Dispersing Obnoxious Facilities on Graphs by Rounding Distances

Tim A. Hartmann
Department of Computer Science, RWTH Aachen University, Germany

Stefan Lendl
Department of Operations and Information Systems, University of Graz, Austria

Abstract

We continue the study of δ-dispersion, a continuous facility location problem on a graph where all edges have unit length and where the facilities may also be positioned in the interior of the edges. The goal is to position as many facilities as possible subject to the condition that every two facilities have distance at least δ from each other.

Our main technical contribution is an efficient procedure to "round-up" distance δ. It transforms a δ-dispersed set S into a $\delta^\star$-dispersed set $S^\star$ of same size where distance $\delta^\star$ is a potentially slightly larger rational $\frac{a}{b}$ with a numerator a upper bounded by the longest (not-induced) path in the input graph.

Based on this rounding procedure and connections to the distance-d independent set problem we derive a number of algorithmic results. When parameterized by treewidth, the problem is in XP. When parameterized by treedepth the problem is FPT and has a matching lower bound on its time complexity under ETH. Moreover, we can also settle the parameterized complexity with the solution size as parameter using our rounding technique: δ-Dispersion is FPT for every $\delta \leq 2$ and W[1]-hard for every $\delta > 2$.

Further, we show that δ-dispersion is NP-complete for every fixed irrational distance δ, which was left open in a previous work.

2012 ACM Subject Classification Theory of computation → Discrete optimization; Mathematics of computing → Graph algorithms; Theory of computation → Parameterized complexity and exact algorithms; Theory of computation → Problems, reductions and completeness

Keywords and phrases facility location, parameterized complexity, packing

Digital Object Identifier 10.4230/LIPIcs.MFCS.2022.55

Related Version *Full Version*: https://arxiv.org/abs/2206.11337

1 Introduction

We study the algorithmic behavior of a continuous dispersion problem. Consider an undirected graph G, whose edges have unit length. Let $P(G)$ be the continuum set of points on all the edges and vertices. For two points $p, q \in P(G)$, we denote by $d(p, q)$ the length of a shortest path containing p and q in the underlying metric space. A subset $S \subseteq P(G)$ is *δ-dispersed* for some positive real number δ, if every distinct points $p, q \in S$ have distance at least $d(p, q) \geq \delta$. Our goal is, for a given graph G and a positive real number δ, to compute a maximum cardinality subset $S \subseteq P(G)$ that is δ-dispersed. We denote by δ-*disp*(G) the maximum size of a δ-dispersed set of G. The decision problem Dispersion asks for a δ-dispersed set of size at least k, where additionally integer $k \geq 0$ is part of the input. When δ is fixed and not part of the input, we refer to the problem as δ-Dispersion.

47th International Symposium on Mathematical Foundations of Computer Science (MFCS 2022).
Editors: Stefan Szeider, Robert Ganian, and Alexandra Silva; Article No. 55; pp. 55:1–55:14
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1.1 Known and Related Results

The area of obnoxious facility location goes back to seminal articles of Goldman & Dearing [6] and Church & Garfinkel [3]. The area includes a wide variety of objectives and models. For example, purely geometric variants have been studied by Abravaya & Segal [1], Ben-Moshe, Katz & Segal [2], and Katz, Kedem & Segal [14]. Recently, van Ee studied the approximability of a generalized covering problem in a metric space that also involves dispersion constraints [19]. Another direction is a graph-theoretic model, where every edge of the given graph G is rectifiable and has some individual length. Tamir discusses the complexity and approximability of several optimization problems. For example, when G is a tree, then a δ-dispersed set can be computed in polynomial time [18]. Another task is to place a single obnoxious facility in a network while maximizing, for example, the smallest distance from the facility to certain clients, as studied by Segal [16].

In a previous work, the complexity of Dispersion was studied for every rational distance δ. When δ is a rational number with numerator 1 or 2, the problem is polynomial time solvable, while it is NP-complete for all other rational values of δ [7, 8]. The complexity when δ is irrational was left as an open problem.

A closely related facility location problem is δ-covering. The objective is to place as few locations as possible on $P(G)$ subject to the condition that any point in $P(G)$ is in distance at most δ to a placed location. This problem is polynomial time solvable whenever δ is a unit fraction, while it is NP-hard for all non unit fractions δ [10]. Furthermore, the parameterized complexity with the parameter solution size k is studied. δ-covering is fixed parameter tractable when $\delta < \frac{3}{2}$, while for $\delta \geq \frac{3}{2}$ the problem is W[2]-complete [10]. Tamir [17] showed that for δ-covering only certain distances δ are of interest. For every amount of points p the distance $\max\{\delta^\star : |\delta^\star\text{-cover}(G)| = p\}$ is of the form $\frac{L'}{2p'}$ where $p' \in \{1, \ldots, p\}$ and L' is roughly at most twice the length of a non-induced path in G.

1.2 Our Contribution

Our main technical contribution is an efficient and constructive rounding procedure. Given a δ-dispersed set S for some distance value $\delta > 0$, it transforms S into a $\delta^\star$-dispersed set $S^\star$ of equal size with a slightly larger well-behaving distance value $\delta^\star \geq \delta$. The new distance $\delta^\star$ is a rational $\frac{a}{b}$ with small numerator a. More precisely, the numerator is upper bounded by the length of the longest (not-induced) path L, hence upper bounded asymptotically by the number of vertices n of the input graph (see Section 5).

Our second technical contribution relates the optimal solution for distance δ and $\frac{\delta}{\delta+1}$ for $\delta \leq 3$. A δ-dispersed set translates to a $\frac{\delta}{\delta+1}$-dispersed set by placing one more point on every edge, and vice versa by removing one point (see Section 3).

Further we explore a connection of Dispersion and an independent set problem (see Section 4). The combination of that connection with our technical contributions yields several algorithmic results for Dispersion (see Section 6 and Section 7):

- Dispersion is NP-hard even for chordal graphs of diameter 4.
- Dispersion is FPT for the graph parameter treedepth $\mathsf{td}(G)$ with a run time matching a lower bound under ETH. We complement this result by showing that δ-Dispersion is W[1]-hard for the slightly more general graph parameter pathwidth $\mathsf{pw}(G)$, even for the combined parameter $\mathsf{pw}(G) + k$. Similarly, δ-Dispersion is W[1]-hard for the graph parameter $\mathsf{fvs}(G)$, the minimum size of a feedback vertex set.
- Dispersion is XP for the parameter treewidth $\mathsf{tw}(G)$, with a running time of $(2L)^{\mathsf{tw}(G)} n^{\mathcal{O}(1)}$, where n is the number of vertices and L is an upper bound on the length

of the longest path in G. We complement this result by the more general lower bound of $n^{o(\mathsf{tw}(G)+\sqrt{k})}$, assuming ETH. It implies the lower bound of $L^{o(\mathsf{tw}(G)+\sqrt{k})}$ since $L \leq n$. Note that a mere lower bound of $L^{o(\mathsf{tw}(G)+\sqrt{k})}$ would not exclude an $n^{o(\mathsf{tw}(G))}$-algorithm.

