

# Shortest Paths in Multimode Graphs

Yael Kirkpatrick   

Massachusetts Institute of Technology, Cambridge, MA, USA

Virginia Vassilevska Williams   

Massachusetts Institute of Technology, Cambridge, MA, USA

---

## Abstract

In this work we study shortest path problems in multimode graphs, a generalization of the min-distance measure introduced by Abboud, Vassilevska W. and Wang in [SODA'16]. A multimode shortest path is the shortest path using one of multiple “modes” of transportation that cannot be combined. This represents real-world scenarios where different modes are not combinable, such as flights operated by different airline alliances. The problem arises naturally in machine learning in the context of learning with multiple embedding. More precisely, a  $k$ -multimode graph is a collection of  $k$  graphs on the same vertex set and the  $k$ -mode distance between two vertices is defined as the minimum among the distances computed in each individual graph.

We focus on approximating fundamental graph parameters on these graphs, specifically diameter and radius. In undirected multimode graphs we first show an elegant linear time 3-approximation algorithm for 2-mode diameter. We then extend this idea into a general subroutine that can be used as a part of any  $\alpha$ -approximation, and use it to construct a 2 and 2.5 approximation algorithm for 2-mode diameter. For undirected radius, we introduce a general scheme that can compute a 3-approximation of the  $k$ -mode radius for any  $k$  and runs in near linear time in the case of  $k = O(1)$ . In the directed case we establish an equivalence between approximating 2-mode diameter on DAGs and approximating the min-diameter, while for general graphs we develop novel techniques and provide a linear time algorithm to determine whether the diameter is finite.

We also develop many conditional fine-grained lower bounds for various multimode diameter and radius approximation problems. We are able to show that many of our algorithms are tight under popular fine-grained complexity hypotheses, including our linear time 3-approximation for 3-mode undirected diameter and radius. As part of this effort we propose the first extension to the Hitting Set Hypothesis [SODA'16], which we call the  $\ell$ -Hitting Set Hypothesis. We use this hypothesis to prove the first parameterized lower bound tradeoff for radius approximation algorithms.

**2012 ACM Subject Classification** Theory of computation  $\rightarrow$  Graph algorithms analysis

**Keywords and phrases** Graph Algorithms, Shortest Paths, Diameter, Radius, Fine-Grained Complexity

**Digital Object Identifier** 10.4230/LIPIcs.MFCS.2025.63

**Related Version** *Full Version*: <http://arxiv.org/abs/2506.22261>

**Funding** *Yael Kirkpatrick*: NSF Graduate Research Fellowship under Grant No. 2141064.

*Virginia Vassilevska Williams*: Supported by NSF Grant CCF-2330048, BSF Grant 2020356 and a Simons Investigator Award.

**Acknowledgements** The authors would like to thank Asaf Etgar, Jenny Kaufmann and Surya Mathialagan for helpful discussions.

## 1 Introduction

The study of shortest paths algorithms in directed graphs has introduced various distance measures. The most standard notion of distance between two vertices  $u$  and  $v$  is the minimum length of a path from  $u$  to  $v$ ,  $d(u, v)$ . This asymmetric one-way notion of distance is not always the most natural. In some applications the *roundtrip* distance  $d(u, v) + d(v, u)$ , defined by Cowen and Wagner [8], is most useful, since unlike its one-way counterpart, it is a metric.



© Yael Kirkpatrick and Virginia Vassilevska Williams;  
licensed under Creative Commons License CC-BY 4.0

50th International Symposium on Mathematical Foundations of Computer Science (MFCS 2025).

Editors: Paweł Gawrychowski, Filip Mazowiecki, and Michał Skrzypczak; Article No. 63; pp. 63:1–63:16

Leibniz International Proceedings in Informatics



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

In directed acyclic graphs (DAGs), however, both the standard notion of distance and the roundtrip notion produce infinite distances for all or a large fraction of the distance pairs. A more useful notion of distance in this case was defined by Abboud, Vassilevska W. and Wang [1]. In this work they introduced the *min-distance*, defined as  $d_{\min}(u, v) = \min\{d(u, v), d(v, u)\}$ . Under one-way distance or roundtrip distance, almost all points in a DAG have infinite eccentricity. Introducing min-distance opened an avenue of research which lead to intricate algorithms for computing or approximating min-diameter, radius and eccentricities in DAGs [1, 9] as well as for general graphs [3, 7, 12].

One real-world motivation for studying the min-distance problem is computing the optimal location for a hospital, as one has the choice of either going to the hospital or having a doctor come to them. The min-distance therefore measures the fastest way to receive care and the point with smallest min-eccentricity in the graph functions as the optimal location for a hospital. Imagine now that the doctor can take a helicopter, whereas a patient can only take a cab to the doctor. Then the two routes are no longer part of the same graph - the routes by which a helicopter can travel are very different from those available to a taxi. The new distance between the patient and the doctor is now the minimum between the distances in two different graphs.

Consider another common scenario. We want to fly from city  $A$  to city  $B$ . We can take any valid itinerary, except that each itinerary needs to stay within the same airline alliance. Now the fastest flight time is the minimum of the distances over several different “graphs”, one for each airline alliance. All of these graphs share a vertex set representing the different cities, while each graph has edges representing the flights of a single airline alliance.

Let us formalize a new shortest paths problem that captures both scenarios. Consider a collection of  $k$  graphs over the same vertex set  $V$ . Let their edge sets be  $E_1, \dots, E_k$ . Together they form a “ $k$ -multimode graph”  $\mathcal{G} = (V, E_1, \dots, E_k)$ . One can associate each  $E_i$  with a “mode of transportation”, where these modes are separate and cannot be combined.

The  $k$ -mode distance between two vertices  $u, v \in V$  in a  $k$ -multimode graph  $\mathcal{G} = (V, E_1, \dots, E_k)$  is defined as

$$d_{\mathcal{G}}(u, v) = \min_i d_{G_i}(u, v),$$

where  $d_{G_i}$  is the distance in the graph  $G_i = (V, E_i)$ . Intuitively,  $d_{\mathcal{G}}(u, v)$  is the length of the shortest path connecting  $u$  and  $v$  using a single mode of transportation.

The  $k$ -mode distance naturally arises in practice in the context of multiple embeddings in machine learning (see [21, 15, 18] and many more). In many scenarios, given a high dimensional data set, there is no clear distance function that accurately captures semantic relationships within the data. To address this, deep learning is now commonly used to obtain a data embedding that encodes semantic information. As research in this field advances, we observe a proliferation of embeddings which can vary due to differences in network architecture, learning paradigms or training data, to name a few. This results in a single set of data points with multiple sets of edges representing distances in the different embeddings spaces. For every pair of points we are interested in the embedding in which they are closest together and often care about finding a pair of points which are furthest apart in all embeddings (i.e. computing the  $k$ -mode diameter) or a point which is closest to all others in some embedding (i.e. computing the  $k$ -mode radius).

Let us first consider the All-Pairs Shortest Paths (APSP) problem in  $k$ -multimode graphs on  $n$  vertices: compute the  $k$ -mode distance between all pairs of vertices of  $\mathcal{G} = (V, E_1, \dots, E_k)$ . One can always run any APSP algorithm on each  $(V, E_i)$  separately, resulting in a running time which is  $k$  times that of the standard APSP algorithm. This approach will also trivially compute the  $k$ -mode radius of the multimode graph. In fact, we show that in general no algorithm can have polynomial improvements over this trivial solution.

► **Theorem 1.** *Under the APSP hypothesis [17], no  $O((kn^3)^{1-\varepsilon})$  time algorithm can solve  $k$ -mode APSP or  $k$ -mode radius for any  $\varepsilon > 0$ .*

While it is clear for APSP that one can compute all  $k$ -multimode distances by solving  $k$  instances of APSP separately on each graph, it is no longer obvious how to solve  $k$ -multimode radius or diameter using algorithms that compute the radius or diameter in a standard graph.

