Information Processing Letters 133 (2018) 10–15

Contents lists available at ScienceDirect

Information Processing Letters

www.elsevier.com/locate/ipl

A note on hardness of diameter approximation

Karl Bringmann [a], Sebastian Krinninger [b,*,1]

[a] Max Planck Institute for Informatics, Saarland Informatics Campus, Campus E1 4, 66123 Saarbrücken, Germany
[b] University of Salzburg, Department of Computer Sciences, Jakob-Haringer-Straße 2, 5020 Salzburg, Austria

ARTICLE INFO

Article history:
Received 12 May 2017
Received in revised form 21 December 2017
Accepted 29 December 2017
Available online 3 January 2018
Communicated by B. Doerr

Keywords:
Distributed computing
Lower bounds
Hardness in P
Graph algorithms

ABSTRACT

We revisit the hardness of approximating the diameter of a network. In the CONGEST model of distributed computing, $\tilde{\Omega}(n)$ rounds are necessary to compute the diameter (Frischknecht et al., 2012 [2]), where $\tilde{\Omega}(\cdot)$ hides polylogarithmic factors. Abboud et al. (2016) [3] 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.

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 [1], 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 [2] even in sparse graphs [3], 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 of the network. In terms of lower bounds, Abboud, Censor-Hillel, and Khoury [3] 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 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 $3\ell + 1$, and more generally between diameter $2\ell + q$ and $3\ell + q$ for any $\ell, q \geq 1$.

The reduction of Abboud et al. [3] is inspired by recent work on 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 between diameter approximation and orthogonal vectors explicit: we consider a communication complexity version of orthogonal vectors that we show to be hard *unconditionally*

* Corresponding author.
E-mail addresses: kbringma@mpi-inf.mpg.de (K. Bringmann), sebastian.krinninger@sbg.ac.at (S. Krinninger).

[1] Work partially done while at Max Planck Institute for Informatics, Saarland Informatics Campus, Germany, and while at University of Vienna, Faculty of Computer Science, Austria.

https://doi.org/10.1016/j.ipl.2017.12.010
0020-0190/

by a reduction from set disjointness and then devise a reduction from orthogonal vectors to diameter approximation.

Additionally, our approach has implications in the RAM model. There, the *Strong Exponential Time Hypothesis (SETH)* [6] states that for every $\delta > 0$ there is an integer $k \geq 3$ such that k-SAT admits no algorithm with running time $O(2^{(1-\delta)N})$ and the *Orthogonal Vectors Hypothesis (OVH)* states that there is no algorithm to decide whether a given set of d-dimensional vectors of length n contains an orthogonal pair in time $O(n^{2-\delta}\operatorname{poly}(d))$ for any constant $\delta > 0$. It is well-known that SETH implies OVH [7]. Prior to our work, the situation in the RAM model was as follows. In their seminal paper [5], Roditty and Vassilevska Williams showed that, for any constants $\epsilon > 0$ and $\delta > 0$ there is no algorithm that computes a $(\frac{3}{2} - \epsilon)$-approximation of the diameter and runs in time $O(m^{2-\delta})$, unless the Strong Exponential Time Hypothesis (SETH) fails. In particular, they show that no algorithm can decide whether a given graph has diameter 2 or 3 in time $O(m^{2-\delta})$, unless the Strong Exponential Time Hypothesis (SETH) fails. The hardness of 2 vs. 3 is already implied by the weaker Orthogonal Vectors Hypothesis (OVH), which in turn is implied by SETH [7] and was popularized after the paper of Roditty and Vassilevska Williams appeared. It has then been shown by Chechik et al. [8] that, for any integer $1 \leq \ell \leq n^{o(1)}$, there is no algorithm that distinguishes between diameter $3(\ell + 1)$ and $4(\ell + 1)$ with running time $O(m^{2-\delta})$ for some constant $\delta > 0$, unless SETH fails. Finally, Cairo, Grossi, and Rizzi [9] showed that, for any integer $1 \leq \ell \leq n^{o(1)}$, there is no algorithm that distinguishes between diameter 2ℓ and 3ℓ with running time $O(m^{2-\delta})$ for some constant $\delta > 0$, unless SETH fails. Our reduction reconstructs the result of Cairo et al. under the weaker hardness assumption OVH, yielding again a more streamlined chain of reductions.

