Brief Announcement: A Note on Hardness of Diameter Approximation*

Karl Bringmann[1] **and Sebastian Krinninger**[2]

1 **Max Planck Institute for Informatics, Saarland Informatics Campus, Germany**
`kbringma@mpi-inf.mpg.de`

2 **Faculty of Computer Science, University of Vienna, Austria**
`sebastian.krinninger@univie.ac.at`

Abstract

We revisit the hardness of approximating the diameter of a network. In the CONGEST model, $\tilde{\Omega}(n)$ rounds are necessary to compute the diameter [Frischknecht et al. SODA'12]. Abboud et al. [DISC 2016] extended this result to sparse graphs and, at a more fine-grained level, showed that, for any integer $1 \leq \ell \leq \text{polylog}(n)$, distinguishing between networks of diameter $4\ell + 2$ and $6\ell + 1$ requires $\tilde{\Omega}(n)$ rounds. We slightly tighten this result by showing that even distinguishing between diameter $2\ell + 1$ and $3\ell + 1$ requires $\tilde{\Omega}(n)$ rounds. The reduction of Abboud et al. is inspired by recent conditional lower bounds in the RAM model, where the orthogonal vectors problem plays a pivotal role. In our new lower bound, we make the connection to orthogonal vectors explicit, leading to a conceptually more streamlined exposition. This is suited for teaching both the lower bound in the CONGEST model and the conditional lower bound in the RAM model.

1998 ACM Subject Classification G.2.2 Graph algorithms

Keywords and phrases diameter, fine-grained reductions, conditional lower bounds

Digital Object Identifier 10.4230/LIPIcs.DISC.2017.44

1 Introduction

In distributed computing, the diameter of a network is arguably the single most important quantity one wishes to compute. In the CONGEST model, where in each round every vertex can send to each of its neighbors a message of size $O(\log n)$, it is known that $\tilde{\Omega}(n)$ rounds are necessary to compute the diameter [3] even in sparse graphs [1], where n is the number of vertices. With this negative result in mind, it is natural that the focus has shifted towards *approximating* the diameter. In this note, we revisit hardness of computing a diameter approximation in the CONGEST model from a *fine-grained* perspective.

The current fastest approximation algorithm [4], which is inspired by a corresponding RAM model algorithm [5], takes $O(\sqrt{n \log n} + D)$ rounds and computes a $\frac{3}{2}$-approximation of the diameter, i.e., an estimate $\hat{D}$ such that $\lfloor \frac{2}{3} D \rfloor \leq \hat{D} \leq D$, where D is the true diameter. In terms of lower bounds, Abboud, Censor-Hillel, and Khoury [1] showed that $\tilde{\Omega}(n)$ rounds are necessary to compute a $(\frac{3}{2} - \epsilon)$-approximation of the diameter for any constant $0 < \epsilon < \frac{1}{2}$. At a more fine-grained level, they show that, for any integer $1 \leq \ell \leq \text{polylog}(n)$, at least $\tilde{\Omega}(n)$ rounds are necessary to decide whether the network has diameter $4\ell + 2$ or $6\ell + 1$, thus ruling out any "relaxed" notions of $(\frac{3}{2} - \varepsilon)$-approximation that additionally allow small

* The work of S. Krinninger was partially done while at Max Planck Institute for Informatics. Full version of this paper available at `https://arxiv.org/abs/1705.02127`.

31st International Symposium on Distributed Computing (DISC 2017).
Editor: Andréa W. Richa; Article No. 44; pp. 44:1–44:3
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

additive error. We tighten this result by showing that, for any integer $\ell \geq 1$, at least $\tilde{\Omega}(n)$ rounds are necessary to distinguish between diameter $2\ell + 1$ and $3\ell + 1$.

