

Distributed Dense Subgraph Detection and Low Outdegree Orientation

Hsin-Hao Su

Boston College, Chestnut Hill, MA, USA

suhx@bc.edu

Hoa T. Vu

San Diego State University, CA, USA

hvu2@sdsu.edu

Abstract

The densest subgraph problem, introduced in the 80s by Picard and Queyranne [Networks 1982] as well as Goldberg [Tech. Report 1984], is a classic problem in combinatorial optimization with a wide range of applications. The lowest outdegree orientation problem is known to be its dual problem. We study both the problem of finding dense subgraphs and the problem of computing a low outdegree orientation in the distributed settings.

Suppose $G = (V, E)$ is the underlying network as well as the input graph. Let D denote the density of the maximum density subgraph of G . Our main results are as follows.

- Given a value $\tilde{D} \leq D$ and $0 < \epsilon < 1$, we show that a subgraph with density at least $(1 - \epsilon)\tilde{D}$ can be identified deterministically in $O((\log n)/\epsilon)$ rounds in the LOCAL model. We also present a lower bound showing that our result for the LOCAL model is tight up to an $O(\log n)$ factor. In the CONGEST model, we show that such a subgraph can be identified in $O((\log^3 n)/\epsilon^3)$ rounds with high probability. Our techniques also lead to an $O(\text{diameter}(G) + (\log^4 n)/\epsilon^4)$ -round algorithm that yields a $1 - \epsilon$ approximation to the densest subgraph. This improves upon the previous $O(\text{diameter}(G)/\epsilon \cdot \log n)$ -round algorithm by Das Sarma et al. [DISC 2012] that only yields a $1/2 - \epsilon$ approximation.
- Given an integer $\tilde{D} \geq D$ and $\Omega(1/\tilde{D}) < \epsilon < 1/4$, we give a deterministic, $\tilde{O}((\log^2 n)/\epsilon^2)$ -round¹ algorithm in the CONGEST model that computes an orientation where the outdegree of every vertex is upper bounded by $(1 + \epsilon)\tilde{D}$. Previously, the best deterministic algorithm and randomized algorithm by Harris [FOCS 2019] run in $\tilde{O}((\log^6 n)/\epsilon^4)$ rounds and $\tilde{O}((\log^3 n)/\epsilon^3)$ rounds respectively and only work in the LOCAL model.

2012 ACM Subject Classification Theory of computation → Distributed algorithms; Networks → Network algorithms; Mathematics of computing → Graph algorithms

Keywords and phrases Distributed Algorithms, Network Algorithms

Digital Object Identifier 10.4230/LIPIcs.DISC.2020.15

Related Version A full version of the paper is available at <https://arxiv.org/abs/1907.12443>.

Acknowledgements We thank David Harris for pointing out Lemma 12, which can be used to improve the running time of the dual rounding algorithm in our previous version.

1 Introduction

The Dense Subgraph Problem. Given a graph $G = (V, E)$, the *maximum density subgraph* problem (or the densest subgraph problem) is to find a subgraph H , where its density $d(H) = |E(H)|/|V(H)|$ is maximized over all subgraphs of G . We denote the density of the maximum density subgraph of G , $\max_{H \subseteq G} d(H)$, by D .

¹ \tilde{O} hides $\log \log n$ factors.

First studied by Picard and Queyranne [37] as well as Goldberg [26], the maximum density subgraph problem has found numerous applications in community detection in social networks [13, 14], link spam identification [8, 25], and computational biology [18, 38]. Faster algorithms [12, 20, 30] have been developed for the problem and its variants since then. Moreover, the problem has been widely studied under different models of computation such as the streaming settings [9, 15, 34], the dynamic setting [9, 40], the massive parallel computation settings [4, 5, 23], and the distributed settings [39].

We study the problem of detecting dense subgraphs in the distributed settings, namely, in both the **LOCAL** and the **CONGEST** models. Let n and m be the number of vertices and the number of edges respectively. Furthermore, let Δ be the maximum degree. In such models, vertices are labeled with unique IDs and they operate in synchronized rounds. In each round, each vertex sends a message to each of its neighbors, receives messages from its neighbors, and performs local computations. The time complexity of an algorithm is defined to be the number of rounds used. In the **LOCAL** model, the message size can be arbitrary. In the **CONGEST** model, the message size is upper bounded by $O(\log n)$ bits. We consider the following parameterized version of the maximum density subgraph problem in the distributed settings, which may capture the computational nature of some applications such as how a dense community can be found by only communicating with the neighbors in social networks.

DENSESUBGRAPH(\tilde{D}, ϵ) : Given a graph $G = (V, E)$, a parameter $\tilde{D} \geq 0$, and $0 < \epsilon < 1$, every vertex u outputs a value $h_u \in \{0, 1\}$ such that $d(H) \geq (1 - \epsilon)\tilde{D}$ where $H = \{u \mid h_u = 1\}$. If $\tilde{D} \leq D$ then H must be non-empty.

The first question we investigate is whether **DENSESUBGRAPH**(\tilde{D}, ϵ) can be solved locally. Intuitively, most dense subgraphs have short diameters because they are well-connected. Thus, they can be detected locally. Our first result justifies this intuition.

► **Theorem 1.** *There exists a deterministic algorithm for **DENSESUBGRAPH**(\tilde{D}, ϵ) that runs in $O((\log n)/\epsilon)$ rounds in the **LOCAL** model.*

We will also present a lower bound showing that the running time of the algorithm is tight up to an $O(\log n)$ factor. The algorithm for the **LOCAL** model uses large message size. This begs the question of whether the problem can be solved in the **CONGEST** model while remaining in the $\text{poly}(1/\epsilon, \log n)$ -round regime. We show that this is indeed possible with randomization.

► **Theorem 2.** *There exists a randomized algorithm that solves **DENSESUBGRAPH**(\tilde{D}, ϵ) w.h.p.², and runs in $O((\log^3 n)/\epsilon^3)$ rounds in the **CONGEST** model.*

Finding the densest subgraph in such distributed settings inevitably requires $\Omega(\text{diameter}(G))$ rounds (e.g., consider two subgraphs of different densities connected by a path of length $\Omega(\text{diameter}(G))$). Das Sarma et al. [39] gave an algorithm for finding a $(1/2 - \epsilon)$ approximation to the densest subgraph in $O(\text{diameter}(G)/\epsilon \cdot \log n)$ rounds in the **CONGEST** model. We show that the approximation factor can be improved and the dependency on $\text{diameter}(G)$ can be made additive:

² W.h.p. denotes with high probability, which means with probability at least $1 - 1/n^c$ for an arbitrarily large constant c .

► **Corollary 3.** *There exists a randomized algorithm that finds a $(1 - \epsilon)$ -approximation to the maximum density subgraph w.h.p. and runs in $O(\text{diameter}(G) + (\log^4 n)/\epsilon^4)$ rounds in the CONGEST model.*

Inspired by web-graphs, Kannan and Vinay [29] defined the notion of density in directed graphs. Suppose that $G = (V, E)$ is a directed graph. The density of a pair of sets $S, T \subseteq V$ is defined as $d(S, T) = \frac{|E(S, T)|}{\sqrt{|V(S)||V(T)|}}$, where $E(S, T)$ denote the set of edges that go from a vertex in S to a vertex in T . Note that we are assuming messages can go in both directions of an edge. Our result in the LOCAL model can be generalized to the directed version of the problem.

The Low Outdegree Orientation Problem. Given an undirected graph $G = (V, E)$, an α -orientation is an orientation of the edges such that the outdegree of every vertex is upper bounded by α . Picard and Queyranne [37] observed that an α -orientation exists if and only if $\alpha \geq \lceil D \rceil$. Charikar [12] formulated the linear program (LP) for the densest subgraph problem, and its dual is the fractional version of the lowest outdegree orientation problem [4].