To illustrate why the  $k$ -mode diameter is not directly related to the standard diameters of the individual graphs  $G_1, \dots, G_k$ , denote the diameter of  $G_1 = (V, E_1)$  by  $D_1 = D(G_1)$  and let  $u_1, v_1 \in V$  be such that  $d_{G_1}(u_1, v_1) = D_1$ . Similarly, denote the diameter of  $G_2 = (V, E_2)$  by  $D_2 = D(G_2)$  and let  $u_2, v_2 \in V$  be such that  $d_{G_2}(u_2, v_2) = D_2$ . We cannot merely take the minimum of  $D_1, D_2$  to get the diameter in the multimode graph  $\mathcal{G} = (V, E_1, E_2)$ . This is because it could be that  $d_{G_2}(u_1, v_1) < D_1$  and  $d_{G_1}(u_2, v_2) < D_2$  so that the multimode distance between  $u_1$  and  $v_1$  is not  $D_1$  but rather  $d_{G_2}(u_1, v_1)$  and similarly for  $u_2$  and  $v_2$ . The multimode diameter can be the distance between some completely different pair of vertices and thus can be arbitrarily smaller than both  $D_1$  and  $D_2$ . It could even be the case that  $D_1, D_2$  are infinite, while the multimode diameter  $D(\mathcal{G})$  is finite.

However, if we wish to compute the exact  $k$ -mode diameter or radius of a  $k$ -multimode undirected graph, we show that we can in fact use any algorithm for computing standard diameter or radius in a blackbox way.

► **Theorem 2.** *If there exists a  $T(n, m)$  time algorithm for computing the diameter/radius of a weighted, undirected graph with  $n$ -nodes and  $m$ -edges, then there exists an algorithm for computing the  $k$ -mode diameter of a  $k$ -multimode undirected graph in time  $T(O(kn), O(m + kn))$ .*

The proofs of both Theorem 1 and Theorem 2 follows from a straightforward application of standard techniques. We therefore defer their proofs to the full version of this work. We instead turn our focus to approximating these parameters.

There exists a plethora of algorithms and conditional lower bounds for approximating diameter and radius under both the standard notion of (directed or undirected) distance and the min-distance measure [2, 1, 7, 9, 10, 11, 6, 16, 4, 13]. As the  $k$ -multimode distance generalizes both these distance measures, we would like to develop approximation algorithms and prove tight conditional lower bounds for approximating diameter and radius in the multimode setting as well.

The known fine-grained hardness results for the min-distance and standard diameter, radius and eccentricities would trivially carry over for the multimode distance case (as it generalizes both). We are thus interested in when stronger hardness is possible, and when similar approximation algorithms can be attained.

Devising approximation algorithms for the min-distance measure is particularly difficult and therefore significantly new techniques had to be developed over the standard diameter ones (see e.g. [12, 7, 9]). The main difficulty lies in the fact that the triangle inequality no longer holds: it could be that  $d_{\min}(u, v) = 1$  and  $d_{\min}(v, w) = 1$  but  $d_{\min}(u, w) = \infty$ , as in the graph with vertices  $\{u, v, w\}$  and edges  $\{(v, u), (v, w)\}$ . This difficulty persists and is even more prominent in the multimode distance setting, as now the various graphs that represent modes of transportation may not share any edges. We can therefore expect that the distance parameters of interest can be even more difficult to approximate.

## 1.1 Our results

We develop various new techniques that work in the multimode distance setting, both for directed and undirected graphs, and for fine-grained conditional lower bounds. Table 1 summarizes our results. Here we present some highlights.

■ **Table 1** Our Results. Algorithms that match their lower bound are in bold. The runtimes marked with \* assume  $\omega = 2$ .

$k$	Approx.	Runtime	Comments	Reference
<b>Undirected <math>k</math>-mode Diameter Upper Bounds</b>				
2	3	$O(m)$	Weighted	Theorem 9
3	3	$\tilde{O}(m)$	<b>Weighted.</b>	<b>Theorem 13</b>
2	$(2, 2M)$	$\tilde{O}(mn^{3/4})^*$	Non negative edge weights $\leq M$ . If $\omega > 2$ the runtime is $m \cdot \left( \frac{n^{1.5}}{(m\sqrt{n})^{1/\omega}} + \left(\frac{m}{n}\right)^{1/(\omega-2)} \right)$ .	See full version
2	$(2.5, 2.5M)$	$\tilde{O}(m\sqrt{n})^*$	Non negative edge weights $\leq M$ . If $\omega > 2$ the runtime is $m \cdot \left( \frac{n}{m^{1/\omega}} + \left(\frac{m}{n}\right)^{1/(\omega-2)} \right)$ .	See full version
<b>Undirected <math>k</math>-mode Diameter Lower Bounds (Under SETH)</b>				
2	$2 - \delta$	$\Omega(m^{2-o(1)})$	Unweighted	See full version
3	$3 - \frac{2}{\ell} - \delta$	$\Omega(m^{1+1/(\ell-1)-o(1)})$	Unweighted.	See full version
$\Omega(\log n)$	Any	$\Omega(m^{2-o(1)})$	Unweighted.	See full version
<b>Directed <math>k</math>-mode Diameter Upper Bounds</b>				
2	$n$	$O(m)$	Weighted.	See full version
2	2	$\tilde{O}(m)$	Weighted DAG.	See full version
2	$(\frac{3}{2}, 1)$	$O(m^{0.414} n^{1.522} + n^{2+o(1)})$	Unweighted DAG.	See full version
<b>Directed <math>k</math>-mode Diameter Lower Bounds (Under SETH)</b>				
2	$2 - \delta$	$\Omega(m^{2-o(1)})$	Unweighted.	See full version
3	Any	$\Omega(m^{2-o(1)})$	Unweighted DAG.	See full version
<b>Undirected <math>k</math>-mode Radius Upper Bounds</b>				
$k$	3	$\tilde{O}(mk!)$	<b>Weighted.</b> <b>Tight for constant <math>k \geq 3</math>.</b>	<b>Theorem 14</b>
<b>Undirected <math>k</math>-mode Radius Lower Bounds (Under the HS Hypothesis)</b>				
2	$2 - \delta$	$\Omega(m^{2-o(1)})$	Unweighted.	See full version
3	$3 - \frac{2}{\ell} - \delta$	$\Omega(m^{1+1/(3\ell-5)-o(1)})$	Weighted. Conditioned on the new $\ell$ -HSC.	See full version
$\Omega(\log n)$	Any	$\Omega(m^{2-o(1)})$	Unweighted.	See full version
<b>Directed <math>k</math>-mode Radius Upper Bounds for DAGs</b>				
2	$n$	$O(m)$	Weighted DAG.	See full version
<b>Directed <math>k</math>-mode Radius Lower Bounds (Under the HS Hypothesis)</b>				
2	Any	$\Omega(m^{2-o(1)})$	Unweighted.	See full version
2	$2 - \delta$	$\Omega(m^{2-o(1)})$	Unweighted DAG.	See full version
3	Any	$\Omega(m^{2-o(1)})$	Unweighted DAG.	See full version

Our hardness results are based on the Strong Exponential Time Hypothesis (SETH) [14, 5] and the Hitting Set Hypothesis [1] (see also [19]). We further extend the Hitting Set Hypothesis to the stronger  $\ell$ -Hitting Set Hypothesis. Based on this new hypothesis we are able to obtain a lower bound tradeoff for approximating the  $k$ -mode radius. As the introduction of the Hitting Set Lemma in [1] allowed for the proof of the first lower bounds on radius approximation, we hope this extension will enable future work on lower bound tradeoffs for radius in various distance frameworks.

► **Theorem 3.** *Let  $k \geq 3$ . Under SETH, there can be no  $O(m^{2-\varepsilon})$  time (for  $\varepsilon > 0$ ) algorithm that achieves any finite multiplicative approximation for the  $k$ -mode diameter in directed unweighted  $m$ -edge multimode graphs, even if all graphs in the multimode graph are DAGs.*

The result appears as Theorem 5.3 in the full version of the text. It implies that, under SETH, to get any approximation algorithm for directed multimode diameter, even for DAGs, one must focus on  $k = 1$  or  $k = 2$ . The case  $k = 1$  is just the standard diameter problem, and so we consider the case where  $k = 2$  next.