The reduction of Abboud et al. [1] is inspired by recent work on conditional lower bounds in the RAM model, where the *orthogonal vectors problem* plays a pivotal role. In particular, the Orthogonal Vectors Hypothesis (OVH) is a weaker "polynomial-time analogue" of the Strong Exponential Time Hypothesis (SETH); it is well-known that SETH implies OVH. In our new lower bound, we make the connection to orthogonal vectors explicit: we consider a communication complexity version of orthogonal vectors that we show to be hard *unconditionally* by a reduction from set disjointness and then devise a reduction from orthogonal vectors to diameter approximation. The latter reduction also has implications in the RAM model. We show that under OVH, for any integer $1 \leq \ell \leq n^{o(1)}$, there is no algorithm that distinguishes between graphs of diameter 2ℓ and 3ℓ with running time $O(m^{2-\delta})$ for some constant $\delta > 0$, where m is the number of edges of the graph. This tightens the result of Cairo, Grossi, and Rizzi [2], who provide the same lower bound under the stronger hardness assumption SETH. To summarize, our approach is more streamlined than in previous works [3, 2, 1], allowing for a more unified view of CONGEST model and RAM model lower bounds.

2 Reduction via Orthogonal Vectors

Set disjointness is a problem in communication complexity between two players, called Alice and Bob, in which Alice is given an n-dimensional bit vector x and Bob is given an n-dimensional bit vector y and the goal for Alice and Bob is to find out whether there is some index k at which both vectors contain a 1, i.e., such that $x[k] = y[k] = 1$. The relevant measure in communication complexity is the number of bits exchanged by Alice and Bob in any protocol that Alice and Bob follow to determine the solution. A classic result states that any such protocol requires Alice and Bob to exchange $\Omega(n)$ bits to solve set disjointness.

In the orthogonal vectors problem, Alice is given a set of bit vectors $L = \{l_1, \ldots, l_n\}$ and Bob is given a set of bit vectors $R = \{r_1, \ldots, r_n\}$, and the goal for them is to find out if there is a pair of orthogonal vectors $l_i \in L$ and $r_j \in R$ (i.e., such that, for every $1 \leq k \leq d$, $l_i[k] = 0$ or $r_j[k] = 0$). We give a reduction from set disjointness to orthogonal vectors.

▶ **Theorem 1.** *Any b-bit protocol for the orthogonal vectors problem in which Alice and Bob each hold n vectors of dimension $d = 2\lceil \log n \rceil + 3$, gives a b-bit protocol for the set disjointness problem where Alice and Bob each hold an n-dimensional bit vector.*

▶ **Corollary 2.** *Any protocol solving the orthogonal vectors problem with n vectors of dimension $d = 2\lceil \log n \rceil + 3$, requires Alice and Bob to exchange $\Omega(n)$ bits.*

We now establish hardness of distinguishing between networks of diameter $2\ell + q$ and $3\ell + q$, where $\ell \geq 1$ and in the CONGEST model $q \geq 1$, whereas in the RAM model $q \geq 0$. To unify the cases of odd and even ℓ, we introduce an additional parameter $p \in \{0, 1\}$ and change the task to distinguishing between networks of diameter $4\ell' - 2p + q$ and $6\ell' - 3p + q$ for integers $\ell' \geq 1$, $q \geq 0$, and $p \in \{0, 1\}$. This covers the original question: if ℓ is even, then set $\ell' := \ell/2$ and $p := 0$ and if ℓ is odd, then set $\ell' := \lceil \ell/2 \rceil$ and $p := 1$.

Given an orthogonal vectors instance $\langle L := \{l_1, \ldots, l_n\}, R := \{r_1, \ldots, r_n\}\rangle$ of d-dimensional vectors and parameters $\ell \geq 1$, $q \geq 0$, and $p \in \{0, 1\}$, we define an unweighted undirected graph $G := G_{L,R,\ell,p,q}$ as follows. The graph G contains the following *exterior* vertices: $u_1^L, \ldots, u_n^L$, $u_1^R, \ldots, u_n^R$, $v_1^L, \ldots, v_n^L$, $v_1^R, \ldots, v_n^R$, $w_1^L, \ldots, w_d^L$, $w_1^R, \ldots, w_d^R$, x^L, x^R, y^L, and y^R. These exterior vertices are connected by paths as depicted in Figure 1, where