An α -orientation can be used to obtain a decomposition of the graph into α *pseudoforests* [37], where a pseudoforest is a collection of disjoint *pseudotrees* (a pseudotree is a connected graph containing at most one cycle). The relation between the pseudoforest decomposition and the maximum density is analogous to that of the forest decomposition and the arboricity shown by Nash-Williams [36].

We consider the low outdegree orientation problem in the distributed setting. The problem can be formally stated as follows.

Given a graph $G = (V, E)$, an integer parameter $\tilde{D} \geq D$, and $0 < \epsilon < 1$, compute a $(1 + \epsilon)\tilde{D}$ -orientation. The orientation of an edge is decided by either of its endpoints.

Our contribution for this problem is a deterministic, $\tilde{O}((\log^2 n)/\epsilon^2)$ -round algorithm in the CONGEST model. Previously, the best deterministic algorithm and randomized algorithm by Harris [28] run in $\tilde{O}((\log^6 n)/\epsilon^4)$ rounds and $\tilde{O}((\log^3 n)/\epsilon^3)$ rounds respectively and only work in the LOCAL model.

► **Theorem 4.** *Given an integer $\tilde{D} \geq D$, for any $32/\tilde{D} \leq \epsilon \leq 1/4$, there exists a deterministic algorithm in the CONGEST model that computes a $(1 + \epsilon)\tilde{D}$ -orientation and runs in*

$$O\left(\frac{\log n}{\epsilon^2} + (\min(\log \log n, \log \Delta) + \log(1/\epsilon))^{2.71} \cdot (1/\epsilon)^{1.71} \cdot \log^2 n\right) \leq \tilde{O}((\log^2 n)/\epsilon^2) \text{ rounds.}$$

► **Table 1** Previous results on the low outdegree orientation problem in the distributed setting.

Reference	Time	Model	Approx.	Rand. or Det.
Barenboim and Elkin [6]	$O((\log n)/\epsilon)$	CONGEST	$2 + \epsilon$	Det.
Ghaffari and Su [24]	$O((\log^4 n)/\epsilon^3)$	LOCAL	$1 + \epsilon$	Rand.
Fischer et al. [17]	$2^{O(\log^2(1/\epsilon \cdot \log n))}$	LOCAL	$1 + \epsilon$	Det.
Ghaffari et al. [21]	$O((\log^{10} n \cdot \log^5 \Delta)/\epsilon^9)$	LOCAL	$1 + \epsilon$	Det.
Harris [28]	$\tilde{O}((\log^6 n)/\epsilon^4)$	LOCAL	$1 + \epsilon$	Det.
Harris [28]	$\tilde{O}((\log^3 n)/\epsilon^3)$	LOCAL	$1 + \epsilon$	Rand.
new	$\tilde{O}((\log^2 n)/\epsilon^2)$	CONGEST	$1 + \epsilon$	Det.

Table 1 summarizes previous algorithms for this problem in the LOCAL and CONGEST models. It is worthwhile to note that $\lceil D \rceil \leq a(G) \leq \lceil D \rceil + 1$ [37], where $a(G)$ is the arboricity of the graph, as several previous results were parameterized in terms of $a(G)$.

Barenboim and Elkin [6] introduced the H-partition algorithm that obtains a $(2 + \epsilon)a(G)$ -orientation as well as a $(2 + \epsilon)a(G)$ forest decomposition. Ghaffari and Su [24] observed that the problem of computing a $(1 + \epsilon)a(G)$ -orientation reduces to computing of maximal independent sets (MIS) on the conflict graphs formed by augmenting paths. The MIS can be computed efficiently by simulating Luby’s randomized MIS on the conflict graph. Fischer et al. [17] initiated the study of computing such MIS (i.e. the maximal matching in hypergraphs) deterministically. They gave a deterministic quasi-polynomial (in r) algorithm for computing the maximal matching in rank r hypergraphs, resulting in a $2^{O(\log^2(1/\epsilon) \cdot \log n)}$ round algorithm for the orientation problem. Later, the dependency on r has been improved to polynomial by [21], which results in a $O((\log^{10} n \cdot \log^5 \Delta)/\epsilon^9)$ rounds algorithm for the orientation problem. Recently, the deterministic running time for the problem is further improved by Harris [28] to $\tilde{O}((\log^6 n)/\epsilon^4)$ as they developed a faster algorithm for the hypergraph maximal matching problem. It is unclear if the above approaches via maximal matching in hypergraphs can be implemented in the CONGEST model without significantly increase on the number of rounds. The low outdegree orientation problem has also been studied in the centralized context by [7, 10, 19, 27, 31, 32].

Our methods and contributions in a nutshell. Our first contribution is a simple yet powerful observation that dense subgraphs have low-diameter approximations. We first give a simple proof via low-diameter decomposition [33, 35] that there exists a subgraph with diameter $O((\log n)/\epsilon)$ that has density at least $(1 - \epsilon)D$ (Lemma 5). Hence, if each vertex examines its local neighborhood up to a small radius, at least one vertex gets a good estimate of the densest subgraph. With appropriate bookkeeping, this leads to our *deterministic* $O((\log n)/\epsilon)$ -round algorithm for detecting dense subgraphs in the LOCAL model. We complement this algorithm with a lower bound showing that $\Omega(1/\epsilon)$ rounds are necessary (Lemma 6).

In the CONGEST model, the starting point for both problems of detecting dense subgraphs and low outdegree orientation is the adaptation of the multiplicative weights update algorithm of Bahmani et al. [4]. Their algorithm solves the dual linear program for $\text{DENSESUBGRAPH}(\tilde{D}, \epsilon)$ (which is the linear program for low outdegree orientation) using the multiplicative weights method. They also showed how to round the dual program’s fractional solution to find a dense subgraph. While it appears that their algorithm can be implemented directly in the CONGEST model, there are a few issues that we need to resolve.

A naive implementation of the algorithm in [4] to solve the dual linear program uses $O(\text{diameter}(G))$ rounds per iteration. Furthermore, for the dense subgraph detection problem $\text{DENSESUBGRAPH}(\tilde{D}, \epsilon)$, the rounding procedures in [4, 12] are inherently global. This is because they require checking whether certain subgraphs have high enough density. These subgraphs may have large diameters. Our contribution here is to remove the dependence on $\text{diameter}(G)$ using ideas from Lemma 5 along with appropriate bookkeeping. This results in a randomized algorithm that runs in $O((\log^3 n)/\epsilon^3)$ rounds instead of $O(\text{diameter}(G) \cdot \text{poly}(\log n, 1/\epsilon))$ rounds. Our adaption also removes the explicit use of real-valued weights and only use integers, which can be transmitted easily in the CONGEST model.

For the low outdegree orientation problem, we still need to develop a procedure to round the dual fractional solution obtained from the multiplicative weights update method into an integral solution efficiently and deterministically in the CONGEST model. Our

contribution here is twofold. We use the idea of recent rounding-type algorithms for distributed matching [1, 16], where we process the fractional solution bit-by-bit and round the solution at each bit scale. We show that each scale reduces to the directed splitting problem where there are known algorithms for the LOCAL model [22, 24]. We will show how to modify these directed splitting algorithms to run in the CONGEST model. We provide an analysis showing that the total error incurred by multiple bit scales is small so that each outdegree is at most $(1 + \epsilon)D$ at the end of the rounding procedure.

Organization. In Section 2, we present our deterministic algorithm for the dense subgraph problem in the LOCAL model. Section 3 presents a randomized algorithm for the dense subgraph problem in the CONGEST model. Section 4 exhibits a deterministic algorithm for the low outdegree orientation problem in the CONGEST model. Due to space constraint, some proofs are omitted; all missing proofs can be found in the full version [41].

