

Inapproximability of Diameter in Super-Linear Time: Beyond the 5/3 Ratio

Édouard Bonnet 

Univ. Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR5668, France

Abstract

We show, assuming the Strong Exponential Time Hypothesis, that for every $\varepsilon > 0$, approximating directed DIAMETER on m -arc graphs within ratio $7/4 - \varepsilon$ requires $m^{4/3 - o(1)}$ time. Our construction uses non-negative edge weights but even holds for sparse digraphs, i.e., for which the number of vertices n and the number of arcs m satisfy $m = \tilde{O}(n)$. This is the first result that conditionally rules out a near-linear time 5/3-approximation for a variant of DIAMETER.

2012 ACM Subject Classification Theory of computation → Graph algorithms analysis

Keywords and phrases Diameter, inapproximability, SETH lower bounds, k-Orthogonal Vectors

Digital Object Identifier 10.4230/LIPIcs.STACS.2021.17

Related Version *Full Version:* <https://arxiv.org/abs/2008.11315>

Funding This work was supported by the grant from French National Agency under PRC program (project Digraphs, ANR-19-CE48-0013-01).

1 Introduction

The diameter of a graph is the largest length of a shortest path between two of its vertices. We denote by DIAMETER the algorithmic task of computing the diameter of an input graph. We will sometimes prefix DIAMETER by the adjectives *undirected/directed* specifying if edges can be oriented (i.e., if they can be arcs), and *unweighted/weighted* specifying if non-negative edge weights can be used. By default, we will assume that both edge orientations and edge weights are allowed. To be clear, the diameter in a digraph (or directed graph) is the maximum taken over all *ordered* pairs of vertices (u, v) of the distance from u to v . Note that it is very possible that the pair (u, v) realizes the distance of the diameter, while there is a much shorter path from v to u (perhaps just an arc).

There is an active line of work aiming to determine the best running time for an algorithm approximating (variants of) DIAMETER within a given ratio (see for instance the survey of Rubinfeld and Vassilevska Williams [12]). We focus here on sparse graphs, for which the number of edges m and the number of vertices n verify $m = \tilde{O}(n)$, where \tilde{O} suppresses the polylogarithmic factors.¹ There is an exact algorithm running in time $\tilde{O}(n^2)$ by computing n shortest-path trees from every vertex of the graph. There is also a folklore 2-approximation running in time $\tilde{O}(n)$ by computing a shortest-path tree from an arbitrary vertex and outputting the largest distance found. There are an $\tilde{O}(n^{3/2})$ time 3/2-approximation for weighted directed DIAMETER [1, 11, 6], and for every non-negative integer k , an $\tilde{O}(n^{1 + \frac{1}{k+1}})$ time $(2 - 2^{-k})$ -approximation for weighted *undirected* DIAMETER [4]. In dense graphs these four algorithms take time $\tilde{O}(mn)$, $\tilde{O}(m)$, $\tilde{O}(m^{3/2})$ and $\tilde{O}(mn^{\frac{1}{k+1}})$, respectively.

¹ Throughout the paper we adopt the convention that n denotes the number of vertices and m , the number of edges of a given graph.



© Édouard Bonnet;

licensed under Creative Commons License CC-BY 4.0

38th International Symposium on Theoretical Aspects of Computer Science (STACS 2021).

Editors: Markus Bläser and Benjamin Monmege; Article No. 17; pp. 17:1–17:13

Leibniz International Proceedings in Informatics

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



There are two competing criteria: minimizing the approximation factor, which is in that case between 1 and 2, and minimizing the exponent of the running time, also a real number between 1 and 2. We now know that the points $(1, 2)$, $(\frac{3}{2}, \frac{3}{2})$, and $(2, 1)$ are feasible for the more general variant of DIAMETER. The question is whether these algorithms can be improved or if conditional lower bounds can be provided instead. The paper is about the latter, so we will now briefly present the relevant framework of fine-grained complexity, as well as its known consequences for DIAMETER.

Fine-grained complexity

The program of fine-grained complexity aims to match fine-grained algorithms (where the precise running time matters more than the membership to some robust complexity class) with tight conditional lower bounds under well-established assumptions. These assumptions are said *problem-centric*. They rely on the fact that we have been collectively unable to “meaningfully” improve over the brute-force or textbook algorithms for some important problems. Then perhaps such improvements are impossible, or at least they are currently out of reach. A fine-grained reduction from one of these central problems to our problem of interest Π tells us that improving on Π would result in a major breakthrough.

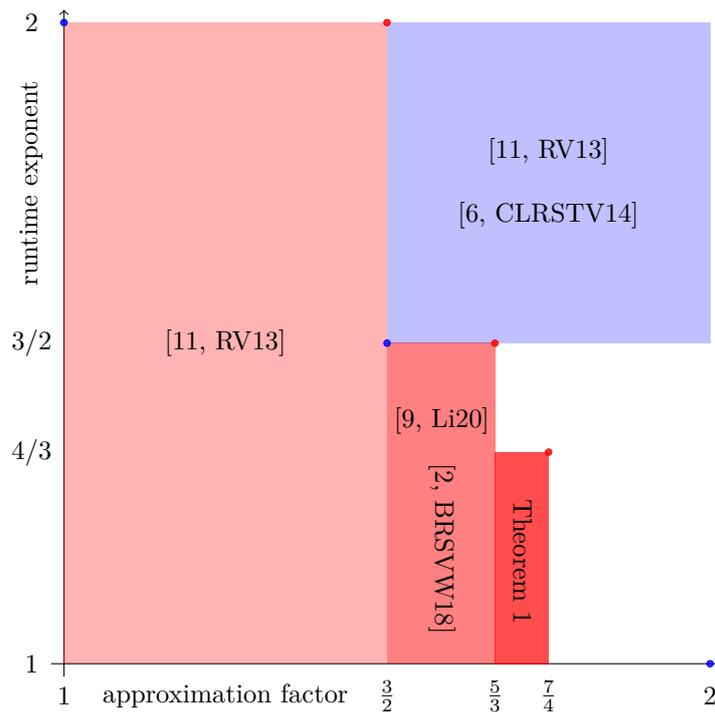
The three main hypotheses are based on SAT, 3-SUM, and ALL-PAIRS SHORTEST-PATHS. One might think that ALL-PAIRS SHORTEST-PATHS is a better starting point for a reduction to DIAMETER. Surprisingly it happens that SAT is. We will now focus on conditional lower bounds for DIAMETER, so we will only define the hypothesis based on SAT. For more on fine-grained complexity, we refer the interested reader to the survey of Vassilevska Williams [14].