In addition, we completely resolve the complexity of δ-dispersion, by showing NP-hardness for irrational δ (see Section 8). We also study the parameterized complexity when parameterized by the solution size k. The problem is W[1]-hard when $\delta > 2$, and FPT otherwise. Thus, there is a sharp threshold at $\delta = 2$ where the complexity jumps from FPT to W[1]-hard (see Section 9).

We mark statements whose proof can be found in the full version of the paper (see [9]) with "($\star$)".

2 Preliminaries

We use the word *vertex* in the graph-theoretic sense, while we use the word *point* to denote the elements of the geometric structure $P(G)$. As an input for δ-dispersion, we consider graphs G that are undirected, connected, and without loops and isolated vertices.

For an edge $\{u, v\} \in E(G)$ and a real number $\lambda \in [0, 1]$, let $p(u, v, \lambda) \in P(G)$ be the point on edge $\{u, v\}$ that has distance λ from u. Note that $p(u, v, 0) = u$, $p(u, v, 1) = v$ and $p(u, v, \lambda) = p(v, u, 1 - \lambda)$. Further, we use $d(p, q)$ for the length of a shortest path between points $p, q \in P(G)$.

For a subset of vertices $V' \subseteq V(G)$ or a subset of edges $E' \subseteq E(G)$, we denote by $G[V']$ and $G[E']$ the subgraph induced by V' and E', respectively. The neighborhood of a vertex u is $N(u) := \{v \in V(G) \mid \{u, v\} \in E(G)\}$. We use n as the number of vertices of G, when G is clear from the context.

For a graph G and integer $c \geq 1$, let the *c-subdivision of G* be the graph G where every edge is replaced by a path of length c.

▶ **Lemma 1** ([8])**.** *Let G be a graph, let $c \geq 1$ be an integer, and let G' be the c-subdivision of G. Then δ-disp$(G) = (c\delta)$-disp(G').*

For integers a and b, we denote the rational number $\frac{a}{b}$ as *b-simple*. A set $S \subseteq P(G)$ is b-simple, if for every point $p(u, v, \lambda)$ in S the edge position λ is *b-simple*.

▶ **Lemma 2** ([8])**.** *Let $\delta = \frac{a}{b}$ with integers a and b, and let G be a graph. Then, there exists an optimal δ-dispersed set $S^\star$ that is $2b$-simple.*

For an introduction into parameterized algorithms, we refer to [4]. We study of the complexity of Dispersion with the natural parameter solution size k, as well as its dependency on structural measures on the input graph. Besides treewidth $\mathsf{tw}(G)$ and pathwidth $\mathsf{pw}(G)$, we also study the parameters "feedback vertex set size" $\mathsf{fvs}(G)$ and treedepth $\mathsf{td}(G)$.

A graph has a feedback vertex set $W \subseteq V(G)$ if G after removing W contains no cycle. The "feedback vertex set size" is the size of a smallest feedback vertex set of G.

The treedepth of a connected graph G can be defined as follows. If G is disconnected, it is the maximum treedepth of its components; If G consists of a single vertex, then $\mathsf{td}(G) = 1$; And else it is one plus the minimum over all $u \in V(G)$ of the treedepth of G without vertex u.

We provide lower bounds for the time-complexity assuming the Exponential Time Hypothesis (ETH): There is no $2^{o(N)}$-time algorithm for 3-SAT with N variables and $\mathcal{O}(N)$ clauses [11]. For more details on ETH, we refer to [4].

3 Translating δ-Dispersion

There is an intriguing relation of the optimal solution for distance δ and $\frac{\delta}{2\delta+1}$ for the similar problem δ-covering [10]. We may analogously expect that an optimal solution for δ-dispersion translates to an optimal solution for $\frac{\delta}{\delta+1}$-dispersion; i.e., that an optimal δ-dispersed set corresponds to an optimal $\frac{\delta}{\delta+1}$-dispersed set of the same size plus one extra point for every edge.

This is not true for $\delta = 3+\varepsilon$ for any $\varepsilon > 0$: Consider a triangle, where a $(3+\varepsilon)$-dispersed set S contains at most one point p. Since $\frac{\delta}{\delta+1} > \frac{3}{4}$, a $\frac{\delta}{\delta+1}$-dispersed set however contains at most $3 < |S| + 3$ points.

Causing trouble is a non-trivial closed walk containing p of length less than δ. The translating lemma may only apply to a variation of dispersion that is sensitive to such walks, a variant which we call *auto-dispersion*. A δ-dispersed set $S \subseteq P(G)$ is δ-*auto-dispersed* if additionally for every point $p \in S$ there is no walk from p to p of length $< \delta$ that is locally-injective. A walk is *locally-injective* if, when interpreted as a continuous mapping $f : [0,1] \to P(G)$ from $f(0) = p$ to $f(1) = p$, has for every pre-image $c \in (0,1)$ a positive range $\varepsilon > 0$ such that f restricted to the interval $(c-\varepsilon, c+\varepsilon)$ is injective.

▶ **Lemma 3.** ($\star$) *Let G be a graph and $\delta > 0$. Then δ-auto-disp$(G) = \frac{\delta}{\delta+1}$-auto-disp$(G) + |E(G)|$.*

Fortunately, this translation lemma is still useful for ordinary δ-dispersion. We have δ-auto-disp$(G) = \delta$-disp(G) for $\delta \le 3$, since there is no such locally-injective walk of length < 3. The threshold of 3 is tight according to the above example with graph K_3.

▶ **Corollary 4.** *Let G be a graph and $\delta \in (0,3]$. Then δ-disp$(G) = \frac{\delta}{\delta+1}$-disp$(G) + |E(G)|$.*

4 Dispersion and Independent Set

To solve Dispersion we can borrow from algorithmic results from a generalized independent set problem. A classical independent set is a set of *vertices* where each two elements have to be at least 2 apart from each other (when we consider that the edges have unit length). In a 2-dispersed set also each two elements need to be at least 2 apart from each other, though the set contains a set of *points* of the graph.

To generalize the independent set problem, we may ask that the vertices are not 2 apart but some integer d apart from each other. Such a generalization for independent set is called a *distance-d independent* set or *d-scattered set*. They have been studied by Eto et al. [5] and Katsikarelis et al. [13].

Let $\alpha_d(G)$ be the maximum size of a distance-d independent set, for a graph G and integer d. We relate δ-dispersion to α_d. We consider the *c-subdivision of a graph G*, denoted as G_c, which is the graph G where every edge is replaced by a path of length c, for some integer $c \ge 1$.

▶ **Lemma 5.** *Consider integers a, b and a $2b$-subdivision G_{2b} of a graph G. Then $\frac{a}{b}$-disp$(G) = \alpha_{2a}(G_{2b})$.*

Proof. Consider the b-subdivision G_b of G. Then G_{2b} is a 2-subdivision of G_b. We know that $\frac{a}{b}$-disp$(G) = 2a$-disp(G_{2b}) from Lemma 1. Hence it remains to show $2a$-disp$(G_{2b}) = \alpha_{2a}(G_{2b})$.