2 Deterministic Dense Subgraph Detection in the LOCAL Model

In this section, we investigate the locality of the dense subgraph detection problem. In particular, we show how to solve $\text{DENSESUBGRAPH}(\tilde{D}, \epsilon)$ *deterministically* in the LOCAL model. Combining with ideas in this section, we show that this problem can also be solved in the CONGEST model with randomization in the next section.

The locality of dense subgraphs. We present an algorithm to solve $\text{DENSESUBGRAPH}(\tilde{D}, \epsilon)$ in $O((\log n)/\epsilon)$ rounds deterministically in the LOCAL model. We first give a structural lemma showing that for some sufficiently large constant K , there exists a subgraph with diameter at most $K/\epsilon \cdot \log n$ that has density at least $(1 - \epsilon)D$.

► **Lemma 5** (Densest subgraph's locality). *For all simple graphs, there exists a subgraph with diameter at most $K/\epsilon \cdot \log n$ for some sufficiently large constant K that has density at least $(1 - \epsilon)D$.*

Proof. It can be shown that for any simple graph G with n vertices and m edges, we can decompose G into disjoint components such that each component has diameter at most $K/\epsilon \cdot \log n$ for some sufficiently large constant K and furthermore the number of inter-component edges is at most ϵm [3, 33, 35] (see also Theorem 11). This is known as the *low-diameter decomposition*.

Consider the densest subgraph $H^* \subseteq G$ with n^* vertices and m^* edges. We apply the low-diameter decomposition to H^* and let the components be H_1^*, \dots, H_t^* . Let n_i^* and m_i^* be the number of vertices and edges in H_i^* respectively. Suppose that $d(H_i^*) < (1 - \epsilon)D$ for all i . Then, $m_i^*/n_i^* < (1 - \epsilon)D$ which implies $m_i^* < (1 - \epsilon)m^*n_i^*/n^*$. Thus,

$$\sum_{i=1}^t m_i^* < (1 - \epsilon) \sum_{i=1}^t \frac{m^* n_i^*}{n^*} = (1 - \epsilon)m^* .$$

This implies that the number of inter-component edges is more than ϵm^* which is a contradiction. Therefore, at least one component H_j^* must have density $d(H_j^*) \geq (1 - \epsilon)D$. ◀

Using Lemma 5, we can design a LOCAL algorithm to solve $\text{DENSESUBGRAPH}(\tilde{D}, \epsilon)$ in $O((\log n)/\epsilon)$ rounds. See Algorithm 1.

► **Theorem 1.** *There exists a deterministic algorithm for $\text{DENSESUBGRAPH}(\tilde{D}, \epsilon)$ that runs in $O((\log n)/\epsilon)$ rounds in the LOCAL model.*

■ **Algorithm 1** Distributed Algorithm for $\text{DENSESUBGRAPH}(\tilde{D}, \epsilon)$ in the LOCAL model.

- 1: Initialize $h_v \leftarrow 0$ for all vertices v .
- 2: Let $r = K/\epsilon \cdot \log n$ for some sufficiently large constant K .
- 3: Each vertex v collects the subgraph induced by its r -neighborhood $N_r(v) = \{u : \text{dist}(v, u) \leq r\}$ and computes the densest subgraph $H(v)$ in $G[N_r(v)]$, i.e., the subgraph induced by $N_r(v)$.
- 4: **if** $d(H(v)) \geq (1 - \epsilon)\tilde{D}$ **then**
- 5: v becomes an active vertex and collects its $2r$ -neighborhood $N_{2r}(v)$.
- 6: **if** an active vertex v has the smallest ID among $N_{2r}(v)$ **then**
- 7: v becomes a black vertex.
- 8: Each black vertex v broadcasts $H(v)$ to $N_r(v)$ and set $h_u \leftarrow 1$ for all u in $V(H(v))$.

Proof. Consider Algorithm 1. The number of rounds is clearly $O(r) = O(1/\epsilon \cdot \log n)$, since it is dominated by the steps in Lines 3, 5, and 8 which is $O(r)$. Appealing to Lemma 5, if $\tilde{D} \leq D$, then there must be a subgraph C with diameter at most r whose density is at least $(1 - \epsilon)D \geq (1 - \epsilon)\tilde{D}$. Therefore, at least one vertex must be active. Among the active vertices, there must be a black vertex, i.e., the active vertex with the smallest ID. We therefore have a non-empty output.

The next observation is that if u and v are black vertices, then $H(v)$ and $H(u)$ are disjoint. Otherwise, there is a path of length at most $2r$ from u to v which leads to a contradiction. This is because u and v cannot both be black vertices if $\text{dist}(u, v) \leq 2r$. Furthermore, it is easy to see that if two subgraphs $G[A]$ and $G[B]$, where $A, B \subset V$ are disjoint, have density at least $(1 - \epsilon)\tilde{D}$, then $d(G[A \cup B]) \geq (1 - \epsilon)\tilde{D}$. To see this,

$$d(G[A \cup B]) \geq \frac{|E(A)| + |E(B)|}{|A| + |B|} \geq \frac{(1 - \epsilon)\tilde{D}(|A| + |B|)}{|A| + |B|} = (1 - \epsilon)\tilde{D}.$$

Let the set of black vertices be B . Since we argued that $H(v)$'s are disjoint for $v \in B$ and the output subgraph is $H = \cup_{v \in B} H(v)$, we deduce that $d(H) \geq (1 - \epsilon)\tilde{D}$.

Finally, we need to argue that the algorithm is correct for when $\tilde{D} > D$. Note that if the output subgraph is non-empty, its density is at least $(1 - \epsilon)\tilde{D}$. Hence, if $\tilde{D} > D$, the algorithm may output an empty subgraph or a subgraph with density $(1 - \epsilon)\tilde{D}$ which are both acceptable. ◀

We show that Theorem 1 is tight in terms of ϵ up to a constant.

► **Lemma 6.** *Let $0 < \epsilon < 0.1$. Any algorithm that solves $\text{DENSESUBGRAPH}(\tilde{D}, \epsilon)$ correctly with probability at least 0.51 requires more than $1/(10\epsilon)$ rounds.*

Proof. Consider deterministic algorithms for $\text{DENSESUBGRAPH}(\tilde{D}, \epsilon)$ where $\tilde{D} = 1 - \epsilon$. Without loss of generality, assume $1/(10\epsilon)$ is an integer and let $\ell = 4/(10\epsilon) + 1$. Consider ℓ vertices $1, \dots, \ell$ where $(i, i+1) \in E$ for $1 \leq i \leq \ell - 1$. We consider two cases. Case 1: $(\ell, 1) \in E$ in which the graph (called G_1) is a cycle. Case 2: $(\ell, 1) \notin E$ in which the graph (called G_2) is a chain of ℓ vertices. Let $v = \lfloor \ell/2 \rfloor + 1$. In both cases, the $(1/(10\epsilon))$ -neighborhood of v is the same and therefore h_v must be the same regardless of whether the network is G_1 or G_2 . We observe that a chain of t vertices has density $1 - 1/t$.

If $h_v = 0$, then let the underlying graph be G_1 . Then $D = 1 > 1 - \epsilon = \tilde{D}$, and therefore the output must be non-empty. The only correct output is when every vertex outputs 1 since otherwise the output subgraph's density is at most $1 - 1/(4/(10\epsilon) + 1) < (1 - \epsilon)^2 = (1 - \epsilon)\tilde{D}$ for $\epsilon < 0.1$. Hence, the algorithm fails.

If $h_v = 1$, then let the underlying graph be G_2 . Then $D = 1 - 1/(4/(10\epsilon) + 1) < (1 - \epsilon)\tilde{D}$. Then, the only correct output is that every vertex outputs 0. Therefore, the algorithm fails.