For 2-mode diameter in directed graphs we provide (1) a near-linear time algorithm that can decide whether the diameter is finite, and (2) a near-linear time 2-approximation algorithm for the case when both graphs in the multimode graph are DAGs. Both of these results are tight, in the sense that for  $k > 2$  such results are impossible under SETH. The results appear in the full version.

For  $k$ -mode diameter and radius in undirected graphs we obtain several tight results as well. Here are two highlights:

► **Theorem 4.** *For every constant  $k$ , there is an  $\tilde{O}(m)$  time<sup>1</sup> 3-approximation algorithm for the radius of  $m$ -edge undirected  $k$ -multimode graphs. This is tight under the  $\ell$ -Hitting Set Hypothesis, in that no  $3 - \varepsilon$ -approximation is possible in  $\tilde{O}(m)$  time, even for  $k = 2$ .*

This algorithmic result is stated in Theorem 14 while the lower bound appears in the full version. It achieves the conditionally best possible result for radius in near-linear time.

► **Theorem 5.** *Under SETH, there can be no  $2 - \delta$ -approximation algorithm for diameter running in  $O(mn^{1-\varepsilon})$  time for  $\varepsilon, \delta > 0$  for  $m$ -edge 2-mode undirected unweighted graphs. If  $\omega = 2$ ,<sup>2</sup> there is an  $\tilde{O}(mn^{3/4})$  time algorithm that achieves a 2-multiplicative, 2-additive approximation to the diameter in  $m$ -edge,  $n$ -node undirected 2-multimode graphs.*

This result appears in the full version. Up to the small additive error, the approximation algorithm is optimal due to the SETH-based lower bound. The running time of the algorithm is still faster than  $mn$ , even with the current bound on  $\omega$ .

## 1.2 Technical Overview

In this subsection we describe an outline of some of our results, and highlight their connections to other, well studied problems.

### Approximating Undirected 2 and 3 Mode Diameter

Our first result is a simple, linear time 3-approximation for 2-mode diameter. Given a threshold  $D$ , we would like to either show that the 2-mode diameter is less than  $D$  or find a pair of points  $u, v$  whose 2-mode distance is greater than  $D/3$ , meaning  $d_{G_1}(u, v) \geq D/3$  and  $d_{G_2}(u, v) \geq D/3$ .

To find such a pair, run BFS from an arbitrary point  $z$ . Let  $X$  be the points within distance  $D/3$  of  $z$  in  $G_1$  and let  $Y$  be the points within distance  $D/3$  of  $z$  in  $G_2$ . If any vertex is contained in neither  $X$  nor  $Y$ , then it has multimode distance greater than  $D/3$  from  $z$  and we are finished.

Next, run BFS from  $X$  (as a set) in  $G_1$  to identify the vertices in  $Y$  that have distance  $\geq D/3$  in  $G_1$  from *all* vertices in  $X$ . Similarly, run BFS from  $Y$  in  $G_2$  to identify all vertices in  $X$  that have distance  $\geq D/3$  in  $G_2$  from all vertices in  $Y$ . If there exists a pair of vertices  $x \in X, y \in Y$  such that  $d_{G_1}(y, X) \geq D/3$  and  $d_{G_2}(x, Y) \geq D/3$  then the multimode distance

<sup>1</sup> As is standard,  $\tilde{O}$  hides polylogarithmic factors.

<sup>2</sup> Here  $\omega$  denotes the fast matrix multiplication exponent,  $\omega < 2.371552$  [20].

of  $x, y$  is greater than  $D/3$  and we are done. Otherwise, we can show that any pair of points has multimode distance  $< D$ , and thus the 2-mode diameter is  $< D$ . Indeed, if  $x, y \in X$  then  $d_{G_1}(x, y) < 2D/3$  using a path through  $z$  in  $G_1$ . Similarly, if  $x, y \in Y$  then  $d_{G_2}(x, y) < 2D/3$ . Otherwise, we can assume w.l.o.g that  $x \in X, y \in Y$  and  $d_{G_1}(y, X) < D/3$ . Then there exists a point  $x' \in X$  such that  $d_{G_1}(x', y) < D/3$ . Now, following a path in  $G_1$  from  $y$  to  $x'$  to  $z$  to  $x$  we have that  $d_G(x, y) < D/3 + D/3 + D/3 = D$ . Thus, for any pair of vertices we have  $d_G(x, y) < D$ .

This completes our 3-approximation for undirected 2-mode diameter in linear time. Next we generalize this approach to a subroutine that can be used for any  $\beta$ -approximation for  $2 \leq \beta \leq 3$ . Informally, we take  $X$  to be the neighborhood in  $G_1$  of some point  $x$  and  $Y$  to be the neighborhood in  $G_2$  of some point  $y$ . We perform a similar search, running BFS from the set  $X$  in  $G_1$  and from  $Y$  in  $G_2$  to find a pair of points in  $X \times Y$  that is far apart in both graphs. In order to conclude that the graph has a small diameter in the case that we do not find a pair of large distance, we need to consider all pairs  $x, y$  where  $x$  is in the  $G_1$  neighborhood of a vertex  $z$  and  $y$  is in the  $G_2$  neighborhood of the same vertex  $z$ . This results in a subroutine that, given a point  $z$  with small  $G_1$  and  $G_2$  neighborhoods, we can compute a  $\beta$ -approximation to the 2-mode diameter. For details see Lemma 12.

We can now use this idea to construct  $\beta$ -approximation algorithms using a classic large-small neighborhoods tradeoff. We show two ways in which to use this subroutine in a complete algorithm, obtaining a 2-approximation and a 2.5-approximation.

For our 2-approximation, if we have a point  $z$  with small  $G_1$  and  $G_2$  neighborhoods, we use the aforementioned subroutine. Otherwise, all vertices have a large  $G_1$  or  $G_2$  neighborhood, including the endpoints of the diameter. We leverage this fact and sample a hitting set, hitting the  $G_1$  or  $G_2$  neighborhood of a diameter endpoint. For every point  $z$  in this hitting set, we compute the set of points  $A$  of distance  $< D/4$  from  $z$  in  $G_1$  and the set of points  $B$  of distance  $> 3D/4$  from  $z$  in  $G_1$ . We show that if the 2-mode diameter is  $\geq D$  and  $z$  hits the  $G_1$  neighborhood of a diameter endpoint, then the ST-diameter of the sets  $A, B$  in  $G_2$  will be  $D$ . We can therefore use an ST-diameter approximation algorithm in a black box way to obtain our desired approximation.

To obtain our 2.5-approximation we use a similar approach, leveraging a blackbox ST-diameter 2.5-approximation. Using a slightly different method, we apply the simple 3-approximation of ST-diameter to turn our linear time 2-mode diameter approximation into an approximation for 3-mode diameter, incurring only poly-logarithmic factors in the running time.

## Connections to ST-Diameter

As hinted to above, the ST-diameter problem turns out to be intricately connected to the problem of computing the multimode diameter. In many cases, one can reduce the problem of computing a multimode diameter to that of computing the ST-diameter of some subsets of the vertices in one of the edge sets. It turns out, this connection between the problems persists in the lower bounds as well.

Consider a lower bound construction showing that an  $\alpha$ -approximation of ST-diameter requires  $T(n)$  time: a graph  $G = (V, E)$  with  $S, T \subseteq V$  such that approximating the ST-diameter of  $S, T$  beyond a factor of  $\alpha$  takes  $\Omega(T(n))$  time. We can now construct a simple lower bound for the problem of approximating 3-mode diameter as follows. Take the 3-multimode graph with  $E$  as its first edge set, using a second set of edges connect the vertices of  $S$  with the vertices of  $V \setminus T$  and using a third set of edges connect  $T$  with  $V \setminus S$ . Now,

the 3-mode diameter of this 3-multimode graph will be the ST-diameter of  $S$  and  $T$  in the original graph. Therefore approximating the 3-mode diameter beyond a factor of  $\alpha$  requires  $\Omega(T(n))$  time as well.