Clearly, a distance-$2a$ independent set $I \subseteq V(G)$ is also a $2a$-dispersed set. For the reverse direction, assume there is a $2a$-dispersed set I_{2a} of G_{2b}. Then I_{2a} corresponds to an a-dispersed set I of G_b of same size, according to Lemma 1. Since a is integer, we may assume that S contains only half-integral points, hence points with edge position from $\{0, \frac{1}{2}, 1\}$, according to Lemma 2. Let G_{2b} result from G_b by replacing each edge $\{u, v\}$ by a path $uw_{u,v}v$. Then let $I \subseteq V(G_{2b})$ consist of vertex $u \in V(G)$ with a point in S and every $w_{u,v}$ for every point $p(u, v, \frac{1}{2}) \in S$. Then I is a distance-$2a$ independent set of G_{2b} of size $|I| = |S|$. ◀

Thus to solve Dispersion for $\delta = \frac{a}{b}$ we can use algorithms for distance-d independent set. For rationals $\frac{a}{b}$ with small values of a and b this possibly leads to efficient algorithms. For example, a distance-d independent set on graphs width treewidth $\mathsf{tw}(G)$ (and a given tree decomposition) can be found in time $d^{\mathsf{tw}(G)} n^{\mathcal{O}(1)}$, see [13]. "Simply" subdivide the edges of the input graph sufficiently often, which does not increase the treewidth of the considered graph. To find a $\frac{a}{b}$-dispersed set in a graph G, we can search for distance $2a$ independent set the $2b$-subdivision of G.

▶ **Corollary 6.** *There is an algorithm that, given a rational distance $\frac{a}{b} > 0$ and a graph G, a tree decomposition of width $\mathsf{tw}(G)$, computes a maximum $\frac{a}{b}$-dispersed set S in time $(2a)^{\mathsf{tw}(G)}(bn)^{\mathcal{O}(1)}$.*

However, in general this constitutes a possibly exponential increase of the input size. While in the input of $\frac{a}{b}$-dispersion encodes a and b in binary, the subdivided graph essentially encodes b in unary. Further, if δ is irrational, we do not have a suitable subdivision at all.

5 Rounding the Distance

For a given graph G and distance δ, we state a rational $\delta^\star \geq \delta$ such that δ-disp$(G) = \delta^\star$-disp(G). Our proof is constructive. We give a procedure that efficiently transforms a δ-dispersed set into a $\delta^\star$-dispersed set. The guaranteed rational $\delta^\star$ has a numerator bounded by the longest path in G (or just n as an upper bound thereof). It is independent of the precise structure of the given graph.

To give some intuition: Generally there is some leeway for δ. For example, in a star $K_{1,k}$, $k \geq 1$ for every $\delta \in (1, 2]$ the optimal solution puts a point on every leaf yielding a δ-dispersed set of size k. Hence for instance $\frac{3}{2}$-disp$(K_{1,k}) = 2$-disp$(K_{1,k})$. However, for $\delta > 2$ only one point can be placed, such that 2-disp$(K_{1,k}) \neq (2 + \varepsilon)$-disp$(K_{1,k})$ for every $\varepsilon > 0$.

So what $\delta^\star$ can be guaranteed such that δ-disp$(G) = \delta^\star$-disp(G)? An illustrative example is a path of length 6. Then $\frac{15}{11}$-disp$(G) = 5 = \frac{3}{2}$-disp(G). For $\delta = \frac{15}{11}$ tightly packing 5 points allows to have a space of size $\frac{6}{11}$ at either end of the path, not enough to place another point. However, placing 5 points in distance $\delta = \frac{3}{2}$ allows no leeway; δ is (already) a divisor of 6, the length of the considered path. Distance $\delta^\star$ relies on L, the length of the longest (not-induced) path in G. We have to take into account that δ might divide any path of length $\leq L$. Our $\delta^\star$ is the smallest rational $\frac{a^\star}{b^\star}$ where the numerator $a^\star \leq 2L$. In other words, the inverse of $\delta^\star$ is the next smaller rational number of the inverse of δ in the Farey sequence of order $2L$.

▶ **Theorem 7.** *Let $\delta \in \mathbb{R}^+$. Let L be an upper bound on the length of paths in G. Let $\delta^\star = \frac{a^\star}{b^\star} \geq \delta$ minimal with $a^\star \leq 2L$ and $b^\star \in \mathbb{N}$. Then δ-disp$(G) = \delta^\star$-disp(G).*

Clearly, a $\delta^\star$-dispersed set $S^\star$, is also δ-dispersed, since $\delta^\star \geq \delta$. We have to show the reverse direction. Consider a δ-dispersed set S (of size $|S| \geq 2$) of a connected graph G that is not $\delta^\star$-dispersed, hence δ is irrational or is equal to $\frac{a}{b}$ for some co-prime a, b with $a > 2L$.

In the following we develop our rounding procedure that shows the reverse direction. Our presentation aims to be accessible by starting from the core algorithmic idea from which we unravel all involved technical concepts piece by piece. The detailed proofs are placed in the appendix.

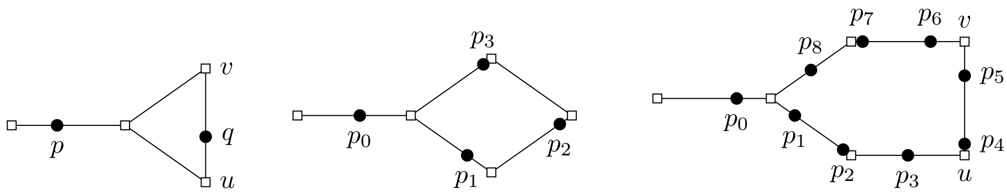

Figure 1 (left) Consider critical points $\{p, q\}$ (points depicted as black dots; vertices as white squares). If we move q away from p by $\varepsilon \geq 0$, their distance increases by ε until q reaches the half-integral point $p(u, v, \frac{1}{2})$. (middle) Let $\{p_i, p_{i-1}\}$ be critical for $i \geq 1$. Consider moving $p_i, i \geq 1$ by $i\varepsilon$ away from p_{i-1}. Once p_2 becomes half-integral, points $\{p_3, p_0\}$ become also critical, hence we cannot continue to move points in the same way. This happens when a point in S becomes half-integral or ... (right) ... a point half-way between two points in S becomes half-integral, as in this example between p_4, p_5. We say p_4, p_5 *witness* the *pivot* $p(u, v, \frac{1}{2})$.

5.1 Overview

Our rounding procedure repeatedly applies a pushing algorithm to the current point set S. We show that each such step strictly decreases a polynomially bounded potential $\Phi : P(G) \to \mathbb{N}$.

