



HAL
open science

Filling Crosswords is Very Hard

Laurent Gourvès, Ararat Harutyunyan, Michael Lampis, Nikolaos Melissinos

► **To cite this version:**

Laurent Gourvès, Ararat Harutyunyan, Michael Lampis, Nikolaos Melissinos. Filling Crosswords is Very Hard. 32nd International Symposium on Algorithms and Computation (ISAAC), Dec 2021, Fukuoka, Japan. 10.4230/LIPIcs.ISAAC.2021.48 . hal-03839913

HAL Id: hal-03839913

<https://hal.science/hal-03839913v1>

Submitted on 4 Nov 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

1 Filling Crosswords is Very Hard

2 **Laurent Gourvès** ✉

3 Université Paris-Dauphine, Université PSL, CNRS, LAMSADE, 75016, Paris, France

4 **Ararat Harutyunyan** ✉

5 Université Paris-Dauphine, Université PSL, CNRS, LAMSADE, 75016, Paris, France

6 **Michael Lampis** ✉ 

7 Université Paris-Dauphine, Université PSL, CNRS, LAMSADE, 75016, Paris, France

8 **Nikolaos Melissinos** ✉ 

9 Université Paris-Dauphine, Université PSL, CNRS, LAMSADE, 75016, Paris, France

10 — Abstract —

11 We revisit a classical crossword filling puzzle which already appeared in Garey&Jonhson’s book.
12 We are given a grid with n vertical and horizontal slots and a dictionary with m words and are
13 asked to place words from the dictionary in the slots so that shared cells are consistent. We attempt
14 to pinpoint the source of intractability of this problem by carefully taking into account the structure
15 of the grid graph, which contains a vertex for each slot and an edge if two slots intersect. Our
16 main approach is to consider the case where this graph has a tree-like structure. Unfortunately, if
17 we impose the common rule that words cannot be reused, we discover that the problem remains
18 NP-hard under very severe structural restrictions, namely, if the grid graph is a union of stars
19 and the alphabet has size 2, or the grid graph is a matching (so the crossword is a collection of
20 disjoint crosses) and the alphabet has size 3. The problem does become slightly more tractable
21 if word reuse is allowed, as we obtain an m^{tw} algorithm in this case, where tw is the treewidth of
22 the grid graph. However, even in this case, we show that our algorithm cannot be improved to
23 obtain fixed-parameter tractability. More strongly, we show that under the ETH the problem cannot
24 be solved in time $m^{o(k)}$, where k is the number of horizontal slots of the instance (which trivially
25 bounds tw).

26 Motivated by these mostly negative results, we also consider the much more restricted case
27 where the problem is parameterized by the number of slots n . Here, we show that the problem does
28 become FPT (if the alphabet has constant size), but the parameter dependence is exponential in
29 n^2 . We show that this dependence is also justified: the existence of an algorithm with running time
30 $2^{o(n^2)}$, even for binary alphabet, would contradict the randomized ETH. Finally, we consider an
31 optimization version of the problem, where we seek to place as many words on the grid as possible.
32 Here it is easy to obtain a $\frac{1}{2}$ -approximation, even on weighted instances, simply by considering only
33 horizontal or only vertical slots. We show that this trivial algorithm is also likely to be optimal,
34 as obtaining a better approximation ratio in polynomial time would contradict the Unique Games
35 Conjecture. The latter two results apply whether word reuse is allowed or not.

36 **2012 ACM Subject Classification** Theory of Computation → Design and Analysis of Algorithms →
37 Parameterized Complexity and Exact Algorithms; Theory of Computation → Design and Analysis
38 of Algorithms → Approximation Algorithms

39 **Keywords and phrases** Crossword Puzzle, Treewidth, ETH

40 **Digital Object Identifier** 10.4230/LIPIcs.ISAAC.2021.48

41 **Related Version** A full version of the paper is available at <https://arxiv.org/abs/2109.11203>

42 **Acknowledgements** We thank Dominique Taton for communicating the puzzle to us.

43 **1** Introduction

44 Crossword puzzles are one-player games where the goal is to fill a (traditionally two-
45 dimensional) grid with words. Since their first appearance more than 100 years ago, crossword



© Laurent Gourvès, Ararat Harutyunyan, Michael Lampis, and Nikolaos Melissinos;
licensed under Creative Commons License CC-BY 4.0


32nd International Symposium on Algorithms and Computation (ISAAC 2021).

Editors: Hee-Kap Ahn and Kunihiko Sadakane; Article No. 48; pp. 48:1–48:16

Leibniz International Proceedings in Informatics



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



■ **Figure 1** Place valid words in this grid. In a possible instance, letters S, U, I, V, R, E, and T have weight 7, 5, 4, 2, 6, 1, and 3, respectively. Any other letter has null weight. Try to obtain at least 330 points.

46 puzzles have rapidly become popular. Nowadays, they can be found in many newspapers and
 47 magazines around the world like the *New York Times* in the USA, or *Le Figaro* in France.
 48 Besides their obvious recreational interest, crossword puzzles are valued tools in education
 49 [2] and medicine. In particular, crossword puzzles participation seems to delay the onset
 50 of memory decline [14]. They are also helpful for developing and testing computational
 51 techniques; see for example [16]. In fact, both the design and the completion of a puzzle
 52 are challenging. In this article, we are interested in the task of solving a specific type of
 53 crossword puzzle.

54 There are different kinds of crossword puzzles. In the most famous ones, some clues are
 55 given together with the place where the answers should be located. A solution contains
 56 words that must be consistent with the given clues, and the intersecting pairs of words are
 57 constrained to agree on the letter they share. *Fill-in* crossword puzzles do not come with
 58 clues. Given a list of words and a grid in which some slots are identified, the objective is to
 59 fill all the slots with the given words. The list of words is typically succinct and provided
 60 explicitly.

61 In a variant of fill-in crossword puzzle currently proposed in a French TV magazine [12],
 62 one has to find up to 14 words and place them in a grid (the grid is the same for every
 63 instance, see Figure 1 for an illustration). The words are not explicitly listed but they must
 64 be *valid* (for instance, belong to the French language). In an instance of the game, some
 65 specified letters have a positive weight; the other letters have weight zero. The objective is
 66 to find a solution whose weight – defined as the total sum of the letters written in the grid –
 67 is at least a given threshold.

68 The present work deals with a theoretical study of this fill-in crossword puzzle (the grid is
 69 not limited to the one of Figure 1). We are mainly interested in two problems: Can the grid
 70 be entirely completed? How can the weight of a solution be maximized? Hereafter, these
 71 problems are called CROSSWORD PUZZLE DECISION and CROSSWORD PUZZLE OPTIMIZATION
 72 (CP-DEC and CP-OPT in short), respectively.

73 CP-DEC is not new; see GP14 in [5]. The proof of NP-completeness is credited to a
 74 personal communication with Lewis and Papadimitriou. Thereafter, an alternative NP-
 75 completeness proof appeared in [4] (see also [10]). Other articles on crossword puzzles exist
 76 and they are mostly empirically validated techniques coming from Artificial Intelligence and
 77 Machine Learning; see for example [6, 13, 11, 1, 16, 15] an references therein.

78 **Our Results** Our goal in this paper is to pinpoint the relevant structural parameters that
 79 make filling crossword puzzles intractable. We begin by examining the structure of the given
 80 grid. It is natural to think that, if the structure of the grid is tree-like, then the problem
 81 should become easier, as the vast majority of problems are tractable on graphs of small
 82 treewidth. We only partially confirm this intuition: by taking into account the structure of a

graph that encodes the intersections between slots (the grid-graph) we show in Section 3 that CP-OPT can be solved in polynomial time on instances of constant treewidth. However, our algorithm is not fixed-parameter tractable and, as we show, this cannot be avoided, even if one considers the much more restricted case where the problem is parameterized by the number of horizontal slots, which trivially bounds the grid-graph's treewidth (Theorem 4). More devastatingly, we show that if we also impose the natural rule that words cannot be reused, the problem already becomes NP-hard when the grid graph is a matching for alphabets of size 3 (Theorem 6), or a union of stars for a binary alphabet (Theorem 5). Hence, a tree-like structure does not seem to be of much help in rendering crosswords tractable.