We use an existing lower bound tradeoff for ST-diameter to establish a lower bound tradeoff for 3-mode diameter approximation. However, for radius approximation no such tradeoff is known. We further extend this result to a lower bound tradeoff for 3-mode radius approximation, establishing the first lower bound tradeoff for any form of radius approximation problem. This result also suggests a potential “radius equivalent” of the ST-diameter problem, which could be an avenue for future work.

## Approximating Directed Multimode Distances

While the problem of approximating diameter and radius of multimode graphs is closely related to the ST-diameter problem, and often ST-diameter techniques allow to “reduce the number of modes” in the graph, the directed case is very different. In directed graphs we cannot hope for techniques like this to work as we show that any approximation to 3-mode diameter or 2-mode radius requires quadratic time.

We are able to show, however, that in directed acyclic graphs, the problem of approximating 2-mode diameter is in fact equivalent to the problem of approximating the min-diameter of each of the 2 graphs making up the multimode graph. In the case of a general directed 2-mode graph, we construct a linear time algorithm which can differentiate between finite and infinite 2-mode diameter. It does so by considering the min-distances in the graphs of the strongly connected components of  $G_1$  and  $G_2$ , and the intersections of these two sets of SCCs with each other.

## 2 Preliminaries

Let  $G = (V, E)$  be a graph with  $n = |V|$  vertices and  $m = |E|$  edges. For a  $k$ -multimode graph  $\mathcal{G} = (V, E_1, \dots, E_k)$ , denote  $n = |V|$  and  $m = |E_1| + \dots + |E_k|$ , we will also use  $e(\mathcal{G})$  to denote the number of edges in  $\mathcal{G}$ . Define  $G_i = (V, E_i)$  and for any  $u, v \in V$  let  $d_{G_i}(u, v)$  be the length of the shortest path from  $u$  to  $v$  in the graph  $G_i$ . When  $\mathcal{G}$  is clear from context, we denote this distance by  $d_i(u, v)$ . The  $k$ -mode distance of  $u, v$  is defined as  $d_{\mathcal{G}}(u, v) = \min_{i \in [k]} d_i(u, v)$ .

Throughout this paper, we will often associate each “mode” with a color. We can assign a color to each set of edges  $E_i$  and think of  $\mathcal{G}$  as the graph  $G = (V, E = \bigcup_i E_i)$  where every edge is assigned at least one color. A path that uses a single mode corresponds to a monochromatic path under this coloring.

Consider the special case of 2-multimode graphs with edge sets  $E_1, E_2$ . We use “red distance” to refer to  $d_{G_1}$ , meaning paths using only edges in  $E_1$  (“red edges”). Likewise, we use “blue distance” to refer to  $d_{G_2}$ . In 3-multimode graphs,  $\mathcal{G} = (V, E_1, E_2, E_3)$ , we use “green distance” to refer to  $d_{G_3}$ .

In this work we present many approximation algorithms to  $k$ -mode diameter, radius and eccentricities of  $k$ -multimode graphs. We define an  $(\alpha, \beta)$ -approximation algorithm for a value  $D$  to be such that outputs  $\tilde{D}$  satisfying  $D \leq \tilde{D} \leq \alpha D + \beta$ . When  $\beta = 0$  we have a multiplicative approximation and call it an  $\alpha$ -approximation.

We call a directed  $k$ -multimode graph  $\mathcal{G} = (V, E_1, \dots, E_k)$  a  $k$ -multimode DAG (Directed Acyclic Graph) if  $G_1, \dots, G_k$  are all DAGs.

Denote the ball around a vertex  $v$  in a graph  $G$ , or the neighborhood of  $v$ , by  $B_G(v, r) := \{u \in V : d_G(u, v) < r\}$ . In a  $k$ -multimode graph, we will use  $B_i(v, r)$  to denote  $B_{G_i}(v, r)$ .



Given a vertex subset  $A \subseteq V$ , denote by  $\mathcal{G}[A]$  the induced subgraph defined by  $\mathcal{G}[A] := (A, E_1 \cap A \times A, \dots, E_k \cap A \times A)$ .

The eccentricity of a vertex  $v$  in a graph  $G$  is defined as  $\text{ecc}_G(v) := \max_{u \in V} d_G(u, v)$ . The diameter of  $G$ , or  $D(G)$ , is defined as the largest eccentricity in the graph and the radius, or  $R(G)$ , is defined as the smallest eccentricity. We naturally extend these definitions to  $k$ -multimode graphs using the  $k$ -mode distance. We refer to points  $u, v$  such that  $d_G(u, v) = D(G)$  as “diameter endpoints” and say that these points “achieve the diameter”. We refer to a point  $c$  such that  $\text{ecc}_G(c) = R(G)$  as a “center” of  $G$ .

The  $k$ -mode distance is a generalization of the well studied min-distance. In a directed graph, the min-distance between two vertices  $u, v$  is defined as  $d_{\min}(u, v) := \min \{d(u, v), d(v, u)\}$ . Approximating the diameter of a graph under the min-distance (or min-diam) has been studied in [1], [7], [9], [12]. In our approximation algorithms we will use the following lemma from [1]:

► **Lemma 6** ([1], Lemma C.2). *There is a  $O(m + n)$ -time algorithm that determines which vertices in a directed graph  $G$  have finite min-eccentricity.*

A variant of the diameter that we use throughout this paper is the  $ST$ -diameter. Given two sets  $S, T \subseteq V$ , the  $ST$ -diameter is defined as  $D_{S,T} := \max_{s \in S, t \in T} d(s, t)$ . Backurs, Roditty, Segal, Vassilevska W. and Wein explored this variant in [2]. We will use many of their results throughout this paper. We will use the following algorithms in approximating the undirected  $k$ -mode diameter.

► **Lemma 7** ([2], Claim 24, Theorem 25). *Given  $G = (V, E)$  and  $S, T \subseteq V$ , there exists an  $\tilde{O}(m\sqrt{n})$  time algorithm that computes a 2-approximation for the  $ST$ -diameter and a  $O(m)$  time algorithm that computes a 3-approximation for the  $ST$ -diameter.*

Backurs et al. [2] also proved conditional lower bounds surrounding the hardness of approximating the  $ST$ -diameter.

► **Lemma 8** ([2], Theorem 7). *Under the Strong Exponential Time Hypothesis (SETH), for every integer  $\ell \geq 2$ , any  $3 - \frac{2}{\ell} - \delta$ -approximation algorithm for  $ST$ -diameter requires  $m^{1+1/(\ell-1)-o(1)}$  time for any  $\delta > 0$ .*

### 3 Approximating Undirected $k$ -mode Diameter and Radius

In the following section we focus on  $k$ -multimode graphs with small values of  $k$ ,  $k \leq \text{poly} \log(n)$ . Under SETH and the Hitting Set Hypothesis, computing the exact  $k$ -mode diameter or radius requires  $\Omega((nm)^{1-o(1)})$  time, as it is at least as hard as computing these values in ordinary graphs [1]. Therefore, we have little hope of computing the exact diameter or radius polynomially faster than the trivial  $\tilde{O}(mn)$  time algorithm. Instead, we focus on approximating these parameters.

In this work we focus on undirected graphs, for directed graphs refer to the full version. First, we present a simple algorithm running in linear time that provides a 3-approximation to the  $k$ -mode diameter when  $k = 2$ . We then show how to extend the idea of this algorithm to use in the context of a general  $\alpha$ -approximation. We use this extended idea to obtain a near 2-approximation and a near 2.5-approximation algorithm for 2-mode diameter in the full version of this paper.



Next, we extend the 3-approximation algorithm for 2-mode diameter to the case of  $k = 3$  and obtain a near linear<sup>3</sup> time 3-approximation algorithm for the 3-mode diameter.

Finally, we turn our attention to approximating the  $k$ -mode radius of a  $k$ -multimode graph. Here we are able to show a 3-approximation algorithm for the  $k$ -mode radius for any  $k$ , with a running time of  $O(mk!)$ .