The Strong Exponential Time Hypothesis (SETH, for short) asserts that for every $\varepsilon > 0$, there is an integer k such that k -SAT cannot be solved in time $(2 - \varepsilon)^n$ on n -variable instances [8]. The first SETH-based lower bound for a polynomial-time solvable graph problem was precisely on unweighted undirected DIAMETER [11]. The authors show that, unless the SETH fails, any $(3/2 - \delta)$ -approximation for sparse unweighted undirected DIAMETER, with $\delta > 0$, requires time $n^{2-o(1)}$. Backurs et al. [2] show under the same assumption that, for every $k \geq 3$, any $((5k - 7)/(3k - 4) - \delta)$ -approximation for sparse unweighted directed DIAMETER, with $\delta > 0$, requires time $n^{1+\frac{1}{k-1}-o(1)}$, and that any $(5/3 - \delta)$ -approximation for sparse weighted undirected DIAMETER requires time $n^{3/2-o(1)}$. Li [9] improves on these results showing that, unless the SETH fails, any $(5/3 - \delta)$ -approximation for unweighted undirected DIAMETER requires time $n^{3/2-o(1)}$.

Since a $5/3$ -approximation of DIAMETER running in near-linear time was consistent with the current knowledge, even in weighted directed graphs, Rubinfeld and Vassilevska Williams [12] and Li [9] ask for such an algorithm or some lower bounds with a ratio closer to 2. We give an evidence that, at least for weighted directed graphs, no such algorithm is possible. More precisely, our contribution is the following.

► **Theorem 1.** *Unless the SETH fails, for any $\varepsilon > 0$, $(7/4 - \varepsilon)$ -approximating DIAMETER on directed n -vertex $\tilde{O}(n)$ -edge graphs where all the edge weights are non-negative integers requires $n^{4/3-o(1)}$ time.*

Figure 1 summarizes what was known for directed DIAMETER and where Theorem 1 fits.



■ **Figure 1** Approximability of sparse directed DIAMETER. The blue region is feasible, as witnessed by algorithms at the bottom-left corners (blue dots). The three algorithms support non-negative edge weights. The red regions would refute the SETH, as witnessed by reductions at top-right corners (red dots). The lower bounds in [11, 9] even hold for the sparse unweighted undirected DIAMETER, and the one in [2], for sparse weighted undirected DIAMETER, while Theorem 1 uses edge weights and orientations.

Techniques

Like all the DIAMETER lower bounds (see also [14, 12]), we reduce from k -ORTHOGONAL VECTORS. In this problem, given a set of N 0, 1-vectors of dimension ℓ , one looks for k vectors such that at every index, at least one of these k vectors has a 0 entry. Unless the SETH fails, k -ORTHOGONAL VECTORS requires time $N^{k-o(1)}$ [13], even when ℓ is polylogarithmic in N .

Here we will more precisely reduce from 4-ORTHOGONAL VECTORS (4-OV, for short). We want to build a digraph on $\tilde{O}(N^3)$ vertices and arcs, with diameter 7 if there is an orthogonal quadruple (that is, a solution to the 4-OV instance), and diameter 4 otherwise. Following a reduction to ST -DIAMETER² by Backurs et al. [2] (arguably also following [11]) all the reductions feature layers $L_0, L_1, \dots, L_{k-1}, L_k$, with only (forward) edges between two consecutive L_i . The vertices within the same layer share the same number of “vector attributes” and “index attributes”. The interplay between vector and index attributes in defining the vertices and edges is adjusted so that if there are no k orthogonal vectors, then there are paths of “optimal” length k between every pair in $L_0 \times L_k$, while if there is set X of k orthogonal vectors, a pair (x, y) in $L_0 \times L_k$ jointly encoding X is far apart (usually and ideally at distance $2k - 1$).

We do not deviate too much from this strategy. Our construction is inspired from and pushes one step forward the elegant reduction of Li [9]. Here we rename L_0, L_1, L_2, L_3, L_4 by ABC, AB, AD_Y, DC, DCB, respectively. Some pairs of vertices are too far apart on 4-OV

² where one seeks the length of a longest shortest path from a vertex of S to a vertex of T

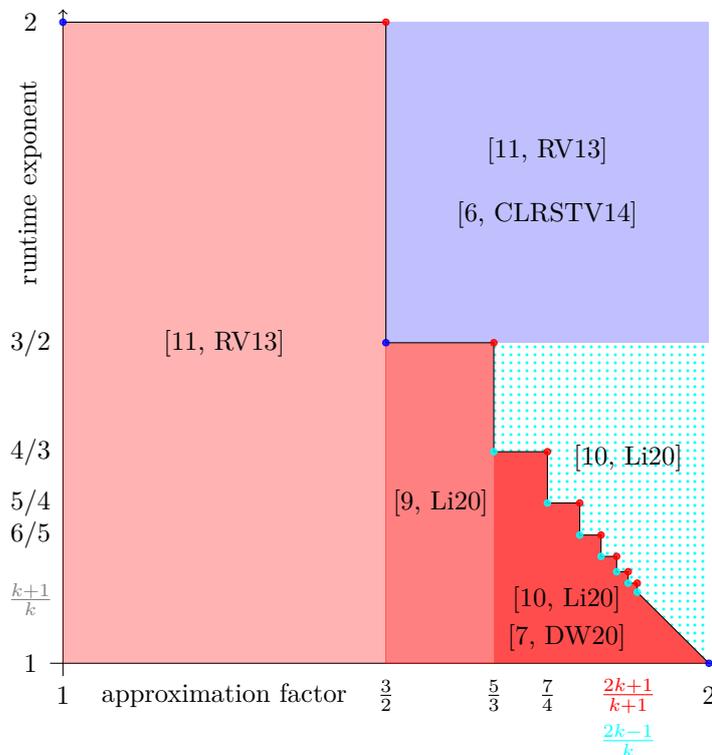
17:4 Inapproximability of Diameter in Super-Linear Time: Beyond the 5/3 Ratio

NO-instances. We thus add two “gates” u and v and link them with weighted arcs to the rest of the graph. This puts many pairs of vertices at distance at most 4 regardless on whether the 4-OV instance is positive or negative.

Of course, we cannot do so for *all* the pairs outside $L_0 \times L_k$. For instance, we do not want the distance from every vertex in ABC to every vertex in AD_Y to be always at most 4. Indeed, that would make the longest path from ABC to DCB of length at most 6. Our main novel contribution is to add a vertex set AD_X (L'_2) only linked to AD_Y (L_2) in a way that gives enough flexibility to make the remaining pairs sufficiently close for 4-OV NO-instances, while not decreasing the distance from x to y .

Recent developments

While the paper was under review, some exciting developments happened. Independently, Wein and Dalirrooyfard [7] and Li [10] showed that, under the SETH, for every integer $k \geq 2$ and real $\delta > 0$, any $(\frac{2k-1}{k} - \delta)$ -approximation for sparse unweighed directed DIAMETER requires time $n^{1+\frac{1}{k-1}-o(1)}$.