For a randomized lower bound, we choose the above two inputs with probability 0.5 each and therefore the probability that a deterministic algorithm is correct is at most 0.5. By Yao's minimax principle, no randomized algorithm outputs correctly with probability more than 0.5. \blacktriangleleft

We remark that an interesting *open question* is whether the $\log n$ factor is necessary for either the lower bound or the upper bound. In our full version [41], we show that similar locality result also holds for directed densest subgraph.

3 Dense Subgraph Detection in the CONGEST Model

Relating the densest subgraph problem and the lowest outdegree orientation problem. We show how to find a) a dense subgraph and b) a low outdegree orientation (in Section 4) by first solving the same feasibility LP. Let us first consider the LP formulations in Figure 1.

(PRIMAL) Maximum Density Subgraph	(DUAL) Lowest Outdegree Orientation
$\text{maximize} \quad \sum_{e \in E} x_e$ <p>subject to</p> $\sum_{v \in V} y_v = 1$ $\forall e = uv \in E, \quad x_e \leq y_u \text{ and } x_e \leq y_v$ $\forall e \in E, u \in V, \quad y_e, x_u \geq 0 .$	$\text{minimize} \quad z$ <p>subject to</p> $\forall e = uv \in E, \quad \alpha_{eu} + \alpha_{ev} \geq 1$ $\forall e \in E, u \in V, \quad \sum_{e \ni u} \alpha_{eu} \leq z$ $\forall e = uv \in E, \quad \alpha_{eu}, \alpha_{ev} \geq 0 .$

■ **Figure 1** Linear programs for densest subgraph and fractional lowest outdegree orientation.

Given a subgraph $H \subseteq G$ of size k , in the primal, we can set $y_v = 1/k$ for all $v \in V(H)$ and $x_e = 1/k$ for all edges $e \in E(H)$ while setting other variables to 0. Then, $\sum_{e \in E} x_e = |E(H)|/k = d(H)$. In fact, the optimal value of the LP is exactly the maximum subgraph density D . Charikar gave a rounding algorithm that recovers the densest subgraph [12].

We observe that the dual models the lowest outdegree orientation problem. In particular, if an edge $e = uv$ is oriented from u to v then we set $\alpha_{eu} = 1$ and $\alpha_{ev} = 0$. By duality, the dual is fractionally feasible if and only if $z \geq D$.

Now we consider the feasibility versions of the programs. We say a primal solution is a solution satisfying PRIMAL(z) if it is a feasible solution whose objective function is at least z . Similarly, we say a dual solution is a solution satisfying DUAL(z) if it is a feasible dual solution whose objective function is at most z .

Adapting the algorithm of Bahmani et al. [4]. Algorithm 2 and Algorithm 3 are adapted from the multiplicative weights update approach of [4] for solving DUAL and PRIMAL. We modify them in a way so that we only have to operate with integers instead of real-valued weights. This is more suitable for the CONGEST model because it may take $\omega(1)$ rounds to transfer a real-valued weight over an edge. Similar to the tree packing method for minimum cuts [42], this is another example of a combinatorial algorithm derived from multiplicative weights update method, where the weights are only used in the analysis but not in the algorithms. However, as mentioned in the first section, this adaptation has a dependence on $\text{diameter}(G)$. We will subsequently show how to remove this dependence when solving DENSESUBGRAPH(\tilde{D}, ϵ).

Algorithm 2 and Algorithm 3 have the same structure except for the outputs. Both algorithms consist of $T = O((\log n)/\epsilon^2)$ iterations. Every edge maintains a load $\ell(e)$ throughout the algorithms. In each iteration t , each vertex u has z units budget. Each vertex u distributes the z units of budget to the $\alpha_{eu}^{(t)}$ variables of the incident edges with the lowest $\lceil z/2 \rceil$ load. It allocates 2 units of budget to the incident edges with the lowest $\lceil z/2 \rceil - 1$ load and the remaining budget to the other edge. Then the load of an edge is updated by adding the allocation from both endpoints. In Algorithm 2, we output the average of the allocations, $\sum_t \alpha_{eu}^{(t)}/T$, multiplied by $(1 + 2\epsilon)$. In Algorithm 3, in each iteration, the loads of the edges will be used as a guide to find a dense subgraph.

We can implement the algorithms by only sending integers across each link. This is because the value of each α_{eu} and $\ell(e)$ is a summation of an integer and an integer multiple of $z - 2(\lceil z/2 \rceil - 1)$. A pair of integers will be enough to express each value involved in the algorithms.

■ **Algorithm 2** FRACTIONAL_DUAL(z, ϵ).

```

1: Let  $T = K \cdot (1/\epsilon^2) \cdot \log n$  for some sufficiently large constant  $K$ .
2: Initialize the load  $\ell(e) \leftarrow 0$  for each edge  $e$ .
3: for  $t = 1, 2, \dots, T$  do
4:   for each vertex  $u$  do
5:     Let  $e_1, e_2, \dots, e_{\deg(u)}$  be the edges adjacent to  $u$  where  $\ell(e_1) \leq \ell(e_2) \dots \leq \ell(e_{\deg(u)})$ .
6:     Set  $\alpha_{e_i u}^{(t)} \leftarrow 2$  for  $i = 1, \dots, \min(\lceil z/2 \rceil - 1, \deg(u))$ .
7:     Set  $\alpha_{e_{\lceil z/2 \rceil u}}^{(t)} \leftarrow z - 2 \cdot (\lceil z/2 \rceil - 1)$  if  $\deg(u) \geq \lceil z/2 \rceil$ .
8:   for each edge  $e = uv$  do
9:     Set  $\ell(e) \leftarrow \ell(e) + \alpha_{eu}^{(t)} + \alpha_{ev}^{(t)}$ .
10:  Return  $\alpha_{eu} = \frac{\sum_{t=1}^T \alpha_{eu}^{(t)}}{T} \cdot (1 + 2\epsilon)$ .

```

■ **Algorithm 3** INTEGRAL_PRIMAL(z, ϵ).

```

1: Let  $T = K \cdot (1/\epsilon^2) \cdot \log n$  for some sufficiently large constant  $K$ .
2: Initialize the load  $\ell(e) \leftarrow 0$  for all  $e$ .
3: for  $t = 1, 2, \dots, T$  do
4:   for each vertex  $u$  do
5:     Let  $e_1, e_2, \dots, e_{\deg(u)}$  be the edges adjacent to  $u$  where  $\ell(e_1) \leq \ell(e_2) \dots \leq \ell(e_{\deg(u)})$ .
6:     Set  $\alpha_{e_i u}^{(t)} \leftarrow 2$  for  $i = 1, \dots, \min(\lceil z/2 \rceil - 1, \deg(u))$ .
7:     Set  $\alpha_{e_{\lceil z/2 \rceil u}}^{(t)} \leftarrow z - 2 \cdot (\lceil z/2 \rceil - 1)$  if  $\deg(u) \geq \lceil z/2 \rceil$ .
8:   Let  $\ell_{\min} = \min_{e \in E} \lfloor \ell(e) \rfloor$  and  $\ell_{\max} = \ell_{\min} + \lceil \frac{1}{\epsilon} \log \frac{2m}{\epsilon} \rceil$ .
9:   for each integer  $\ell \in [\ell_{\min}, \ell_{\max}]$  do
10:    Let  $V'_\ell$  be the set of vertices with at least  $\lceil z/2 \rceil$  incident edges  $e$  with  $\lceil \ell(e) \rceil \leq \ell$ .
11:    Test if  $G[V'_\ell]$  is a graph of density at least  $(1 - 3\epsilon)z$ .
12:    If yes, output  $G[V'_\ell]$  and terminate.
13:   for each edge  $e = uv$  do
14:     Set  $\ell(e) \leftarrow \ell(e) + \alpha_{eu}^{(t)} + \alpha_{ev}^{(t)}$ .