▶ **Theorem 8.** *Suppose that there is an algorithm, that given a δ-dispersed set S with $\delta < \delta^\star$ computes an $\varepsilon > 0$ and a $(\delta+\varepsilon)$-dispersed set S_ε of size $|S_\varepsilon| = |S|$ that satisfies $\Phi(S) > \Phi(S_\varepsilon)$ for some polynomially bounded potential $\Phi : P(G) \to \mathbb{N}$. Then Theorem 7 follows.*

Proof. Let S be a δ-dispersed set. Apply the assumed algorithm to obtain a $\varepsilon > 0$ and a $(\delta + \varepsilon)$-dispersed set S_ε of size $|S_\varepsilon| = |S|$. If $\delta = \delta^\star$, we reached our goal. Else we apply the assumed algorithm again. Since the potential $\Phi : P(G) \to \mathbb{N}$ decreases for S_ε compared to S and Φ is polynomially bounded, we have to reach $\delta^\star$ in polynomial many steps. ◀

In the remainder of this section we will develop such an algorithm. It pushes the points of point set S away from each other such that their pairwise distance increases from "at least δ" to "at least $\delta + \varepsilon$". We choose $\varepsilon \geq 0$ as large as possible limited by some events. Either we already reach $\delta + \varepsilon = \delta^\star$, hence we reached our goal, or at least one of three events occurs. We will specify these events in the course of this section. These events mean that one pushing step, i.e., one step for Theorem 8 terminated. All the following preparations for such a pushing step start anew.

We make sure that our potential $\Phi : P(G) \to \mathbb{N}$ decreases when an event occurs. Each of the three events has a corresponding partial potential $\Phi_1(S)$, $\Phi_2(S)$ and $\Phi_3(S)$. They define the overall potential as $\Phi(S) := \Phi_1(S) + \Phi_2(S) + \Phi_3(S)$. Each part never increases. Whenever event i occurs, $\Phi_i(S)$ strictly decreases.

We denote a pair of points $\{p, q\}$ from our given point set S as δ-*critical*, if they have distance exactly δ. Hence the critical pairs of points are exactly those that we need to push away from each other. At the same time we make sure that, once $\{p, q\}$ are δ-critical, they never turn uncritical again, i.e., they are $(\delta + \varepsilon)$-critical in the next step. An uncritical pair of points $\{p, q\}$ might become critical, hence we have to take care of $\{p, q\}$ in future steps. This constitutes our first event. The corresponding partial potential is $\Phi_1(S)$, the number of uncritical pairs of points $\{p, q\}$.

(Event 1) A δ-uncritical pair of points $\{p, q\}$ becomes $(\delta + \varepsilon)$-critical.

$$\Phi_1(S) := \left|\left\{\{p, q\} \in \tbinom{S}{2} \mid \{p, q\} \text{ are not } \delta\text{-critical}\right\}\right| \leq |S|^2$$

5.2 Coordination of Movement

We need to coordinate the movement of all critical pairs of points. To this end, we will fix some set of *root points* R. Our movement will be locally prescribed for sequences of points $p_0, p_1, \dots, p_s$ that originate in $p_0 \in R$ and where each $\{p_0, p_1\}, \dots, \{p_{s-1}, p_s\}$ is critical. The overall movement will be uniquely defined by movement defined for these sequences.

For now, consider such a sequence of points $p_0, p_1, p_2, \dots$. Our idea is to do not move p_0, to move p_1 by distance ε away from p_0, point p_2 by distance 2ε away from p_1 and so on. We have to stop pushing in this way as soon as one of the points, say p_i, becomes half-integral, i.e., p_i is moved onto a vertex or the midpoint of an edge. See Figure 1 for examples. This constitutes the second event.

(Event 2) A non-half-integral $p \in S$ becomes half-integral.

$$\Phi_2(S) \coloneqq \left|\{p \in S \mid p \text{ is not half-integral}\}\right| \leq |S|.$$

The next pushing step will choose p_i as one of the root point R and will move the points away from p_i instead of p_0. Very similarly, we stop when a point $r \in P(G)$ that is "half-way" between two points p_i, p_{i-1} becomes half-integral. Formally, we denote such a point r as an (S, δ)-*pivot*, or simply a pivot, if it is half-integral and there is a (critical) pair of points $\{p, q\} \in \binom{S}{2}$, the witnesses, that have equal distances to r, which means $d(p, r) = d(q, r) = \frac{\delta}{2}$. Let $\mathsf{pivots}(S, \delta)$ be the set of (S, δ)-pivots, and let $W(S, \delta) \subseteq \binom{S}{2}$ be the family of pairs of points from S, that witness some (S, δ)-pivot. This leads to the third and final event.

(Event 3) A non-pivot point $r \in P(G)$ becomes a pivot.

$$\Phi_3(S) \coloneqq \left|\{r \in P(G) \mid r \text{ is half-integral}\} \setminus \mathsf{pivots}(S, \delta)\right| \leq |V(G)|^2.$$

Hence a root point R may not only be a point $p \in S$ but also come from the set of pivots. We will later properly define R as a superset of half-integral points $p_i \in S$ and the (S, δ)-pivots.

We use an *auxiliary graph* G_S for the current δ-dispersed set S. Its vertex set is $S \cup \mathsf{pivots}(S, \delta)$. Essentially we make all pairs of critical $\{p, q\}$ adjacent unless they witness a pivot; If they do witness a pivot, we make them adjacent to the pivot:

- For $\{p, q\} \in W(S, \delta)$ and for every pivot $r \in \mathsf{pivots}(S, \delta)$ they witness, add edges $\{p, r\}, \{r, q\}$; and
- for every critical pair of points $\{p, q\} \in \binom{S}{2} \setminus W(S, \delta)$ add edge $\{p, q\}$.

Note that, for every edge $\{r, p\}$ with $r \in \mathsf{pivots}(S, \delta)$, there is at least one other edge $\{r, q\}$ such that p, q witness r as a pivot.

Now we define the sequence of points which serve as the structure to state the movement. A path $P = (p_0, p_1, \dots, p_s)$ in the auxiliary graph G_S of length $s \geq 1$ is a *spine* if $p_1, \dots, p_s$ are not half-integral. Note that any sub-sequence $(p_0, \dots, p_i)$ for $1 \leq i \leq s$ is also a spine.

5.3 Velocities

We assign velocities vel_P to the points $p_0, \dots, p_s$ of a spine P that specify their movement speed. The point p_i for $i \in \{1, \dots, s\}$ is moved by $\mathsf{vel}(p_i)\varepsilon$. Thus setting $\mathsf{vel}(p_1) = 1$ makes the δ-critical $\{p_0, p_1\}$ become $(\delta + \varepsilon)$-critical, as desired. Setting $\mathsf{vel}(p_i) = i$ for $i \geq 1$, however, can make consecutive points $\{p_{i-1}, p_i\}$ uncritical. Figure 2 provides an example. To see this, fix some shortest p_{i-1}, p_i-path P_i and some shortest p_i, p_{i+1}-path P_{i+1}. The paths P_i and P_{i+1} can have a trivial intersection of only $\{p_i\}$ or their intersection may contain more than one point. We denote this bit of information as $\mathsf{flip}_P(p_i) \in \{-1, 1\}$. We set $\mathsf{flip}_P(p_i) = 1$