■ **Figure 2** The new results for sparse unweighed directed DIAMETER. The blue region is feasible, as witnessed by an algorithm at the bottom-left corner (blue dot). The red regions would refute the SETH, as witnessed by reductions at top-right corners (red dots). A red dot in the interior of the dotted cyan region would refute the NSETH. The lower bounds in [11, 9] even hold for the sparse unweighed undirected DIAMETER.

In the same paper, Li gives a piece of evidence that this could be as far as the hardness of unweighed directed DIAMETER goes. This evidence is based on the NSETH (for Non-deterministic SETH), a strengthening of SETH introduced by Carmosino et al. [5]. NSETH asserts that for every $\varepsilon > 0$, there is an integer k such that the k -TAUT problem cannot

be solved in *non-deterministic* time $(2 - \varepsilon)^n$, where k -TAUT asks, given a k -DNF formula whether every truth assignment satisfies it. Li shows, for all four variants of DIAMETER *but* the weighted directed one that improving on any of these (deterministic) SETH lower bounds would refute the NSETH (see dotted cyan region in Figure 2).

The construction of Dalirrooyfard and Wein [7] works for $k \geq 5$. It makes a more intricate use of “parallel layers” (such as AD_X). For $k = 2$ and 3 , the authors cite the existing lower bounds for unweighted undirected DIAMETER, and for $k = 4$ they tune the construction presented in this paper, in order to remove the edge weights. Namely, there is a simpler and weight-free way of connecting the “gates” u, v to the rest of the graph. The construction of Li [10] effectively combines *index-changing* “back” edges (which re-implements and extends the *skew edges* of this paper) with the usual *vector-changing* “forward” edges.

These generalizations crucially rely on edge orientations, while in the particular case of $k = 4$, our construction seems to only accidentally require arcs. In a subsequent work [3], we show how to obtain the lower bound for $k = 4$ in the most constrained case of unweighted *undirected* DIAMETER. This makes a non-trivial use of additional sets of vertices without “vector fields”.

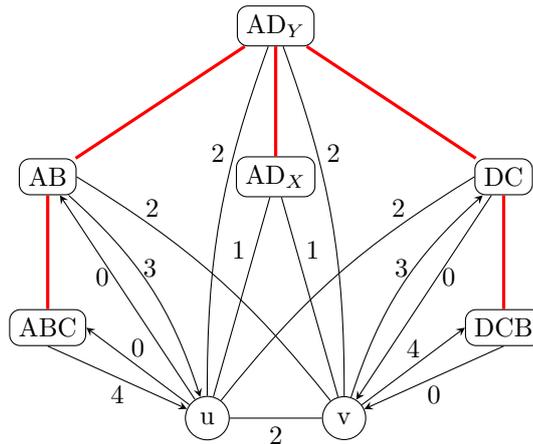
Preliminaries

We use the standard graph-theoretic notations. If G is a graph, $V(G)$ denotes its vertex set. If $S \subseteq V(G)$, $G[S]$ denotes the subgraph of G induced by S , and $G - S$ is a short-hand for $G[V(G) \setminus S]$. For $u, v \in V(G)$, $d_G(u, v)$ denotes the distance from u to v in G , that is the length of a shortest path from u to v , or equivalently, the minimum sum of weights on the edges on a path from u to v . Note that, in a directed graph, $d_G(u, v)$ and $d_G(v, u)$ may well be different values. We drop the subscript, if the graph G is clear from the context. We denote by $\text{diam}(G)$ the diameter of G , that is, $\max_{u, v \in V(G)} d_G(u, v)$. Note that both the pairs (u, v) and (v, u) are considered in this maximum. If ℓ is positive integer, $[\ell]$ denotes the set $\{1, 2, \dots, \ell\}$.

2 Reduction from 4-Orthogonal Vectors to 4 vs 7 Diameter

For every fixed positive integer k , the k -ORTHOGONAL VECTORS (k -OV for short) problem is as follows. It asks, given a set S of 0,1-vectors in $\{0, 1\}^\ell$, if there are k vectors $v_1, \dots, v_k \in S$ such that for every $i \in [\ell]$, $\prod_{h \in [k]} v_h[i] = 0$ or equivalently that $v_1[i] = v_2[i] = \dots = v_k[i] = 1$ does not hold. Williams [13] showed that, assuming the SETH, k -OV requires $N^{k-o(1)}$ time with $N := |S|$. Here we will leverage this lower bound for $k = 4$. This is a usual opening step: for example, Roditty and Vassilevska Williams [11] uses this lower bound for $k = 2$, and Li [9] uses it for $k = 3$.

From any set S of N vectors in $\{0, 1\}^\ell$, we build a directed weighted graph $G := \rho(S)$ (without negatively-weighted arcs) with $O(N^3 + N^2\ell^3)$ vertices and $O(N^3\ell^3 + N^2\ell^6)$ arcs such that if S admits an orthogonal quadruple then the diameter of G is (at least) 7, whereas if S has no orthogonal quadruple then the diameter of G is (at most) 4. There is a large enough constant c such that 4-OV requires $N^{4-o(1)}$ time, unless the SETH fails, even when $\ell = c \lceil \log N \rceil$ [13]. In that case, the graph G has $O(N^3)$ vertices and $\tilde{O}(N^3)$ edges. Hence any algorithm approximating sparse, weighted, directed DIAMETER within ratio better than $7/4$ in time $n^{4/3-\delta}$, with $\delta > 0$, would refute the SETH.



■ **Figure 3** The part of the reduction *not* depending on the 4-OV instance. The edges represented without arrow are double-arcs with the indicated weight. The black arc between, say, ABC and u , symbolizes that every vertex of ABC is linked by an arc of weight 4 to vertex u . Thick red edges represent *some* double-arcs of weight 1. Not every double-arc (or edge) is present between two sets linked by a red edge. This will be specified in the rest of the construction.

2.1 Constant part

We start by describing the part of the construction which does *not* depend on the 4-OV instance. Its purpose is to make many pairs of vertices at distance at most 4 regardless on whether the 4-OV instance is positive or negative. The vertex set of the eventually-built graph G consists of two special vertices u and v , and six (disjoint) sets ABC, AB, AD_X , AD_Y , DC, and DCB. Vertices u and v are unconditionally linked to these sets (and to each other) by weighted arcs as specified in Figure 3. As we wrote in the introduction, there is a simpler way, that does not require edge weights, of realizing the constant part (see [7, Section 6]). We keep our construction for the sake of consistency but invite the reader to have a look at [7, Figure 9].

In this figure, a black arc between a vertex $x \in \{u, v\}$ and a set $Z \in \{ABC, AB, AD_X, AD_Y, DC, DCB\}$ (or vice versa) indicates that x is linked to *every* vertex of Z by such an arc. Note that edges represented without arrow are double-arcs. Double-arcs will sometimes simply be called *edges*. The only arcs *not* incident to $\{u, v\}$ are edges (double-arcs) of weight 1. These edges are only present between ABC and AB, AB and AD_Y , AD_X and AD_Y , AD_Y and DC, and finally DC and DCB. We will describe them later. At this point, one just needs to know that every vertex in ABC (resp. DCB) has at least one neighbor in AB (resp. DC), and that these edges have weight 1.

