

FPT Algorithms for Diverse Collections of Hitting Sets*

Julien Baste¹, Lars Jaffke², Tomáš Masařík^{3,4},
Geevarghese Philip⁵, and Günter Rote⁶

¹Institute of Optimization and Operations Research, Ulm University, Germany

`julien.baste@uni-ulm.de`

²University of Bergen, Norway

`lars.jaffke@uib.no`

³Charles University, Prague, Czech Republic

⁴University of Warsaw, Poland

`masarik@kam.mff.cuni.cz`

⁵Chennai Mathematical Institute, Chennai, India and UMI ReLaX

`gphilip@cmi.ac.in`

⁶Freie Universität Berlin

`rote@inf.fu-berlin.de`

Abstract

In this work, we study the d -HITTING SET and FEEDBACK VERTEX SET problems through the paradigm of finding diverse collections of r solutions of size at most k each, which has recently been introduced to the field of parameterized complexity [Baste et al., 2019]. This paradigm is aimed at addressing the loss of important *side information* which typically occurs during the abstraction process which models real-world problems as computational problems. We use two measures for the diversity of such a collection: the sum of all pairwise Hamming distances, and the minimum pairwise Hamming distance. We show that both problems are FPT in $k + r$ for both diversity measures. A key ingredient in our algorithms is a (problem independent) network flow formulation that, given a set of ‘base’ solutions, computes a maximally diverse collection of solutions. We believe that this could be of independent interest.

1 Introduction

The typical approach in modeling a real-world problem as a computational problem has, broadly speaking, two steps: (i) abstracting the problem into a mathematical formulation which captures the crux of the real-world problem, and (ii) asking for a best solution to the mathematical problem.

*Tomáš Masařík received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme Grant Agreement 714704, and from Charles University student grant SVV-2017-260452. Lars Jaffke is supported by the Bergen Research Foundation (BFS). Geevarghese Philip’s work was supported by BFS (Bergens Forsknings Stiftelse) ”Putting Algorithms Into Practice” Grant Number 810564 and NFR (Norwegian Research Foundation) grant number 274526d ”Parameterized Complexity for Practical Computing”.

38 Consider the following scenario. Dr. \mathcal{O} organizes a panel discussion, and has a shortlist of
 39 candidates to invite. From that shortlist, Dr. \mathcal{O} wants to invite as many candidates as possible,
 40 such that each of them will bring an individual contribution to the panel. Given two candidates
 41 A and B , it may not be beneficial to invite both A and B , for various reasons: their areas of
 42 expertise or opinions may be too similar for both to make a distinguishable contribution, or it
 43 may be preferable not to invite more than one person from each institution. It may even be the
 44 case that A and B do not see eye-to-eye on some issues which could come up at the discussion,
 45 and Dr. \mathcal{O} wishes to avoid a confrontation.

46 A natural mathematical model to resolve Dr. \mathcal{O} 's dilemma is as an instance of the VERTEX
 47 COVER problem: each candidate on the shortlist corresponds to a vertex, and for each pair of
 48 candidates A and B , we add the edge between A and B if it is *not* beneficial to invite both of
 49 them. Removing a smallest *vertex cover* in the resulting graph results in a largest possible set of
 50 candidates such that each of them may be expected to individually contribute to the appeal of
 51 the event.

52 Formally, a *vertex cover* of an undirected graph G is any subset $S \subseteq V(G)$ of the vertex set of
 53 G such that every edge in G has at least one end-point in S . The VERTEX COVER problem asks
 54 for a vertex cover of the smallest size:

VERTEX COVER

55 **Input:** Graph G .
Solution: A vertex cover S of G of the smallest size.

56 While the above model does provide Dr. \mathcal{O} with a set of candidates to invite that is *valid* in
 57 the sense that each invited candidate can be expected to make a unique contribution to the panel,
 58 a vast amount of *side information* about the candidates is lost in the modeling process. This side
 59 information could have helped Dr. \mathcal{O} to get more out of the panel discussion. For instance, Dr. \mathcal{O}
 60 may have preferred to invite more well-known or established people over ‘newcomers’, if they
 61 wanted the panel to be highly visible and prestigious; or they may have preferred to have more
 62 ‘newcomers’ in the panel, if they wanted the panel to have more outreach. Other preferences that
 63 Dr. \mathcal{O} may have had include: to have people from many different cultural backgrounds, to have
 64 equal representation of genders, or preferential representation for affirmative action; to have a
 65 variety in the levels of seniority among the attendants, possibly skewed in one way or the other.
 66 Other factors, such as the total carbon footprint caused by the participants’ travels, may also be
 67 of interest to Dr. \mathcal{O} . This list could go on and on.

68 Now, it is possible to plug in some of these factors into the mathematical model, for instance
 69 by including weights or labels. Thus a vertex weight could indicate ‘how well-established’ a
 70 candidate is. However, the complexity of the model grows fast with each additional criterion. The
 71 classic field of multicriteria optimization [46, 39, 29, 36] addresses the issue of bundling multiple
 72 factors into the objective function, but it is seldom possible to arrive at a balance in the various
 73 criteria in a way which captures more than a small fraction of all the relevant side information.
 74 Moreover, several side criteria may be conflicting or incomparable (or both); consider in Dr. \mathcal{O} 's
 75 case ‘maximizing the number of different cultural backgrounds’ vs. ‘minimizing total carbon
 76 footprint.’

77 While Dr. \mathcal{O} 's story is admittedly a made-up one, the VERTEX COVER problem is in fact used
 78 to model *conflict resolution* in far more realistic settings. In each case there is a *conflict graph* G
 79 whose vertices correspond to entities between which one wishes to avoid a conflict of some kind.
 80 There is an edge between two vertices in G if and only if they could be in conflict, and finding
 81 and deleting a smallest vertex cover of G yields a largest conflict-free subset of entities. We describe
 82 three examples to illustrate the versatility of this model. In each case it is intuitively clear, just

83 like in Dr. \mathcal{O} 's problem, that formulating the problem as VERTEX COVER results in a lot of
84 significant side information being thrown away, and that while finding a smallest vertex cover
85 in the conflict graph will give a *valid* solution, it may not really help in finding a *best* solution,
86 *or even a reasonably good solution*. We list some side information that is lost in the modeling
87 process; the reader should find it easy to come up with any amount of other side information
88 that would be of interest, in each case.

89 **Air traffic control.** Conflict graphs are used in the design of decision support tools for aiding
90 Air Traffic Controllers (ATCs) in preventing untoward incidents involving aircraft [44, 45, 25].
91 Each node in the graph G in this instance is an aircraft, and there is an edge between two
92 nodes if the corresponding aircraft are at risk of interfering with each other. A vertex cover
93 of G corresponds to a set of aircraft which can be issued *resolution commands* which ask
94 them to change course, such that afterwards there is no risk of interference.

95 In a situation involving a large number of aircraft it is unlikely that *every* choice of ten
96 aircraft to redirect is *equally* desirable. For instance, in general it is likely that (i) it is
97 better to ask smaller aircraft to change course in preference to larger craft, and (ii) it is
98 better to ask aircraft which are cruising to change course, in preference to those which are
99 taking off or landing.

100 **Wireless spectrum allocation.** Conflict graphs are a standard tool in figuring out how to
101 distribute wireless frequency spectrum among a large set of wireless devices so that no
102 two devices whose usage could potentially interfere with each other are allotted the same
103 frequencies [22, 24]. Each node in G is a user, and there is an edge between two nodes if
104 (i) the users request the same frequency, and (ii) their usage of the same frequency has the
105 potential to cause interference. A vertex cover of G corresponds to a set of users whose
106 requests can be denied, such that afterwards there is no risk of interference.

107 When there is large collection of devices vying for spectrum it is unlikely that *every* choice
108 of ten devices to deny the spectrum is *equally* desirable. For instance, it is likely that
109 denying the spectrum to a remote-controlled toy car on the ground is preferable to denying
110 the spectrum to a drone in flight.

111 **Managing inconsistencies in database integration.** A database constructed by integrating
112 data from different data sources may end up being inconsistent (that is, violating specified
113 integrity constraints) even if the constituent databases are individually consistent. Handling
114 these inconsistencies is a major challenge in database integration, and conflict graphs are
115 central to various approaches for restoring consistency [37, 26, 12, 3, 13]. Each node in G is
116 a database item, and there is an edge between two nodes if the two items together form an
117 inconsistency. A vertex cover of G corresponds to a set of database items in whose *absence*
118 the database achieves consistency.