```

We will show that if we run the algorithms on the same input with the same parameters z and ϵ , then at least one of the following must be true:

- Algorithm 2 returns a solution satisfying DUAL($(1 + 2\epsilon)z$).
- Algorithm 3 outputs a subgraph of density at least $(1 - 3\epsilon)z$, i.e., an integral solution for PRIMAL($(1 - 3\epsilon)z$).

Define $w_e = (1 - \epsilon)^{\ell(e)}$ to be the weight associated with the edge e . We will use $\ell^{(t)}(e)$ and $w_e^{(t)}$ if we are referring to the load or the weight of an edge e at the beginning of iteration t . We often use u and v to denote the endpoints of an edge e when the context is clear.

If it is the case that $\sum_{e \in E} w_e^{(t)}(\alpha_{eu}^{(t)} + \alpha_{ev}^{(t)}) \geq \sum_{e \in E} w_e^{(t)}$ for every iteration t , then we show that Algorithm 2 returns a feasible solution for $\text{DUAL}((1 + 2\epsilon)z)$. The following lemma is a standard derivation of the multiplicative weights update method [2, 43].

► **Lemma 7.** *Let $0 < \epsilon < 1/4$. Suppose that $\sum_{e \in E} w_e^{(t)}(\alpha_{eu}^{(t)} + \alpha_{ev}^{(t)}) \geq \sum_{e \in E} w_e^{(t)}$ for every iteration t in Algorithm 2. Then the dual variables $\{\alpha_{eu}\}$ returned by Algorithm 2 must be a feasible solution satisfying $\text{DUAL}((1 + 2\epsilon)z)$.*

Note that if we provide the algorithms with $z \geq D$, the condition $\sum_{e \in E} w_e^{(t)}(\alpha_{eu}^{(t)} + \alpha_{ev}^{(t)}) \geq \sum_{e \in E} w_e^{(t)}$ holds for every iteration t . This is because if $z \geq D$ then we know there exists a feasible solution such that $(\alpha_{eu} + \alpha_{ev}) \geq 1$ for every edge $e = uv$. This implies $\sum_{e \in E} w_e^{(t)}(\alpha_{eu} + \alpha_{ev}) \geq \sum_{e \in E} w_e^{(t)}$. Note that

$$\sum_{e \in E} w_e^{(t)}(\alpha_{eu}^{(t)} + \alpha_{ev}^{(t)}) = \sum_{u \in V} \sum_{e \ni u} \alpha_{eu}^{(t)} w_e^{(t)} = \sum_{u \in V} \sum_{e \ni u} \alpha_{eu}^{(t)} (1 - \epsilon)^{\ell^{(t)}(e)}.$$

Hence, the way we assign $\{\alpha_{eu}^{(t)}\}$ maximizes $\sum_{e \in E} w_e^{(t)}(\alpha_{eu}^{(t)} + \alpha_{ev}^{(t)})$ over the set of feasible solutions $S(z)$, where

$$S(z) = \{\alpha : \sum_{e \ni u} \alpha_{eu} \leq z \text{ for all } u \text{ and } 0 \leq \alpha_{eu} \leq 2 \text{ for all } \alpha_{eu}\}.$$

Therefore, $\sum_{e \in E} w_e^{(t)}(\alpha_{eu}^{(t)} + \alpha_{ev}^{(t)}) \geq \sum_{e \in E} w_e^{(t)}(\alpha_{eu} + \alpha_{ev}) \geq \sum_{e \in E} w_e^{(t)}$. From the above lemma, this implies Algorithm 2 will output a fractional solution satisfying $\text{DUAL}(z)$.

Next, we need to show that if the opposite holds, i.e., $\sum_{e \in E} w_e^{(t)}(\alpha_{eu}^{(t)} + \alpha_{ev}^{(t)}) < \sum_{e \in E} w_e^{(t)}$ for some iteration t , then Algorithm 3 outputs a subgraph of density at least $(1 - 3\epsilon)z$. First, we rely on the following lemma which is a paraphrase of [4].

► **Lemma 8 (Paraphrase of [4]).** *Let $1/2 < F < 1$. Suppose there exists a set of (non-negative) weights $\{w'_e\}$ such that $\sum_{e \in E} w'_e > F \cdot \max_{\alpha \in S(z)} (\sum_{e \in E} w'_e(\alpha_{eu} + \alpha_{ev}))$. Let V''_λ denote the set of vertices that are incident to at least $\lceil z/2 \rceil$ edges e with $w'_e \geq \lambda$. Then, there exists λ such that $d(G[V''_\lambda]) > F \cdot z$.*

Hence, if there is an iteration t where $\sum_{e \in E} w_e^{(t)}(\alpha_{eu}^{(t)} + \alpha_{ev}^{(t)}) < \sum_{e \in E} w_e^{(t)}$, we can apply the above lemma with $F = 1$ and $w'_e = w_e^{(t)}$ to find the desired dense subgraph. However, there may be $\Omega(m)$ potential values for λ to check. We circumvent this by discretizing the weights. Since we are only testing the densities of $G[V'_\ell]$ for integer loads ℓ , we are effectively discretizing the weights. We will show that this discretization only introduces a small error on the inequality so that we can apply Lemma 8 with $F = 1 - O(\epsilon)$. This, in turn, gives us the following.

► **Lemma 9.** *Suppose that there exists an iteration t in Algorithm 3 where $\sum_{e \in E} w_e^{(t)}(\alpha_{eu}^{(t)} + \alpha_{ev}^{(t)}) < \sum_{e \in E} w_e^{(t)}$. Then for $0 < \epsilon < 1/4$, Algorithm 3 will output a subgraph of density at least $(1 - 3\epsilon)z$ in that iteration.*

► **Lemma 10.** *Suppose $0 < \epsilon < 1/4$ and $\tilde{D} \leq D$. Then, there exists a deterministic algorithm that solves $\text{DENSESUBGRAPH}(\tilde{D}, \epsilon)$ in $O(\text{diameter}(G) + 1/\epsilon^3 \cdot \log^2 n)$ rounds in the CONGEST model.*

Proof. Assume that $\tilde{D} \leq D$. Let $z = (1 - \epsilon/2)\tilde{D}$ and $\epsilon' = \epsilon/8$. If we were to run Algorithm 2 with parameters z and ϵ' , then it cannot possibly output a feasible solution for $\text{DUAL}(z(1 + 2\epsilon'))$, since $\text{DUAL}(z(1 + 2\epsilon'))$ is infeasible. This is because $z(1 + 2\epsilon') = (1 - \epsilon/2)(1 + \epsilon/4)\tilde{D} < \tilde{D} \leq D$. This implies that Algorithm 3 has to output a subgraph with density of at least $(1 - 3\epsilon')z \geq (1 - \epsilon/2)(1 - \epsilon/2)\tilde{D} \geq (1 - \epsilon)\tilde{D}$.

We now argue that Algorithm 3 can be implemented in $O(\text{diameter}(G) + 1/\epsilon^3 \cdot \log^2 n)$ rounds in the CONGEST model. Each step in the main loop of Algorithm 3 uses $O(1)$ rounds except for Line 8 and the inner loop (Line 9 – Line 12). Line 8 can be done in $O(\text{diameter}(G))$ rounds. The inner loop tests the density of $O((1/\epsilon) \log n)$ subgraphs. In total, there are $O((1/\epsilon) \log n) \cdot O((1/\epsilon^2) \log n)$ subgraphs to be tested. Each subgraph and its density can be computed in $O(\text{diameter}(G))$ rounds. These steps can be pipelined to run in $O(\text{diameter}(G) + (\log^2 n)/\epsilon^3)$ rounds. Finally, we argue that each $\alpha_{eu}^{(t)}$ can be represented using $O(\log \log n + \log(1/\epsilon))$ bits (Lemma 12) and therefore in each iteration, Line 14 in Algorithm 3 (and Line 9 in Algorithm 2) can be done in one CONGEST round. \blacktriangleleft

Using the above lemma, we now present the following results:

- There exists an algorithm that solves $\text{DENSESUBGRAPH}(\tilde{D}, \epsilon)$ in $O(\text{poly}(\log n, 1/\epsilon))$ rounds w.h.p. In particular, we are able to avoid $\text{diameter}(G)$ rounds in Lemma 10.
- There exists an algorithm that finds a $1 - \epsilon$ approximation of the densest subgraph (instead of solving the parameterized version) in $O(\text{diameter}(G) + \text{poly}(\log n, 1/\epsilon))$ rounds w.h.p.