We check that many pairs of vertices are at distance at most 4 (even without the knowledge of the edges symbolized in red). For the sake of conciseness, when we write, say, “ $\mathbf{u} \leftrightarrow \mathbf{DCB}$ ”, we intend to provide paths from u to every vertex in DCB, and from every vertex in DCB to u . Similarly, the paragraph “ $\mathbf{DCB} \rightarrow \mathbf{ABC}$ ” gives a path (of length 2) from every vertex of DCB to every vertex of ABC.

$\mathbf{u} \leftrightarrow \mathbf{v}$. There is an edge of weight 2 between u and v .

$\mathbf{u} \leftrightarrow \mathbf{Z} \in \{ABC, AB, AD_X, AD_Y, DC\}$. There are double-arcs with weight at most 4 between these pairs.

$\mathbf{u} \leftrightarrow \mathbf{DCB}$. There is a path of double-arcs of total weight 3: From a vertex of DCB, take any (weight-1) edge to DC, followed by the weight-2 edge to u . Recall that every vertex in DCB will have at least one neighbor in DC (via a weight-1 edge).

The next two cases are symmetric.

$\mathbf{v} \leftrightarrow \mathbf{Z} \in \{\mathbf{AB}, \mathbf{AD}_X, \mathbf{AD}_Y, \mathbf{DC}, \mathbf{DCB}\}$. There are double-arcs with weight at most 4 between these pairs.

$\mathbf{v} \leftrightarrow \mathbf{ABC}$. There is a path of double-arcs of total weight 3: From a vertex of \mathbf{ABC} , take any (weight-1) edge to \mathbf{AB} , followed by the weight-2 edge to v . Recall that every vertex in \mathbf{ABC} will have at least one neighbor in \mathbf{AB} (via a weight-1 edge).

So far, we have seen that for each $x \in \{u, v\}$ and $y \in V(G)$, $d(x, y) \leq 4$ and $d(y, x) \leq 4$.

For each $\mathbf{Z} \in \{\mathbf{ABC}, \mathbf{AB}, \mathbf{AD}_X, \mathbf{AD}_Y, \mathbf{DC}, \mathbf{DCB}\}$, $\mathbf{Z} \leftrightarrow \mathbf{Z}$. Each of these six sets Z has a double-arc to u or to v (or both) whose sum of weights is at most 4.

$\mathbf{AD}_X \rightarrow \mathbf{Z} \in \{\mathbf{ABC}, \mathbf{AB}, \mathbf{AD}_Y, \mathbf{DC}\}$. These pairs are at distance at most 3. There is an edge of weight 1 from every vertex of \mathbf{AD}_X to u , and an arc of weight at most 2 from u to every vertex of $Z \in \{\mathbf{ABC}, \mathbf{AB}, \mathbf{AD}_Y, \mathbf{DC}\}$.

$\mathbf{AD}_X \rightarrow \mathbf{DCB}$. There is a path of length 4, via u and \mathbf{DC} . Again recall that every vertex of \mathbf{DCB} has at least one neighbor in \mathbf{DC} (via a double-arc of weight 1).

The next two cases are symmetric.

$\mathbf{AD}_X \leftarrow \mathbf{Z} \in \{\mathbf{AB}, \mathbf{AD}_Y, \mathbf{DC}, \mathbf{DCB}\}$. These pairs are at distance at most 3. There is an arc of weight at most 2 from every vertex of $Z \in \{\mathbf{AB}, \mathbf{AD}_Y, \mathbf{DC}, \mathbf{DCB}\}$ to v , and an edge of weight 1 from v to every vertex of \mathbf{AD}_X .

$\mathbf{AD}_X \leftarrow \mathbf{ABC}$. There is path of length 4, via \mathbf{AB} and v .

We have now established that for every $x \in \mathbf{AD}_X$ and $y \in V(G)$, $d(x, y) \leq 4$ and $d(y, x) \leq 4$.

$\mathbf{AD}_Y \leftrightarrow \mathbf{AB}$. There is a path of two double-arcs of weight 2, via v .

$\mathbf{AD}_Y \leftrightarrow \mathbf{DC}$. There is a path of two double-arcs of weight 2, via u .

$\mathbf{AD}_Y \rightarrow \mathbf{ABC}$. There is a path of length 2 via u .

$\mathbf{AD}_Y \leftarrow \mathbf{DCB}$. There is a path of length 2 via v .

$\mathbf{ABC} \leftrightarrow \mathbf{AB}$. Via u , there is a path of length 4 from every vertex of \mathbf{ABC} to every vertex of \mathbf{AB} , and a path of length 3 from every vertex of \mathbf{AB} to every vertex of \mathbf{ABC} .

$\mathbf{DCB} \leftrightarrow \mathbf{DC}$. Via v , there is a path of length 3 from every vertex of \mathbf{DCB} to every vertex of \mathbf{DC} , and a path of length 4 from every vertex of \mathbf{DC} to every vertex of \mathbf{DCB} .

$\mathbf{DCB} \rightarrow \mathbf{ABC}$. There is a path of length 2 via v and u .

$\mathbf{DCB} \rightarrow \mathbf{AB}$. There is a path of length 2 via v .

$\mathbf{DC} \rightarrow \mathbf{Z} \in \{\mathbf{ABC}, \mathbf{AB}\}$. There is a path of length 2 via u .

To summarize, we have obtained that for every pair $x, y \in V(G)$, $d(x, y) > 4$ implies that $(x, y) \in P$ for some $P \in \{\mathbf{ABC} \times \mathbf{AD}_Y, \mathbf{ABC} \times \mathbf{DC}, \mathbf{ABC} \times \mathbf{DCB}, \mathbf{AB} \times \mathbf{DC}, \mathbf{AB} \times \mathbf{DCB}, \mathbf{AD}_Y \times \mathbf{DCB}\}$.

2.2 Variable part

We think of the vector set S as having four copies A, B, C, D with $S = A = B = C = D$. Equivalently, one looks for $a \in A$, $b \in B$, $c \in C$, and $d \in D$ such that a, b, c, d are orthogonal.

Vertex set

We first describe the vertex set $V(G) \setminus \{u, v\}$.

- For every $(a, b, c) \in A \times B \times C$, we add vertex $(a, b, c)_{ABC}$ to ABC.
- Similarly for every $(d, c, b) \in D \times C \times B$, we add vertex $(d, c, b)_{DCB}$ to DCB.
- For every $(a, b) \in A \times B$ and every triple $i, j, k \in [\ell]$ such that $a[i] = a[j] = a[k] = 1$ and b takes value 1 on at least two indices of $\{i, j, k\}$, we add vertex $(a, b, i, j, k)_{AB}$ to AB.
- The set of vertices DC is defined analogously with D and C playing the roles of A and B (recall that actually $A = B = C = D$).
- For every $(a, d) \in A \times D$ and every triple of indices $i, j, k \in [\ell]$ such that $a[i] = a[j] = a[k] = 1$ and $d[i] = d[j] = d[k] = 1$, we add vertex $(a, d, i, j, k)_{AD_Y}$ to AD_Y .
- For every $(a, d) \in A \times D$ and every $i, j, k \in [\ell]$ such that at most one of $a[i], a[j], a[k], d[i], d[j], d[k]$ is equal to 0, we add vertex $(a, d, i, j, k)_{AD_X}$ to AD_X .