We then go on to consider CP-OPT parameterized by the total number of slots n . This is arguably a very natural parameterization of the problem, as in real-life crosswords, the size of the grid can be expected to be significantly smaller than the size of the dictionary. We show that in this case the problem does become fixed-parameter tractable (Corollary 9), but the running time of our algorithm is exponential in n^2 . Our main result is to show that this disappointing dependence is likely to be best possible: even for a binary alphabet, an algorithm solving CP-DEC in time $2^{o(n^2)}$ would contradict the randomized ETH (Theorem 12). Note that all our positive results up to this point work for the more general CP-OPT, while our hardness results apply to CP-DEC.

Finally, in Section 5 we consider the approximability of CP-OPT. Here, it is easy to obtain a $\frac{1}{2}$ -approximation by only considering horizontal or vertical slots. We are only able to slightly improve upon this, giving a polynomial-time algorithm with ratio $\frac{1}{2} + O(\frac{1}{n})$. Our main result in this direction is to show that this is essentially best possible: obtaining an algorithm with ratio $\frac{1}{2} + \epsilon$ would falsify the Unique Games Conjecture (Theorem 15).

2 Problem Statement and Preliminaries

We are given a dictionary $\mathcal{D} = \{d_1, \dots, d_m\}$ whose words are constructed on an alphabet $\mathcal{L} = \{l_1, \dots, l_\ell\}$, and a two-dimensional grid consisting of horizontal and vertical slots. A slot is composed of consecutive cells. Horizontal slots do not intersect each other; the same goes for vertical slots. However horizontal slots can intersect vertical slots. A cell is *shared* if it lies at the intersection of two slots. Unless specifically stated, n , m and ℓ denote the total number of slots, the size of \mathcal{D} , and the size of \mathcal{L} , respectively. Finally, let us mention that we consider only instances where the alphabet is of constant size, i.e., $\ell = O(1)$.

In a feasible solution, each slot S receives either a word of \mathcal{D} of length $|S|$, or nothing (we sometimes say that a slot receiving nothing gets an *empty word*). Each cell gets at most one letter, and the words assigned to two intersecting slots must agree on the letter placed in the shared cell. All filled horizontal slots get words written from left to right (across) while all vertical slots get words written from top to bottom (down).

There is a weight function $w : \mathcal{L} \rightarrow \mathbb{N}$. The weight of a solution is the total sum of the weights of the letters placed in the grid. Observe that, for a given solution, the total weight of all filled-in words is not the same as the weight of this solution as, in the latter, the letters of the shared cells are counted only once.

The two main problems studied in this article are the following. Given a grid, a dictionary \mathcal{D} on alphabet \mathcal{L} , and a weight function $w : \mathcal{L} \rightarrow \mathbb{N}$, the objective of CROSSWORD PUZZLE OPTIMIZATION (CP-OPT in short) is to find a feasible solution of maximum weight. Given a grid and a dictionary \mathcal{D} on alphabet \mathcal{L} , the question posed by CROSSWORD PUZZLE DECISION (CP-DEC in short) is whether the grid can be completely filled or not?

Two cases will be considered: whether each word is used at most once, or if each word

129 can be assigned multiple times. In this article, we will sometimes suppose that some cells are
 130 pre-filled with some elements of \mathcal{L} . In this case, a solution is feasible if it is consistent with
 131 the pre-filled cells. Below we propose a first result when all the shared cells are pre-filled.

132 ► **Proposition 1.** *CP-DEC and CP-OPT can be solved in polynomial time if all the shared
 133 cells in the grid are pre-filled, whether word reuse is allowed or not.*

134 **Proof.** If word reuse is allowed, then for each combination of letters placed in these cells, we
 135 greedily fill out the rest of each slot with the maximum value word that can still be placed
 136 there. This is guaranteed to produce the optimal solution. On the other hand, if word reuse
 137 is not allowed, we construct a bipartite graph, with elements of \mathcal{D} on one side and the slots
 138 on the other, and place an edge between a word and a slot if the word can still be placed
 139 in the slot. If we give each edge weight equal to the value of its incident word reduced by
 140 the weight of the letters imposed by the shared cells of the slot, then an optimal solution
 141 corresponds to a maximum weight matching. ◀

142 One can associate a bipartite graph, hereafter called the *grid graph*, with each grid: each
 143 slot is a vertex and two vertices share an edge if the corresponding slots overlap. The grid
 144 (and then, the grid graph) is not necessarily connected.

145 Let us also note that so far we have been a bit vague about the encoding of the problem.
 146 Concretely, we could use a simple representation which lists for each slot the coordinates of
 147 its first cell, its size, and whether the slot is horizontal or vertical; and then supplies a list of
 148 all words in the dictionary and an encoding of the weight function. Such a representation
 149 would allow us to perform all the basic operations needed by our algorithms in polynomial
 150 time, such as deciding if it is possible to place a word d in a slot S , and which letter would
 151 then be placed in any particular cell of S . However, one drawback of this encoding is that its
 152 size may not be polynomially bounded in $n + m$, as some words may be exponentially long.
 153 We can work around this difficulty by using a more succinct representation: we are given
 154 the same information as above regarding the n slots; for each word we are given its total
 155 weight; and for each slot S and word d , we are told whether d fits exactly in S , and if yes,
 156 which letters are placed in the cells of S which are shared with other slots. Since the number
 157 of shared cells is $O(n^2)$ this representation is polynomial in $n + m$ and it is not hard to see
 158 that we are still able to perform any reasonable basic operation in polynomial time and that
 159 we can transform an instance given in the simple representation to this more succinct form.
 160 Hence, in the remainder, we will always assume that the size of the input is polynomially
 161 bounded in $n + m$.

162 We will rely on the Exponential Time Hypothesis (ETH) of Impagliazzo, Paturi, and
 163 Zane [8], which states the following:

164 ► **Conjecture 2.** *Exponential Time Hypothesis: there exists an $\epsilon > 0$, such that 3-SAT on
 165 instances with n variables and m clauses cannot be solved in time $2^{\epsilon(n+m)}$.*

166 Note that it is common to use the slightly weaker formulation which states the ETH as
 167 the assumption that 3-SAT cannot be solved in time $2^{o(n+m)}$. This is known to imply that
 168 k -INDEPENDENT SET cannot be solved in time $n^{o(k)}$ [3]. We use this fact in Theorem 4. In
 169 Section 4 we will rely on the randomized version of the ETH, which has the same statement
 170 as Conjecture 2 but for randomized algorithms with expected running time $2^{\epsilon(n+m)}$.

171 **3 When the Grid Graph is Tree-like**

172 In this section we are considering instances of CP-DEC and CP-OPT where the grid graph is
 173 similar to a tree. First, we give an algorithm for both problems in cases where the grid graph

174 has bounded treewidth and we are allowed to reuse words and we show that this algorithm
 175 is essentially optimal. Then, we show that CP-DEC and CP-OPT are much harder to deal
 176 with, in the case where we are not allowed to reuse words, by proving that the problems are
 177 NP-hard even for instances where the grid graph is just a matching. For the instances such
 178 that CP-DEC is NP-hard, we know that CP-OPT is NP-hard. That happens because we
 179 can assume that all the letters have weight equal to 1 hence a solution for CP-DEC is an
 180 optimal solution for CP-OPT.

181 3.1 Word Reuse

182 We propose a dynamic programming algorithm for CP-OPT and hence also for CP-DEC.
 183 Note that it can be extended to the case where some cells of the instance are pre-filled.

184 ► **Theorem 3.** *If we allow word reuse, then CP-OPT can be solved in time $(m+1)^{tw}(n+m)^{O(1)}$
 185 on inputs where tw is the treewidth of the grid graph.*

186 **Proof.** As the techniques we are going to use are standard we are sketching some details. For
 187 more details on tree decomposition (definition and terminology) see [3, Chap. 7]. Assuming
 188 that we have a rooted nice tree decomposition of the grid graph, we are going to perform
 189 dynamic programming on the nodes of this tree decomposition. For a node B_t of the given
 190 tree decomposition of the grid graph we denote by B_t^\downarrow the set of vertices of the grid graph
 191 that appears in the nodes of the subtree with B_t as a root. Since each vertex of the grid
 192 graph corresponds to a slot, we interchangeably mention a vertex of the grid graph and its
 193 corresponding slot. In particular, we say that a solution σ assigns words to the vertices of
 194 the grid graph, and $\sigma(v)$ denotes the word assigned to v .