119 In a database of large size it is unlikely that all data are created equal; some database items
120 are likely to be of better relevance or usefulness than others, and so it is unlikely that *every*
121 choice of ten items to delete is *equally* desirable.

122 Getting back to our first example, it seems difficult to help Dr. \mathcal{O} with their decision by
123 employing the 'traditional' way of modeling computational problems, where one looks for one
124 best solution. If on the other hand, Dr. \mathcal{O} was presented with a *small set of good solutions* that
125 in some sense are *far apart*, then they might hand-pick the list of candidates that they consider
126 the best choice for the panel and make a more informed decision. Moreover, several forms of
127 side-information may *only become apparent once Dr. \mathcal{O} is presented some concrete alternatives*,
128 and are more likely to be retrieved from alternatives that look very different. That is, a bunch of

129 good quality, dissimilar solutions may end up capturing a lot of the “lost” side information. And
130 this applies to each of the other three examples as well. In each case, finding one best solution
131 could be of little utility in solving the original problem, whereas finding a *small set of solutions*,
132 *each of good quality, which are not too similar to one another* may offer much more help.

133 To summarize, real-world problems typically have complicated side constraints, and the
134 optimality criterion may not be clear. Therefore, the abstraction to a mathematical formulation
135 is almost always a simplification, omitting important side information. There are at least two
136 obstacles to simply adapting the model by incorporating these secondary criteria into the objective
137 function or taking into account the side constraints: (i) they make the model complicated and
138 unmanageable, and (ii) more importantly, these criteria and constraints are often not precisely
139 formulated, potentially even unknown a priori. There may even be no sharp distinction between
140 optimality criteria and constraints (the so-called “soft constraints”).

141 One way of dealing with this issue is to present a small number r of *good* solutions and let the
142 *user* choose between them, based on all the experience and additional information that the user
143 has and that is ignored in the mathematical model. Such an approach is useful even when the
144 objective can be formulated precisely, but is difficult to optimize: After generating r solutions,
145 each of which is *good enough* according to some quality criterion, they can be compared and
146 screened in a second phase, evaluating their exact objective function or checking additional side
147 constraints. In this context, it makes little sense to generate solutions that are very similar to
148 each other and differ only in a few features. It is desirable to present a *diverse* variety of solutions.

149 It should be clear that the issue is scarcely specific to VERTEX COVER. Essentially *any*
150 computational problem motivated by practical applications likely has the same issue: the modeling
151 process throws out so much relevant side information that any algorithm which finds just one
152 optimal solution to an input instance may not be of much use in solving the original problem in
153 practice. One scenario where the traditional approach to modeling computational problems fails
154 completely is when computational problems may combined with a human sense of aesthetics or
155 intuition to solve a task, or even to stimulate inspiration. Some early relevant work is on the
156 problem of designing a tool which helps an architect in creating a floor plan which satisfies a
157 specified set of constraints. In general, the number of feasible floor plans—those which satisfy
158 constraints imposed by the plot on which the building has to be erected, various regulations
159 which the building should adhere to, and so on—would be too many for the architect to look at
160 each of them one by one. Further, many of these plans would be very similar to one another, so
161 that it would be pointless for the architect to look at more than one of these for inspiration. As
162 an alternative to optimization for such problems, Galle proposed a “Branch & Sample” algorithm
163 for generating a “limited, representative sample of solutions, uniformly scattered over the entire
164 solution space” [21].

165 **The Diverse X Paradigm.** Mike Fellows has proposed *the Diverse X Paradigm* as a solution
166 for these issues and others [19]. In this paradigm “ X ” is a placeholder for an optimization problem,
167 and we study the complexity—specifically, the fixed-parameter tractability—of the problem of
168 finding a few different good quality solutions for X . Contrast this with the traditional approach
169 of looking for just one good quality solution. Let X denote an optimization problem where one
170 looks for a minimum-size subset of some set; VERTEX COVER is an example of such a problem.
171 The generic form of X is then:

	X
Input:	An instance I of X .
Solution:	A solution S of I of the smallest size.

172

173 Here the form that a “solution S of I ” takes is dictated by the problem X ; compare this with
 174 the earlier definition of VERTEX COVER.

175 The *diverse* variant of problem X , as proposed by Fellows, has the form

DIVERSE X	
Input:	An instance I of X , and positive integers k, r, t .
Parameter:	(k, r)
Solution:	A set \mathcal{S} of r solutions of I , each of size at most k , such that a <i>diversity measure</i> of \mathcal{S} is at least t .

176 Note that one can construct diverse variants of other kinds of problems as well, following this
 177 model: it doesn’t have to be a minimization problem, nor does the solution have to be a subset of
 178 some kind. Indeed, the example about floor plans described above has neither of these properties.
 179 What is relevant is that one should have (i) some notion of “good quality” solutions (for X , this
 180 equates to a small size) and (ii) some notion of a set of solutions being “diverse”.

181 **Diversity measures.** The concept of diversity appears also in other fields, and there are many
 182 different ways to measure the diversity of a collection. For example, in ecology, the diversity of a
 183 set of species (“biodiversity”) is a topic that has become increasingly important in recent times,
 184 see for example Solow and Polasky [41].

185 Another possible viewpoint, in the context of multicriteria optimization, is to require that the
 186 sample of solutions should try to represent the *whole solution space*. This concept can be quantified
 187 for example by the geometric *volume* of the represented space [28, 10], or by the *discrepancy* [34].
 188 See [43, Section 3] for an overview of diversity measures in multicriteria optimization.

189 In this paper, we follow the simple possibility of looking for a collection of good solutions
 190 that have large *distances* from each other, in a sense that will be made precise below (1)–(2).
 191 Direction (2), i.e., taking the pairwise sum of all Hamming distances, has been taken by many
 192 practical papers in the area of genetic algorithms, see e.g. [20, 33]. This now classical approach
 193 can be traced as far back as 1992 [32]. In [47], it has been boldly stated that this measure (and
 194 its variations) is one of the most broadly used measures in describing population diversity within
 195 genetic algorithms. One of its advantages is that it can be computed very easily and efficiently
 196 unlike many other measures, e.g., some geometry or discrepancy based measures.

197 1.1 Our problems and results.

198 In this work we focus on diverse versions of two minimization problems, d -HITTING SET and
 199 FEEDBACK VERTEX SET, whose solutions are subsets of a finite set. d -HITTING SET is in fact a
 200 *class* of such problems which includes VERTEX COVER, as we describe below. We will consider
 201 two natural diversity measures for these problems: the minimum Hamming distance between any
 202 two solutions, and the sum of pairwise Hamming distances of all the solutions.

203 The *Hamming distance* between two sets S and S' , or *the size of their symmetric difference*, is

$$204 \quad d_H(S, S') := |(S \setminus S') \cup (S' \setminus S)|.$$

205 We use

$$206 \quad \text{div}_{\min}(S_1, \dots, S_r) := \min_{1 \leq i < j \leq r} d_H(S_i, S_j) \quad (1)$$

207 to denote the minimum Hamming distance between any pair of sets in a collection of finite sets,
 208 and

$$209 \quad \text{div}_{\text{total}}(S_1, \dots, S_r) := \sum_{1 \leq i < j \leq r} d_H(S_i, S_j) \quad (2)$$

246 **Related Work.** The parameterized complexity of finding a diverse collection of good-quality
247 solutions to algorithmic problems seems to be largely unexplored. To the best of our knowledge,
248 the only existing work in this area consists of: (i) a privately circulated manuscript by Fellows [19]
249 which introduces the Diverse X Paradigm and makes a forceful case for its relevance, and (ii) a
250 manuscript by Baste et al. [5] which applies the Diverse X Paradigm to *vertex-problems* with
251 the *treewidth* of the input graph as an extra parameter. In this context a *vertex-problem* is any
252 problem in which the input contains a graph G and the solution is some subset of the vertex set
253 of G which satisfies some problem-specific properties. Both VERTEX COVER and FEEDBACK
254 VERTEX SET are vertex-problems in this sense, as are many other graph problems. The *treewidth*
255 of a graph is, informally put, a measure of how tree-like the graph is. See, e.g., [14, Chapter 7]
256 for an introduction of the use of the treewidth of a graph as a parameter in designing FPT
257 algorithms. The work by Baste et al. [5] shows how to convert essentially any treewidth-based
258 dynamic programming algorithm for solving a vertex-problem, into an algorithm for computing a
259 diverse set of r solutions for the problem, with the diversity measure being the sum $\text{div}_{\text{total}}$ of
260 Hamming distances of the solutions. This latter algorithm is FPT in the combined parameter
261 (r, w) where w is the treewidth of the input graph. As a special case, they obtain a running time
262 of $\mathcal{O}((2^{k+2}(k+1))^r kr^2 n)$ for DIVERSE VERTEX COVER. Further, they show that the r -DIVERSE
263 versions (i.e., where the diversity measure is $\text{div}_{\text{total}}$) of a handful of problems have polynomial
264 kernels. In particular, they show that DIVERSE VERTEX COVER has a kernel with $\mathcal{O}(k(k+r))$
265 vertices, and that DIVERSE d -HITTING SET has a kernel with a universe size of $\mathcal{O}(k^d + kr)$.