As observed in [7], the definition of AD_X could be simpler. We keep it as is, again for the sake of consistency.

Edge Set

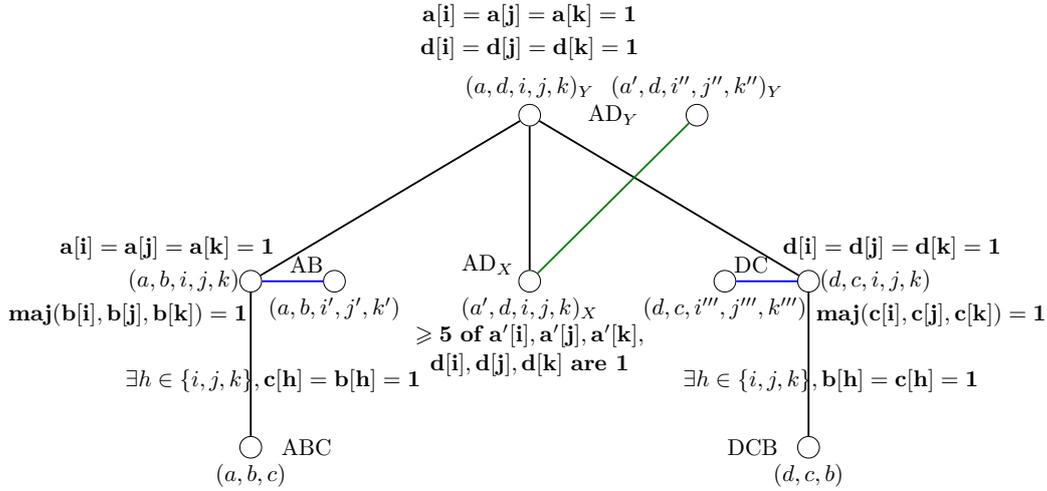
We now describe the edge set on $V(G) \setminus \{u, v\}$. All these edges are double-arcs (we do not need to orient them) of weight 1 (we do not need to put weights).

- We add an edge (double-arc) of weight 1 between every pair $(a, b, c)_{ABC}$ and $(a, b, i, j, k)_{AB}$ if c takes value 1 on at least one index of $\{i, j, k\}$ where b also takes value 1. In our construction, the existence of an edge is implicitly conditional to the existence of both of its endpoints. The edge exists only if $(a, b, i, j, k)_{AB}$ is indeed a vertex of AB. The edges between DCB and DC are defined similarly.
- We add every edge between $(a, b, i, j, k)_{AB}$ and $(a, b, i', j', k')_{AB}$, with $a \in A, b \in B$, and $i, j, k, i', j', k' \in [\ell]$. Similarly we add every edge between $(d, c, i, j, k)_{DC}$ and $(d, c, i', j', k')_{DC}$, with $d \in D, c \in C$, and $i, j, k, i', j', k' \in [\ell]$. We call these edges *index-switching*. AB and DC are the only two sets among $\{ABC, AB, AD_X, AD_Y, DC, DCB\}$ which are *not* independent sets.
- We link every pair $(a, b, i, j, k)_{AB}$ and $(a, d, i, j, k)_{AD_Y}$ by an edge, as well as every pair $(a, d, i, j, k)_{AD_Y}$ and $(d, c, i, j, k)_{DC}$.
- We add an edge between every pair $(a, d, i, j, k)_{AD_X}$ and $(a, d', i, j, k)_{AD_Y}$ (such that $a \in A, d \neq d' \in D$, and $i, j, k \in [\ell]$), and $(a, d, i, j, k)_{AD_X}$ and $(a', d, i, j, k)_{AD_Y}$ (such that $a \neq a' \in A, d \in D$, and $i, j, k \in [\ell]$).
- Finally we add an edge between every pair $(a, d, i, j, k)_{AD_X}$ and $(a, d, i', j', k')_{AD_Y}$ with $a \in A, d \in D$, and $i, j, k, i', j', k' \in [\ell]$. We call this type of edge *skew*. Skew edges are the only way to change the indices while also moving from one set to another (it is not internal to AB or to DC).

All the edges defined in this section that are not index-switching or skew are called *regular*. This ends the construction of $G = \rho(S)$. See Figure 4 for an illustration.

Simplified vertex notations, vector fields, and index fields

Henceforth we will drop the indices in the vertex labels. The set a vertex belongs to will be implicit by the choice of the variable labels. For instance (a, b', i, j, k) is in AB, and (d', c, b') is in DCB. This does not allow us to distinguish vertices of AD_X and AD_Y . We will denote



■ **Figure 4** The rules for the existence of vertices and edges in $G - \{u, v\}$. We removed the subscripts in the vertex labels as discussed in the second-to-last paragraph of Section 2.2. Conditions to the existence of a vertex appear in bold next to the vertex. Conditions to the existence of an edge appear in bold along the edge. All edges (double-arcs) have weight 1, so we omit their weight. Regular edges are represented in black. Index-switching edges are represented in blue (they are only present in AB and DC). Skew edges are represented in green (they are only present between AD_X and AD_Y). Note that the regular edge $(a, d, i, j, k)_Y(a', d, i, j, k)_X$ could also be of the form $(a, d, i, j, k)_Y(a', d', i, j, k)_X$.

by $(a, d, i, j, k)_Y = (a, d, i, j, k)_{AD_Y}$ a vertex in AD_Y and by $(a, d, i, j, k)_X$ the “same” vertex in AD_X . Note that it is possible that $(a, d, i, j, k)_X$ exists but not $(a, d, i, j, k)_Y$, if exactly five of $a[i], a[j], a[k], d[i], d[j], d[k]$ are equal to 1.

We call *vector fields* the first three coordinates of every vertex in $ABC \cup DCB$, and the first two coordinates of every vertex in $AB \cup AD_X \cup AD_Y \cup DC$. We call *index fields* the last three coordinates of every vertex in $AB \cup AD_X \cup AD_Y \cup DC$. We can assume that the 4-OV instance does not have an orthogonal *triple* (this can be checked in time $\tilde{O}(N^3)$). Thus every vertex $(a, b, c) \in ABC$ (resp. $(d, c, b) \in DCB$) indeed has at least one neighbor in AB (resp. DC), namely (a, b, i, i, i) (resp. (d, c, i, i, i)) where $i \in [\ell]$ is such that $a[i] = b[i] = c[i] = 1$ (resp. $d[i] = c[i] = b[i] = 1$).