In the full version we prove that for  $k \geq 3$ , any  $3 - \delta$  approximation algorithm for  $k$ -mode diameter must run in super linear time, showing that our near linear time 3-approximation algorithm for 3-mode diameter is in fact tight. Similarly, in the full version we also show that for  $k \geq 3$ , any  $3 - \delta$  approximation algorithm for  $k$ -mode radius must run in super linear time. Therefore, for any constant  $k \geq 3$  our 3-approximation algorithm for  $k$ -mode radius runs in near linear time and is tight.

### 3.1 Linear Time 3-Approximation for 2-mode Diameter

In this section we show our first approximation algorithm for undirected 2-mode diameter, running in linear time and providing a multiplicative 3-approximation. We prove this result for unweighted graphs but note that by replacing BFS with Dijkstra's algorithm we obtain the same result for weighted graphs while increasing the runtime complexity by only a  $\log n$  factor.

► **Theorem 9.** *There exists an  $O(m)$  time algorithm that computes a 3-approximation of the 2-mode diameter of an unweighted, undirected 2-multimode graph.*

**Proof.** Given a 2-multimode graph  $\mathcal{G} = (V, E_1, E_2)$  with 2-mode diameter  $D$ , our algorithm will output a pair of points  $a, b$  such that  $\frac{D}{3} \leq d_{\mathcal{G}}(a, b) \leq D$ .

Start by running BFS in the red graph  $G_1$  and the blue graph  $G_2$  from an arbitrary node  $z \in V$ . Let  $X = \{v \in V \mid d_1(v, z) < d_2(v, z)\}$ , the points closer to  $z$  in red than in blue. Let  $Y = V \setminus X$ . Denote by  $\alpha_0 := \max_{v \in X} d_1(z, v)$ , the largest distance between  $z$  and a node in  $X$ . Likewise define  $\beta_0 := \max_{v \in Y} d_2(z, v)$ .

Run BFS in  $G_1$  from the set  $X$  by contracting the set  $X$  into a single vertex and running BFS from it. Let  $y \in Y$  be the point furthest away from  $X$  in red. Similarly, run BFS in  $G_2$  from the set  $Y$  and find the point  $x \in X$  furthest away from  $Y$  in blue.

Return  $\tilde{D} = \max(d_{\mathcal{G}}(x, y), \alpha_0, \beta_0)$ . We claim that  $D/3 \leq \tilde{D} \leq D$ .

Clearly  $\tilde{D} \leq D$ . We are left to show  $\tilde{D} \geq D/3$ . If  $\alpha_0 \geq D/3$  or  $\beta_0 \geq D/3$ , we are finished. Otherwise, consider a pair of diameter endpoints  $s, t$ . Since any two points in  $X$  are within distance  $\leq 2\alpha_0 < D$ , the points  $s, t$  cannot both be in  $X$ . Likewise, they cannot both be in  $Y$ . Thus, we can assume without loss of generality  $s \in Y, t \in X$ .

Denote by  $\alpha_1 := d_1(y, X)$ . By our choice of  $y$ , we have  $d_1(s, X) \leq \alpha_1$ . Therefore, there is some point  $x' \in X$  such that  $d_1(s, x') \leq \alpha_1$ . Thus,

$$D = d_1(s, t) \leq d_1(s, x') + d_1(x', z) + d_1(z, t) \leq \alpha_1 + 2\alpha_0.$$

Since  $\alpha_0 < D/3$  we conclude  $\alpha_1 \geq D/3$ . By a similar argument, if we define  $\beta_1 = d_2(x, Y)$ , we conclude that  $\beta_1 \geq D/3$ . Therefore,  $d_1(x, y) \geq d_1(y, X) = \alpha_1 \geq D/3$  and  $d_2(x, y) \geq d_2(x, Y) = \beta_1 \geq D/3$ . The pair of points  $x, y$  have both red and blue distance greater than  $D/3$  and so their 2-multimode distance is  $d_{\mathcal{G}}(x, y) \geq D/3$ .

We conclude that  $\tilde{D} \geq D/3$ . ◀

---


<sup>3</sup> Running in  $\tilde{O}(m)$  time.

### 3.2 Subquadratic Time $< 3$ -Approximation Technique

Still considering a 2-multimode graph, our next goal is to obtain a better-than-3 approximation for the 2-mode diameter. The main tool we will use in this section is a generalization of our linear time 3-approximation algorithm. For the following algorithms we will be given a value of  $D$  and determine whether  $D(\mathcal{G}) < D$  or  $D(\mathcal{G}) \geq \alpha D$ . Thus, binary searching over  $D$  will give a  $\frac{1}{\alpha}$ -approximation to the diameter. For simplicity, we will assume  $\frac{1-\alpha}{2}D$  is an integer and  $\mathcal{G}$  is unweighted. If that is not the case, we get an additional additive error to our approximation, which we address in the full version of this work.

Given  $D$ , we can rephrase the above 3-approximation algorithm as follows: pick an arbitrary node  $z$  and run BFS from it to compute its “red neighborhood”  $X_1 := B_1(z, \frac{D}{3})$  and its “blue neighborhood”  $Y_1 := B_2(z, \frac{D}{3})$ . Find the points in the blue neighborhood  $Y_1$  that are within red distance  $< \frac{D}{3}$  of some point in the red neighborhood, call these points  $X_2$ . Similarly, take the points in  $X_1$  that are within blue distance  $< \frac{D}{3}$  of some point in the blue neighborhood and call them  $Y_2$ . If we take a pair of points  $a \in X_1 \setminus Y_2, b \in Y_1 \setminus X_2$  we know that their distance from each other in both red and blue is greater than  $\frac{D}{3}$  and thus the pair  $(a, b)$  provides a 3-approximation to the diameter. Furthermore, we know any point in  $X_2$  is within red distance  $< D$  of all points in  $X_1$  and within blue distance  $< \frac{2D}{3}$  of all points in  $Y_1$ . Thus, if the diameter is  $\geq D$ , no diameter endpoint can be in  $X_2$  and similarly no diameter endpoint can be in  $Y_2$ . We conclude that if  $X_1 \setminus Y_2$  and  $Y_1 \setminus X_2$  are not empty we can find points  $a, b$  that provide a 3-approximation to the 2-mode diameter, and otherwise  $D(\mathcal{G}) < D$ .

For a general approximation factor  $\alpha$ , instead of considering the red and the blue neighborhoods of a single point  $z$ , consider a pair of points  $x, y$  and take the red neighborhood of one and the blue neighborhood of the other. Define the red neighborhood of  $x$ ,  $X_1 = B_1(x, \frac{1-\alpha}{2}D)$  and the blue neighborhood of  $y$ ,  $Y_1 = B_2(y, \frac{1-\alpha}{2}D)$ . Consider the points in  $Y_1$  that are within red distance  $\alpha D$  of some point in  $X_1$ ,  $X_2 = B_1(x, \frac{1+\alpha}{2}D) \cap Y_1$ . Similarly, define  $Y_2 = B_1(y, \frac{1+\alpha}{2}D) \cap X_1$ . The sets are illustrated in figure 1, where  $d_1$  distance is illustrated in red and  $d_2$  distance is illustrated in blue.



■ **Figure 1** Defining the sets  $X_1, X_2, Y_1, Y_2$ .

As with the 3-approximation, these sets have two useful properties.

▷ **Claim 10.** If  $X_1 \setminus Y_2 \neq \emptyset$  and  $Y_1 \setminus X_2 \neq \emptyset$  then we can find a pair of points  $a, b$  with  $d_{\mathcal{G}}(a, b) \geq \alpha D$ .

*Proof.* Take an arbitrary  $a \in X_1 \setminus Y_2$  and  $b \in Y_1 \setminus X_2$ . If  $d_1(a, b) < \alpha D$ , then

$$d_1(x, b) \leq d_1(x, a) + d_1(a, b) < \alpha D + \frac{1-\alpha}{2}D = \frac{1+\alpha}{2}D$$

and so  $b \in X_2$ , a contradiction. Thus  $d_1(a, b) \geq \alpha D$  and similarly  $d_2(a, b) \geq \alpha D$ . Therefore  $d_{\mathcal{G}}(a, b) \geq \alpha D$ .  $\triangleleft$

▷ **Claim 11.** Let  $s, t$  have distance  $d_G(s, t) \geq D$ . If  $s \in X_1$  and  $t \in Y_1$ , then  $X_1 \setminus Y_2 \neq \emptyset$  and  $Y_1 \setminus X_2 \neq \emptyset$ .