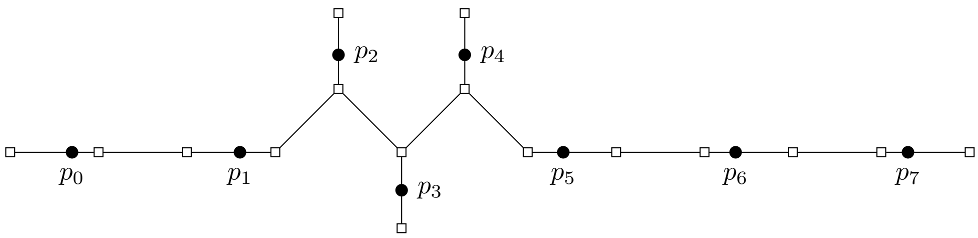

Figure 2 A spine $P = (p_0, \ldots, p_7)$. The shortest path between p_0, p_1 and the shortest path between p_1, p_2 have the trivial intersection of $\{p_1\}$, hence $\mathsf{flip}_P(p_1) = 1$. In turn, the shortest path between p_1, p_2 and the shortest path between p_2, p_3 have a non-trivial intersection, hence $\mathsf{flip}_P(p_1) = -1$. Also $\mathsf{flip}_P(p_3) = \mathsf{flip}_P(p_4) = -1$ while the other values are positive. Consequently $(\mathsf{sgn}_P(p_1), \ldots, \mathsf{sgn}_P(p_7)) = (1, 1, -1, 1, -1, -1, -1)$. Thus $(\mathsf{vel}_P(p_1), \ldots, \mathsf{vel}_P(p_7)) = (1, 2, 1, 2, 1, 0, -1)$. In particular, $\mathsf{sgn}_P(p_5)$ is negative such that it is moved towards p_4. In turn, $\mathsf{sgn}_P(p_7)$ and $\mathsf{vel}_P(p_7)$ are negative such that p_7 has a net movement away from p_0. Under this movement all $\{p_0, p_1\}, \ldots, \{p_6, p_7\}$ remain critical.

if and only if the path P_i and P_{i-1} have a trivial intersection. (The definition of flip_P is independent on the exact considered shortest paths and we will define it properly in the next subsection.)

The easy case is when P_i and P_{i+1} have a trivial intersection, i.e., $\mathsf{flip}_P(p_i) = 1$. Then we increase the velocity of the next point p_{i+1}. The first time we encounter the other case, that $\mathsf{flip}_P(p_i) = -1$, we decrease the velocity of the next point p_{i+1}. Further we move p_{i+1} *towards* and not away of p_i. Hence we also specify a sgn of the velocity that records whether a point p_{i+1} is pushed towards or away from its predecessor p_i. All these changes are relative to whether the movement of predecessor p_i is away from p_{i-1}, i.e., whether $\mathsf{sgn}(p_i)$ is positive. For example, the second time we encounter a point p_j with $\mathsf{flip}_P(p_j) = -1$, point p_{j+1} is again moved away from its predecessor.

This leads to the following definition of vel_P and sgn_P for a spine P. We define its half-integral *velocities* $\mathsf{vel}_P : \{p_0, \ldots, p_s\} \to \{\frac{z}{2} \mid z \in \mathbb{Z}\}$ depending on *signs* $\mathsf{sgn}_P : \{p_1, \ldots, p_s\} \to \{-1, 1\}$, which in turn depend on flip_P. We may drop the subscript P, if it is clear from the context. Let $\mathsf{vel}(p_0) = 0$. Let $\mathsf{vel}(p_1) = \frac{1}{2}$, if $p_0 \in \mathsf{pivots}(S, \delta)$, and let $\mathsf{vel}(p_1) = 1$, if $p_0 \in S$. For $i \geq 1$, let

$$\mathsf{vel}(p_{i+1}) \;\coloneqq\; \mathsf{vel}(p_i) + \mathsf{sgn}(p_{i+1}).$$

Thus $\mathsf{sgn} \in \{-1, 1\}$ indicates whether the velocity increases or decreases. The current sgn is unchanged unless flip is negative. Let $\mathsf{sgn}(p_1) = 1$. For $2 \leq i \leq s$, let

$$\mathsf{sgn}(p_i) \;\coloneqq\; \mathsf{flip}(p_{i-1})\,\mathsf{sgn}(p_{i-1}) \;=\; \prod_{0<j<i} \mathsf{flip}(p_j).$$

The movement step of a point p_i in a spine $P = (p_0, \ldots, p_i)$ is now as follows. We push the point p_i by the (possibly negative) distance $\mathsf{sgn}_P(p_i)\,\mathsf{vel}_P(p_i)\varepsilon$ away from its predecessor p_{i-1}. In other words, the point $p_i = p(u, v, \lambda)$ is replaced by the point $p(u, v, \lambda + \mathsf{sgn}_P(p_i)\,\mathsf{vel}_P(p_i)\varepsilon)$ assuming that vertex u compared to v is in some sense closer to the predecessor point p_{i-1}. We make this notion formal in the next subsection.

5.4 Directions

We formalize the notion the direction of a point p towards another point q. The *direction* $\mathsf{dir}(p \to q) \in \{u, v\}$ for distinct points $p = p(u, v, \lambda) \in P(G)$ and $q \in P(G)$ is defined as follows:

- For points $p = p(u, v, \lambda_p)$ and $q = p(u, v, \lambda_q)$ on a common edge $\{u, v\} \in E(G)$ with $\lambda_p < \lambda_q$, let $\mathsf{dir}(p \to q) = v$. Let $\overline{\mathsf{dir}}(p \to q) = u$.
- For points $p = p(u_p, v_p, \lambda_p)$ and $q = p(u_q, v_q, \lambda_q)$ on distinct edges $\{u_p, v_p\} \neq \{u_q, v_q\}$, let $\mathsf{dir}(p \to q)$ be the unique vertex of $\{u_p, v_p\}$ that is contained in *every* shortest path between p and q, if such a vertex exists. If $\mathsf{dir}(p \to q)$ is defined, let $\overline{\mathsf{dir}}(p \to q)$ be the unique vertex in $\{u_p, v_p\} \setminus \{\mathsf{dir}(p \to q)\}$.

▶ **Lemma 9.** ($\star$) *For distinct points $p, q \in V(G_S)$, $\mathsf{dir}(p \to q)$ is well-defined, unless p is half-integral.*

Hence we can properly define $\mathsf{flip}(p_i)$ of a point p_i of a spine $(p_0, \ldots, p_s)$ with $1 \leq i \leq s-1$, since p_i with $i \geq 1$ is non-half-integral. Let

$$\mathsf{flip}(p_i) := \begin{cases} 1, & \mathsf{dir}(p_i \to p_{i-1}) \neq \mathsf{dir}(p_i \to p_{i+1}), \\ -1, & \text{else.} \end{cases}$$

Further, for a non-half-integral point $p = p(u, v, \lambda)$ we have $\{u, v\} = \{\mathsf{dir}(p \to q), \overline{\mathsf{dir}}(p \to q)\}$. By symmetry assume that $\overline{\mathsf{dir}}(p \to q) = u$. We can equivalently specify point p as $p(\overline{\mathsf{dir}}(p \to q), \mathsf{dir}(p \to q), \lambda)$. Conveniently, we may write $p = p(\cdot, \overline{\mathsf{dir}}(p \to q), \lambda)$ since the missing entry is clear from the context. Doing so, the edge position λ measures a part of the length of any shortest p, q-path, specifically the part using the edge of p (assuming q is on another edge).