266 **Organization of the rest of the paper.** In Section 2 we list some definitions which we use
267 in the rest of the paper. In Section 3 we describe a generic framework which can be used for
268 computing solution families of maximum diversity for a variety of problems whose solutions form
269 subsets of some finite set. We prove Theorem 1 in Section 3.3 and Theorem 2 in Section 4. In
270 Section 5 we discuss some potential pitfalls in using $\text{div}_{\text{total}}$ as a measure of diversity. In Section 6
271 we prove Theorem 3 and Theorem 4. We conclude in Section 7.

272 2 Preliminaries

273 Given two integers p and q , we denote by $[p, q]$ the set of all integers r such that $p \leq r \leq q$ holds.
274 Given a graph G , we denote by $V(G)$ (resp. $E(G)$) the set of *vertices* (resp. *edges*) of G . For
275 a subset $S \subset V(G)$ we use $G[S]$ to denote the subgraph of G induced by S , and $G \setminus S$ for the
276 graph $G[V(G) \setminus S]$. A set $S \subseteq V(G)$ is a vertex cover (resp. a feedback vertex set) if $G \setminus S$ has
277 no edge (resp. no cycle). Given a graph G and a vertex v such that v has exactly two neighbors,
278 say w and w' , *contracting* v consists in removing the edges $\{v, w\}$ and $\{v, w'\}$, removing v and
279 adding the edge $\{w, w'\}$. Given a graph G and a vertex $v \in V(G)$, we denote by $\delta_G(v)$ the *degree*
280 of v in G . For two vertices u, v in a connected graph G we use $\text{dist}_T(u, v)$ to denote the *distance*
281 between u and v in G , which is the length of a shortest path in G between u and v .

282 A *deepest leaf* in a tree T is a vertex $v \in V(T)$ such that there exists a root $r \in V(T)$
283 satisfying $\text{dist}_T(r, v) = \max_{u \in V(T)} \text{dist}_T(r, u)$. A *deepest leaf* in a forest F is a deepest leaf in
284 some connected component of F . A deepest leaf v has the property that there is another leaf in
285 the tree at distance at most 2 from v unless v is an isolated vertex or v 's neighbor has degree 2.

286 The objective function $\text{div}_{\text{total}}$ in (2) has an alternative representation in terms of frequencies
287 of occurrence [5]: If y_v is the number of sets of $\{S_1, \dots, S_r\}$ in which v appears, then

$$288 \text{div}_{\text{total}}(S_1, \dots, S_r) = \sum_{v \in U} y_v (r - y_v). \quad (3)$$

3 A Framework for Maximally Diverse Solutions

290
291
292
293

In this section we describe a framework for computing solution families of maximum diversity for a variety of hitting set problems. This framework requires that the solutions form a family of subsets of a ground set U which is upward closed: Any superset $T \supseteq S$ of a solution S is also a solution.

294
295
296
297
298

The approach is as follows: In a first phase, we enumerate the class \mathcal{S} of all *minimal solutions* of size at most k . (A larger class \mathcal{S} is also fine as long as it is guaranteed to contain all minimal solutions of size at most k .) Then we form all r -tuples $(S_1, \dots, S_r) \in \mathcal{S}^k$. For each such family (S_1, \dots, S_r) , we try to *augment* it to a family (T_1, \dots, T_r) under the constraints $T_i \supseteq S_i$ and $|T_i| \leq k$, for each $i \in [1, r]$, in such a way that $\text{div}_{\text{total}}(T_1, \dots, T_r)$ is maximized.

299
300
301
302
303

For this augmentation problem, we propose a network flow model that computes an optimal augmentation in polynomial time, see Section 3.1. This has to be repeated for each family, $O(|\mathcal{S}|^r)$ times. The first step, the generation of \mathcal{S} , is problem-specific. Section 3.3 shows how to solve it for d -HITTING SET. In Section 4, we will adapt our approach to deal with FEEDBACK VERTEX SET.

304

3.1 Optimal Augmentation

305
306
307

Given a universe U and a set \mathcal{S} of subsets of U , the problem $\text{diverse}_{r,k}(\mathcal{S})$ consists in finding an r -tuple (S_1, \dots, S_r) that maximizes $\text{div}_{\text{total}}(S_1, \dots, S_r)$, over all r -tuples (S_1, \dots, S_r) such that for each $i \in [1, r]$, $|S_i| \leq k$ and there exists $S \in \mathcal{S}$ such that $S \subseteq S_i \subseteq U$.

308
309

Theorem 6. *Let U be a finite universe, r and k be two integers and \mathcal{S} be a set of s subsets of U . $\text{diverse}_{r,k}(\mathcal{S})$ can be solved in time $r^2 s^r \cdot |U|^{O(1)}$.*

310
311
312
313
314

Proof. The algorithm that proves Theorem 6 starts by enumerating all r -tuples $(S_1, S_2, \dots, S_r) \in \mathcal{S}^r$ of elements from \mathcal{S} . For each of these s^r r -tuples we try to augment each S_i , using elements of U , in such a way that the diversity d of the resulting tuple (T_1, \dots, T_r) is maximized and such that for each $i \in [1, r]$, $S_i \subseteq T_i \subseteq U$ and $|T_i| \leq k$. It is clear that this algorithm will find the solution to $\text{diverse}_{r,k}(\mathcal{S})$.

315
316
317

We show how to model this problem as a maximum-cost network flow problem with piecewise linear concave costs. This problem can be solved in polynomial time. (See for example [42] for basic notions about network flows.)

318
319
320

Without loss of generality, let $U = \{1, 2, \dots, n\}$. We use a variable $0 \leq x_{ij} \leq 1$ to decide whether element j of U should belong to set T_i . In an optimal flow, these values are integral. Some of these variables are already fixed because T_i must contain S_i :

321

$$x_{ij} = 1 \text{ for } j \in S_i \quad (4)$$

322

The size of T_i must not exceed k :

323

$$\sum_{j=1}^n x_{ij} \leq k, \text{ for } i = 1, \dots, r \quad (5)$$

324

Finally, we can express the number y_j of sets T_i in which an element j occurs:

325

$$y_j = \sum_{i=1}^r x_{ij}, \text{ for } j = 1, \dots, n \quad (6)$$




Figure 1: The network. The middle layer between the vertices T_i and V_j is a complete bipartite graph, but only a few selected arcs are shown. A potential augmenting path is highlighted.

326 These variables y_j are the variables in terms of which the objective function (3) is expressed:

327
$$\text{maximize } \sum_{j=1}^n y_j(r - y_j) \tag{7}$$

328 These constraints can be modeled by a network as shown in Figure 1. There are nodes T_i
 329 representing the sets T_i and a node V_j for each element $j \in U$. In addition, there is a source s
 330 and a sink t . The arcs emanating from s have capacity k . Together with the flow conservation
 331 equations at the nodes T_i , this models the constraints (5). Flow conservation at the nodes V_j
 332 gives rise to the flow variables y_j in the arcs leading to t according to (6). The arcs with fixed
 333 flow (4) could be eliminated from the network, but for ease of notation, we leave them in the
 334 model. The only arcs that carry a cost are the arcs leading to t , and the costs are given by the
 335 concave function (7).

336 There is now a one-to-one correspondence between integral flows from s to t in the network
 337 and solutions (T_1, \dots, T_r) , and the cost of the flow is equal to the diversity (2) or (3). We are
 338 thus looking for a flow of maximum cost. The *value* of the flow (to total flow out of s) can be
 339 arbitrary. (It is equal to the sum of the sizes of the sets T_i .)