Vertex and edge count

Before tackling the correctness of the reduction, we check that G has $O(N^3)$ vertices and $\tilde{O}(N^3)$ arcs. The number of vertices of G is bounded by $2 + 2 \cdot N^3 + 4 \cdot N^2 \ell^3 = O(N^3)$ since $\ell = O(\log N)$. The number of arcs of G is bounded by $4 \cdot O(N^3) + 4 \cdot N^2 \ell^6 + 10 \cdot N^3 \ell^3 + 2 \cdot N^2 \ell^6 = \tilde{O}(N^3)$, where the first term accounts for the arcs incident with $\{u, v\}$, the second for the index-switching arcs, the third for the regular arcs, and the fourth for the skew arcs.

2.3 No orthogonal quadruple implies diameter at most 4

We exhibit in this section short paths (of length at most 4) between every pair of vertices of G . For that we extensively use that, as there is no orthogonal quadruple in S , for every $u, v, w, x \in S$, $\text{ind}(u, v, w, x) := \min\{i \in [\ell] \mid u[i] = v[i] = w[i] = x[i] = 1\}$ is a well-defined index in $[\ell]$. We only take the minimum index to have a deterministic notation. There will

17:10 Inapproximability of Diameter in Super-Linear Time: Beyond the 5/3 Ratio

not be anything particular with the minimum, and any index of the non-empty $\{i \in [\ell] \mid u[i] = v[i] = w[i] = x[i] = 1\}$ would work as well. We will also use $\text{ind}(u, v, w)$ as a short-hand for $\text{ind}(u, v, w, w)$.

In Section 2.1, we have reduced the task of showing that $\text{diam}(G) \leq 4$ to considering only six pairs of sets.

ABC \rightarrow AD_Y. Let (a, b, c) and $(a', d, i', j', k')_Y$ be two vertices in ABC and AD_Y respectively. We define the indices $i := \text{ind}(a, b, c, d)$, $j := \text{ind}(a, a', b, d)$, and $k := \text{ind}(a, a', d)$. Then, $(a, b, c) \rightarrow (a, b, i, j, k) \rightarrow (a, d, i, j, k)_Y \rightarrow (a', d, i, j, k)_X \rightarrow (a', d, i', j', k')_Y$ is a path of length 4 in G .

We first justify the existence of the inner vertices of this path (i.e., all but the endpoints). Indeed the endpoints exist by assumption. Vertex $(a, b, i, j, k) \in \text{AB}$ is present in G since $a[i] = a[j] = a[k] = 1$, and $b[i] = b[j] = 1$. Vertex $(a, d, i, j, k)_Y \in \text{AD}_Y$ exists since $a[i] = a[j] = a[k] = 1 = d[i] = d[j] = d[k]$. Finally $(a', d, i, j, k)_X \in \text{AD}_X$ is indeed a vertex of G since $a'[j] = a'[k] = 1$ and $d[i] = d[j] = d[k] = 1$. Recall that in AD_X (contrary to AD_Y) it is fine if at most one of the six values obtained by evaluating one of the two vectors at one of the three indices is 0.

We now justify the existence of the edges. The arc $(a, b, c) \rightarrow (a, b, i, j, k)$ exists since $c[i] = b[i] = 1$ (and both its endpoints exist). The arcs $(a, b, i, j, k) \rightarrow (a, d, i, j, k)_Y$ and $(a, d, i, j, k)_Y \rightarrow (a', d, i, j, k)_X$ are regular edges of G : one vector field and the three index fields remain unchanged. Finally the arc $(a', d, i, j, k)_X \rightarrow (a', d, i', j', k')_Y$ is a skew edge of G : it is between AD_X and AD_Y, and both vector fields remain the same (while the indices are allowed to change).

ABC \rightarrow DC. Let (a, b, c) and (d, c', i', j', k') be two vertices in ABC and DC respectively. We define the indices $i := \text{ind}(a, b, c, d)$, $j := \text{ind}(a, b, c', d)$, and $k := \text{ind}(a, c', d)$. Then, $(a, b, c) \rightarrow (a, b, i, j, k) \rightarrow (a, d, i, j, k)_Y \rightarrow (d, c', i, j, k) \rightarrow (d, c', i', j', k')$ is a path of length 4 in G .

As in the previous case, the existence of vertices $(a, b, i, j, k) \in \text{AB}$, $(a, d, i, j, k)_Y \in \text{AD}_Y$, $(d, c', i, j, k) \in \text{DC}$ is ensured by the fact that $a[i] = a[j] = a[k] = 1 = d[i] = d[j] = d[k]$ and $b[i] = b[j] = 1 = c'[j] = c'[k]$. The arc $(a, b, c) \rightarrow (a, b, i, j, k)$ exists since $c[i] = b[i] = 1$, and the arcs $(a, b, i, j, k) \rightarrow (a, d, i, j, k)_Y \rightarrow (d, c', i, j, k)$ are two (existing) regular arcs. Finally $(d, c', i, j, k) \rightarrow (d, c', i', j', k')$ is an index-switching arc internal to DC (note that the two vector fields remain the same, as they should).

ABC \rightarrow DCB. Let (a, b, c) and (d, c', b') be two vertices in ABC and DCB respectively. We define the indices $i := \text{ind}(a, b, c, d)$, $j := \text{ind}(a, b, c', d)$, and $k := \text{ind}(a, b', c', d)$. Then, $(a, b, c) \rightarrow (a, b, i, j, k) \rightarrow (a, d, i, j, k)_Y \rightarrow (d, c', i, j, k) \rightarrow (d, c', b')$ is a path of length 4 in G .

The vertices $(a, b, i, j, k) \in \text{AB}$, $(a, d, i, j, k)_Y \in \text{AD}_Y$, $(d, c', i, j, k) \in \text{DC}$ exist since $a[i] = a[j] = a[k] = 1 = d[i] = d[j] = d[k]$ and $b[i] = b[j] = 1 = c'[j] = c'[k]$. The arc $(a, b, c) \rightarrow (a, b, i, j, k)$ is in G since $c[i] = b[i] = 1$. The arcs $(a, b, i, j, k) \rightarrow (a, d, i, j, k)_Y \rightarrow (d, c', i, j, k)$ are two regular arcs in G . The arc $(d, c', i, j, k) \rightarrow (d, c', b')$ exists since $b'[k] = c'[k] = 1$.

AB \rightarrow DC. Let (a, b, i', j', k') and (d, c, i'', j'', k'') be two vertices in AB and DC respectively. We define the index $i := \text{ind}(a, b, c, d)$. Then, $(a, b, i', j', k') \rightarrow (a, b, i, i, i) \rightarrow (a, d, i, i, i)_Y \rightarrow (d, c, i, i, i) \rightarrow (d, c, i'', j'', k'')$ is a path of length 4 in G .