Now we can also properly define one pushing step for a point p_i of a spine $P = (p_0, \ldots, p_s)$ and for $\varepsilon > 0$. Let λ_i be such that $p_i = p(\cdot, \overline{\mathsf{dir}}(p_i \to p_{i-1}), \lambda_i)$. Then the new point is

$$(p_i)_{P,\varepsilon} := p\big(\cdot,\ \overline{\mathsf{dir}}(p_i \to p_{i-1}),\ \lambda_i + \mathsf{sgn}_P(p_i)\, \mathsf{vel}_P(p_i)\varepsilon\big).$$

▶ **Lemma 10.** ($\star$) *For a spine $P = (p_0, \ldots, p_s)$ and $i \in \{0, \ldots, s-1\}$, points $(p_i)_{P,\varepsilon}, (p_{i+1})_{P,\varepsilon}$ are $(\delta + \varepsilon)$-critical for the maximal $\varepsilon \leq \delta^\star - \delta$ that is limited by the Events 1,2,3.*

5.5 Root Points

We formally define the set of root points R. Let R_0 be the set of half-integral points in G_S. There may be some components of the auxiliary graph G_S without a point in R_0. Let R result from R_0 by adding exactly one point from every component that has no point in R_0.

We consider only spines $P = (p_0, \ldots, p_i)$ where $p_0 \in R$. Clearly every point G_S is part of at least one spine and hence has some movement prescribed. We also claim that the prescribed movement is uniquely defined. In other words, there are no spines $P = (p_0, \ldots, p_i)$ and $Q = (q_0, \ldots, q_j)$ with $p_0, q_0 \in R$ that terminate at the same point $p_i = q_j$ and contradict in their prescribed movement for $p_i = q_j$.

▶ **Lemma 11.** ($\star$) *Let $\delta < \delta^\star$. Consider spines $P = (p_0, \ldots, p_i)$ and $Q = (q_0, \ldots, q_j)$ with $p_0, q_0 \in R$ and $p_i = q_j$. Then (1) $\mathsf{vel}_P(p_i) = \mathsf{vel}_Q(q_j)$; and (2) $\overline{\mathsf{dir}}(p_i \to p_{i-1}) = \overline{\mathsf{dir}}(q_j \to q_{j-1})$ if and only if $\mathsf{sgn}_P(p_i) = \mathsf{sgn}_Q(q_j)$.*

Therefore we can uniquely define the moved version of a point p as $p_\varepsilon := (p)_{P,\varepsilon}$ where we may choose P to be an arbitrary spine starting in R and containing p. This defines the set of pushed points $S_\varepsilon := \{\, p_\varepsilon \mid p \in S \,\}$.

For the proof of Lemma 11, we use that $\delta^\star = \frac{a^\star}{b^\star} \geq \delta$ is minimal with $a^\star \leq 2L$. Since spines P, Q meet at the point $p_i = q_j$ their roots p_0, q_0 must be from the same component in the auxiliary graph G_S; in other words p_0, q_0 are both half-integral or are the same point.

Our proof by contradiction considers these two options and whether P, Q reach p_i and q_j from the same vertex relatively to if they agree on the sgn of $p_i = q_j$. An example is that $p_0 = q_0$ and P, Q reach $p_i = q_j$ from the same vertex (formally with $\overline{\mathsf{dir}}(p_i \to p_{i-1}) \neq \overline{\mathsf{dir}}(q_j \to q_{j-1})$) while they agree on the sgn (formally $\mathsf{sgn}_P(p_i) = \mathsf{sgn}_Q(q_j)$). Then we can glue P and Q together forming a walk starting in $p_0 = q_0$ and returning to $p_0 = q_0$. This walk then has half-integral length at most $2L$ but is made up of hops of length δ. That implies the contradiction that already $\delta = \delta^\star$.

5.6 Summery

With the previous observations we can assemble the algorithm for Theorem 8. Our potential counts how many elements can still trigger Events 1,2,3. That is

$$\Phi(S) \;\coloneqq\; \Phi_1(S) + \Phi_2(S) + \Phi_3(S) \;\leq\; 2|S|^2 + |V(G)|^2.$$

We define $\varepsilon^\star \geq 0$ as the maximal $\varepsilon \leq \delta^\star - \delta$ limited by the Events 1,2,3. We claim that $\varepsilon^\star \geq 0$ is defined. This is due to that the above events depend on continuous functions in ε, which are the distance of p_ε to its closest half-integral point, and the distance between points p_ε and q_ε for $p, q \in S$.

To show termination, we prove that no such element can trigger its according event more than once. Lemma 10 already implies that a δ-critical pair of points $\{p, q\}$ stays $(\delta + \varepsilon^\star)$-critical. It remains to show the following monotonicities:

▶ **Lemma 12.** ($\star$) *Let S be a δ-dispersed set for $\delta < \delta^\star$ and $\varepsilon^\star$ defined as above. Then:*

- (1) *$S_{\varepsilon^\star}$ is a $(\delta + \varepsilon^\star)$-dispersed set of size $|S|$.*
- (2) *If $\{p, q\} \in \binom{S}{2}$ is δ-critical, then $\{p_{\varepsilon^\star}, q_{\varepsilon^\star}\}$ is $(\delta + \varepsilon^\star)$-critical.*
- (3) *If $r \in \mathsf{pivots}(S, \delta)$, then $r \in \mathsf{pivots}(S_{\varepsilon^\star}, \delta + \varepsilon^\star)$.*

Now we have all the tools to show Theorem 8. Determine $\varepsilon^\star$ and the $(\delta + \varepsilon^\star)$-dispersed set $S_{\varepsilon^\star}$ as defined before. The resulting set $S_{\varepsilon^\star}$ is a $(\delta + \varepsilon^\star)$-dispersed set of the same size, according to Lemma 12. If $\delta + \varepsilon^\star = \delta^\star$ then $S_{\varepsilon^\star}$ is already the desired $\delta^\star$-dispersed set. Else one of the Events 1,2,3 occurred. We observe that the potential strictly decreases, that is $\Phi(S_{\varepsilon^\star}) < \Phi(S)$. Because of the monotonicities of Lemma 10 and Lemma 12 the partial potentials Φ_1, Φ_2 and Φ_3 do not increase. If Event 1 occurs then Φ_1 strictly decreases. If Event 2 occurs then Φ_2 strictly decreases. If Event 3 occurs then Φ_3 strictly decreases. All in all at least one part strictly decreases and so does Φ. This completes the proof of Theorem 8 and hence of Theorem 7.

6 Algorithmic Implications

Based on the rounding procedure from Section 5, the translation result from Section 3 and connections to distance-d independent set we derive a number of algorithmic results.