Solving DenseSubgraph(\tilde{D}, ϵ). Our first goal is to avoid $O(\text{diameter}(G))$ rounds to solve the parameterized version of the densest subgraph problem. The main idea is to apply the low diameter decomposition and solve the problem in each component. First, we recall that the low-diameter decomposition can be implemented efficiently in the CONGEST model. See also [11].

► **Lemma 11** ([35]). *There exists an algorithm that decomposes the graph into disjoint components such that: 1) Each edge is inter-component with probability at most ϵ , 2) Each component has diameter $O(1/\epsilon \cdot \log n)$ w.h.p. and 3) Runs in $O(1/\epsilon \cdot \log n)$ rounds in the CONGEST model w.h.p.*

In the decomposition given by Miller et al. [35], the rough idea is that each vertex v draws δ_v from the exponential distribution $\text{Exp}(\epsilon)$. Let $\delta = K/\epsilon \cdot \log n$ for some sufficiently large constant K . At time step $\lfloor \delta - \delta_v \rfloor$, v wakes up and starts a breadth first search (BFS) if it has not been covered by another vertex's BFS. At the end of the decomposition, $v \in \text{cluster}(u)$ if $u = \min_{y \in V} (\text{dist}(y, v) + \lfloor \delta - \delta_y \rfloor)$. This algorithm can be simulated in the CONGEST model in $O(1/\epsilon \cdot \log n)$ rounds. Now, we are ready to prove our next main result.

► **Theorem 2.** *There exists a randomized algorithm that solves $\text{DENSESUBGRAPH}(\tilde{D}, \epsilon)$ w.h.p. and runs in $O((\log^3 n)/\epsilon^3)$ rounds in the CONGEST model.*

Proof. We apply the decomposition above with parameter $\epsilon/2$ to the graph to obtain low-diameter components C_1, \dots, C_t . Let H^* be the densest subgraph with n^* vertices and m^* edges. We condition on the event ξ that at most $\epsilon|E(H^*)|$ edges in $E(H^*)$ are inter-component. This happens with probability at least $1/2$ according to Markov's inequality.

First, consider the case $\tilde{D} \leq D$. Using a similar argument as in the proof of Lemma 5, we can show that the densest subgraph of at least one component has density at least $(1 - \epsilon)D \geq (1 - \epsilon)\tilde{D}$. Specifically, let $H_i^* = H^* \cap C_i$ and let $C_i^* \subseteq C_i$ be the densest subgraph in C_i . Furthermore, let $|V(H_i^*)| = n_i^*$ and $|E(H_i^*)| = m_i^*$.

If $d(C_i^*) < (1 - \epsilon)D$ for all $1 \leq i \leq t$. Then, $d(H_i^*) \leq d(C_i^*) < (1 - \epsilon)D$. This implies $m_i^* < (1 - \epsilon)m^*n_i^*/n^*$. Therefore, $\sum_{i=1}^t m_i^* < (1 - \epsilon)m^*$ which means that more than ϵm^* edges in $E(H^*)$ are inter-component. This is a contradiction since we condition on ξ .

For each low-diameter component C_i , using the algorithm in Lemma 10, we can solve DENSESUBGRAPH(\tilde{D}, ϵ) in C_i . As argued above, since $\tilde{D} \leq D$, we must have a non-empty output in some component i . Observe that since the diameter of each C_i is $O(1/\epsilon \cdot \log n)$, this takes $O(1/\epsilon^3 \cdot \log^2 n)$ rounds.

We ensure ξ happen w.h.p. by repeating $O(\log n)$ trials. The total number of rounds becomes $O(1/\epsilon^3 \cdot \log^3 n)$. However, there is a catch regarding the consistency of the output. Recall that we want the subgraph induced by the marked vertices $\{v : h_v = 1\}$ has density at least $(1 - \epsilon)\tilde{D}$. It is possible that for different trials j and j' , v might be marked and unmarked respectively. We need to address the issue of how to decide the final output.

This can be done with proper bookkeeping. Originally, all vertices are unmarked. In each trial, compute a low-diameter decomposition C_1, C_2, \dots, C_t . For each low-diameter component C_i , if it does not contain a marked vertex, check if there exists a subgraph $H_i \subseteq C_i$ with density at least $(1 - \epsilon)\tilde{D}$ using the algorithm in Lemma 10. If there exists such subgraph, mark every vertex in H_i . In the end, return the subgraph induced by the marked vertices. The output must be non-empty w.h.p. by considering the first time the event ξ occurs. Since all the output subgraphs among the trials are disjoint and they have density at least $(1 - \epsilon)\tilde{D}$, their union must have density at least $(1 - \epsilon)\tilde{D}$ as argued in the proof of Theorem 1.

In the case $\tilde{D} > D$, the algorithm either returns a subgraph with density at least $(1 - \epsilon)\tilde{D}$ or an empty subgraph which are both acceptable. \blacktriangleleft

Approximating the densest subgraph. We can also apply the ideas above to find a $1 - \epsilon$ approximation of the densest subgraph. This is done by simply running the algorithm in Theorem 2 on different guesses for the maximum subgraph density D .

► **Corollary 12.** *There exists a randomized algorithm that finds a $(1 - \epsilon)$ -approximation to the maximum density subgraph w.h.p. and runs in $O(\text{diameter}(G) + (\log^4 n / \epsilon^4))$ in the CONGEST model.*

Proof. For each $\tilde{D} = 1, (1 + \epsilon), (1 + \epsilon)^2, \dots, (1 + \epsilon)^{\lceil \log_{1+\epsilon} n \rceil}$, we run the algorithm in Theorem 2. This requires $O(1/\epsilon^4 \cdot \log^4 n)$ rounds. We then identify the largest \tilde{D} in which we have a non-empty output. In particular, we refer to when $\tilde{D} = (1 + \epsilon)^i$ as phase i . Each vertex v sets $\psi(v) = i$ where i is the largest phase in which it is marked. In the end, we can broadcast $j = \max_{v \in V} \psi(v)$ to all vertices v in $O(\text{diameter}(G))$ rounds. Then, for every vertex v , if v is marked in phase j , set $h_v = 1$. \blacktriangleleft

We include the following lemma, which will be useful in bounding the running time for rounding the fractional dual solution into an integer dual solution in the next section.

► **Lemma 12.** *Suppose that we run Algorithm 2 with parameters z and ϵ , where z is an integer and ϵ is a (negative) power of 2. Moreover, assume, $T = \Theta((\log n) / \epsilon^2)$, in the main loop of Algorithm 2, is a power of 2. Each α_{eu} in the solution returned by Algorithm 2 contains at most $O(\log \log n + \log(1/\epsilon))$ bits.*

4 Deterministic Algorithms for Low Outdegree Orientation

Recall that in the low outdegree orientation problem, we are given an integer parameter $\tilde{D} \geq D$. The goal is to find an orientation of the edges such that for every vertex, the number of outgoing edges is upper bounded by $(1 + \epsilon)\tilde{D}$.

In Section 3, we showed that Algorithm 2 can be used to obtain a solution for the fractional version of the low outdegree orientation problem. In this section, we will show how to round a fractional solution to an integral solution deterministically in the CONGEST model. We first present the framework and then describe the subroutine for the directed splitting procedure adapted from [22] in Section 4.1.