Vertices $(a, b, i, i, i) \in \text{AB}$, $(a, d, i, i, i)_Y \in \text{AD}_Y$, $(d, c, i, i, i) \in \text{DC}$ exist since $a[i] = b[i] = c[i] = d[i] = 1$. The arc $(a, b, i', j', k') \rightarrow (a, b, i, i, i)$ is a legal index-switching arc, internal to AB. The arcs $(a, b, i, i, i) \rightarrow (a, d, i, i, i)_Y \rightarrow (d, c, i, i, i)$ are two regular arcs in G . Finally the arc $(d, c, i, i, i) \rightarrow (d, c, i'', j'', k'')$ is an index-switching arc, internal to DC.

The next two cases are symmetric to $\mathbf{ABC} \rightarrow \mathbf{DC}$ and $\mathbf{ABC} \rightarrow \mathbf{AD}_Y$, respectively. We spell them out since it is not much longer than making the symmetry explicit.

AB \rightarrow DCB. Let (a, b, i', j', k') and (d, c, b') be two vertices in AB and DCB respectively. We define the indices $i := \text{ind}(a, b, c, d)$, $j := \text{ind}(a, b, d)$, and $k := \text{ind}(a, b', c, d)$. Then, $(a, b, i', j', k') \rightarrow (a, b, i, j, k) \rightarrow (a, d, i, j, k)_Y \rightarrow (d, c, i, j, k) \rightarrow (d, c, b')$ is a path of length 4 in G .

Vertices $(a, b, i, j, k) \in \mathbf{AB}$, $(a, d, i, j, k)_Y \in \mathbf{AD}_Y$, $(d, c, i, j, k) \in \mathbf{DC}$ exist since $a[i] = a[j] = a[k] = 1 = d[i] = d[j] = d[k]$ and $b[i] = b[j] = 1 = c[i] = c[k]$. The arc $(a, b, i', j', k') \rightarrow (a, b, i, j, k)$ is index-switching in AB. The arcs $(a, b, i, j, k) \rightarrow (a, d, i, j, k)_Y \rightarrow (d, c, i, j, k)$ are two regular arcs present in G . Finally the arc $(d, c, i, j, k) \rightarrow (d, c, b')$ exists since $b'[k] = c[k] = 1$.

AD_Y \rightarrow DCB. Let $(a, d, i', j', k')_Y$ and (d', c, b) be two vertices in AD_Y and DCB respectively. We define the indices $i := \text{ind}(a, b, c, d')$, $j := \text{ind}(a, c, d, d')$, and $k := \text{ind}(a, d, d')$. Then, $(a, d, i', j', k')_Y \rightarrow (a, d, i, j, k)_X \rightarrow (a, d', i, j, k)_Y \rightarrow (d', c, i, j, k) \rightarrow (d', c, b)$ is a path of length 4 in G .

The vertices $(a, d, i, j, k)_X \in \mathbf{AD}_X$, $(a, d', i, j, k)_Y \in \mathbf{AD}_Y$, $(d', c, i, j, k) \in \mathbf{DC}$ are present in G since $a[i] = a[j] = a[k] = 1 = d'[i] = d'[j] = d'[k]$ and $d[j] = d[k] = 1 = c[i] = c[j]$. Recall that $a[i] = a[j] = a[k] = 1 = d[j] = d[k]$ suffices for the existence of $(a, d, i, j, k)_X$ (but not for the one of $(a, d, i, j, k)_Y$). The arc $(a, d, i', j', k')_Y \rightarrow (a, d, i, j, k)_X$ is a skew arc: it is between AD_Y and AD_X, and the two vector fields remain unchanged. The arcs $(a, d, i, j, k)_X \rightarrow (a, d', i, j, k)_Y \rightarrow (d', c, i, j, k)$ are two regular arcs of G . Finally the arc $(d', c, i, j, k) \rightarrow (d', c, b)$ exists since $b[i] = c[i] = 1$.

We have proved that there is a path of length at most 4 between every (ordered) pair of vertices in G , when there is no orthogonal quadruple. Thus the diameter of G is then (at most) 4.

2.4 An orthogonal quadruple implies two vertices at distance at least 7

We now suppose that S admits at least one orthogonal quadruple, say, a, b, c, d . We show that G has diameter at least 7, by arguing that there is no path of length at most 6 from $(a, b, c) \in \mathbf{ABC}$ to $(d, c, b) \in \mathbf{DCB}$.

The first observation is that there is no path of length at most 6 from (a, b, c) to (d, c, b) intersecting $\{u, v\}$. Indeed one can check that $d((a, b, c), u) = 4$ and $d(u, (d, c, b)) = 3$, and that $d((a, b, c), v) = 3$ and $d(v, (d, c, b)) = 4$. We can now rule out the existence of a path P of length 6 from (a, b, c) and (d, c, b) in $G - \{u, v\}$.

We distinguish two cases:

- (a) P does not intersect AD_X, or
- (b) P intersects AD_X.

Case (a). We further distinguish two cases: either (a1) P contains no index-switching arc, or (a2) P contains at least one index-switching arc. In case (a1), the three index fields cannot change at all in P (recall that the skew edges are between AD_X and AD_Y). Thus the first and penultimate vertices of P are $(a, b, i, j, k) \in \mathbf{AB}$ and $(d, c, i, j, k) \in \mathbf{DC}$ for some $i, j, k \in [\ell]$. The existence of these vertices imply that $a[i] = a[j] = a[k] = 1 = d[i] = d[j] = d[k]$, and vectors b and c both take value 1 on at least two indices among $\{i, j, k\}$. Therefore there exists an index $h \in \{i, j, k\}$ such that $a[h] = b[h] = c[h] = d[h] = 1$. This contradicts the fact that a, b, c, d are orthogonal.

17:12 Inapproximability of Diameter in Super-Linear Time: Beyond the 5/3 Ratio

We now tackle case (a2). Since the removal of AD_Y separates $ABC \cup AB$ from $DCB \cup DC$ in $G - \{u, v\}$, $G[AD_Y]$ is an independent set, and there are no edges between AD_Y and $ABC \cup DCB$, such a path P has to contain a subpath $x \rightarrow y \rightarrow z$ with $x \in AB$, $y \in AD_Y$, and $z \in DC$. As P contains at least one index-switching arc, it cannot also contain a back-and-forth along $AB \rightarrow ABC \rightarrow AB$, $AD_Y \rightarrow AB \rightarrow AD_Y$, $DC \rightarrow AD_Y \rightarrow DC$, or $DCB \rightarrow DC \rightarrow DCB$. Indeed that would amount to at least three additional arcs (at least one index-switching plus two for the back-and-forth) to the mandatory four arcs $ABC \rightarrow AB \rightarrow AD_Y \rightarrow DC \rightarrow DCB$, hence a path of length at least 7. Therefore, there are three indices $i, j, k \in [\ell]$ such that $x = (a, b, i, j, k)$, $y = (a, d, i, j, k)_Y$, and $z = (d, c, i, j, k)$. Indeed every path in $G[ABC \cup AB]$ and every path in $G[DCB \cup DC]$ preserve the first two vector fields. Again the existence of (a, b, i, j, k) (in AB) and (d, c, i, j, k) (in DC) contradicts that a, b, c, d are orthogonal.