▶ **Theorem 13.** *There is an algorithm that, given distance $\delta \geq 0$, a graph G, a tree decomposition and an upper bound $L \in \mathbb{N}$ on the length of the longest path in G, computes a maximum δ-dispersed set S in time $(2L)^{\mathsf{tw}(G)} n^{\mathcal{O}(1)}$.*

Proof. According to Theorem 7, we may consider the rounded up distance, that is a rational $\frac{a}{b} \geq \delta$ with $a \leq 2L$, instead of δ. Notice that $\frac{a}{b}$ is polynomial time computable. As long as $\frac{a}{b} \leq \frac{3}{4}$, we may repeatedly apply Corollary 4 such that eventually we obtain that $\frac{3}{4}b < a \leq 2L$. Let G_{2b} be a $2b$-subdivision of G. Observe that $\mathsf{tw}(G) = \mathsf{tw}(G_{2b})$ and the number of vertices increases only by a factor of $\mathcal{O}(n^2L)$. According to Lemma 5 $\frac{a}{b}$-disp$(G) = \alpha_{2a}(G_{2b})$. Thus, to answer the original δ-Dispersion-instance we may find a maximum distance-$2a$ independent set in G_{2b}, which is possible in time $(2a)^{\mathsf{tw}(G)}n^{\mathcal{O}(1)}$, according to [12]. ◀

This result immediately yields parameterized complexity results for the parameters treedepth and treewidth. Regarding the treewidth, note that n is an upper bound on L. Thus the above algorithm is an XP-algorithm for the parameter treewidth. When a treewidth decomposition is given, Dispersion can be solved in time $2n^{\mathsf{tw}(G)}n^{\mathcal{O}(1)}$.

▶ **Corollary 14.** *Dispersion can be solved in time $2n^{\mathsf{tw}(G)}n^{\mathcal{O}(1)}$, assuming a tree decomposition is given.*

Similarly we obtain an FPT algorithm for treedepth $\mathsf{td}(G)$ of the input graph. The treedepth $\mathsf{td}(G)$ implies a bound on L, which is $L \leq 2^{\mathsf{td}(G)}$. Since also $\mathsf{td}(G) \geq \mathsf{tw}(G)$, we obtain an $2^{\mathcal{O}(\mathsf{td}(G)^2)}n^{\mathcal{O}(1)}$-time algorithm, assuming a treedepth decomposition is given.

▶ **Corollary 15.** *Dispersion can be solved in time $2^{\mathcal{O}(\mathsf{td}(G)^2)}n^{\mathcal{O}(1)}$, assuming a treedepth decomposition is given.*

7 Parameterized Hardness Results

We complement the positive results by hardness results. These results borrow ideas from hardness-reductions for the similar problem Distance Independent Set (DIS), see Section 4.

A natural generalization of treedepth is the maximum diameter of graph G, which is the maximum distance between any vertices $u, v \in V(G)$ (since we only consider connected graphs G). We show NP-hardness for graphs of any diameter ≥ 3 even for chordal graphs by a reduction from Independent Set, similarly as NP-hardness for DIS is shown by Eto et al. [5]. Our reduction also shows W[1]-hardness with respect to the solution size k.

▶ **Lemma 16.** ($\star$) *For every $\delta > 3$, δ-Dispersion is NP-complete and W[1]-hard with parameter solution size, even for connected chordal graphs of diameter $\leq \lceil\delta\rceil$.*

Another direct generalization of treedepth is pathwidth of the input graph G. We show W[1]-hardness even for the combined parameters pathwidth and solution size $\mathsf{pw}(G) + k$. With the same reduction also W[1]-hardness for the combined parameters "feedback vertex set size" $\mathsf{fvs}(G)$ and solution size k follows. We can essentially use the same reduction as used by Katsikarelis et al. to show W[1]-hardness of DIS when parameterized by $\mathsf{fvs}(G) + k$ by reducing from Multi-Colored-Independent-Set [12].

▶ **Theorem 17.** ($\star$) *Dispersion is W[1]-hard parameterized by $\mathsf{pw}(G)+k$. Further, there is no $n^{o(\sqrt{\mathsf{pw}(G)}+\sqrt{k})}$-time algorithm unless ETH fails. Dispersion is W[1]-hard parameterized by $\mathsf{fvs}(G) + k$. Further, there is no $n^{o(\mathsf{fvs}(G)+\sqrt{k})}$-time algorithm unless ETH fails.*

Since $\mathsf{fvs}(G)$ is a linear upper bound for the treewidth of G, we also obtain: Dispersion is W[1]-hard parameterized by $\mathsf{tw}(G) + k$. Further, there is no $n^{o(\mathsf{tw}(G)+\sqrt{k})}$-time algorithm unless ETH fails. Similarly as in [12] we obtain a lower bound for treedepth.

▶ **Theorem 18.** ($\star$) *Assuming ETH, there is no $2^{o(\mathsf{td}(G)^2)}$-time algorithm for Dispersion.*

8 NP-hardness for Irrational Distance

We show NP-hardness of δ-Dispersion for every irrational distance $\delta > 0$. Thus together with earlier results [8] the complexity for every real $\delta > 0$ is resolved: For rational distance $\delta = \frac{a}{b}$ where $a \in \{1, 2\}$ the problem is polynomial time solvable, while it is NP-complete for every other distance $\delta > 0$.

▶ **Theorem 19.** *For every irrational $\delta > 0$, δ-Dispersion is NP-complete.*

The key step is a reduction from Independent Set which shows NP-hardness not only for a single distance δ but for the whole interval $\delta \in (2, 3]$.

Construction. Given a graph G and integer $k \in \mathbb{N}$, we construct an input for δ-Dispersion consisting of a graph G' and integer $k' = k$ as follows. For every vertex $u \in V(G)$ introduce vertices u_1, u_2 and edge $\{u_1, u_2\}$. For every edge $\{u, v\} \in E(G)$ introduce edges $\{u_i, v_j\}$ for every $i, j \in \{1, 2\}$.

▶ **Lemma 20.** *For every $\delta \in (2, 3]$, δ-Dispersion is NP-hard and W[1]-hard when parameterized by solution size.*

Proof. Clearly, this construction is polynomial time computable. Further, the reduction is parameter preserving such that W[1]-hardness of Independent Set implies W[1]-hardness of Dispersion, assuming correctness of the reduction.

Hence, it remains to show the correctness, that G has an independent set of size k if and only if G' has a δ-dispersed set of size k.

($\Rightarrow$) Let $I \subseteq V(G)$ be an independent set of graph G. We define $S := \{p(u_1, u_2, \frac{1}{2}) \mid u \in I\} \subseteq P(G)$, which has size $|S| = |I|$. We claim that S is δ-dispersed in G' for $\delta \in (2, 3]$. Since any vertices $u, v \in V(G)$ have distance at least 2 in G, their corresponding points $p(u_1, u_2, \frac{1}{2})$ and $p(v_1, v_2, \frac{1}{2})$ have distance at least 3 in $P(G)$. Thus they are δ-dispersed for $\delta \in (2, 3]$.

($\Leftarrow$) Let $S \subseteq P(G)$ be a δ-dispersed set for some $\delta \in (2, 3]$. We define the *ball* B_u for $u \in V(G)$ as the points in $P(G)$ with distance at most $\frac{1}{2}$ to u_1 or u_2, which is $B_u := \{p(u_i, v, \lambda) \mid i \in \{1, 2\}, \{u_i, v\} \in E(G'), \lambda \in [0, \frac{1}{2}]\}$. Every ball B_u for $u \in V(G)$ contains at most one point from S since points $p, q \in B_u$ can be at most $2 < \delta$ apart. Every union $B_u \cup B_v$ for adjacent $\{u, v\} \in E(G)$ contains at most one point from S since points $p, q \in B_u \cup B_v$ can also be at most $2 < \delta$ apart.