Let $\epsilon_1, \epsilon_2 \in \Theta(\epsilon)$ be error control parameters which will be determined later. First, we will run Algorithm 2 with parameters \tilde{D} and $\epsilon_1/2$ to obtain a fractional solution for $\text{DUAL}((1 + \epsilon_1)\tilde{D})$. Let $\{\alpha'_{eu}\}_{e \in E, u \in V}$ be the output fractional solution. They satisfy the conditions that $\alpha'_{eu} + \alpha'_{ev} \geq 1$ for every edge $e = uv$ and $\sum_{e \ni u} \alpha'_{eu} \leq (1 + \epsilon_1)\tilde{D}$ for every vertex u . We show how to round the α' -values to $\{0, 1\}$ bit-by-bit deterministically and incur bounded errors on the constraints.

We may also assume without loss of generality that both $1/\epsilon_1$ and the number of iterations, $T = \Theta((\log n)/\epsilon_1^2)$, in the main loop of Algorithm 2, is a power of 2. We can apply Lemma 12 to obtain an upper bound, $B = O(\log \log n + \log(1/\epsilon_1))$, on the number of bits needed to store each α'_{eu} . In the case where the maximum degree Δ is very small, we can even truncate the bits without creating much error. Let $t = \min(B, \lceil \log_2(\Delta/\epsilon_2) \rceil)$. We round the α' -values up to the t 'th bit after the decimal point. In other words, we set $\alpha_{eu}^{(0)} = \lceil 2^t \cdot \alpha'_{eu} \rceil / 2^t$ for every variable α_{eu} . If $t = B$, we are just setting $\alpha_{eu}^{(0)} = \alpha'_{eu}$.

Note that in the case $t = \lceil \log_2(\Delta/\epsilon_2) \rceil$, because $\tilde{D} \geq 1$, we have

$$\sum_{e \ni u} \alpha_{eu}^{(0)} \leq \sum_{e \ni u} \alpha'_{eu} + \deg(u) \cdot 2^{-\lceil \log_2(\Delta/\epsilon_2) \rceil} \leq (1 + \epsilon)\tilde{D} + \epsilon_2 \leq (1 + \epsilon_1)(1 + \epsilon_2)\tilde{D}.$$

The algorithm (Algorithm 4) consists of t iterations. It processes the α -values bit-by-bit, from the t 'th bit to the first bit after the decimal point. When it is processing the k 'th bit, for each α_{eu} , we will round its k 'th bit either up or down. Therefore, after we have processed the first bit in the last iteration, all the α -values are integers. Let $\alpha_{eu}(i)$ be the $(t - i + 1)$ 'th bit of α_{eu} after the decimal point (i.e. the i 'th rightmost bit after the initial rounding).

Let $\alpha_{eu}^{(k)}$ denote the value α_{eu} at the end of iteration k . During iteration k , if $\alpha_{eu}^{(k-1)}(k) = 1$, we will either need to round it up (set $\alpha_{eu}^{(k)} = \alpha_{eu}^{(k-1)} + 2^{-(t-k+1)}$) or round it down (set $\alpha_{eu}^{(k)} = \alpha_{eu}^{(k-1)} - 2^{-(t-k+1)}$) so that $\alpha_{eu}^{(k)}(k) = 0$.

Consider an edge $e = uv$. If $\alpha_{eu}^{(k-1)}(k) = 0$ or $\alpha_{ev}^{(k-1)}(k) = 0$, we will round both of their k 'th bit down so that $\alpha_{eu}^{(k)}(k) = 0$ and $\alpha_{ev}^{(k)}(k) = 0$. All the remaining edges must be contained in the graph $G_k = (V, E_k)$, where $E_k = \{uv \mid \alpha_{eu}^{(k-1)}(k) = 1 \text{ and } \alpha_{ev}^{(k-1)}(k) = 1\}$. We will perform a deterministic directed splitting algorithm on G_k which we adapt from [22] to the CONGEST model. Let $\deg_k(u)$ denote the degree of u in G_k . The outcome of the algorithm is an orientation of the edges in E_k such that for each vertex u , $|\text{outdeg}_k(u) - \text{indeg}_k(u)| \leq \epsilon_3 \deg_k(u) + 12$.

Suppose that $e = uv$ is oriented from u to v . We will round $\alpha_{eu}^{(k-1)}$ up and round $\alpha_{ev}^{(k-1)}$ down. We do the opposite if it is oriented from v to u .

► **Lemma 13.** *Suppose that $\epsilon_3 \leq 1/4$. The $\{\alpha_{ev}\}$ values produced by Algorithm 4 satisfy the following properties. (1) $\alpha_{eu} \in \{0, 1\}$ (2) For every edge $e = uv$, $\alpha_{eu} + \alpha_{ev} \geq 1$. (3) For every vertex u , $\sum_{e \ni u} \alpha_{eu} \leq (1 + \epsilon_1)(1 + \epsilon_2)(1 + \epsilon_3)^t \tilde{D} + 16$.*

Proof. For (1), since we either round $\alpha_{ev}^{(k-1)}(k)$ up or down during iteration k , we have $\alpha_{ev}^{(k)}(k) = 0$ at the end of iteration k . Moreover, once $\alpha_{ev}^{(k)}(k)$ becomes 0, $\alpha_{ev}^{(k')}(k)$ remains 0 for $k' \geq k$. Therefore, at the end of iteration t , we have $\alpha_{ev}^{(t)}(i) = 0$ for $1 \leq i \leq t$. This implies $\alpha_{ev}^{(t)}$ is a non-negative integer. Since the final output, α_{ev} , is the minimum of 1 and $\alpha_{ev}^{(t)}$, we have $\alpha_{ev} \in \{0, 1\}$.

■ **Algorithm 4** Deterministic Rounding for Low Outdegree Orientation.

```

1: Obtain a feasible  $\{\alpha'_{eu}\}_{eu}$  for DUAL( $(1 + \epsilon_1)\tilde{D}$ ).
2: Set  $\alpha_{eu}^{(0)} = \lceil 2^t \cdot \alpha'_{eu} \rceil / 2^t$  for every  $\alpha_{eu}$ .
3: for  $k = 1 \dots t$  do
4:   for every edge  $e = uv$  s.t.  $\alpha_{eu}^{(k-1)}(k) = 0$  or  $\alpha_{ev}^{(k-1)}(k) = 0$  do
5:     Set  $\alpha_{eu}^{(k)} = \alpha_{eu}^{(k-1)}$  and then set  $\alpha_{eu}^{(k)}(k) = 0$ .                                ▷ round the  $k$ 'th bit down
6:     Set  $\alpha_{ev}^{(k)} = \alpha_{ev}^{(k-1)}$  and then set  $\alpha_{ev}^{(k)}(k) = 0$ .
7:   Let  $G_k = (V, E_k)$ , where  $E_k = \{uv \mid \alpha_{eu}^{(k-1)}(k) = 1 \text{ and } \alpha_{ev}^{(k-1)}(k) = 1\}$ .
8:   Obtain a directed splitting of  $G_k$  whose discrepancy is at most  $\epsilon_3 \deg_k(u) + 12$  for each
   vertex  $u$ .
9:   for every edge  $e = uv$  where  $u$  is oriented toward  $v$  do
10:    Set  $\alpha_{eu}^{(k)} = \alpha_{eu}^{(k-1)} + 2^{-(t-k+1)}$ .                                ▷ round up
11:    Set  $\alpha_{ev}^{(k)} = \alpha_{ev}^{(k-1)} - 2^{-(t-k+1)}$ .                                ▷ round down
12: for each  $e = uv$  do
13:   Set  $\alpha_{eu} = \min(\alpha_{eu}^{(t)}, 1)$ .

