G(L), (L \cap A, L \cap B), k', T', M \cap E(G(L)))$ . For each  $f' : T' \rightarrow \{X, Y\}$  and  $0 \leq k'' \leq k'$ , it returns a partition  $(X_{f', k''}, Y_{f', k''})$  of  $L$  that is (1) consistent with  $f'$ , (2) consistent with  $M \cap E(G(L))$ , and (3)  $k''$ -close to  $(L \cap A, L \cap B)$  (or nothing).

Let  $f^* : T' \rightarrow \{X, Y\}$  be the correct guessing with respect to the good partition  $(X, Y)$  and  $k^*$  be the correct closeness parameter  $|C(A \cap L, B \cap L) \Delta C(X \cap L, Y \cap L)|$  in  $L$ . We update  $(X, Y)$  with  $X \leftarrow (X \cap R) \cup X_{f^*, k^*}$  and  $Y \leftarrow (Y \cap R) \cup Y_{f^*, k^*}$ . Note that the new  $(X, Y)$  is still a good cut since  $(X_{f^*, k^*}, Y_{f^*, k^*})$  is good within  $L$  and the interaction between  $L$  and  $R$  only depends on the terminals  $T'$ , where the previous and new solutions agree. By the guarantee (2) and (3) in the above paragraph, the new  $(X, Y)$  is still  $k'$ -close to  $(A, B)$  and consistent with  $M$ .

Since  $L$  has at least  $q = k^2 2^{2k+2}$  non-marked edges and there are at most  $2^{2k}(k+1)$  possible values for  $f'$  and  $k''$ , we have  $|E(L) \setminus (\cup_{f', k''} (C(A \cap L, B \cap L) \Delta C(X_{f', k''}, Y_{f', k''})))| \geq \ell - q/2$ . These edges do not belong to  $C(A, B) \Delta C(X, Y) = E(A, B) \Delta E(X, Y)$  for sure. For each such edge  $e$ , perform the following operation.

- If  $e \notin E(A, B)$ , which means that  $e$  is not cut in the current solution, this means that  $e$  is not cut in the good solution too. Contract it.
- If  $e \in E(A, B)$ , mark it. To ensure that the set of marked edges  $M$  is a matching, if  $e = (u, v)$ ,  $f = (v, w)$  are both in  $M$  with  $u \neq w$ , then we know that  $u, w$  will be in the same side in the good solution, so we can contract them.

By doing this, the number of unmarked edges is decreased by at least  $\ell - q/2$ .  $\triangleleft$

With the three main claims, we finish the proof of Theorem 17.

**Proof of Theorem 17.** From the definition of the terminal version, given graph  $G$  and the current solution  $(A, B)$ , running the terminal version with  $T = M = \emptyset$  and  $k' = k$  we can find a good cut given the promise  $|C(A, B) \Delta C(X, Y)| \leq k$  for some (possibly different) good cut  $(X, Y)$ .

Let us analyze the running time. As defined previously, let  $T(m)$  be the running time of our algorithm with  $m$  unmarked edges. Since the graph is connected and marked edges form a matching, we have  $m \geq \frac{n}{2}$ . Our algorithm first uses Claim 20 to find a  $(k, q)$ -cut in time  $2^{O(k \log q)} n^{O(1)}$ . If there is none, it uses Claim 19 to solve the terminal version in time  $2^{O(k \log q)} n^{O(1)}$ . Otherwise, it uses Claim 21 to reduce the number of unmarked edges by  $\ell - q/2$  in time  $T(\ell) + n^{O(1)}$ , for some  $\ell \in [q, m - q]$ . Therefore, we have the following recurrence relation:

$$T(m) \leq 2^{O(k \log q)} n^{O(1)} + \max_{\ell \in [q, m - q]} \left( T(\ell) + T(m - (\ell - q/2)) + n^{O(1)} \right).$$

Using  $m \in [\frac{n}{2}, n^2]$ ,  $T(m) \leq 2^{O(k \log q)} n^{O(1)}$  satisfies the above.  $\blacktriangleleft$

---

## References

- 1 Aditya Anand, Euiwoong Lee, and Amartya Sharma. Min-CSPs on complete instances. In *Proceedings of the 2025 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2025.
- 2 Edouard Bonnet, László Egri, and Dániel Marx. Fixed-parameter approximability of boolean mincsp. In *24th European Symposium on Algorithms (ESA 2016)*, page 246, 2016.