2. Reduction from set disjointness to 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 $x[k] = y[k] = 1$ (meaning the sets represented by x and y are not disjoint). 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 [10,11] 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 $L = \{l_1, \dots, l_n\}$ and Bob is given a set of bit vectors $R = \{r_1, \dots, r_n\}$, 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 $l_i[k] = 0$ or $r_j[k] = 0$ in each dimension k). We give a reduction from set disjointness to orthogonal vectors.

Theorem 2.1. *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.*

Proof. We show that, without any communication, Alice and Bob can transform a set disjointness instance $\langle x, y \rangle$ with n-dimensional bit vectors into an orthogonal vectors instance $\langle L, R \rangle$ such that x and y are not disjoint if and only if $\langle L, R \rangle$ contains an orthogonal pair. For every integer $1 \leq i \leq n$, let s_i denote the binary representation of i with $\lceil \log n \rceil$ bits. For every bit b, let $\bar{b}$ be the result of 'flipping' bit b, i.e., $\bar{1} = 0$, and $\bar{0} = 1$. Similarly, for a bit vector b, let $\bar{b}$ be the result of flipping each bit of b. For every $1 \leq i \leq n$, let l_i be the vector obtained from concatenating $x[i]$, $\bar{x}[i]$, $\bar{x}[i]$, s_i, and $\bar{s}_i$. For every $1 \leq j \leq n$, let r_i be the vector obtained from concatenating $\bar{y}[i]$, $y[i]$, $\bar{y}[i]$, $\bar{s}_i$, and s_i.

We now claim that the vectors x and y are not disjoint if and only if $\langle L, R \rangle$ contains an orthogonal pair. If the vectors x and y are not disjoint, then there is some i such that $x[i] = y[i] = 1$. Clearly, s_i and $\bar{s}_i$ are orthogonal and, as the vectors $(x[i], \bar{x}[i], \bar{x}[i])$ and $(\bar{y}[i], y[i], \bar{y}[i])$ are equal to $(1, 0, 0)$ and $(0, 1, 0)$, respectively, they are also orthogonal. It follows that l_i and r_i are orthogonal.

Now assume that $\langle L, R \rangle$ contains an orthogonal pair $l_i \in L$ and $r_j \in R$. We first show that $i = j$. Suppose for the sake of contradiction that $i \neq j$. Then the binary representations s_i and s_j differ in at least one bit, say $s_i[k] \neq s_j[k]$. If $s_i[k] = 0$ and $s_j[k] = 1$, then $\bar{s}_i$ and s_j are not orthogonal and thus l_i and r_j are not orthogonal, contradicting the assumption. If $s_i[k] = 1$ and $s_j[k] = 0$, then s_i and $\bar{s}_j$ are not orthogonal and thus l_i and r_j are not orthogonal, contradicting the assumption. It follows that $i = j$ and thus the vectors $(x[i], \bar{x}[i], \bar{x}[i])$ and $(\bar{y}[i], y[i], \bar{y}[i])$ are orthogonal. Orthogonality of $x[i]$ and $\bar{y}[i]$ rules out $x[i] = 1$ and $y[i] = 0$, orthogonality of $\bar{x}[i]$ and $y[i]$ rules out $x[i] = 0$ and $y[i] = 1$, and orthogonality of $\bar{x}[i]$ and $\bar{y}[i]$ rules out $x[i] = 0$ and $y[i] = 0$. It follows that $x[i] = y[i] = 1$, making x and y not disjoint. $\square$

The hardness of set disjointness now directly transfers to orthogonal vectors.

Corollary 2.2. *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.*

3. Reduction from orthogonal vectors to diameter

We now establish hardness of distinguishing between networks of diameter $2\ell + q$ and $3\ell + q$ for any $\ell \geq 1$ and $q \geq 1$ in the CONGEST model and for any $\ell \geq 1$ and $q \geq 0$ in the RAM model, respectively. 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 $p := 1$.