340 The concave arc costs (7) on the arcs leading to t can be modeled in a standard way by
 341 multiple arcs. Denote the concave cost function by $f_y := y(r - y)$, for $y = 0, 1, \dots, r$. Then each
 342 arc (V_i, t) in the last layer is replaced by r parallel arcs of capacity 1 with costs $f_1 - f_0, f_2 - f_1,$
 343 $\dots, f_r - f_{r-1}$. This sequence of values $f_y - f_{y-1} = r - 2y + 1$ is decreasing, starting out with
 344 positive values and ending with negative values. If the total flow along such a bundle is y , the
 345 maximum-cost way to distribute this flow is to fill the first y arcs to capacity, for a total cost of
 346 $(f_1 - f_0) + (f_2 - f_1) + \dots + (f_y - f_{y-1}) = f_y - f_0 = f_y$, as desired.

347 An easy way to compute a maximum-cost flow is the longest augmenting path method.
 348 (Commonly it is presented as the *shortest* augmenting path method for the *minimum*-cost flow.)
 349 This holds for the classical flow model where the cost on each arc is a linear function of the flow.
 350 An augmenting path is a path in the residual network with respect to the current flow, and the
 351 cost coefficient of an arc in such a path must be taken with opposite sign if it is traversed in the
 352 direction opposite to the original graph.

353 **Proposition 1** (The shortest augmenting path algorithm, cf. [42, Theorem 8.12]). *Suppose a*
 354 *maximum-cost flow among all flows of value v from s to t is given. Let P be a maximum-cost*

355 *augmenting path from s to t . If we augment the flow along this path, this results in a new flow, of*
 356 *some value v' . Then the new flow is a maximum-cost flow among all flows of value v' from s to t .*

357 Let us apply this algorithm to our network. We initialize the constrained flow variables x_{ij}
 358 according to (4) to 1 and all other variables x_{ij} to 0. This corresponds to the original solution
 359 (S_1, S_2, \dots, S_r) , and it is clearly the optimal flow of value $\sum_{i=1}^r |S_i|$ because it is the only feasible
 360 flow of this value.

361 We can now start to find augmenting paths. Our graph is bipartite, and augmenting paths
 362 have a very simple structure: They start in s , alternate back and forth between the T -nodes and
 363 the V -nodes, and finally make a step to t . Moreover, in our network, all costs are zero except in
 364 the last layer, and an augmenting path contains precisely one arc from this layer. Therefore, *the*
 365 *cost of an augmenting path is simply the cost of the final arc.*

366 The flow variables in the final layer are never decreased. The resulting algorithm has therefore
 367 a simple greedy-like structure. Starting from the initial flow, we first try to saturate as many of
 368 the arcs of cost $f_1 - f_0$ as possible. Next, we try to saturate as many of the arcs of cost $f_2 - f_1$
 369 as possible, and so on. Once the incremental cost $f_{y+1} - f_y$ becomes negative, we stop.

370 Trying to find an augmenting path whose last arc is one of the arcs of cost $f_{y+1} - f_y$, for fixed
 371 y , is a reachability problem in the residual graph, and it can be solved by graph search in $O(nr)$
 372 time because the network has $O(nr)$ vertices. Every augmentation increases the flow value by 1
 373 unit. Thus, there are at most kr augmentations, for a total runtime of $O(kr^2n)$. \square

374 3.2 Faster Augmentation

375 We can obtain faster algorithms by using more advanced network algorithms from the literature.
 376 We will derive one such algorithm here. The best choice depends on the relation between n , k , and
 377 r . We will apply the following result about *b-matchings*, which are generalizations of matchings:
 378 Each node v has a given *supply* $b(v)$, specifying that v should be incident to at most v edges.

379 **Proposition 2** ([1]). *A maximum-weight b-matching in a bipartite graph with $N_1 + N_2$ nodes on*
 380 *the two sides of the bipartition and M edges that have integer weights between 0 and W can be*
 381 *found in time $O(N_1 M \log(2 + \frac{N_1^2}{M} \log(N_1 W)))$.*

382 We will describe below how the network flow problem from above can be converted into
 383 a *b-matching* problem with $N_1 = r + 1$ plus $N_2 = n$ nodes and $M = 2rn$ edges of weight at
 384 most $W = 2r$. Plugging these values into Proposition 2 gives a running time of $O(r^2n \log(2 +$
 385 $\frac{r}{n} \log(r^2))) = O(r^2n \max\{1, \log \frac{r \log r}{n}\})$ for finding an optimal augmentation. This improves over
 386 the run time $O(r^2nk)$ from the previous section unless r is extremely large (at least 2^k).

387 From the network of Figure 1, we keep the two layers of nodes T_i and V_j . Each vertex T_i gets
 388 a supply of $b(T_i) := k$, and each vertex V_j gets a supply of $b(V_j) := r$. To mimic the piecewise
 389 linear costs on the arcs (V_j, t) in the original network, we introduce r parallel *slack edges*
 390 from a new source vertex s' to each vertex V_i . The costs are as follows. Let $g_1 > g_2 > \dots > g_r$ with
 391 $g_y = f_y - f_{y-1}$ denote the costs in the last layer of the original network, and let $\hat{g} := r$. Since
 392 $g_1 = r - 1$, this is larger than all costs. Then every edge (T_i, V_j) from the original network gets a
 393 weight of \hat{g} , and the r new slack edges entering each V_j get positive weights $\hat{g} - g_1, \hat{g} - g_2, \dots, \hat{g} - g_r$.
 394 We set the supply of the extra source node to $b(s') := rn$, which imposes no constraint on the
 395 number of incident edges.

396 Now suppose that we have a solution for the original network in which the total flow into vertex
 397 V_j is y . In the corresponding *b-matching*, we can then use $b(V_j) - y = r - y$ of the slack edges
 398 incident to V_j . The $r - y$ maximum-weight slack edges have weights $\hat{g} - g_r, \hat{g} - g_{r-1}, \dots, \hat{g} - g_{y+1}$.

399 The total weight of the edges incident to V_j is therefore

$$400 \quad r\hat{g} - g_r - g_{r-1} - \cdots - g_{y+1} = r\hat{g} + (g_1 + g_2 + \cdots + g_y),$$

401 using the equation $g_1 + g_2 + \cdots + g_r = f_r - f_0 = 0$. Thus, up to an addition of the constant $nr\hat{g}$,
402 the maximum weight of a b -matching agrees with the maximum cost of a flow in the original
403 network.

404 3.3 Diverse Hitting Set

405 In this section we show how to use the optimal augmentation technique developed in Section 3
406 to solve DIVERSE d -HITTING SET. For this we use the following folklore lemma about minimal
407 hitting sets.

408 **Lemma 7.** *Let (U, \mathcal{F}) be an instance of d -HITTING SET, and let k be an integer. There are at
409 most d^k inclusion-minimal hitting sets of \mathcal{F} of size at most k , and they can all be enumerated in
410 time $d^k|U|^2$.*

411 Combining Lemma 7 and Theorem 6, we obtain the following result.

412 **Theorem 1.** DIVERSE d -HITTING SET can be solved in time $r^2d^{kr} \cdot |U|^{O(1)}$.

413 *Proof.* Using Lemma 7, we can construct the set \mathcal{S} of all inclusion-minimal hitting sets of \mathcal{F} , each
414 of size at most k . Note that the size of \mathcal{S} is bounded by d^k . As every superset of an element of \mathcal{S}
415 is also a hitting set, the theorem follows directly from Theorem 6. \square

416 4 Diverse Feedback Vertex Set

417 A *feedback vertex set* (FVS) (also called a *cycle cutset*) of a graph G is any subset $S \subseteq V(G)$ of
418 vertices of G such that every cycle in G contains at least one vertex from S . The graph $G - S$
419 obtained by deleting S from G is thus an acyclic graph. Finding an FVS of small size is an
420 NP-hard problem with a number of applications in Artificial Intelligence, many of which stem
421 from the fact that many hard problems become easy to solve in acyclic graphs.

422 The Propositional Model Counting (or #SAT) problem asks for the number of satisfying assign-
423 ments for a given CNF formula, and has a number of applications, for instance in planning [35, 17]
424 and in probabilistic inference problems such as Bayesian reasoning [4, 11, 23, 15, 30, 40, 2]. A
425 popular approach to solving #SAT consists of first finding a small FVS S of the CNF formula. As-
426 signing values to all the variables in S results in an acyclic instance of CNF. The algorithm assigns
427 all possible sets of values to the variables in S , computes the number of satisfying assignments of
428 the resulting acyclic instances, and returns the sum of these counts [16]. Other applications of
429 finding a small FVS include faster sampling for Bayesian networks, solving constraint satisfaction
430 problems, credulous and skeptical acceptance problems in abstract argumentation, and learning
431 and inference in graphical models [8, 6, 9, 7, 18, 31].