195 For each node B_t of the tree decomposition we are going to keep all the triplets (σ, W, W_t)
 196 such that:

- 197 ■ σ is an assignment of words to the vertices of B_t ;
- 198 ■ W is the weight of σ restricted to the vertices appearing in B_t ;
- 199 ■ and W_m is the maximum weight, restricted to the vertices appearing in B_t^\downarrow , of an
 200 assignment consistent with σ .

201 In order to create all the possible triplets for all the nodes of the tree decomposition we are
 202 going to explore the nodes from leaves to the root. Therefore, each time we visit a node we
 203 assume that we have already created the triplets for all its children. Let us explain how we
 204 deal with the different types of nodes.

205 In the Leaf nodes we have no vertices so we keep an empty assignment (σ does not assign
 206 any word) and the weights W and W_m are equal to 0.

207 For an Introduce node B_t we need to take in consideration its child node. Assume that u
 208 is the introduced vertex; for each triplet (σ, W, W_m) of the child node we are going to create
 209 all the triplets (σ', W', W'_m) for the new node as follows. First we find all the words $d \in \mathcal{D}$
 210 that fit in the corresponding slot of u and respect the assignment σ (i.e., if there are cells
 211 that are already filled under σ and d uses these cells then it must have the same letters). We
 212 create one triplet (σ', W', W'_m) for each such a d as follows:

- 213 ■ We set $\sigma'(u) := d$ and $\sigma'(v) := \sigma(v)$ for all $v \in B_t \setminus \{u\}$.
- 214 ■ We can easily calculate the total weight, W' , of the words in B_t where the shared letters
 215 are counted only once under the assignment σ' .
- 216 ■ For the maximum weight W'_m we know that it is increased by the same amount as W ; so
 217 we set $W'_m = W_m + W' - W$.

218 Observe that we do not need to consider the intersection with slots whose vertices appear in
 219 $B_t^\downarrow \setminus B_t$ as each node of a tree decomposition is a cut set.

220 Finally, we need to take in consideration that we can leave a slot empty. For this case we
 221 create a new word d_* which, we assume that, fits in all slots and d_* has weight 0. Because
 222 the empty word has weight 0, W' and W'_m are identical to W and W_m so for each triplet of
 223 the child node, we only need to extend σ by assigning d_* to u . In the case we assign the
 224 empty word somewhere we will consider that the cells of this slot are empty unless another
 225 word $d \neq d_*$ uses them.

226 For the Forget nodes we need to restrict the assignments of the child node to the vertex
 227 set of the Forget node, as it has been reduced by one vertex (the forgotten vertex), and
 228 reduce the weight W (which we can calculate easily). The maximum weight is not changed
 229 by the deletion.

230 However, if we restrict the assignments we may end up with several triplets (σ, W, W_m)
 231 with identical assignments σ . In that case we are keeping only the triplet with maximum
 232 W_m . Observe that we are allowed to keep only triplets with the maximum W_m because each
 233 node of a tree decomposition is a cut set so the same holds for the Forget nodes. Specifically,
 234 the vertices that appear in the nodes higher than a Forget node B_t of the tree decomposition
 235 do not have edges incident to vertices in $B_t^{\downarrow} \setminus B_t$ so we only care for the assignment in B_t .

236 Finally, we need to consider the Join nodes. Each Join node has exactly two children.
 237 For each possible assignment σ on the vertices of this Join node, we create a triplet iff this σ
 238 appears in a triplet of both children of the Join node.

239 Because W is related only to the assignment σ , it is easy to see that it will be the same
 240 as in the children of the Join node. So we need to find the maximum weight W_m . Observe
 241 that between the vertices that appear in the subtrees of two children of a Join node there are
 242 no edges except those incident to the vertices of the Join node. Therefore, we can calculate
 243 the maximum weight W_m as follows: first we consider the maximum weight of each child of
 244 the Join node reduced by W , we add all these weights and, in the end, we add again the W .
 245 It is easy to see that this way we consider the weight of the cells appearing in each subtree
 246 without those of the slots of the Join node and we add the weight of the words assigned to
 247 the vertices of the Join node in the end.

248 For the running time we need to observe that the number of nodes of a nice tree
 249 decomposition is $O(\text{tw} \cdot n)$ and all the other calculations are polynomial in $n + m$ so we only
 250 need to consider the different assignments for each node. Because for each vertex we have
 251 $|\mathcal{D}| + 1$ choices, the number of different assignments for a node is at most $(|\mathcal{D}| + 1)^{\text{tw}+1}$. ◀

252 It seems that the algorithm we propose for CP-DEC is essentially optimal, even if we
 253 consider a much more restricted case.

254 ► **Theorem 4.** *CP-DEC with word reuse is $W[1]$ -hard parameterized by the number of
 255 horizontal slots of the grid, even for alphabets with two letters. Furthermore, under the ETH,
 256 no algorithm can solve this problem in time $m^{o(k)}$, where k is the number of horizontal slots.*

257 **Proof.** We perform a reduction from k -INDEPENDENT SET, where we are given a graph
 258 $G = (V, E)$ with $|V|$ vertices and $|E|$ edges and are looking for an independent set of size k .
 259 This problem is well-known to be $W[1]$ -hard and not solvable in $|V|^{o(k)}$ time under the ETH
 260 [3]. We assume without loss of generality that $|E| \neq k$. Furthermore, we can safely assume
 261 that G has no isolated vertices.

262 We first describe the grid of our construction which fits within an area of $2k - 1$ lines
 263 and $2|E| - 1$ columns. We construct:

- 264 1. k horizontal slots, each of length $2|E| - 1$ (so each of these slots is as long horizontally as
 265 the whole grid). We place these slots in the unique way so that no two of these slots are
 266 in consecutive lines. We number these horizontal slots $1, \dots, k$ from top to bottom.

267 2. $|E|$ vertical slots, each of length $2k - 1$ (so each of these slots is long enough to cover the
 268 grid top to bottom). We place these slots in the unique way so that no two of them are
 269 in consecutive columns. We number them $1, \dots, |E|$ from left to right.

270 Before we describe the dictionary, let us give some intuition about the grid. The main
 271 idea is that in the k horizontal slots we will place k words that signify which vertices we
 272 selected from the original graph. Each vertical slot represents an edge of E , and we will be
 273 able to place a word in it if and only if we have not placed words representing two of its
 274 endpoints in the horizontal slots.

275 Our alphabet has two letters, say 0, 1. In the remainder, we assume that the edges of the
 276 original graph are numbered, that is, $E = \{e_1, \dots, e_{|E|}\}$. The dictionary is as follows:

- 277 1. For each vertex v we construct a word of length $2|E| - 1$. For each $i \in \{1, \dots, |E|\}$, if
 278 the edge e_i is incident on v , then the letter at position $2i - 1$ of the word representing v
 279 is 1. All other letters of the word representing v are 0. Observe that this means that if e_i
 280 is incident on v and we place the word representing v on a horizontal slot, the letter i
 281 will appear on the i -th vertical slot. Furthermore, the word representing v has a number
 282 of 1s equal to the degree of v .
- 283 2. We construct $k + 1$ words of length $2k - 1$. One of them is simply 0^{2k-1} . The remaining
 284 are $0^{2j-2}10^{2k-2j}$, for $j \in \{1, \dots, k\}$, that is, the words formed by placing a 1 in an
 285 odd-numbered position and 0s everywhere else. Observe that if we place one of these k
 286 words on a vertical slot, a 1 will be placed on exactly one horizontal slot.

287 This completes the construction. We now observe that the k horizontal slots correspond
 288 to a vertex cover of the grid-graph. Therefore, if the reduction preserves the answer, the
 289 hardness results for k -INDEPENDENT SET transfer to our problem, since we preserve the
 290 value of the parameter.