Now we define an independent set $I \subseteq V(G)$. Add vertex $u \in V(G)$ for every point $p \in S \cap B_u$ except when $p \in B_u \cap B_v$ for some $v \in P(G)$. If $p \in S \cap B_u \cap B_v$, add either the point u or v to I. Then $|I| = |S|$ since the union of B_u for $u \in V(G)$ is the whole point space $P(G)$. Further, no adjacent vertices u, v are in I since $B_u \cup B_v$ contain at most one point from S. Thus $I \subseteq V(G)$ is an independent set of size $|S|$. ◀

Because $\delta \leq 3$ we may apply Lemma 3 to obtain NP-hardness for δ in the interval $(\frac{2}{2x+1}, \frac{3}{3x+1}]$ for every integer $x \geq 0$. Applying Lemma 1 yields NP-hardness for δ in the interval $(\frac{2c}{2x+1}, \frac{3c}{3x+1}]$ for every integer $c \geq 1$.

Now, consider any irrational distance $\delta > 0$. Consider $F := \{c\delta^{-1} - \lfloor c\delta^{-1} \rfloor \mid c \geq 1\}$, the set of fractional parts of multiples of δ^{-1}. Since δ^{-1} is irrational, F is a dense subset of the interval $[0, 1]$. Let integer $c \geq 1$ be such that $\frac{1}{3} \leq c\delta^{-1} - \lfloor c\delta^{-1} \rfloor < \frac{1}{2}$. Thus there is a non-negative x such that $x + \frac{1}{3} \leq c\delta^{-1} < x + \frac{1}{2}$. This implies that $\frac{2c}{2x+1} < \delta \leq \frac{3c}{3x+1}$ and hence NP-hardness for δ-dispersion. This finishes the proof of Theorem 19.

9 Parameter Solution Size

δ-Dispersion parameterized by the solution size k is W[1]-hard when $\delta > 2$: When $\delta \in (2,3]$ Lemma 20 shows W[1]-hardness, while for $\delta > 3$ Lemma 16 implies W[1]-hardness even when the input graph is chordal. It remains to consider $\delta \le 2$. Observe that for $\delta \le 2$, every graph G satisfies $\delta\text{-disp}(G) \ge \nu(G)$ [8], where $\nu(G)$ is the maximum matching size of G. Thus, a win-win situation occurs. Determine $\nu(G)$ in polynomial time. If $k \le \nu(G)$, we may immediately answer "yes". Otherwise $k > \nu(G) \ge \frac{\mathsf{vc}(G)}{2}$, where $\mathsf{vc}(G)$ is the minimum size of a vertex cover in G. The size of a vertex cover upper bounds the treedepth. A treedepth decomposition of size $\mathsf{td}(G)$ is computable in FPT-time [15]. Thus we may apply the FPT algorithm for parameter treedepth from Theorem 13.

▶ **Theorem 21.** *δ-Dispersion parameterized by solution size k is FPT if $\delta \le 2$; and $W[1]$-hard if $\delta > 2$.*

References

1. Shimon Abravaya and Michael Segal. Maximizing the number of obnoxious facilities to locate within a bounded region. *Comput. Oper. Res.*, 37(1):163–171, 2010. doi:10.1016/j.cor.2009.04.004.
2. Boaz Ben-Moshe, Matthew J. Katz, and Michael Segal. Obnoxious facility location: complete service with minimal harm. *International Journal of Computational Geometry and Applications*, 10(6):581–592, 2000.
3. Richard L. Church and Robert S. Garfinkel. Locating an obnoxious facility on a network. *Transportation Science*, 12:107–118, 1978.
4. Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015. doi:10.1007/978-3-319-21275-3.
5. Hiroshi Eto, Fengrui Guo, and Eiji Miyano. Distance-d independent set problems for bipartite and chordal graphs. *J. Comb. Optim.*, 27(1):88–99, 2014. doi:10.1007/s10878-012-9594-4.
6. Alan J. Goldman and Perino M. Dearing. Concepts of optimal location for partially noxious facilities. *ORSA Bulletin*, 23:B–31, 1975.
7. Alexander Grigoriev, Tim A. Hartmann, Stefan Lendl, and Gerhard J. Woeginger. Dispersing obnoxious facilities on a graph. In *36th International Symposium on Theoretical Aspects of Computer Science, STACS 2019, March 13-16, 2019, Berlin, Germany*, pages 33:1–33:11, 2019. doi:10.4230/LIPIcs.STACS.2019.33.
8. Alexander Grigoriev, Tim A Hartmann, Stefan Lendl, and Gerhard J Woeginger. Dispersing obnoxious facilities on a graph. *Algorithmica*, 83(6):1734–1749, 2021.
9. Tim A. Hartmann and Stefan Lendl. Dispersing obnoxious facilities on graphs by rounding distances, 2022. doi:10.48550/ARXIV.2206.11337.
10. Tim A. Hartmann, Stefan Lendl, and Gerhard J. Woeginger. Continuous facility location on graphs. *Math. Program.*, 192(1):207–227, 2022. doi:10.1007/s10107-021-01646-x.
11. Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *Journal of Computer and System Sciences*, 63(4):512–530, 2001. doi:10.1006/jcss.2001.1774.
12. Ioannis Katsikarelis, Michael Lampis, and Vangelis Th. Paschos. Structurally parameterized d-scattered set. In *Graph-Theoretic Concepts in Computer Science - 44th International Workshop, WG 2018, Cottbus, Germany, June 27-29, 2018, Proceedings*, pages 292–305, 2018. doi:10.1007/978-3-030-00256-5_24.
13. Ioannis Katsikarelis, Michael Lampis, and Vangelis Th. Paschos. Structurally parameterized d-scattered set. *Discret. Appl. Math.*, 308:168–186, 2022. doi:10.1016/j.dam.2020.03.052.

14 Matthew J. Katz, Klara Kedem, and Michael Segal. Improved algorithms for placing undesirable facilities. *Comput. Oper. Res.*, 29(13):1859–1872, 2002. doi:10.1016/S0305-0548(01)00063-6.

15 Felix Reidl, Peter Rossmanith, Fernando Sánchez Villaamil, and Somnath Sikdar. A faster parameterized algorithm for treedepth. In *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*, pages 931–942, 2014. doi:10.1007/978-3-662-43948-7_77.

16 Michael Segal. Placing an obnoxious facility in geometric networks. *Nordic Journal of Computing*, 10:224–237, 2003.

17 Arie Tamir. On the solution value of the continuous p-center location problem on a graph. *Math. Oper. Res.*, 12(2):340–349, 1987. doi:10.1287/moor.12.2.340.

18 Arie Tamir. Obnoxious facility location on graphs. *SIAM J. Discret. Math.*, 4(4):550–567, 1991. doi:10.1137/0404048.

19 Martijn van Ee. Approximability of the dispersed p→-neighbor k-supplier problem. *Discret. Appl. Math.*, 289:219–229, 2021. doi:10.1016/j.dam.2020.10.007.