**Proof.** If  $X_1 = Y_2$ , then  $s \in X_1 = Y_2$  and therefore  $d_2(s, y) < \frac{1+\alpha}{2}D$ . Thus,  $d_2(s, t) \leq d_2(s, y) + d_2(y, t) < \frac{1+\alpha}{2}D + \frac{1-\alpha}{2}D = D$ , in contradiction to  $d_G(s, t) \geq D$ . We arrive at a similar contradiction if  $Y_1 = X_2$ .  $\triangleleft$

Therefore, our goal is now to find a pair of points  $x, y$  such that  $s \in X_1$  and  $t \in Y_1$ . Using the above claims, this will allow us to find a pair of points with  $d_G(a, b) \geq \alpha D$  in linear time.

To do so, consider running BFS from an arbitrary point  $z$  in both the red and the blue graphs  $G_1, G_2$ . If the algorithm found no point  $p$  such that  $d_G(z, p) \geq \alpha D$ , then  $z$  is close to  $s$  and  $t$  in either red or blue. Since  $\alpha \leq \frac{1}{2}$ , we can assume without loss of generality that  $d_1(z, s) < \alpha D$  and  $d_2(z, t) < \alpha D$ . Now, if  $d_1(z, a) \geq \frac{1-\alpha}{2}D$ , since we assumed  $\frac{1-\alpha}{2}D \in \mathbb{N}$ , there is some point  $x$  on the shortest red path between  $z$  and  $s$  such that  $d_1(x, s) = \frac{1-\alpha}{2}D$ , and so

$$d_1(z, x) = d_1(z, s) - d_1(x, s) < \left( \alpha - \frac{1-\alpha}{2} \right) D = \frac{3\alpha-1}{2}D.$$

If  $d_1(z, a) < \frac{1-\alpha}{2}D$  then there exists a point  $x$  on this shortest path such that  $d_1(x, s) \leq \frac{1-\alpha}{2}D$  and  $d_1(z, x) < \frac{3\alpha-1}{2}D$ .

Similarly, there exists a point  $y$  such that  $d_2(y, t) \leq \frac{1-\alpha}{2}D$  and  $d_2(z, y) < \frac{3\alpha-1}{2}D$ . Therefore, by considering all pairs of points  $x \in B_1(z, \frac{3\alpha-1}{2}D), y \in B_2(z, \frac{3\alpha-1}{2}D)$  we can find a pair that satisfy the conditions of claim 11 and thus also the conditions of claim 10. Now, by running a linear time algorithm for every pair of points  $x \in B_1(z, \frac{3\alpha-1}{2}D), y \in B_2(z, \frac{3\alpha-1}{2}D)$  we can find a pair of points of distance  $\geq \alpha D$ . In fact, we don't need to run the above algorithm for every pair of potential points  $x, y$  independently. We can improve our running time using fast rectangular matrix multiplication. We use the standard notation  $M(J, K, L)$  to indicate the time it takes to multiply a  $J \times K$  matrix by a  $K \times L$  matrix using fast rectangular matrix multiplication.

We are now ready to claim the following generalization to our 2-approximation algorithm.

► **Lemma 12.** Let  $\mathcal{G}$  be a 2-multimode graph with 2-mode diameter  $\geq D$ . Given  $\frac{1}{3} < \alpha \leq \frac{1}{2}$  and a point  $z$  with  $|B_1(z, \frac{3\alpha-1}{2}D)| \leq n^\delta$  and  $|B_2(z, \frac{3\alpha-1}{2}D)| \leq n^\delta$ , we can find a pair of points  $a, b$  with  $d_G(a, b) \geq \alpha D$  in time  $O(n^\delta m) + M(n^\delta, n, n^\delta)$ .

**Proof.** Define  $X = B_1(z, \frac{3\alpha-1}{2}D)$  and  $Y = B_2(z, \frac{3\alpha-1}{2}D)$ . From every  $x \in X$  run BFS in  $d_1$  to compute  $X_1 := B_1(x, \frac{1-\alpha}{2}D)$  and  $X_2 := B_1(x, \frac{1+\alpha}{2}D)$  as illustrated in figure 1 (without taking an intersection with  $Y_1$ ). Let  $M_X^1$  be a  $n^\delta \times n$  matrix with rows indexed by  $X$  and columns indexed by  $V$  and let  $M_X^2$  be a  $n \times n^\delta$  matrix with rows indexed by  $V$  and columns indexed by  $X$ . Define  $M_X^1$  and  $M_X^2$  as follows,

$$M_X^1[x, u] = \begin{cases} 1 & \text{if } u \in X_1, \\ 0 & \text{otherwise.} \end{cases}$$

$$M_X^2[u, x] = \begin{cases} 1 & \text{if } u \notin X_2, \\ 0 & \text{otherwise.} \end{cases}$$

Similarly, run BFS from every  $y \in Y$  in  $d_2$  and define  $M_Y^1, M_Y^2$  by,

$$M_Y^1[y, u] = \begin{cases} 1 & \text{if } u \in Y_1 := B_2(y, \frac{1-\alpha}{2}D), \\ 0 & \text{otherwise.} \end{cases}$$

$$M_Y^2[u, y] = \begin{cases} 1 & \text{if } u \notin Y_2 := B_2(y, \frac{1+\alpha}{2}D), \\ 0 & \text{otherwise.} \end{cases}$$

Finally, compute the boolean matrix product

$$Z = M_X^1 \cdot M_Y^2 \wedge M_Y^1 \cdot M_X^2.$$

For any  $x \in X, y \in Y$ , we have  $Z[x, y] = 1 \iff X_1 \setminus Y_2 \neq \emptyset$  and  $Y_1 \setminus X_2 \neq \emptyset$ . Therefore, by claim 10, if  $Z[x, y] = 1$  we can find a pair of points with  $d_G(a, b) \geq \alpha D$ . If  $D(\mathcal{G}) \geq D$ , there exist some  $x, y$  for which  $s \in X_1$  and  $t \in Y_1$ , and so by claim 11 at least one pair will have  $Z[x, y] = 1$ .

To summarize, in order to find a pair of points with  $d_G(a, b) \geq \alpha D$  we first run BFS from every point in  $X$  and  $Y$ . We construct the matrices  $M_X^1, M_X^2, M_Y^1, M_Y^2$  and use them to compute  $Z$ . We then find a pair  $x, y$  such that  $Z[x, y] = 1$ . For this pair, we compute the sets  $X_1, X_2, Y_1, Y_2$  and return  $a \in X_1 \setminus Y_2$  and  $b \in Y_1 \setminus X_2$ . These points satisfy  $d_G(a, b) \geq \alpha D$ . If we are unable to find such points we conclude that  $D(\mathcal{G}) < D$ .

The final running time is dominated by  $O(n^\delta m)$  for the BFS searches and  $M(n^\delta, n, n^\delta)$  to perform the matrix multiplications.  $\blacktriangleleft$

### 3.3 Linear Time 3-Approximation for 3-mode Diameter

In this section we consider a 3-multimode graph  $\mathcal{G} = (V, E_1, E_2, E_3)$  and attempt to approximate its 3-mode diameter. We extend our result for 2-multimode graphs and provide a near linear time 3-approximation algorithm for the 3-mode diameter.

To find a 3-approximation to the 3-mode diameter we will look for a pair of points with  $d_G(a, b) \geq \frac{D}{3}$ . We will use a similar approach to the 2-mode case, finding two sets in which all points are far apart in both the red and blue graphs ( $G_1, G_2$ ). We can then compute the ST-diameter of the two sets in green ( $G_3$ ), finding two points with large  $d_3$  distance as well. As in previous sections, we provide an algorithm for unweighted graphs and note that by replacing BFS with Dijkstra's algorithm the algorithm works for weighted graphs, while adding a  $\log n$  factor to the running time.