432 In this section, we focus on the DIVERSE FEEDBACK VERTEX SET problem and prove the
433 following theorem.

434 **Theorem 2.** DIVERSE FEEDBACK VERTEX SET can be solved in time $2^{7kr} \cdot n^{O(1)}$.

435 In order to solve r -DIVERSE k -FEEDBACK VERTEX SET, one natural way would be to generate
436 every feedback vertex set of size at most k and then check which set of k solutions provide the
437 required sum of Hamming distances. Unfortunately, the number of feedback vertex set is not

438 FPT parameterized by k . Indeed, one can consider a graph containing k cycle of size $\frac{n}{k}$, leading
 439 to $\left(\frac{n}{k}\right)^k$ different feedback vertex sets of size k .

440 We avoid this problem by generating all such small feedback vertex sets up to some equivalence
 441 of degree two vertices. We obtain an exact and efficient description of all feedback vertex sets
 442 of size at most k , which is formally captured by Lemma 8. A *class of solutions* of a graph G ,
 443 is a pair (S, ℓ) such that $S \subseteq V(G)$ and $\ell : S \rightarrow 2^{V(G)}$ is a function such that for each $u \in S$,
 444 $u \in \ell(u)$, and for each $u, v \in S$, $u \neq v$, $\ell(u) \cap \ell(v) = \emptyset$. Given a class of solutions (S, ℓ) , we
 445 define $\text{sol}(S, \ell) = \{S' : |S'| = |S| \text{ and } \forall v \in S, |S' \cap \ell(v)| = 1\}$. A *class of FVS solutions* is a
 446 class of solutions (S, ℓ) such that each $S' \in \text{sol}(S, \ell)$ is a feedback vertex set of G . Moreover, if
 447 $S' \in \text{sol}(S, \ell)$ and $S' \subseteq S'' \subseteq V(G)$, we say that S'' is *described* by (S, ℓ) . Note that S'' is also a
 448 feedback vertex set. In a class of FVS solutions (S, ℓ) , the meaning of the function ℓ is that, for
 449 each cycle C in G , there exists $v \in S$ such that each element of $\ell(v)$ hits C . This allows us to
 450 group related solutions into only one set $\text{sol}(S, \ell)$.

451 **Lemma 8.** *Let G be a n -vertex graph. There exists a set \mathcal{S} of classes of FVS solutions of G of*
 452 *size at most 2^{7k} such that each feedback vertex set of size at most k is described by an element of*
 453 *\mathcal{S} . Moreover, \mathcal{S} can be constructed in time $2^{7k} \cdot n^{O(1)}$.*

454 *Proof.* Let G be a n -vertex graph. We start by generating a feedback vertex set $F \subseteq V$ of size
 455 at most k . The current best deterministic algorithm for this by Kociumaka and Pilipczuk [27]
 456 finds such a set in time $3.62^k \cdot n^{O(1)}$. In the following, we use the ideas used for the iterative
 457 compression approach [38].

458 For each subset $F' \subseteq F$, we initiate a branching process by setting $A := F'$, $B := F - F'$, and
 459 $G' := G$. Observe that initially, as $B \subseteq F$ and $|F| \leq k$, the graph $G[B]$ has at most k components.
 460 In the branching process, we will add more vertices to A and B , and we will remove vertices and
 461 edges from G' , but we will maintain the property that $A \subseteq V(G')$ and $B \subseteq V(G')$. The set C
 462 will always denote the vertex set $V(G') \setminus (A \cup B)$. Note that $G'[C]$ is initially a forest; we ensure
 463 that it always remains a forest.

464 We also initialize a function $\ell : V(G) \rightarrow 2^{V(G)}$ by setting $\ell(v) = \{v\}$ for each $v \in V(G)$. This
 465 function will keep information about vertices that are deleted from G . While searching for a
 466 feedback vertex set, we consider only feedback vertex sets that contain all vertices of A but no
 467 vertex of B . Vertices in C are still undecided. The function ℓ will maintain the invariant that for
 468 each $v \in V(G')$, $\ell(v) \cap V(G') = \{v\}$, and for each $v \in C$, all vertices of $\ell(v)$ intersect exactly the
 469 same cycles in $G \setminus A$. Moreover, for each $v \in A$, the value $\ell(v)$ is fixed and will not be modified
 470 anymore in the branching process. During the branching process, we will progressively increase
 471 the size of A , B , and the sets $\ell(v)$, $v \in V(G)$.

472 By *reducing* (G', A, B, ℓ) we mean that we apply the following rules exhaustively.

- 473 - If there is a $v \in C$ such that $\delta_{G'[B \cup C]}(v) \leq 1$, we delete v from G' .
 474 - If there is an edge $\{u, v\} \in E(G'[C])$ such that $\delta_{G'[B \cup C]}(u) = \delta_{G'[B \cup C]}(v) = 2$, we contract
 475 u in G' and set $\ell(v) := \ell(v) \cup \ell(u)$.

476 These are classical preprocessing rules for the FEEDBACK VERTEX SET problem, see for in-
 477 stance [14, Section 9.1]. Indeed, vertices of degree one cannot appear in a cycle, and consecutive
 478 vertices of degree 2 hit exactly the same cycles. After this preprocessing, there are no adjacent
 479 degree-two vertices and no degree-one vertices in C . (Degrees are measured in $G'[B \cup C]$.)

480 We start to describe the branching procedure. We work on the tuple (G', A, B, ℓ) . After
 481 each step, the value $|A| - \text{cc}(B)$ will increase, where $\text{cc}(B)$ denotes the number of connected
 482 components of $G'[B]$.

483 At each step of the branching we do the following. If $|A| > k$ or if $G'[B]$ contains a cycle, we
484 immediately stop this branch as there is no solution to be found in it. If A is a feedback vertex
485 set of size at most k , then $(A, \ell|_A)$ is a class of FVS solutions, we add it to \mathcal{S} and stop working
486 on this branch. Otherwise, we reduce (G', A, B, ℓ) . We pick a deepest leaf v in $G'[C]$ and apply
487 one of the two following cases, depending of the vertex v .

488 - **Case 1:** The vertex v has at least two neighbors in B (in the graph G').
489 If there is a path in B between two neighbors of v , then we have to put v in A , as otherwise
490 this path together with v will induce a cycle. If there is no such path, we branch on both
491 possibilities, inserting v either into A or into B .

492 - **Case 2:** The vertex v has at most one neighbor in B .
493 Since v is a leaf in $G'[C]$, it has at most one neighbor also in C . On the other hand, we
494 know that v has degree at least 2 in $G'[B \cup C]$. Thus, v has exactly one neighbor in B and
495 one neighbor in C , for a degree of 2 in $G'[B \cup C]$. Let p be the neighbor in C . Again, as
496 we have reduced (G', A, B, ℓ) , the degree of p in $G'[B \cup C]$ is at least 3. So either it has a
497 neighbor in B , or, as v is a deepest leaf, it has another child, say w , that is also a leaf in
498 $G'[C]$, and w has therefore a neighbor in B . We branch on the at most $2^3 = 8$ possibilities
499 to allocate v , p , and w if considered, between A and B , taking care not to produce a cycle
500 in B .

501 In both cases, either we put at least one vertex in A , and so $|A|$ increases by one, or all
502 considered vertices are added to B . In the latter case, the considered vertices are connected, at least
503 two of them have a neighbor in B , and no cycles were created; therefore, the number of components
504 in B drops by one. Thus $|A| - \text{cc}(B)$ increases by at least one. As $-k \leq |A| - \text{cc}(B) \leq k$, there
505 can be at most $2k$ branching steps.

506 Since we branch at most $2k$ times and at each branch we have at most 2^3 possibilities, the
507 branching tree has at most 2^{6k} leaves. So, for each of the at most 2^k subsets F' of F , we add at
508 most 2^{6k} elements to \mathcal{S} .

509 It is clear that we have obtained all solutions of FVS and they are described by the classes of
510 FVS solutions in \mathcal{S} , which is of size 2^{7k} . \square

511 *Proof of Theorem 2.* We generate all 2^{7kr} r -tuples of the classes of solutions given by Lemma 8,
512 with repetition allowed.