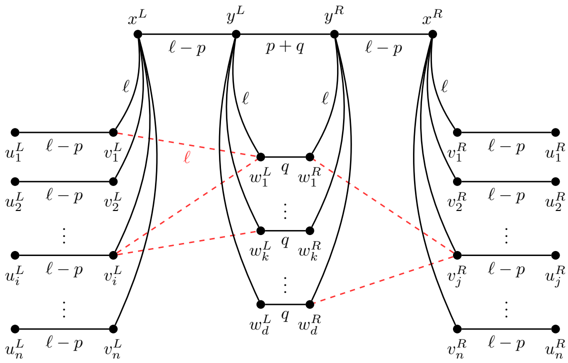

Figure 1 Visualization of the graph $G := G_{L,R,\ell,p,q}$ used in our reduction from orthogonal vectors to diameter distinction. The red, dashed edges encode the orthogonal vectors instance.

each path introduces a separate set of interior vertices. In particular, the instance $\langle L, R\rangle$ is encoded as follows: for every $1 \leq i \leq n$ and every $1 \leq k \leq d$, if $l_i[k] = 1$, then add a path from v_i^L to w_k^L of length ℓ, and if $r_i[k] = 1$, then add a path from v_i^R to w_k^R of length ℓ.

▶ **Theorem 3.** *Let $\langle L, R\rangle$ be an orthogonal vectors instance of two sets of d-dimensional vectors of size n each and let $\ell \geq 1$, $p \in \{0, 1\}$, and $q \geq 0$ be integer parameters. Then the unweighted, undirected graph $G := G_{L,R,\ell,p,q}$ has $O(nd\ell + dq)$ vertices and edges and its diameter D has the following property: if $\langle L, R\rangle$ contains an orthogonal pair, then $D = 6\ell - 3p + q$, and if $\langle L, R\rangle$ contains no orthogonal pair, then $D = 4\ell - 2p + q$.*

For the CONGEST model, observe that G has a small cut of size $d + 1$ between its left hand side and its right hand side. A standard simulation argument, where communication between Alice and Bob is limited to messages sent along the small cut, yields our main result.

▶ **Corollary 4.** *In the CONGEST model, any algorithm distinguishing between graphs of diameter $2\ell + q$ and $3\ell + q$ when $\ell \geq 1$ and $q \geq 1$ requires $\Omega(n/((\ell + q)\log^3 n))$ rounds.*

In the RAM model, the Orthogonal Vectors Hypothesis (OVH) states that there is no algorithm that decides whether a given orthogonal vectors instance contains an orthogonal pair in time $O(n^{2-\delta}\,\mathrm{poly}(d))$ for some constant $\delta > 0$. Under this hardness assumption, our reduction has the following straightforward implication.

▶ **Corollary 5.** *In the RAM model, under OVH, there is no algorithm distinguishing between graphs of diameter $2\ell + q$ and graphs of diameter $3\ell + q$ when $\ell \geq 1$ and $q \geq 0$ in time $O(m^{2-\delta}/(\ell + q)^{2-\delta})$ for any constant $\delta > 0$.*

References

1 Amir Abboud, Keren Censor-Hillel, and Seri Khoury. Near-linear lower bounds for distributed distance computations, even in sparse networks. In *DISC*, pages 29–42, 2016.

2 Massimo Cairo, Roberto Grossi, and Romeo Rizzi. New bounds for approximating extremal distances in undirected graphs. In *SODA*, pages 363–376, 2016.

3 Silvio Frischknecht, Stephan Holzer, and Roger Wattenhofer. Networks cannot compute their diameter in sublinear time. In *SODA*, pages 1150–1162, 2012.

4 Stephan Holzer, David Peleg, Liam Roditty, and Roger Wattenhofer. Brief announcement: Distributed 3/2-approximation of the diameter. In *DISC*, pages 562–564, 2014.

5 Liam Roditty and Virginia Vassilevska Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In *STOC*, pages 515–524, 2013.