► **Theorem 13.** *Given a 3-multimode undirected, unweighted graph  $\mathcal{G}$  and integer  $D$ , there exists an  $O(m)$  time algorithm that does one of the following with high probability:*

1. *Finds a pair of points  $a, b$  with  $d_G(a, b) \geq \frac{D}{3}$ .*
2. *Determines that the 3-mode diameter of  $G$  is  $< D$ .*

By performing a binary search over the possible values of  $D$  we obtain our desired result.

**Proof.** Choose an arbitrary vertex  $p \in V$  and run BFS from it in all three graphs,  $G_1, G_2, G_3$ . Let  $X = B_1(p, D/3), Y = B_2(p, D/3) \setminus X$  and  $Z = V \setminus (X \cup Y)$ . If there exists a vertex  $z \in Z$  such that  $d_3(p, z) \geq D/3$ , then  $z \notin \bigcup_{i=1}^3 B_i(p, D/3)$  and so  $d_G(p, z) \geq D/3$ . We can therefore return  $(p, z)$  and finish. Otherwise  $Z \subseteq B_3(p, D/3)$ .

As was the case for 2-mode diameter, we note that for any pair of points  $u, u' \in X$ ,  $d_1(u, u') \leq d_1(u, p) + d_1(p, u') < 2D/3$  and so  $d_G(u, u') < 2D/3$ . Likewise, for any pair  $v, v' \in Y$  or  $w, w' \in Z$  we have  $d_G(v, v') < 2D/3$  and  $d_G(w, w') < 2D/3$ .

Thus, any pair of points within one of the sets  $X, Y, Z$  has 3-mode distance  $< D$ . Next, consider pairs of points  $x \in X$  and  $y \in Y$ . Run BFS in red from the set  $X$  and run BFS in blue from the set  $Y$ . Let  $X_0 = Y \cap B_1(X, D/3)$  be the points in  $Y$  of red distance  $< D/3$  from the set  $X$ . We claim that any pair of points  $x' \in X, y' \in X_0$  have 3-mode distance  $< D$ . For any point  $y \in X_0$  there exists  $x' \in X$  with  $d_1(x', y) < D/3$ . Therefore,  $d_1(x, y) \leq d_1(x, x') + d_1(x', y) < 2D/3 + D/3 = D$ .

Now define  $Y_0 = X \cap B_2(Y, D/3)$ . Using a symmetric argument we can show every pair of points  $y \in Y, x \in Y_0$  have  $d_2$  distance  $< D$ . Therefore, if  $X = Y_0$  or  $Y = X_0$ , then every pair of points  $x \in X, y \in Y$  have distance  $d_G(x, y) < D$ .

Denote  $X_1 = X \setminus X_0$  and  $Y_1 = Y \setminus Y_0$ . We are left to handle the case when both  $X_1, Y_1$  are non-empty. Every pair of points  $x \in X_1, y \in Y_1$  has  $d_1(x, y) \geq D/3$  and  $d_2(x, y) \geq D/3$  so in order to determine if  $d_G(x, y) \geq D/3$  we only need to consider  $d_3(x, y)$ .

To do so, preform a simple approximation to the ST-diameter of the sets  $X_1, Y_1$  in  $G_3$  as follows. Pick two arbitrary nodes  $\hat{x} \in X_1, \hat{y} \in Y_1$  and run BFS from each one of them in  $G_3$ . If  $Y_1 \not\subseteq B_3(\hat{x}, D/3)$  then we will have found a point  $y \in Y_1$  such that  $d_3(\hat{x}, y) \geq D/3$  and so  $d_G(\hat{x}, y) \geq D/3$ . Similarly, if  $X_1 \not\subseteq B_3(\hat{y}, D/3)$ , we will have found a point  $x \in X$  such that  $d_G(x, \hat{y}) \geq D/3$ .

Otherwise,  $X_1 \subseteq B_3(\hat{y}, D/3)$  and  $Y_1 \subseteq B_3(\hat{x}, D/3)$ . Thus, for every pair of points  $x \in X_1, y \in Y_1$ ,

$$d_G(x, y) \leq d_3(x, y) \leq d_3(x, \hat{y}) + d_3(\hat{y}, \hat{x}) + d_3(\hat{x}, y) < D/3 + D/3 + D/3 = D.$$

And so we conclude that any pair of points  $x \in X, y \in Y$  have  $d_G(x, y) < D$ .

We can now repeat this algorithm for the pairs of sets  $X, Z$  and  $Y, Z$  and their respective pairs of colors. If the algorithms finds no pair of points with  $d_G(u, v) \geq D/3$ , we conclude that the 3-mode diameter of  $\mathcal{G}$  is  $< D$ .  $\blacktriangleleft$

### 3.4 Linear Time 3-Approximation for $k$ -mode Radius

In the following section we consider the problem of approximating the  $k$ -mode radius of a  $k$ -multimode graph. Unlike the algorithms we constructed for approximating the  $k$ -mode diameter, in this case we provide a general algorithm for all  $k$ , with running time parameterized by  $k$  and  $m$ . Our algorithm provides a 3-approximation for the  $k$ -mode radius and runs in time  $\tilde{O}(mk!)$ . Thus, for a constant  $k$  we have a near linear time algorithm. We again present our algorithm for unweighted graphs and note that by replacing BFS with Dijkstra's algorithm we can extend the result to weighted graphs while adding  $\log n$  to the running time.

► **Theorem 14.** *There exists an  $\tilde{O}(mk!)$  time algorithm that computes a 3-approximation of the  $k$ -mode radius of an unweighted, undirected  $k$ -multimode graph.*

**Proof.** Our algorithm receives a threshold  $R$  and either finds a vertex  $v$  with  $\text{ecc}(v) \leq 3R$  or determines that  $R(\mathcal{G}) > R$ . Performing a binary search over the values of  $R$  will give the desired approximation.

First we consider the 2-mode radius ( $k = 2$ ) and then we will show how to generalize this algorithm to any  $k$ . Assuming  $R(\mathcal{G}) \leq R$ , our approach will be to find a point that is close to the center  $c$  of the graph in both colors. If we find a vertex  $x$  with  $d_1(c, x) \leq 2R$  and  $d_2(x, c) \leq 2R$ , then for any vertex  $v$  we have  $d_i(c, v) \leq R$  for some  $i \in \{1, 2\}$  and so  $d_i(x, v) \leq 3R$ .

To find such a point, run BFS from an arbitrary node  $z$ . If  $\text{ecc}(z) \leq 3R$  we are finished, otherwise we have a point  $w$  with  $d_G(z, w) > 3R$ . The points  $z, w$  are both within distance  $R$  of the center  $c$  in either red or blue, if  $z, w$  were both within distance  $R$  of the center in the same color we would have a contradiction. Thus, if  $d_G(z, c) = d_1(z, c) \leq R$ , then we must have  $d_G(w, c) = d_2(w, c) \leq R$  and if  $d_G(z, c) = d_2(z, c) \leq R$ , then we must have  $d_G(w, c) = d_1(w, c) \leq R$ .

Let  $X := B_1(z, R) \cap B_2(w, R)$  and  $Y := B_2(z, R) \cap B_1(w, R)$ . By the above observation,  $c$  is either in  $X$  or in  $Y$ . If  $c \in X$ , then any point  $x \in X$  will have  $d_1(x, c) \leq d_1(x, z) + d_1(z, c) \leq 2R$  and  $d_2(x, c) \leq d_2(x, w) + d_2(w, c) \leq 2R$ . This would imply  $\text{ecc}(x) \leq 3R$ . Similarly, if  $c \in Y$  then any point  $y \in Y$  will have  $\text{ecc}(y) \leq 3R$ . Therefore, by choosing arbitrary  $x \in X, y \in Y$  and running BFS from each, we will have found a vertex of eccentricity  $\leq 3R$ .

This algorithm runs a constant number of BFSs and if  $R(\mathcal{G}) \leq R$  it finds a point with  $\text{ecc}(v) \leq 3R$ . Otherwise, we can conclude  $R(\mathcal{G}) > R$ . Therefore, in  $\tilde{O}(m)$  time our algorithm provides a 3-approximation to the 2-mode radius. We can now extend this idea to any value of  $k$ .