291 We claim that if there exists an independent set of size k in G , then it is possible to fill
 292 the grid. Indeed, take such a set S and for each $v \in S$ we place the word representing v in a
 293 horizontal slot. Consider the i -th vertical slot. We will place in this slot one of the $k + 1$
 294 words of length $2k - 1$. We claim that the vertical slot at this moment contains the letter 1
 295 at most once, and if 1 appears it must be at an odd position (since these are the positions
 296 shared with the horizontal slots). If this is true, clearly there is a word we can place. To see
 297 that the claim is true, recall that since S is an independent set of k distinct vertices, there
 298 exists at most one vertex in S incident on e_i .

299 For the converse direction, recall that $|E| \neq k$. This implies that if there is a way to fill
 300 out the whole grid, then words representing vertices must go into horizontal slots and words
 301 of length $2k - 1$ must go into vertical slots. By looking at the words that have been placed
 302 in the horizontal slots we obtain a collection of k (not necessarily distinct) vertices of G .
 303 We will prove that these vertices must actually be an independent set of size exactly k . To
 304 see this, consider the i -th vertical slot. If our collection of vertices contained two vertices
 305 incident on e_i , it would have been impossible to fill out the i -th vertical slot, since we would
 306 need a word with two 1s. Observe that the same argument rules out the possibility that
 307 our collection contains the same vertex v twice, as the column corresponding to any edge e_i
 308 incident on v would have been impossible to fill. ◀

309 3.2 No Word Reuse

310 If a word cannot be reused, then CP-DEC looks more challenging. Indeed, in the following
 311 theorem we prove that if reusing words is not allowed, then the problem becomes NP-hard

312 even if the grid graph is acyclic and the alphabet size is 2. (Note that if the alphabet size is
313 1, the problem is trivial, independent of the structure of the graph).

314 ► **Theorem 5.** *CP-DEC is NP-hard, even for instances where all of the following restrictions*
315 *apply: (i) the grid graph is a union of stars (ii) the alphabet contains only two letters (iii)*
316 *words cannot be reused.*

317 **Proof.** We show a reduction from 3-PARTITION. Recall that in 3-PARTITION we are given a
318 collection of $3n$ distinct positive integers x_1, \dots, x_{3n} and are asked if it is possible to partition
319 these integers into n sets of three integers (triples), such that all triples have the same sum.
320 This problem has long been known to be strongly NP-hard [5] and NP-hardness when the
321 integers are distinct was shown by Hulett et al. [7]. We can assume that $\sum_{i=1}^{3n} x_i = nB$ and
322 that if a partition exists each triple has sum B . Furthermore, we can assume without loss
323 of generality that $x_i > 6n$ for all $i \in \{1, \dots, 3n\}$ (otherwise, we can simply add $6n$ to all
324 numbers and adjust B accordingly without changing the answer).

325 Given an instance of 3-PARTITION as above, we construct a crossword instance as follows.
326 First, the alphabet only contains two letters, say the letters $*$ and $!$. To construct our
327 dictionary we do the following:

- 328 1. For each $i \in \{1, \dots, 3n\}$, we add to the dictionary one word of length x_i that begins with
329 $!$ and $n - 1$ words of length x_i that begin with $*$. The remaining letters of these words
330 are chosen in an arbitrary way so that all words remain distinct.
- 331 2. For each $i, j, k \in \{1, \dots, 3n\}$ with $i < j < k$ we check if $x_i + x_j + x_k = B$. If this is the
332 case, we add to the dictionary the word $*^{2i-2}!_*^{2j-2i-1}!_*^{2k-2j-1}!_*^{6n-2k}$. In other words,
333 we constructed a word that has $*$ everywhere except in positions $2i - 1, 2j - 1,$ and $2k - 1$.
334 The length of this word is $6n - 1$. Let f be the number of words added to the dictionary
335 in this step. We have $f \leq \binom{3n}{3} = O(n^3)$.

336 We now also need to specify our grid. We first construct f horizontal slots, each of length
337 $6n - 1$. Among these f slots, we select n , which we call the “interesting” horizontal slots.
338 For each interesting horizontal slot, we construct $3n$ vertical slots, such that the i -th of these
339 slots has length x_i and its first cell is the cell in position $2i - 1$ of the interesting horizontal
340 slot. This completes the construction, which can clearly be carried out in polynomial time.
341 Observe that the first two promised restrictions are satisfied as we have an alphabet with
342 two letters and each vertical slot intersects at most one horizontal slot (so the grid graph is
343 a union of stars).

344 We claim that if there exists a partition of the original instance, then we can place all
345 the words of the dictionary on the grid. Indeed, for each $i, j, k \in \{1, \dots, 3n\}$ such that
346 $\{x_i, x_j, x_k\}$ is one of the triples of the partition, we have constructed a word of length $6n - 1$
347 corresponding to the triple (i, j, k) , because $x_i + x_j + x_k = B$. We place each of these n
348 words on an interesting horizontal slot and we place the remaining words of length $6n - 1$ on
349 the non-interesting horizontal slots. Now, for every $i \in \{1, \dots, 3n\}$ we have constructed n
350 words, one starting with $!$ and $n - 1$ starting with $*$. We observe that among the interesting
351 horizontal slots, there is one that contains the letter $!$ at position $2i - 1$ (the one corresponding
352 to the triple containing x_i in the partition) and $n - 1$ containing the letter $*$ at position $2i - 1$.
353 By construction, the vertical slots that begin in these positions have length x_i . Therefore,
354 we can place all n words corresponding to x_i on these vertical slots. Proceeding in this way
355 we fill the whole grid, fulfilling the third condition.

356 For the converse direction, suppose that there is a way to fill the whole grid. Then, vertical
357 slots must contain words that were constructed in the second step and represent integers x_i ,

358 while horizontal slots must contain words constructed in the first step (this is a consequence
 359 of the fact that $x_i > 6n$ for all $i \in \{1, \dots, 3n\}$). We consider the n interesting horizontal
 360 slots. Each such slot contains a word that represents a triple (i, j, k) with $x_i + x_j + x_k = B$.
 361 We therefore collect these n triples and attempt to construct a partition from them. To do
 362 this, we must prove that each x_i must belong to exactly one of these triples. However, recall
 363 that we have exactly n words of length x_i (since all integers of our instance are distinct)
 364 and exactly n vertical slots of this length. We conclude that exactly one vertical slot must
 365 have ! as its first letter, therefore x_i appears in exactly one triple and we have a proper
 366 partition. ◀

367 Actually, the problem remains NP-hard even in the case where the grid graph is a
 368 matching and the alphabet contains three letters. This is proved for grid graphs composed
 369 of \mathcal{T} s, where a \mathcal{T} is a horizontal slot solely intersected by the first cell of a vertical slot.

370 ▶ **Theorem 6.** *CP-DEC is NP-hard, even for instances where all of the following restrictions*
 371 *apply: (i) each word can be used only once (ii) the grid is consisted only by \mathcal{T} s and (iii) the*
 372 *alphabet contains only three letters.*

373 ▶ **Remark 7.** Theorem 4 can be adjusted to work also for the case where word reuse is not
 374 allowed. We simply need to add a suffix of length $\log m$ to all words of length $2k - 1$ and add
 375 rows to the grid accordingly. Hence, under the ETH, no algorithm can solve this problem in
 376 time $m^{o(k)}$, where k is the number of horizontal slots.

377 Finally, observe that by filling the slots of a vertex cover of the grid graph, all the shared
 378 cells are pre-filled. Since there are at most m^k (where k is the size of the vertex cover) ways
 379 to assign words to these slots, by Proposition 1, we get the following corollary.

380 ▶ **Corollary 8.** *Given a vertex cover of size k of the grid graph we can solve CP-DEC and*
 381 *CP-OPT in time $m^k(n + m)^{O(1)}$. Furthermore, as vertex cover we can take the set of*
 382 *horizontal slots.*

383 Therefore, the bound given in Remark 7 for the parameter vertex cover is tight.

384 4 Parameterized by Total Number of Slots

385 In this section we consider a much more restrictive parameterization of the problem: we
 386 consider instances where the parameter is n , the total number of slots. Recall that in
 387 Theorem 4 (and Remark 7) we already considered the complexity of the problem parameterized
 388 by the number of *horizontal* slots of the instance. We showed that this case of the problem
 389 cannot be solved in $m^{o(k)}$ and that an algorithm with running time roughly m^k is possible
 390 whether word reuse is allowed or not.