513 We now consider each r -tuple $((S_1, \ell_1), (S_2, \ell_2), \dots, (S_r, \ell_r)) \in \mathcal{S}^r$ and try to pick an ap-
514 propriate solution T_i from each class of solutions (S_i, ℓ_i) , $i \in [1, k]$, in such a way that the
515 diversity of the resulting tuple of feedback vertex sets (T_1, \dots, T_r) is maximized. The network
516 of Section 3.1 must be adapted to model the constraints resulting from solution classes. Let
517 (S, ℓ) be a solution class, with $|S| = b$. For our construction, we just need to know the family
518 $\{\ell(v) \mid v \in S\} = \{L_1, L_2, \dots, L_b\}$ of disjoint nonempty vertex sets. The solutions that are
519 described by this class are all sets that can be obtained by picking at least one vertex from each
520 set L_q . Figure 2 shows the necessary adaptations for *one* solution $T = T_i$. In addition to a single
521 node T that is either directly or indirectly connected to all nodes V_1, \dots, V_n , like in Figure 1, we
522 have additional nodes representing the sets L_q . For each vertex j that appears in one of the sets
523 L_q , there is an additional node U_j in an intermediate layer of the network. The flow from s to L_q
524 is forced to be equal to 1, and this ensures that at least one element of the set L_q is chosen in the
525 solution. Here it is important that the sets L_q are disjoint.

526 A similar structure must be built for each set T_1, \dots, T_r , and all these structures share the
527 vertices s and V_1, \dots, V_n . The rightmost layer of the network is the same as in Figure 1.




Figure 2: Part of the modified network for a solution T which is specified by $b = 3$ sets $L_1 = \{1, 2\}$, $L_2 = \{3\}$, and $L_3 = \{4, 5, 6\}$.

528 The initial flow is not so straightforward as in Section 3.1 but is still easy to find. We simply
 529 saturate the arc from s to each of the nodes L_q in turn by a shortest augmenting path. Such a
 530 path can be found by a simple reachability search in the residual network, in $O(rn)$ time. The
 531 total running time $O(kr^2n)$ from Section 3.1 remains unchanged. \square

532 5 Modeling Aspects: Discussion of the Objective Function

533 In Sections 3 and 4, we have used the sum of the Hamming distances, $\text{div}_{\text{total}}$, as the measure of
 534 diversity. While this metric is of natural interest, it appears that in some specific cases, it may
 535 not be a useful choice. We present a simple example where the *most diverse* solution according
 536 to $\text{div}_{\text{total}}$ is not what one might expect.

537 Let r be an even number. We consider the path with $2r - 2$ vertices, and we are looking for r
 538 vertex covers of size at most $r - 1$, of maximum diversity. Figure 3 shows an example with $r = 6$.
 539 The smallest size of a vertex cover is indeed $r - 1$, and there are r different solutions. One would
 540 hope that the “maximally diverse” selection of r solutions would pick all these different solutions.
 541 But no, the selection that maximizes $\text{div}_{\text{total}}$ consists of $r/2$ copies of just *two* solutions, the “odd”
 542 vertices and the “even” vertices (the first and last solution in Figure 3).

543 This can be seen as follows. If the selected set contains in total n_i copies of the first i solutions
 544 in the order of Figure 3, then the objective can be written as

$$545 \quad 2n_1(r - n_1) + 2n_2(r - n_2) + \cdots + 2n_{r-1}(r - n_{r-1}).$$

546 Here, each term $2n_i(r - n_i)$ accounts for two consecutive vertices $2i - 1, 2i$ of the path in the
 547 formulation (3). The unique way of maximizing each term individually is to set $n_i = r/2$ for
 548 all i . This corresponds to the selection of $r/2$ copies of the first solution and $r/2$ copies of the
 549 last solution, as claimed.




Figure 3: The $r = 6$ different vertex covers of size $r - 1 = 5$ in a path with $2(r - 1) = 10$ vertices

550 In a different setting, namely the distribution of r points inside a square, an analogous
551 phenomenon has been observed [43, Figure 1]: Maximizing the sum of pairwise Euclidean
552 distances places all points at the corners of the square. In fact it is easy to see that, in this
553 geometric setting, any locally optimal solution must place all points on the boundary of the feasible
554 region. By contrast, for our combinatorial problem, we don't know whether this pathological
555 behavior is typical or rare in instances that are not specially constructed. Further research is
556 needed. A notion of diversity which is more robust in this respect is the *smallest* difference
557 between two solutions, which we consider in Section 6.

558 6 Maximizing the Smallest Hamming distance

559 The undesired behavior highlighted in Section 5 is the fact that the collection that maximizes the
560 sum of the Hamming distances uses several copies of the same set. In this section we explore
561 how to handle this unexpected behavior by changing the distance to the minimal Hamming
562 distance between two sets of the collection. This modification naturally removes the possibility
563 of selecting the same solution twice. We show how to solve MIN-DIVERSE d -HITTING SET and
564 r -MIN-DIVERSE k -FEEDBACK VERTEX SET for this metric.

565 **Theorem 3.** MIN-DIVERSE d -HITTING SET can be solved in time

566 - $2^{kr^2} \cdot (kr)^{O(1)}$ if $|U| < kr$ and

567 - $d^{kr} \cdot |U|^{O(1)}$ otherwise.

568 *Proof.* Let $(U, \mathcal{F}, k, r, t)$ be an instance of MIN-DIVERSE d -HITTING SET where $|U| = n$. If
569 $n < kr$, we solve the problem by complete enumeration: There are trivially at most 2^n hitting
570 sets of size at most k . We form all r -tuples (T_1, \dots, T_r) of them and select the one that maximizes
571 $\text{div}_{\min}(T_1, \dots, T_r)$. The running time is at most $O((2^n)^r r^2 n) = O(2^{kr^2} kr^3)$.

572 We now assume that $n \geq kr$. We use the same strategy as in Section 3: We generate all r -tuples
573 (S_1, \dots, S_r) of *minimal solutions* and try to augment each one to a r -tuple (T_1, \dots, T_r) such that
574 for each $i \in [1, r]$, $|T_i| \leq k$ and $S_i \subseteq T_i \subseteq V(G)$ hold. The difference is that we try to maximize
575 $\text{div}_{\min}(T_1, \dots, T_r)$ instead of $\text{div}_{\text{total}}(T_1, \dots, T_r)$ in the augmentation. Given that we have a large
576 supply of $n \geq kr$ elements in U , this is easy. To each set S_i we add $k - |S_i|$ new elements, taking
577 care that we pick different elements for each S_i with are not in any of the other sets S_j . The

578 Hamming distance between two resulting sets is then $d_H(T_i, T_j) = d_H(S_i, S_j) + (k - |S_i|) + (k - |S_j|)$,
579 and it is clear that this is the largest possible distance that two sets $T'_i \supseteq S_i$ and $T'_j \supseteq S_j$ with
580 $|T'_i|, |T'_j| \leq k$ can achieve. Thus, since our choice of augmentation individually maximizes each
581 pairwise Hamming distance, it also maximizes the smallest Hamming distance. This procedure
582 can be carried out in $O(kr + n) = O(n)$ time. In addition, we need $O(kr^2) = O(n^2)$ time to
583 compute the smallest distance.

584 Using Lemma 7, we construct the set \mathcal{S} of all minimal solutions of the d -HITTING SET instance
585 (U, \mathcal{F}) , each of size at most k . We then go through every r -tuple $(S_1, \dots, S_r) \in \mathcal{S}^r$ and augment
586 it optimally, as just described. The running time is $d^{kr} \cdot O(n^2)$. \square

587 **Theorem 4.** MIN-DIVERSE FEEDBACK VERTEX SET *can be solved in time* $2^{kr \cdot \max(r, 7 + \log_2(kr))} \cdot$
588 $(nr)^{O(1)}$.

589 *Proof.* Let G be a n -vertex graph. If $n < kr$, we again solve the problem by complete enumeration:
590 There are trivially at most 2^n feedback vertex sets of size at most k . We form all r -tuples
591 (T_1, \dots, T_r) of them and select the one that maximizes $\text{div}_{\min}(T_1, \dots, T_r)$. The running time is
592 at most $O((2^n)^r r^2 n) = O(2^{kr^2} r^2 n)$.

593 We assume now that $n \geq kr$. As in Section 4, we construct a set \mathcal{S} of at most 2^{7k} classes
594 of FVS solutions of G , using Lemma 8. Then we go through all $(2^{7k})^r$ r -tuples of classes $S =$
595 $((S_1, \ell_1), \dots, (S_r, \ell_r)) \in \mathcal{S}^r$. For each such r -tuple, we look for the r -tuple (T_1, \dots, T_r) of feedback
596 vertex sets such that each T_i is described by (S_i, ℓ_i) , and the objective value $\text{div}_{\min}(T_1, \dots, T_r)$
597 is maximized. So far, the procedure is completely analogous to the algorithm of Theorem 2 in
598 Section 4 for maximizing $\text{div}_{\text{total}}(T_1, \dots, T_r)$.