For a general  $k$  we can define a recursive algorithm, where at each node we “guess” the color (or mode) in which it achieves its distance to the center (in fact we try all possible colors). We run algorithm 1, RADIUS-3-APPROX, on input  $\mathcal{G}, R$ , a subset of colors  $C$  and a subset of vertices  $W \subseteq V$ . We would like  $\mathcal{G}, R, C, W$  to have the following property:

**(P1)** If  $R(\mathcal{G}) \leq R$ , then  $c \in W$  and  $d_i(w, c) \leq 2R$  for  $w \in W, i \in C$ .

■ **Algorithm 1** RADIUS-3-APPROX.

---

```

1: Input: Graph  $G = (V, E)$ , threshold  $R$ , color subset  $C \subseteq [k]$  and vertex subset  $W \subseteq V$ .
2: Output: Node  $x$  with  $\text{ecc}(x) \leq 3R$  or determine  $R(G) > R$ .
3: Select arbitrary  $x \in W$ .
4: if  $\text{ecc}(x) \leq 3R$  then return  $x$ .
5: end if
6: Let  $y$  be a point with  $d_G(x, y) > 3R$ .
7: for  $i \in [k] \setminus C$  do
8:    $C' \leftarrow C \cup \{i\}, W' \leftarrow W \cap B_i(y, R)$ .
9:   RADIUS-3-APPROX( $G, R, C', W'$ ).
10: end for

```

---

**Correctness.** We run algorithm 1 on  $W = V$  and  $C = \emptyset$  at first, at which point (P1) trivially holds. Subsequently, we claim that if  $C, W$  satisfy (P1), then in at least one of the recursive calls in line 9, the sets  $C', W'$  satisfy (P1).

Let  $x \in W$  and  $y$  be such that  $d_G(x, y) > 3R$  (as chosen in lines 3,6). Since  $d_j(x, c) \leq 2R$  for any  $j \in C$ , we must have  $d_j(c, y) > R$ . Thus, there exists  $i \in [k] \setminus C$  such that  $d_i(c, y) \leq R$ . For this iteration of the for the loop in line 7, we have  $c \in B_i(y, R)$  and so any point  $z \in B_i(y, R)$  has  $d_i(z, c) \leq 2R$ . Hence  $C', W'$  satisfy property (P1).

Therefore, there exists at least one path of length  $k$  in the recursion tree for which (P1) is maintained at all levels of the recursion. Since each call adds a color to  $C$ , after  $k$  calls we will have  $C = [k]$ . In this case,  $W$  is a nonempty set which has  $d_i(c, x) \leq 2R$  for any  $x \in W, i \in [k]$ . Thus, any  $x \in W$  will have  $\text{ecc}(x) \leq 3R$  and will be returned in line 4. If no  $x$  is returned by the algorithm we can conclude that  $R(\mathcal{G}) > R$ .

**Runtime.** Every call to RADIUS-3-APPROX( $G, R, C, W$ ) performs a constant number of BFS searches and  $k - |C|$  calls to RADIUS-3-APPROX with  $|C'| = |C| + 1$ . Therefore, the total running time of the algorithm is  $O(mk!)$ . ◀



## References

- 1 Amir Abboud, Virginia Vassilevska Williams, and Joshua Wang. Approximation and fixed parameter subquadratic algorithms for radius and diameter in sparse graphs. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '16, pages 377–391, USA, 2016. Society for Industrial and Applied Mathematics. doi:10.1137/1.9781611974331.CH28.
- 2 Arturs Backurs, Liam Roditty, Gilad Segal, Virginia Vassilevska Williams, and Nicole Wein. Toward tight approximation bounds for graph diameter and eccentricities. *SIAM Journal on Computing*, 50(4):1155–1199, 2021. doi:10.1137/18M1226737.
- 3 Aaron Berger, Jenny Kaufmann, and Virginia Vassilevska Williams. Approximating Min-Diameter: Standard and Bichromatic. In Inge Li Gørtz, Martin Farach-Colton, Simon J. Puglisi, and Grzegorz Herman, editors, *31st Annual European Symposium on Algorithms (ESA 2023)*, volume 274 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 17:1–17:14, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ESA.2023.17.
- 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 Chris Calabro, Russell Impagliazzo, and Ramamohan Paturi. On the exact complexity of evaluating quantified  $k$ -cnf. In Venkatesh Raman and Saket Saurabh, editors, *Parameterized and Exact Computation - 5th International Symposium, IPEC 2010, Chennai, India, December 13-15, 2010. Proceedings*, volume 6478 of *Lecture Notes in Computer Science*, pages 50–59. Springer, 2010. doi:10.1007/978-3-642-17493-3\_7.
- 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 Shiri Chechik and Tianyi Zhang. Constant approximation of min-distances in near-linear time. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 896–906. IEEE, 2022. doi:10.1109/FOCS54457.2022.00089.
- 8 Lenore Cowen and Christopher G. Wagner. Compact roundtrip routing for digraphs. In Robert Endre Tarjan and Tandy J. Warnow, editors, *Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms, 17-19 January 1999, Baltimore, Maryland, USA*, pages 885–886. ACM/SIAM, 1999. URL: <http://dl.acm.org/citation.cfm?id=314500.315068>.
- 9 Mina Dalirrooyfard and Jenny Kaufmann. Approximation Algorithms for Min-Distance Problems in DAGs. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*, volume 198 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 60:1–60:17, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ICALP.2021.60.
- 10 Mina Dalirrooyfard, Ray Li, and Virginia Vassilevska Williams. Hardness of approximate diameter: Now for undirected graphs. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 1021–1032. IEEE, 2021. doi:10.1109/FOCS52979.2021.00102.
- 11 Mina Dalirrooyfard, Virginia Vassilevska Williams, Nikhil Vyas, and Nicole Wein. Tight approximation algorithms for bichromatic graph diameter and related problems. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019*,



- Patras, Greece*, volume 132 of *LIPICs*, pages 47:1–47:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.ICALP.2019.47.
- 12 Mina Dalirrooyfard, Virginia Vassilevska Williams, Nikhil Vyas, Nicole Wein, Yinzhan Xu, and Yuancheng Yu. Approximation Algorithms for Min-Distance Problems. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, volume 132 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 46:1–46:14, Dagstuhl, Germany, 2019. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPICs.ICALP.2019.46.
  - 13 Mina Dalirrooyfard and Nicole Wein. Tight conditional lower bounds for approximating diameter in directed graphs. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21–25, 2021*, pages 1697–1710. ACM, 2021. doi:10.1145/3406325.3451130.
  - 14 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-sat. *Journal of Computer and System Sciences*, 62(2):367–375, 2001. doi:10.1006/jcss.2000.1727.
  - 15 Lung-Hao Lee and Yi Lu. Multiple embeddings enhanced multi-graph neural networks for chinese healthcare named entity recognition. *IEEE Journal of Biomedical and Health Informatics*, 25(7):2801–2810, 2021. doi:10.1109/JBHI.2020.3048700.
  - 16 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.
  - 17 Liam Roditty and Uri Zwick. On dynamic shortest paths problems. In *Algorithms-ESA 2004: 12th Annual European Symposium, Bergen, Norway, September 14–17, 2004. Proceedings 12*, pages 580–591. Springer, 2004. doi:10.1007/978-3-540-30140-0\_52.
  - 18 Fei Tian, Hanjun Dai, Jiang Bian, Bin Gao, Rui Zhang, Enhong Chen, and Tie-Yan Liu. A probabilistic model for learning multi-prototype word embeddings. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 151–160, 2014. URL: <https://aclanthology.org/C14-1016/>.
  - 19 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.
  - 20 Zixuan Xu Virginia Vassilevska Williams, Yinzhan Xu and Refei Zhou. New bounds for matrix multiplication: from alpha to omega. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3792–3835, 2024. doi:10.1137/1.9781611977912.134.
  - 21 Canwen Xu, Feiyang Wang, Jialong Han, and Chenliang Li. Exploiting multiple embeddings for chinese named entity recognition. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM '19*, pages 2269–2272, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3357384.3358117.