- 3 Nicolas Bousquet, Jean Daligault, and Stéphan Thomassé. Multicut is fpt. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 459–468, 2011. doi:10.1145/1993636.1993698.
- 4 Rajesh Chitnis, Marek Cygan, MohammadTaghi Hajiaghayi, Marcin Pilipczuk, and Michał Pilipczuk. Designing fpt algorithms for cut problems using randomized contractions. *SIAM Journal on Computing*, 45(4):1171–1229, 2016. doi:10.1137/15M1032077.
- 5 Marek Cygan, Daniel Lokshtanov, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. Minimum bisection is fixed parameter tractable. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 323–332, 2014. doi:10.1145/2591796.2591852.
- 6 Marek Cygan, Marcin Pilipczuk, Michał Pilipczuk, and Jakub Onufry Wojtaszczyk. On multiway cut parameterized above lower bounds. *ACM Transactions on Computation Theory (TOCT)*, 5(1):1–11, 2013. doi:10.1145/2462896.2462899.
- 7 Konrad K Dabrowski, Peter Jonsson, Sebastian Ordyniak, George Osipov, and Magnus Wahlström. Almost consistent systems of linear equations. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3179–3217. SIAM, 2023. doi:10.1137/1.9781611977554.CH121.
- 8 Michael R. Fellows, Fedor V. Fomin, Daniel Lokshtanov, Frances Rosamond, Saket Saurabh, and Yngve Villanger. Local search: Is brute-force avoidable? *Journal of Computer and System Sciences*, 78(3):707–719, 2012. In Commemoration of Amir Pnueli. doi:10.1016/j.jcss.2011.10.003.
- 9 Serge Gaspers, Eun Jung Kim, Sebastian Ordyniak, Saket Saurabh, and Stefan Szeider. Don't be strict in local search! In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 26, pages 486–492, 2012. doi:10.1609/AAAI.V26I1.8128.
- 10 Ken-ichi Kawarabayashi and Mikkel Thorup. The minimum k-way cut of bounded size is fixed-parameter tractable. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 160–169. IEEE, 2011. doi:10.1109/FOCS.2011.53.
- 11 Eun Jung Kim, Stefan Kratsch, Marcin Pilipczuk, and Magnus Wahlström. Solving hard cut problems via flow-augmentation. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 149–168. SIAM, 2021. doi:10.1137/1.9781611976465.11.
- 12 Eun Jung Kim, Stefan Kratsch, Marcin Pilipczuk, and Magnus Wahlström. Directed flow-augmentation. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 938–947, 2022. doi:10.1145/3519935.3520018.
- 13 Eun Jung Kim, Stefan Kratsch, Marcin Pilipczuk, and Magnus Wahlström. Flow-augmentation iii: Complexity dichotomy for boolean csps parameterized by the number of unsatisfied constraints. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3218–3228. SIAM, 2023. doi:10.1137/1.9781611977554.CH122.
- 14 Andrei Krokhin and Dániel Marx. On the hardness of losing weight. *ACM Transactions on Algorithms (TALG)*, 8(2):1–18, 2012. doi:10.1145/2151171.2151182.
- 15 Tommaso Lanciano, Atsushi Miyauchi, Adriano Fazzino, and Francesco Bonchi. A survey on the densest subgraph problem and its variants. *ACM Computing Surveys*, 56(8):1–40, 2024. doi:10.1145/3653298.
- 16 Daniel Lokshtanov, Saket Saurabh, and Vaishali Surianarayanan. A parameterized approximation scheme for min k-cut. *SIAM Journal on Computing*, (0):FOCS20–205, 2022.
- 17 Dániel Marx. Parameterized complexity of constraint satisfaction problems. *computational complexity*, 14:153–183, 2005. doi:10.1007/S00037-005-0195-9.
- 18 Dániel Marx and Igor Razgon. Fixed-parameter tractability of multicut parameterized by the size of the cutset. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 469–478, 2011. doi:10.1145/1993636.1993699.
- 19 Dániel Marx and Ildikó Schlotter. Stable assignment with couples: Parameterized complexity and local search. *Discrete Optimization*, 8(1):25–40, 2011. Parameterized Complexity of Discrete Optimization. doi:10.1016/j.disopt.2010.07.004.

- 20 Moni Naor, Leonard J Schulman, and Aravind Srinivasan. Splitters and near-optimal derandomization. In *Proceedings of IEEE 36th Annual Foundations of Computer Science*, pages 182–191. IEEE, 1995. doi:10.1109/SFCS.1995.492475.
- 21 NS Narayanaswamy, Venkatesh Raman, MS Ramanujan, and Saket Saurabh. Lp can be a cure for parameterized problems. In *29th International Symposium on Theoretical Aspects of Computer Science (STACS 2012)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2012. doi:10.4230/LIPIcs.STACS.2012.338.
- 22 Venkatesh Raman, MS Ramanujan, and Saket Saurabh. Paths, flowers and vertex cover. In *Algorithms–ESA 2011: 19th Annual European Symposium, Saarbrücken, Germany, September 5-9, 2011. Proceedings 19*, pages 382–393. Springer, 2011.
- 23 Igor Razgon and Barry O’Sullivan. Almost 2-sat is fixed-parameter tractable. In *Automata, Languages and Programming: 35th International Colloquium, ICALP 2008, Reykjavik, Iceland, July 7-11, 2008, Proceedings, Part I 35*, pages 551–562. Springer, 2008.
- 24 Neil Robertson and Paul D Seymour. Graph minors. xiii. the disjoint paths problem. *Journal of combinatorial theory, Series B*, 63(1):65–110, 1995. doi:10.1006/JCTB.1995.1006.
- 25 Stefan Szeider. The parameterized complexity of k-flip local search for sat and max sat. *Discrete Optimization*, 8(1):139–145, 2011. doi:10.1016/J.DISOPT.2010.07.003.