599 Now, in going from a class (S_i, ℓ_i) to T_i , we have to select a vertex from every set $\ell_i(v)$, for
600 $v \in S_i$, and we may add an arbitrary number of additional vertices, up to size k . We make this
601 selection as follows: Whenever $|\ell_i(v)| < kr$, we simply try all possibilities of choosing an element
602 of $\ell_i(v)$ and putting it into T_i . If $|\ell_i(v)| \geq kr$, we defer the choice for later. In this way, we have
603 created at most $(kr)^{kr}$ “partial” feedback vertex sets (T_1^0, \dots, T_r^0)

604 For each such (T_1^0, \dots, T_r^0) , we now add the remaining elements. In each list $\ell_i(v)$ which
605 has been deferred, we greedily pick an element that is distinct from all other chosen elements.
606 This is always possible since the list is large enough. Finally, we fill up the sets to size k , again
607 choosing fresh elements each time. Each such choice is an optimal choice, because it increases the
608 Hamming distance between the concerned set T_i and *every* other set T_j by 1, which is the best
609 that one can hope for. As we proceed to this operation for each $S \in \mathcal{S}^r$, where $|\mathcal{S}| \leq 2^{7k}$, and
610 that for each such S , we create at most $(kr)^{kr}$ r -tuples, we obtain an algorithm running in time
611 $2^{7kr} \cdot (kr)^{kr} \cdot n^{O(1)}$. The theorem follows. \square

612 7 Conclusions and Open Problems

613 In this work, we have considered the paradigm of finding small diverse collections of reasonably
614 good solutions to combinatorial problems, which has recently been introduced to the field of
615 fixed-parameter tractability theory [5].

616 We have shown that finding diverse collections of d -hitting sets and feedback vertex sets
617 can be done in FPT time. While these problems can be classified as FPT via the kernels and
618 a treewidth-based meta-theorem proved in [5], the methods proposed here are of independent
619 interest. We introduced a method of generating a maximally diverse set of solutions from a set
620 that either contains all minimal solutions of bounded size (d -HITTING SET) or from a collection
621 of structures that in some way *describes* all solutions of bounded size (FEEDBACK VERTEX SET).
622 In both cases, the maximally diverse collection of solutions is obtained via a network flow model,

623 which does not rely on any specific properties of the studied problems. It would be interesting
624 to see if this strategy can be applied to give FPT-algorithms for diverse problems that are not
625 covered by the meta-theorem or the kernels presented in [5].

626 While the problems in [5] as well as the ones in Sections 3 and 4, seek to maximize the *sum* of
627 all pairwise Hamming distances, we also studied the variant that asks to maximize the *minimum*
628 Hamming distance, taken over each pair of solutions. This was motivated by an example where
629 the former measure does not perform as intended (Section 5). We showed that also under this
630 objective, the diverse variants of d -HITTING SET and FEEDBACK VERTEX SET are FPT. It
631 would be interesting to see whether this objective also allows for a (possibly treewidth-based)
632 meta-theorem.

633 In [5], the authors ask whether there is a problem that is in FPT parameterized by solution
634 size whose r -diverse variant becomes $W[1]$ -hard upon adding r as another component of the
635 parameter. We restate this question here.

636 **Question 9** ([5]). *Is there a problem Π with solution size k , such that Π is FPT parameterized*
637 *by k , while DIVERSE Π , asking for r solutions, is $W[1]$ -hard parameterized by $k + r$?*

638 To the best of our knowledge, this problem is still wide open. We believe that the div_{\min}
639 measure is more promising to obtain such a result rather than the $\text{div}_{\text{total}}$ measure. A possible
640 way to tackle both measures at once might be a parameterized (and strengthened) analogue
641 of the following approach that is well-studied in classical complexity. Yato and Seta propose a
642 framework [48] to prove NP-completeness of finding a *second* solution to an NP-complete problem.
643 In other words, there are some problems where given one solution it is still NP-hard to determine
644 whether the problem has a different solution.

645 From a different perspective, one might want to identify problems where obtaining one solution
646 is polynomial-time solvable, but finding a diverse collection of r solutions becomes NP-hard. The
647 targeted running time should be FPT parameterized by r (and maybe t , the diversity target) only.
648 We conjecture that this is most probably NP- or $W[-]$ hard in general. However, we believe it is
649 interesting to search for well-known problems where it is not the case.

650 **Acknowledgements.** The second, third and fourth authors would like to thank Mike Fellows
651 for introducing them to the notion of diverse FPT algorithms and sharing the manuscript “The
652 Diverse X Paradigm” [19].

653 References

- 654 [1] R. K. Ahuja, J. B. Orlin, C. Stein, and R. E. Tarjan. Improved algorithms for bipartite
655 network flow. *SIAM Journal on Computing*, 23:906–933, 1994.
- 656 [2] Udi Apsel and Ronen I. Brafman. Lifted MEU by weighted model counting. In *Proceedings*
657 *of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, pages 1861–1867. AAAI
658 Press, 2012.
- 659 [3] Marcelo Arenas, Leopoldo Bertossi, Jan Chomicki, Xin He, Vijay Raghavan, and Jeremy
660 Spinrad. Scalar aggregation in inconsistent databases. *Theoretical Computer Science*,
661 296(3):405–434, 2003.
- 662 [4] Fahiem Bacchus, Shannon Dalmao, and Toniann Pitassi. Algorithms and complexity results
663 for #SAT and Bayesian inference. In *44th Annual IEEE Symposium on Foundations of*
664 *Computer Science, 2003. Proceedings.*, pages 340–351. IEEE, 2003.

- 665 [5] Julien Baste, Michael R. Fellows, Lars Jaffke, Tomáš Masařík, Mateus de Oliveira Oliveira,
666 Geevarghese Philip, and Frances A. Rosamond. Diversity in combinatorial optimization,
667 2019. [arXiv:1903.07410](https://arxiv.org/abs/1903.07410).
- 668 [6] Bozhena Bidyuk and Rina Dechter. An anytime scheme for bounding posterior beliefs.
669 In *Proceedings of the 21st National Conference on Artificial Intelligence, Volume 2*, pages
670 1095–1100. AAAI Press, 2006.
- 671 [7] Bozhena Bidyuk and Rina Dechter. Cutset sampling with likelihood weighting. In *Proceedings*
672 *of the Twenty-Second Conference on Uncertainty in Artificial Intelligence*, pages 39–46. AUAI
673 Press, 2006.
- 674 [8] Bozhena Bidyuk and Rina Dechter. Cutset sampling for Bayesian networks. *Journal of*
675 *Artificial Intelligence Research*, 28:1–48, 2007.
- 676 [9] Bozhena Petrovna Bidyuk. *Exploiting graph cutsets for sampling-based approximations in*
677 *Bayesian networks*. PhD thesis, University of California, Irvine, 2006.
- 678 [10] Karl Bringmann, Sergio Cabello, and Michael T. M. Emmerich. Maximum volume subset se-
679 lection for anchored boxes. In Boris Aronov and Matthew J. Katz, editors, *33rd International*
680 *Symposium on Computational Geometry (SoCG 2017)*, volume 77 of *Leibniz International*
681 *Proceedings in Informatics (LIPIcs)*, pages 22:1–22:15, Dagstuhl, Germany, 2017. Schloss
682 Dagstuhl–Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SoCG.2017.22.
- 683 [11] Mark Chavira and Adnan Darwiche. On probabilistic inference by weighted model counting.
684 *Artificial Intelligence*, 172(6–7):772–799, 2008.
- 685 [12] Jan Chomicki and Jerzy Marcinkowski. Minimal-change integrity maintenance using tuple
686 deletions. *Information and Computation*, 197(1–2):90–121, 2005.
- 687 [13] Jan Chomicki and Jerzy Marcinkowski. On the computational complexity of minimal-change
688 integrity maintenance in relational databases. In *Inconsistency Tolerance*, pages 119–150.
689 Springer, 2005.
- 690 [14] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin
691 Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.
- 692 [15] Adnan Darwiche. The quest for efficient probabilistic inference. *Invited Talk, IJCAI-05*,
693 2005.
- 694 [16] Rina Dechter and David Cohen. *Constraint Processing*. Morgan Kaufmann, 2003.
- 695 [17] Carmel Domshlak and Jörg Hoffmann. Fast probabilistic planning through weighted model
696 counting. In *Proceedings of the Sixteenth International Conference on Automated Planning*
697 *and Scheduling, ICAPS 2006*, pages 243–252. AAAI Press, 2006.
- 698 [18] Wolfgang Dvořák, Sebastian Ordyniak, and Stefan Szeider. Augmenting tractable fragments
699 of abstract argumentation. *Artificial Intelligence*, 186:157–173, 2012.
- 700 [19] Michael Ralph Fellows. The diverse X paradigm. Manuscript, November 2018.
- 701 [20] Thomas Gabor, Lenz Belzner, Thomy Phan, and Kyrill Schmid. Preparing for the unexpected:
702 Diversity improves planning resilience in evolutionary algorithms. In *2018 IEEE International*
703 *Conference on Autonomic Computing, ICAC 2018, Trento, Italy, September 3-7, 2018*, pages
704 131–140, 2018. doi:10.1109/ICAC.2018.00023.