Case (b). We can now assume that P intersects AD_X . Thus P has length exactly 6 and is of the form $(a, b, c) \in ABC \rightarrow AB \rightarrow AD_Y \rightarrow AD_X \rightarrow AD_Y \rightarrow DC \rightarrow DCB \ni (d, c, b)$. In particular, P cannot contain an index-switching edge. If P contains no skew edge too, the index fields cannot change. So the second and sixth vertices of P are some $(a, b, i, j, k) \in AB$ and $(d, c, i, j, k) \in DC$, and we can conclude as in case (a).

Thus P has to contain at least one skew edge. Let us show that P has to contain exactly one skew edge. We recall that the skew edges are only present between AD_X and AD_Y .

We first argue that the third vertex of P is $(a, d', i, j, k)_Y$ for some $d' \neq d \in D$ and $i, j, k \in [\ell]$. The first vector field cannot change in a path of the form $ABC \rightarrow AB \rightarrow AD_Y$, so we only have to show that d' cannot be equal to d . Indeed, otherwise $a[i] = a[j] = a[k] = 1 = d[i] = d[j] = d[k]$ by the existence of $(a, d, i, j, k)_Y$. Furthermore, the existence of the arc $(a, b, c) \rightarrow (a, b, i, j, k)$ (which has to be the first arc of P) implies that there is an index $h \in \{i, j, k\}$ such that $b[h] = c[h] = 1$. This index thus contradicts the orthogonality of a, b, c, d .

As the third vertex of P is $(a, d', i, j, k)_Y$ with $d' \neq d$, two skew edges $AD_Y \rightarrow AD_X \rightarrow AD_Y$ would lead to a vertex $(a, d', i', j', k')_Y$. This latter vertex is linked in DC to vertices of the form (d', c', i', j', k') (where c' can be c). This cannot lead to (d, c, b) since the first vector field does not change in an arc from DC to DCB .

We have established that from vertex $(a, d', i, j, k)_Y$, P takes exactly one skew edge, either from AD_Y to AD_X , or from AD_X to AD_Y . In both cases, by the previous remark, the second vector field of the fifth vertex of P should be d . This implies that the fifth vertex of P is of the form $(a, d, i', j', k')_Y$ for some indices $i', j', k' \in [\ell]$. Indeed the skew edge of P in $AD_Y \rightarrow AD_X \rightarrow AD_Y$ preserves both vector fields, whereas the regular edge of P in $AD_Y \rightarrow AD_X \rightarrow AD_Y$ can only change one vector field, and has to change d' ($\neq d$) to d .

The end of P is thus $(a, d, i', j', k')_Y \rightarrow (d, c, i', j', k') \rightarrow (d, c, b)$, since the first two vector fields cannot be changed by an arc from DC to DCB . The existence of vertex $(a, d, i', j', k')_Y$ implies that $a[i'] = a[j'] = a[k'] = 1 = d[i'] = d[j'] = d[k']$. The existence of the arc $(d, c, i', j', k') \rightarrow (d, c, b)$ implies that there is an index $h \in \{i', j', k'\}$ such that $c[h] = b[h] = 1$. This yields $a[h] = b[h] = c[h] = d[h] = 1$, contradicting the orthogonality of a, b, c, d .

We have ruled out the existence of a path in G of length at most 6 between (a, b, c) and (d, c, b) , when a, b, c, d are orthogonal. Hence the diameter of G is at least 7 when there is an orthogonal quadruple.

References

- 1 Donald Aingworth, Chandra Chekuri, Piotr Indyk, and Rajeev Motwani. Fast Estimation of Diameter and Shortest Paths (Without Matrix Multiplication). *SIAM J. Comput.*, 28(4):1167–1181, 1999. doi:10.1137/S0097539796303421.
- 2 Arturs Backurs, Liam Roditty, Gilad Segal, Virginia Vassilevska Williams, and Nicole Wein. Towards tight approximation bounds for graph diameter and eccentricities. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 267–280. ACM, 2018. doi:10.1145/3188745.3188950.
- 3 Édouard Bonnet. 4 vs 7 sparse undirected unweighted Diameter is SETH-hard at time $n^{4/3}$, 2021. arXiv:2101.02312.
- 4 Massimo Cairo, Roberto Grossi, and Romeo Rizzi. New Bounds for Approximating Extremal Distances in Undirected Graphs. In Robert Krauthgamer, editor, *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 363–376. SIAM, 2016. doi:10.1137/1.9781611974331.ch27.
- 5 Marco L. Carmosino, Jiawei Gao, Russell Impagliazzo, Ivan Mihajlin, Ramamohan Paturi, and Stefan Schneider. Nondeterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility. In Madhu Sudan, editor, *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, Cambridge, MA, USA, January 14-16, 2016*, pages 261–270. ACM, 2016. doi:10.1145/2840728.2840746.
- 6 Shiri Chechik, Daniel H. Larkin, Liam Roditty, Grant Schoenebeck, Robert Endre Tarjan, and Virginia Vassilevska Williams. Better Approximation Algorithms for the Graph Diameter. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 1041–1052. SIAM, 2014. doi:10.1137/1.9781611973402.78.
- 7 Mina Dalirrooyfard and Nicole Wein. Tight Conditional Lower Bounds for Approximating Diameter in Directed Graphs. *CoRR*, abs/2011.03892, 2020. arXiv:2011.03892.
- 8 Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which Problems Have Strongly Exponential Complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001. doi:10.1006/jcss.2001.1774.
- 9 Ray Li. Improved SETH-hardness of unweighted Diameter. *CoRR*, abs/2008.05106, 2020. arXiv:2008.05106.
- 10 Ray Li. Settling SETH vs. Approximate Sparse Directed Unweighted Diameter (up to (NU)NSETH). *CoRR*, abs/2008.05106, 2020. arXiv:2008.05106.
- 11 Liam Roditty and Virginia Vassilevska Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013*, pages 515–524. ACM, 2013. doi:10.1145/2488608.2488673.
- 12 Aviad Rubinfeld and Virginia Vassilevska Williams. SETH vs Approximation. *SIGACT News*, 50(4):57–76, 2019. doi:10.1145/3374857.3374870.
- 13 Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theor. Comput. Sci.*, 348(2-3):357–365, 2005. doi:10.1016/j.tcs.2005.09.023.
- 14 Virginia Vassilevska Williams. On some fine-grained questions in algorithms and complexity. In *Proceedings of the ICM*, volume 3, pages 3431–3472. World Scientific, 2018.