391 Since parameterizing by the number of horizontal slots is not sufficient to render the
 392 problem FPT, we therefore consider our parameter to be the total number of slots. This is,
 393 finally, sufficient to obtain a simple FPT algorithm.

394 ▶ **Corollary 9.** *There is an algorithm that solves CP-DEC and CP-OPT in time $O^*(\ell^{n^2/4})$,*
 395 *where n is the total number of slots and ℓ the size of the alphabet, whether word reuse is*
 396 *allowed or not.*

397 Even though the running time guaranteed by Corollary 9 is FPT for parameter n , we
 398 cannot help but observe that the dependence on n is rather disappointing, as our algorithm is
 399 exponential *in the square* of n . It is therefore a natural question whether an FPT algorithm
 400 for this problem can achieve complexity $2^{o(n^2)}$, assuming the alphabet size is bounded. The
 401 main result of this section is to establish that this is likely to be impossible.

48:10 Filling Crosswords is Very Hard

402 **Overview** Our hardness proof consists of two steps. In the first step we reduce 3-SAT to
403 a version of the same problem where variables and clauses are partitioned into $O(\sqrt{n+m})$
404 groups, which we call SPARSE 3-SAT. The key property of this intermediate problem is that
405 interactions between groups of variables and groups of clauses are extremely limited. In
406 particular, for each group of variables V_i and each group of clauses C_j , at most one variable
407 of V_i appears in a clause of C_j . We obtain this rather severe restriction via a randomized
408 reduction that runs in expected polynomial time. The second step is to reduce SPARSE
409 3-SAT to CP-DEC. Here, every horizontal slot will represent a group of variables and every
410 vertical slot a group of clauses, giving $O(\sqrt{n+m})$ slots in total. Hence, an algorithm for
411 CP-DEC whose dependence on the total number of slots is subquadratic in the exponent will
412 imply a sub-exponential time (randomized) algorithm for 3-SAT. The limited interactions
413 between groups of clauses and variables will be key in allowing us to execute this reduction
414 using a *binary* alphabet.

415 Let us now define our intermediate problem.

416 **► Definition 10.** *In SPARSE 3-SAT we are given an integer n which is a perfect square and
417 a 3-SAT formula ϕ with at most n variables and at most n clauses, such that each variable
418 appears in at most 3 clauses. Furthermore, we are given a partition of the set of variables V
419 and the set of clauses C into \sqrt{n} sets $V_1, \dots, V_{\sqrt{n}}$ and $C_1, \dots, C_{\sqrt{n}}$ of size at most \sqrt{n} each,
420 such that for all $i, j \in [\sqrt{n}]$ the number of variables of V_i which appear in at least one clause
421 of C_j is at most one.*

422 Now, we are going to prove the hardness of SPARSE 3-SAT, which is the first step of our
423 reduction.

424 **► Lemma 11.** *Suppose the randomized ETH is true. Then, there exists an $\epsilon > 0$ such that
425 SPARSE 3-SAT cannot be solved in time $2^{\epsilon n}$.*

426 We are now ready to prove the main theorem of this section.

427 **► Theorem 12.** *Suppose the randomized ETH is true. Then, there exists an $\epsilon > 0$ such that
428 CP-DEC on instances with a binary alphabet cannot be solved in time $2^{\epsilon n^2} \cdot m^{O(1)}$. This
429 holds also for instances where all slots have distinct sizes (so words cannot be reused).*

430 **Proof.** Suppose for the sake of contradiction that for any fixed $\epsilon > 0$, CP-DEC on instances
431 with a binary alphabet can be solved in time $2^{\epsilon n^2} \cdot m^{O(1)}$. We will then contradict Lemma 11.
432 In particular, we will show that for any ϵ' we can solve SPARSE 3-SAT in time $2^{\epsilon' N}$, where
433 N is the upper bound on the number of variables and clauses. Fix some $\epsilon' > 0$ and suppose
434 that ϕ is an instance of SPARSE 3-SAT with at most N variables and at most N clauses,
435 where N is a perfect square. Recall that the variables are given partitioned into \sqrt{N} sets,
436 $V_1, \dots, V_{\sqrt{N}}$ and the clauses partitioned into \sqrt{N} sets $C_1, \dots, C_{\sqrt{N}}$. In the remainder, when
437 we write $V(C_j)$ we will denote the set of variables that appear in a clause of C_j . Recall
438 that the partition satisfies the property that for all $i, j \in [\sqrt{N}]$ we have $|V_i \cap V(C_j)| \leq 1$.
439 Suppose that the variables of ϕ are ordered x_1, x_2, \dots, x_N .

440 We construct a grid as follows: for each group V_i we construct a horizontal slot and for
441 each group C_j we construct a vertical slot, in a way that all slots have distinct lengths. More
442 precisely, the i -th horizontal slot, for $i \in [\sqrt{N}]$ is placed on row $2i - 1$, starts in the first
443 column and has length $2\sqrt{N} + 2i$. The j -th vertical slot is placed in column $2j - 1$, starts
444 in the first row and has length $5\sqrt{N} + 2j$. (As usual, we number the rows and columns
445 top-to-bottom and left-to-right). Observe that all horizontal slots intersect all vertical slots;
446 in particular, the cell in row $2i - 1$ and column $2j - 1$ is shared between the i -th horizontal
447 and j -th vertical slot, for $i, j \in [\sqrt{N}]$. We define \mathcal{L} to contain two letters $\{0, 1\}$.

448 What remains is to describe the dictionary.

- 449 ■ For each $i \in [\sqrt{N}]$ and for each assignment function $\sigma : V_i \rightarrow \{0, 1\}$ we construct a word
450 w_σ of length $2\sqrt{N} + 2i$. The word w_σ has the letter 0 in all positions, except positions
451 $2j - 1$, for $j \in [\sqrt{N}]$. For each such j , we consider σ restricted to $V_i \cap V(C_j)$. By the
452 properties of SPARSE 3-SAT, we have $|V_i \cap V(C_j)| \leq 1$. If $V_i \cap V(C_j) = \emptyset$ then we place
453 letter 0 in position $2j - 1$; otherwise we set in position $2j - 1$ the letter that corresponds
454 to the value assigned by σ to the unique variable of $V_i \cap V(C_j)$.
- 455 ■ For each $j \in [\sqrt{N}]$ and for each *satisfying* assignment function $\sigma : V(C_j) \rightarrow \{0, 1\}$, that
456 is, every assignment function that satisfies all clauses of C_j , we construct a word w'_σ of
457 length $5\sqrt{N} + 2j$. The word w'_σ has the letter 0 in all positions, except positions $2i - 1$,
458 for $i \in [\sqrt{N}]$. For each such i , we consider σ restricted to $V_i \cap V(C_j)$. If $V_i \cap V(C_j) = \emptyset$
459 then we place letter 0 in position $2i - 1$; otherwise we set in position $2i - 1$ the letter
460 that corresponds to the value assigned by σ to the unique variable of $V_i \cap V(C_j)$.

461 The construction is now complete. We claim that if ϕ is satisfiable, then it is possible to
462 fill out the grid we have constructed. Indeed, fix a satisfying assignment σ to the variables of
463 ϕ . For each $i \in [\sqrt{N}]$ let σ_i be the restriction of σ to V_i . We place in the i -th horizontal slot
464 the word w_{σ_i} . Similarly, for each $j \in [\sqrt{N}]$ we let σ'_j be the restriction of σ to $V(C_j)$ and
465 place $w'_{\sigma'_j}$ in the j -th vertical slot. Now if we examine the cell shared by the i -th horizontal
466 and j -th vertical slot, we can see that it contains a letter that represents σ restricted to (the
467 unique variable of) $V_i \cap V(C_j)$ or 0 if $V_i \cap V(C_j) = \emptyset$, and both the horizontal and vertical
468 word place the same letter in that cell.