- 705 [21] Per Galle. Branch & sample: A simple strategy for constraint satisfaction. *BIT Numerical*
706 *Mathematics*, 29(3):395–408, 1989.
- 707 [22] Sorabh Gandhi, Chiranjeeb Buragohain, Lili Cao, Haitao Zheng, and Subhash Suri. A general
708 framework for wireless spectrum auctions. In *2007 2nd IEEE International Symposium on*
709 *New Frontiers in Dynamic Spectrum Access Networks*, pages 22–33. IEEE, 2007.
- 710 [23] Carla P Gomes, Ashish Sabharwal, and Bart Selman. Model counting: a new strategy
711 for obtaining good bounds. In *Proceedings of the 21st National Conference on Artificial*
712 *Intelligence-Volume 1*, pages 54–61. AAAI Press, 2006.
- 713 [24] Martin Hoefer, Thomas Kesselheim, and Berthold Vöcking. Approximation algorithms for
714 secondary spectrum auctions. *ACM Transactions on Internet Technology (TOIT)*, 14(2–3):16,
715 2014.
- 716 [25] M. Idan, G. Iosilevskii, and L. Ben-Yishay. Efficient air traffic conflict resolution by minimizing
717 the number of affected aircraft. *International Journal of Adaptive Control and Signal*
718 *Processing*, 24(10):867–881, 2010.
- 719 [26] Ekaterini Ioannou and Slawek Staworko. Management of inconsistencies in data integration.
720 In *Dagstuhl Follow-Ups*, volume 5. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2013.
- 721 [27] Tomasz Kociumaka and Marcin Pilipczuk. Faster deterministic feedback vertex set. *Informa-*
722 *tion Processing Letters*, 114(10):556–560, 2014. doi:10.1016/j.ipl.2014.05.001.
- 723 [28] Tobias Kuhn, Carlos M. Fonseca, Luís Paquete, Stefan Ruzika, Miguel M. Duarte, and
724 José Rui Figueira. Hypervolume subset selection in two dimensions: Formulations and
725 algorithms. *Evolutionary Computation*, 24(3):411–425, 2016. doi:10.1162/EVCO_a_00157.
- 726 [29] G. Leitmann and A. Marzollo, editors. *Handbook of Multicriteria Analysis*. Springer, Berlin,
727 2010.
- 728 [30] Michael L. Littman, Stephen M. Majercik, and Toniann Pitassi. Stochastic Boolean satisfiability.
729 *Journal of Automated Reasoning*, 27(3):251–296, 2001. doi:10.1023/A:1017584715408.
- 730 [31] Ying Liu and Alan Willsky. Learning Gaussian graphical models with observed or latent
731 FVSS. In *Advances in Neural Information Processing Systems*, pages 1833–1841, 2013.
- 732 [32] Sushil J. Louis and Gregory J. E. Rawlins. Syntactic analysis of convergence in genetic algo-
733 rithms. In *Proceedings of the Second Workshop on Foundations of Genetic Algorithms. Vail,*
734 *Colorado, USA, July 26-29 1992.*, pages 141–151, 1992. doi:10.1016/b978-0-08-094832-4.
735 50015-5.
- 736 [33] Ronald W. Morrison and Kenneth A. De Jong. Measurement of population diversity.
737 In *Artificial Evolution, 5th International Conference, Evolution Artificielle, EA 2001, Le*
738 *Creusot, France, October 29-31, 2001, Selected Papers*, pages 31–41, 2001. doi:10.1007/
739 3-540-46033-0_3.
- 740 [34] Aneta Neumann, Wanru Gao, Carola Doerr, Frank Neumann, and Markus Wagner.
741 Discrepancy-based evolutionary diversity optimization. In *Proceedings of the Genetic and*
742 *Evolutionary Computation Conference, GECCO '18*, pages 991–998, New York, NY, USA,
743 2018. ACM. doi:10.1145/3205455.3205532.

- 744 [35] Héctor Palacios, Blai Bonet, Adnan Darwiche, and Héctor Geffner. Pruning conformant
745 plans by counting models on compiled d-DNNF representations. In *Proceedings of the*
746 *Fifteenth International Conference on Automated Planning and Scheduling, ICAPS 2005*,
747 pages 141–150. AAAI Press, 2005.
- 748 [36] Jason Papathanasiou and Nikolaos Ploskas. *Multiple Criteria Decision Aid: Methods,*
749 *Examples and Python Implementations*. Springer Nature, 2018.
- 750 [37] Enela Pema, Phokion G. Kolaitis, and Wang-Chiew Tan. On the tractability and intractability
751 of consistent conjunctive query answering. In *Proceedings of the 2011 Joint EDBT/ICDT*
752 *Ph. D. Workshop*, pages 38–44. ACM, 2011.
- 753 [38] Bruce Reed, Kaleigh Smith, and Adrian Vetta. Finding odd cycle transversals. *Operations*
754 *Research Letters*, 32(4):299–301, 2004.
- 755 [39] Bernard Roy. *Multicriteria Methodology for Decision Aiding*. Kluwer, Dordrecht, 1996.
- 756 [40] Tian Sang, Paul Bearne, and Henry Kautz. Performing Bayesian inference by weighted model
757 counting. In *Proceedings of the 20th National Conference on Artificial Intelligence-Volume 1*,
758 pages 475–481. AAAI Press, 2005.
- 759 [41] Andrew R. Solow and Stephen Polasky. Measuring biological diversity. *Environmental and*
760 *Ecological Statistics*, 1(2):95–103, Jun 1994. doi:10.1007/BF02426650.
- 761 [42] Robert Endre Tarjan. *Data Structures and Network Algorithms*. SIAM, Philadelphia, 1983.
- 762 [43] Tamara Ulrich, Johannes Bader, and Lothar Thiele. Defining and optimizing indicator-based
763 diversity measures in multiobjective search. In Robert Schaefer, Carlos Cotta, Joanna
764 Kolodziej, and Günter Rudolph, editors, *Parallel Problem Solving from Nature, PPSN XI*,
765 pages 707–717, Berlin, Heidelberg, 2010. Springer. doi:10.1007/978-3-642-15844-5_71.
- 766 [44] Adan Vela, John-Paul Clarke, Eric Feron, Nicolas Durand, and William Singhose. Determining
767 the value of information for minimizing controller taskload: a graph-based approach. In
768 *ATM Seminar 2011, 9th USA/Europe Seminar on ATM R&D*, 2011.
- 769 [45] Adan Ernesto Vela. *Understanding conflict-resolution taskload: implementing advisory*
770 *conflict-detection and resolution algorithms in an airspace*. PhD thesis, Georgia Institute of
771 Technology, 2011.
- 772 [46] Philippe Vincke. *Multicriteria Decision-aid*. Wiley, Chichester, 1992.
- 773 [47] Mark Wineberg and Franz Oppacher. The underlying similarity of diversity measures used
774 in evolutionary computation. In *Genetic and Evolutionary Computation - GECCO 2003,*
775 *Genetic and Evolutionary Computation Conference, Chicago, IL, USA, July 12-16, 2003.*
776 *Proceedings, Part II*, pages 1493–1504, 2003. doi:10.1007/3-540-45110-2_21.
- 777 [48] Takayuki Yato and Takahiro Seta. Complexity and completeness of finding another so-
778 lution and its application to puzzles. *IEICE transactions on fundamentals of electronics,*
779 *communications and computer sciences*, 86(5):1052–1060, 2003.