469 For the converse direction, if the grid is filled, we can extract an assignment σ for the
470 variables of ϕ as follows: for each $x \in V_i$ we find a C_j such that x appears in some clause of
471 C_j (we can assume that every variable appears in some clause). We then look at the cell
472 shared between the i -th horizontal and the j -th vertical slot. The letter we have placed
473 in that cell gives an assignment for the variable contained $V_i \cap V(C_j)$, that is x . Having
474 extracted an assignment to all the variables, we claim it must satisfy ϕ . If not, there is a
475 group C_j that contains an unsatisfied clause. Nevertheless, in the j -th vertical slot we have
476 placed a word that corresponds to a *satisfying* assignment for the clauses of C_j , call it σ_j .
477 Then σ_j must disagree with σ in a variable x that appears in C_j . Suppose this variable is
478 part of V_i . Then, this would contradict the fact that we extracted an assignment for x from
479 the word placed in the i -th horizontal slot.

480 Observe that the new instance has $n = 2\sqrt{N}$ slots. If there exists an algorithm that
481 solves CP-DEC in time $2^{\epsilon n^2} m^{O(1)}$ for any $\epsilon > 0$, we set $\epsilon = \epsilon'/8$ (so ϵ only depends on ϵ')
482 and execute this algorithm on the constructed instance. We observe that $m \leq 2\sqrt{N} \cdot 7^{\sqrt{N}}$,
483 and that $2^{\epsilon n^2} \leq 2^{\epsilon' N/2}$. Assuming that N is sufficiently large, using the supposed algorithm
484 for CP-DEC we obtain an algorithm for SPARSE 3-SAT with complexity at most $2^{\epsilon' N}$. Since
485 we can do this for arbitrary ϵ' , this contradicts the randomized ETH. ◀

486 5 Approximability of CP-Opt

487 This section begins with a $(\frac{1}{2} + O(\frac{1}{n}))$ -approximation algorithm which works when words
488 can, or cannot, be reused. After that, we prove that under the unique games conjecture, an
489 approximation algorithm with a significantly better ratio is unlikely.

490 ▶ **Theorem 13.** CP-OPT is $(\frac{1}{2} + \frac{1}{2(\epsilon n + 1)})$ -approximable in polynomial time, for all $\epsilon \in (0, 1]$.

48:12 Filling Crosswords is Very Hard

491 **Proof.** Fix some $\varepsilon \in (0, 1]$. Let $k_v := \min(\lceil \frac{1}{\varepsilon} \rceil, n - h)$ and $r_v := \lceil \frac{n-h}{k_v} \rceil$, where h is the
 492 number of horizontal slots in the grid. Create r_v groups of vertical slots G_1, \dots, G_{r_v} such
 493 that $|G_i| \leq k_v$ for all $i \in [r_v]$ and $G_1 \cup \dots \cup G_{r_v}$ covers the entire set of vertical slots. For
 494 each G_i , guess an optimal choice of words, i.e., identical to a global optimum, and complete
 495 this partial solution by filling the horizontal slots (use the aforementioned matching technique
 496 where the words selected for G_i are excluded from \mathcal{D}). Each slot of $\bigcup_{j \neq i} G_j$ gets the empty
 497 word.

498 Since $|G_i| \leq k_v$, guessing an optimal choice of words for G_i by brute force requires
 499 at most m^{k_v} combinations. This is done r_v times (once for each G_i). The maximum
 500 matching runs in time $O((m+n)^2 \cdot mn)$. In all, the time complexity of the algorithm is
 501 $O(m^{k_v} \cdot r_v \cdot (m+n)^2 \cdot mn) \leq O(m^{1/\varepsilon} \cdot \varepsilon n \cdot (m+n)^2 \cdot mn)$.

502 Assume that, given an optimal solution, W_H^* and W_V^* are the total weight of the words
 503 assigned to the horizontal and vertical slots, respectively, both including the shared cells.
 504 Furthermore, let W_S^* be the weight of the letters assigned to the shared cells in the optimal
 505 solution. Observe that the weight of the optimal solution is $W_H^* + W_V^* - W_S^*$ and the weight
 506 of our solution is at least $W_H^* + \frac{1}{r_v}(W_V^* - W_S^*)$.

507 We repeat the same process, but the roles of vertical and horizontal slots are interchanged.
 508 Fix a parameter $k_h := \min(\lceil \frac{1}{\varepsilon} \rceil, h)$. Create $r_h := \lceil \frac{h}{k_h} \rceil$ groups of horizontal slots G_1, \dots, G_{r_h}
 509 such that $|G_i| \leq k_h$ for all $i \in [r_h]$ and $G_1 \cup \dots \cup G_{r_h}$ covers the entire set of horizontal slots.
 510 For each G_i , guess an optimal choice of words and complete this partial solution by filling
 511 the vertical slots. Each slot of $\bigcup_{j \neq i} G_j$ gets the empty word.

512 Using the same arguments as above, we can conclude that the time complexity is
 513 $O(m^{1/\varepsilon} \cdot \varepsilon n \cdot (m+n)^2 \cdot mn)$ and that we return a solution of weight at least $W_V^* + \frac{1}{r_h}(W_H^* - W_S^*)$.

514 Finally, between the two solutions, we return the one with the greater weight. It remains
 515 to argue about the approximation ratio. We need to consider two cases: $W_H^* \geq W_V^*$ and
 516 $W_V^* > W_H^*$.

517 Suppose $W_H^* \geq W_V^*$. The first approximate solution has value $W_H^* + \frac{1}{r_v}(W_V^* - W_S^*) \geq$
 518 $\frac{1+1/r_v}{2}(W_H^* + W_V^* - W_S^*)$. If $k_v = n-h$ then $r_v = 1$ and our approximation ratio is 1. Otherwise,
 519 $k_v = \lceil \frac{1}{\varepsilon} \rceil$ and $r_v = \lceil \frac{n-h}{\lceil 1/\varepsilon \rceil} \rceil \leq \frac{n-h}{\lceil 1/\varepsilon \rceil} + 1 = \frac{n-h + \lceil 1/\varepsilon \rceil}{\lceil 1/\varepsilon \rceil}$. It follows that $\frac{1}{r_v} \geq \frac{\lceil 1/\varepsilon \rceil}{n-h + \lceil 1/\varepsilon \rceil}$. Use
 520 $n-h + \lceil 1/\varepsilon \rceil \leq n + \frac{1}{\varepsilon}$ and $\lceil 1/\varepsilon \rceil \geq 1/\varepsilon$ to get that $\frac{1}{r_v} \geq \frac{1/\varepsilon}{n+1/\varepsilon} = \frac{1}{\varepsilon n + 1}$. Our approximation
 521 ratio is at least $\frac{1+1/(\varepsilon n + 1)}{2}$.

522 Suppose $W_V^* > W_H^*$. The second approximate solution has value $W_V^* + \frac{1}{r_h}(W_H^* - W_S^*) >$
 523 $\frac{1+1/r_h}{2}(W_H^* + W_V^* - W_S^*)$. If $k_h = h$, then our approximation ratio is 1. Otherwise, $k_h = \lceil \frac{1}{\varepsilon} \rceil$
 524 and, using the same arguments, our approximation ratio is at least $\frac{1+1/(\varepsilon n + 1)}{2}$.

525 Note that $\frac{1+1/(\varepsilon n + 1)}{2} \leq 1$. In all, we have a $\frac{1+1/(\varepsilon n + 1)}{2}$ -approximate solution in $O(m^{1/\varepsilon} \cdot$
 526 $\varepsilon n \cdot (m+n)^2 \cdot mn)$ for all $\varepsilon \in (0, 1]$. ◀

527 The previous approximation algorithm only achieves an approximation ratio of $\frac{1}{2} + O(\frac{1}{n})$,
 528 which tends to $\frac{1}{2}$ as n increases. At first glance this is quite disappointing, as someone can
 529 observe that a ratio of $\frac{1}{2}$ is achievable simply by placing words only on the horizontal or the
 530 vertical slots of the instance. Nevertheless, we are going to show that this performance is
 531 justified, as improving upon this trivial approximation ratio would falsify the Unique Games
 532 Conjecture (UGC).

533 Before we proceed, let us recall some relevant definitions regarding Unique Games. The
 534 UNIQUE LABEL COVER problem is defined as follows: we are given a graph $G = (V, E)$, with
 535 some arbitrary total ordering \prec of V , an integer R , and for each $(u, v) \in E$ with $u \prec v$ a
 536 1-to-1 constraint $\pi_{(u,v)}$ which can be seen as a permutation on $[R]$. The vertices of G are

537 considered as variables of a constraint satisfaction problem, which take values in $[R]$. Each
 538 constraint $\pi_{(u,v)}$ defines for each value of u a unique value that must be given to v in order
 539 to satisfy the constraint. The goal is to find an assignment to the variables that satisfies
 540 as many constraints as possible. The Unique Games Conjecture states that for all $\epsilon > 0$,
 541 there exists R , such that distinguishing instances of UNIQUE LABEL COVER for which it is
 542 possible to satisfy a $(1 - \epsilon)$ -fraction of the constraints from instances where no assignment
 543 satisfies more than an ϵ -fraction of the constraints is NP-hard. In this section we will need a
 544 slightly different version of this conjecture, which was defined by Khot and Regev as the
 545 Strong Unique Games Conjecture. Despite the name, Khot and Regev showed that this
 546 version is implied by the standard UGC. The precise formulation is the following:

547 ► **Theorem 14.** [Theorem 3.2 of [9]] *If the*

548 *Unique Games Conjecture is true, then for all $\epsilon > 0$ it is NP-hard to distinguish between*
 549 *the following two cases of instances of UNIQUE LABEL COVER $G = (V, E)$:*

550 ■ *(Yes case): There exists a set $V' \subseteq V$ with $|V'| \geq (1 - \epsilon)|V|$ and an assignment for V'*
 551 *such that all constraints with both endpoints in V' are satisfied.*

552 ■ *(No case): For any assignment to V , for any set $V' \subseteq V$ with $|V'| \geq \epsilon|V|$, there exists a*
 553 *constraint with both endpoints in V' that is violated by the assignment.*

554 Using the version of the UGC given in Theorem 14 we are ready to present our hardness
 555 of approximation argument for the crossword puzzle.

556 ► **Theorem 15.** *Suppose that the Unique Games Conjecture is true. Then, for all ϵ with*
 557 $\frac{1}{4} > \epsilon > 0$, *there exists an alphabet Σ_ϵ such that it is NP-hard to distinguish between the*
 558 *following two cases of instances of the crossword problem on alphabet Σ_ϵ :*

559 ■ *(Yes case): There exists a valid solution that fills a $(1 - \epsilon)$ -fraction of all cells.*

560 ■ *(No case): No valid solution can fill more than a $(\frac{1}{2} + \epsilon)$ -fraction of all cells.*

561 *Moreover, the above still holds if all slots have distinct lengths (and hence reusing words*
 562 *is trivially impossible).*

563 **Proof.** Fix an $\epsilon > 0$. We will later define an appropriately chosen value $\epsilon' \in (0, \epsilon)$ whose
 564 value only depends on ϵ . We present a reduction from a UNIQUE LABEL COVER instance,
 565 as described in Theorem 14. In particular, suppose we have an instance $G = (V, E)$, with
 566 $|V| = n$, alphabet $[R]$, such that (under UGC) it is NP-hard to distinguish if there exists a
 567 set V' of size $(1 - \epsilon')n$ that satisfies all its induced constraints, or if all sets V' of size $\epsilon'n$
 568 induce at least one violated constraint for any assignment. Throughout this proof we assume
 569 that n is sufficiently large (otherwise the initial instance is easy). In particular, let $n > \frac{20}{\epsilon}$.

570 We construct an instance of the crossword puzzle that fits in an $N \times N$ square, where
 571 $N = 4n + n^2$. We number the rows $1, \dots, N$ from top to bottom and the columns $1, \dots, N$
 572 from left to right. The instance contains n horizontal and n vertical slots. For $i \in [n]$, the
 573 i -th horizontal slot is placed in row $2i$, starting at column 1, and has length $2n + n^2 + i$.
 574 For $j \in [n]$, the j -th vertical slot is placed in column $2j$, starts at row 1 and has length
 575 $3n + n^2 + j$. Observe that all horizontal slots intersect all vertical slots and in particular, for
 576 all $i, j \in [n]$ the cell in row $2i$, column $2j$ belongs to the i -th horizontal slot and the j -th
 577 vertical slot. Furthermore, each slot has a distinct length, as the longest horizontal slot has
 578 length $3n + n^2$ while the shortest vertical slot has length $3n + n^2 + 1$.

579 We define the alphabet as $\Sigma_\epsilon = [R] \cup \{*\}$. Before we define our dictionary, let us give some
 580 intuition. Let $V = \{v_1, \dots, v_n\}$. The idea is that a variable $v_i \in V$ of the original instance
 581 will be represented by both the i -th horizontal slot and the i -th vertical slot. In particular,

48:14 Filling Crosswords is Very Hard

582 we will define, for each $\alpha \in [R]$ a pair of words that we can place in these slots to represent
 583 the fact that v_i is assigned with the value α . We will then ensure that if we place words
 584 on both the i -th horizontal slot and the j -th horizontal slot, where $(v_i, v_j) \in E$, then the
 585 assignment that can be extracted by reading these words will satisfy the constraint $\pi_{(v_i, v_j)}$.
 586 The extra letter $*$ represents an indifferent assignment (which we need if $(v_i, v_j) \notin E$).

587 Armed with this intuition, let us define our dictionary.

- 588 ■ For each $i \in [n]$, for each $\alpha \in [R]$ we define a word $d_{(i, \alpha)}$ of length $2n + n^2 + i$. The word
 589 $d_{(i, \alpha)}$ has the character $*$ everywhere except at position $2i$ and at positions $2j$ for $j \in [n]$
 590 and $(v_i, v_j) \in E$. In these positions the word $d_{(i, \alpha)}$ has the character α .
- 591 ■ For each $j \in [n]$, for each $\alpha \in [R]$ we define a word $d'_{(j, \alpha)}$ of length $3n + n^2 + j$. The word
 592 $d'_{(j, \alpha)}$ has the character $*$ everywhere except at position $2j$ and at positions $2i$ for $i \in [n]$
 593 and $(v_i, v_j) \in E$. In position $2j$ we have the character α . In position $2i$ with $(v_i, v_j) \in E$,
 594 we place the character $\beta \in [R]$ such that the constraint $\pi_{(v_i, v_j)}$ is satisfied by assigning
 595 β to v_i and α to v_j . (Note that β always exists and is unique, as the constraints are
 596 permutations on $[R]$, that is, for each value α of v_j there exists a unique value β of v_i
 597 that satisfies the constraint).

598 This completes the construction. Suppose now that $V = \{v_1, \dots, v_n\}$ and that we started
 599 from the Yes case of UNIQUE LABEL COVER, that is, there exists a set $V' \subseteq V$ such that
 600 $|V'| \geq (1 - \epsilon')n$ and all constraints induced by V' can be simultaneously satisfied. Fix an
 601 assignment $\sigma : V' \rightarrow [R]$ that satisfies all constraints induced by V' . For each $i \in [n]$ such
 602 that $v_i \in V'$ we place in the i -th horizontal slot (that is, in row $2i$) the word $d_{(i, \sigma(v_i))}$. For
 603 each $j \in [n]$ such that $v_j \in V'$ we place in the j -th vertical slot the word $d'_{(j, \sigma(v_j))}$. We leave
 604 all other slots empty. We claim that this solution is valid, that is, no shared cell is given
 605 different values from its horizontal and vertical slot. To see this, examine the cell in row $2i$
 606 and column $2j$. If both of the slots that contain it are filled, then $v_i, v_j \in V'$. If $(v_i, v_j) \notin E$
 607 and $i \neq j$, then the cell contains $*$ from both words. If $i = j$, then the cell contains $\sigma(v_i)$
 608 from both words. If $i \neq j$ and $(v_i, v_j) \in E$, then the cell contains $\sigma(v_i)$. This is consistent
 609 with the vertical word, as the constraint $\pi_{(v_i, v_j)}$ is assumed to be satisfied by σ . We now
 610 observe that this solution covers at least $2(1 - \epsilon')n^3$ cells, as we have placed $2(1 - \epsilon')n$ words,
 611 each of length at least $n^2 + 2n$, that do not pairwise intersect beyond their first $2n$ characters.

612 Suppose now we started our construction from a No instance of UNIQUE LABEL COVER.
 613 We claim that the optimal solution in the new instance cannot cover significantly more than
 614 half the cells. In particular, suppose a solution covers at least $(1 + \epsilon')n^3 + 10n^2$ cells. We
 615 claim that the solution must have placed at least $(1 + \epsilon')n$ words. Indeed, if we place at most
 616 $(1 + \epsilon')n$ words, as the longest word has length $n^2 + 4n$, the maximum number of cells we
 617 can cover is $(1 + \epsilon')n(n^2 + 4n) \leq (1 + \epsilon')n^3 + 4(1 + \epsilon')n^2 < (1 + \epsilon')n^3 + 10n^2$. Let x be the
 618 number of indices $i \in [n]$ such that the supposed solution has placed a word in both the i -th
 619 horizontal slot and the i -th vertical slot. We claim that $x \geq \epsilon'n$. Indeed, if $x < \epsilon'n$, then
 620 the total number of words we might have placed is at most $(n - x) + 2x < (1 + \epsilon')n$, which
 621 contradicts our previous observation that we placed at least $(1 + \epsilon')n$ words. Let $V' \subseteq V$
 622 be defined as the set of $v_i \in V$ such that the solution places words in the i -th horizontal
 623 and vertical slot. Then $|V'| \geq \epsilon'n$. We claim that it is possible to satisfy all the constraints
 624 induced by V' in the original instance, obtaining a contradiction. Indeed, we can extract an
 625 assignment for each $v_i \in V'$ by assigning to v_i value α if the i -th horizontal slot contains the
 626 word $d_{(i, \alpha)}$. Note that the i -th horizontal slot must contain such a word, as these words are
 627 the only ones that have an appropriate length. Observe that in this case the i -th vertical
 628 slot must also contain $d'_{(i, \alpha)}$. Now, for $v_i, v_j \in V'$, with $(v_i, v_j) \in E$ we see that $\pi_{(v_i, v_j)}$ is

629 satisfied by our assignment, otherwise we would have a conflict in the cell in position $(2i, 2j)$.
 630 Therefore, in the No case, it must be impossible to fill more than $(1 + \epsilon')n^3 + 10n^2$ cells.

631 The only thing that remains is to define ϵ' . Let C be the total number of cells in the
 632 instance. Recall that we proved that in the Yes case we cover at least $2(1 - \epsilon')n^3$ cells
 633 and in the No case at most $(1 + \epsilon')n^3 + 10n^2$ cells. So we need to define ϵ' such that
 634 $2(1 - \epsilon')n^3 \geq (1 - \epsilon)C$ and $(1 + \epsilon')n^3 + 10n^2 \leq (\frac{1}{2} + \epsilon)C$. To avoid tedious calculations, we
 635 observe that $2n^3 \leq C \leq 2n^3 + 8n^2$. Therefore, it suffices to have $2(1 - \epsilon')n^3 \geq 2(1 - \epsilon)(n^3 + 4n^2)$
 636 and $(1 + \epsilon')n^3 + 10n^2 \leq (1 + 2\epsilon)n^3$. The first inequality is equivalent to $(\epsilon - \epsilon')n \geq 4(1 - \epsilon)$
 637 and the second inequality is equivalent to $(2\epsilon - \epsilon')n \geq 10$. Since we have assumed that
 638 $n \geq 20/\epsilon$, it is sufficient to set $\epsilon' = \epsilon/2$. ◀

639 6 Conclusion

640 We studied the parameterized complexity of some crossword puzzles under several different
 641 parameters and we gave some positive results followed by proofs which show that our
 642 algorithms are essentially optimal. Based on our results the most natural questions that
 643 arise are: What is the complexity of CP-DEC when the grid graph is a matching and the
 644 alphabet has size 2? Can Theorem 12 be strengthened by starting from ETH instead of
 645 randomized ETH? Can we beat the $1/2$ approximation ratio of CP-OPT if we restrict our
 646 instances? Can Theorem 14 be strengthened by dropping the UGC? Furthermore, it would
 647 be interesting to investigate if there exist non trivial instances of the problem that can be
 648 solved in polynomial time.

649 Finally, we could consider a variation of the crossword puzzle problems where each word
 650 can be used a given number of times. This would be an intermediate case between word
 651 reuse and no word reuse.

652 References

- 653 1 Anbulagan and Adi Botea. Crossword puzzles as a constraint problem. In *Principles and Prac-*
 654 *tice of Constraint Programming, 14th International Conference, CP 2008, Sydney, Australia,*
 655 *September 14-18, 2008. Proceedings*, pages 550–554, 2008.
- 656 2 Edward K. Crossman and Sharyn M. Crossman. The crossword puzzle as a teaching tool.
 657 *Teaching of Psychology*, 10(2):98–99, 1983.
- 658 3 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin
 659 Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.
- 660 4 Jakob Engel, Markus Holzer, Oliver Ruepp, and Frank Sehnke. On computer integrated
 661 rationalized crossword puzzle manufacturing. In *Fun with Algorithms - 6th International*
 662 *Conference, FUN 2012, Venice, Italy, June 4-6, 2012. Proceedings*, pages 131–141, 2012.
- 663 5 Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory*
 664 *of NP-Completeness*. W. H. Freeman, 1979.
- 665 6 Matthew L. Ginsberg, Michael Frank, Michael P. Halpin, and Mark C. Torrance. Search
 666 lessons learned from crossword puzzles. In *Proceedings of the 8th National Conference on*
 667 *Artificial Intelligence. Boston, Massachusetts, USA, July 29 - August 3, 1990, 2 Volumes*,
 668 pages 210–215, 1990.
- 669 7 Heather Hulett, Todd G. Will, and Gerhard J. Woeginger. Multigraph realizations of degree
 670 sequences: Maximization is easy, minimization is hard. *Oper. Res. Lett.*, 36(5):594–596, 2008.
- 671 8 Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly
 672 exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001.
- 673 9 Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within 2-epsilon.
 674 *J. Comput. Syst. Sci.*, 74(3):335–349, 2008.

48:16 Filling Crosswords is Very Hard

- 675 **10** Michael Lampis, Valia Mitsou, and Karolina Soltys. Scrabble is PSPACE-complete. *J. Inf.*
676 *Process.*, 23(3):284–292, 2015.
- 677 **11** Michael L. Littman, Greg A. Keim, and Noam M. Shazeer. Solving crosswords with PROVERB.
678 In *Proceedings of the Sixteenth National Conference on Artificial Intelligence and Eleventh*
679 *Conference on Innovative Applications of Artificial Intelligence, July 18-22, 1999, Orlando,*
680 *Florida, USA*, pages 914–915, 1999.
- 681 **12** Télé Magazine. Publications Grand Public.
- 682 **13** Gary Meehan and Peter Gray. Constructing crossword grids: Use of heuristics vs constraints.
683 In *In: Proceedings of Expert Systems 97: Research and Development in Expert Systems XIV,*
684 *SGES*, pages 159–174, 1997.
- 685 **14** Jagan A. Pillai, Charles B. Hall, Dennis W. Dickson, Herman Buschke, Richard B. Lipton,
686 and Joe Verghese. Association of crossword puzzle participation with memory decline in
687 persons who develop dementia. *Journal of the International Neuropsychological Society*,
688 17(6):1006–1013, 2011.
- 689 **15** Leonardo Rigutini, Michelangelo Diligenti, Marco Maggini, and Marco Gori. Automatic
690 generation of crossword puzzles. *Int. J. Artif. Intell. Tools*, 21(3), 2012.
- 691 **16** Christopher D. Rosin. Nested rollout policy adaptation for Monte Carlo tree search. In
692 *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence,*
693 *Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 649–654, 2011.