```

We show (2) inductively. In the beginning, since $\alpha_{eu}^{(0)} \geq \alpha'_{eu}$, we must have $\alpha_{eu}^{(0)} + \alpha_{ev}^{(0)} \geq 1$. Suppose that $\alpha_{eu}^{(k-1)} + \alpha_{ev}^{(k-1)} \geq 1$. During iteration k , if $\alpha_{eu}^{(k-1)}(k) = 0$ and $\alpha_{ev}^{(k-1)}(k) = 0$, then $\alpha_{eu}^{(k)} = \alpha_{eu}^{(k-1)}$ and $\alpha_{ev}^{(k)} = \alpha_{ev}^{(k-1)}$ and so $\alpha_{eu}^{(k)} + \alpha_{ev}^{(k)} \geq 1$.

If $\alpha_{eu}^{(k-1)}(k) = 1$ and $\alpha_{ev}^{(k-1)}(k) = 0$, then it must be the case that $\alpha_{eu}^{(k-1)} + \alpha_{ev}^{(k-1)} \geq 1 + 2^{-(t-k+1)}$. After rounding $\alpha_{eu}^{(k-1)}$ down, we still have $\alpha_{eu}^{(k)} + \alpha_{ev}^{(k)} \geq 1$. The case for $\alpha_{eu}^{(k-1)}(k) = 0$ and $\alpha_{ev}^{(k-1)}(k) = 1$ is symmetric.

The remaining case is when $\alpha_{eu}^{(k-1)}(k) = 1$ and $\alpha_{ev}^{(k-1)}(k) = 1$. In this case, $e \in G_k$. We must have $\alpha_{eu}^{(k)} + \alpha_{ev}^{(k)} = \alpha_{eu}^{(k-1)} + \alpha_{ev}^{(k-1)} \geq 1$, since one of them is rounded up and the other is rounded down.

For (3), let $D_0 = (1 + \epsilon_1)(1 + \epsilon_2)\tilde{D}$ and $D_k = (1 + \epsilon_3)D_{k-1} + 12 \cdot 2^{-(t-k+1)}$ for $k \geq 1$. We will show by induction that $\sum_{e \ni u} \alpha_{eu}^{(k)} \leq D_k$. For the base case, initially, we argued that $\sum_{e \ni u} \alpha_{eu}^{(0)} \leq (1 + \epsilon_1)(1 + \epsilon_2)\tilde{D}$. For $k \geq 1$, note that the increase on the quantity $\sum_{e \ni u} \alpha_{eu}$ during iteration k is at most $2^{-(t-k+1)} \cdot (\epsilon_3 \cdot \deg_k(u) + 12)$. Therefore,

$$\begin{aligned} \sum_{e \ni u} \alpha_{eu}^{(k)} &\leq 2^{-(t-k+1)} \cdot (\epsilon_3 \cdot \deg_k(u) + 12) + \sum_{e \ni u} \alpha_{eu}^{(k-1)} \\ &\leq \epsilon_3 D_{k-1} + 12 \cdot 2^{-(t-k+1)} + D_{k-1} \leq (1 + \epsilon_3)D_{k-1} + 12 \cdot 2^{-(t-k+1)} = D_k. \end{aligned}$$

The second inequality follows since $2^{-(t-k+1)} \deg_k(u) \leq \sum_{e \ni u} \alpha_{eu}^{(k-1)} \leq D_{k-1}$. This completes the induction. Since $\epsilon_3 \leq 1/4$, in the end, we have

$$\begin{aligned} \sum_{e \ni u} \alpha_{eu} &\leq \sum_{e \ni u} \alpha_{eu}^{(t)} \leq D_t \leq (1 + \epsilon_1)(1 + \epsilon_2)(1 + \epsilon_3)^t \tilde{D} + 12 \cdot (1/2) \cdot \sum_{k=0}^{t-1} \left(\frac{1 + \epsilon_3}{2} \right)^k \\ &\leq (1 + \epsilon_1)(1 + \epsilon_2)(1 + \epsilon_3)^t \tilde{D} + 16. \end{aligned}$$

◀

► **Theorem 4.** Given an integer $\tilde{D} \geq D$, for any $32/\tilde{D} \leq \epsilon \leq 1/4$, there exists a deterministic algorithm in the CONGEST model that computes a $(1 + \epsilon)\tilde{D}$ -orientation and runs in

$$O\left(\frac{\log n}{\epsilon^2} + (\min(\log \log n, \log \Delta) + \log(1/\epsilon))^{2.71} \cdot (1/\epsilon)^{1.71} \cdot \log^2 n\right) \leq \tilde{O}((\log^2 n)/\epsilon^2) \text{ rounds.}$$

Proof. We set parameters $\epsilon_1 = \epsilon_2 = \epsilon/8$ and $\epsilon_3 = \epsilon/(4t)$. Run Algorithm 4 to obtain integral $\{\alpha_{eu}\}$ that satisfy $\alpha_{eu} + \alpha_{ev} \geq 1$ for every $e = uv$ and $\sum_{e \ni u} \alpha_{eu} \leq (1 + \epsilon_1)(1 + \epsilon_2)(1 + \epsilon_3)^t \tilde{D} + 16$ for every u . For each edge $e = uv$, if $\alpha_{eu} = 1$ then we orient e from u to v . Otherwise, we

orient e from v to u . Since $16 \leq \epsilon \tilde{D}/2$, the out-degree of each vertex is upper bounded by:

$$\sum_{e \ni u} \alpha_{eu} \leq (1 + \epsilon_1)(1 + \epsilon_2)(1 + \epsilon_3)^t \tilde{D} + 16 \leq (1 + \epsilon/2) \tilde{D} + 16 \leq (1 + \epsilon) \tilde{D}.$$

The running time for Algorithm 4 consists of the following. The number of rounds to obtain a fractional solution is $O((\log n)/\epsilon_1^2) = O((\log n)/\epsilon^2)$. Then, it consists of $t = \min(\log(\Delta/\epsilon_2), B) = O(\min(\log \log n, \log \Delta) + \log(1/\epsilon))$ iterations. Each iteration invokes a directed splitting procedure that runs in $O((1/\epsilon_3)^{1.71} \cdot \log^2 n)$ rounds by Theorem 16. Recall that $\epsilon_3 = \epsilon/(4t)$, the total number of rounds is therefore:

$$O\left(\frac{\log n}{\epsilon^2} + t \cdot (1/\epsilon_3)^{1.71} \cdot \log^2 n\right) = O\left(\frac{\log n}{\epsilon^2} + t^{2.71} \cdot (1/\epsilon)^{1.71} \cdot \log^2 n\right). \quad \blacktriangleleft$$

4.1 Distributed Splitting in the CONGEST Model

Given a graph $G = (V, E)$, a *weak $f(v)$ -orientation* of G is an orientation of the edges in G such that there are at least $f(v)$ outgoing edges for each $v \in V$. In order to adapt the algorithm of [22] for directed splitting, we need an algorithm for weak $\lfloor \deg(v)/3 \rfloor$ -orientation in the CONGEST model. The previous algorithms [22, 24] for weak $\lfloor \deg(v)/3 \rfloor$ -orientation requires finding short cycles for containing each edge. They are not adaptable to the CONGEST model. We use an augmenting path approach for finding a weak $\lfloor \deg(v)/3 \rfloor$ -orientation instead.

► **Lemma 14.** *There exists a deterministic distributed algorithm that computes a weak $\lfloor \deg(v)/3 \rfloor$ -orientation in $O(\log^2 n)$ rounds in the CONGEST model.*

Proof. First we construct a new graph G' as follows: Split every vertex v into $\lceil \deg(v)/3 \rceil$ copies. Attach evenly the edges to each copy of the vertex so every copy except possibly the last gets 3 edges and the last copy gets $\deg(v) - 3 \cdot (\lceil \deg(v)/3 \rceil - 1)$ edges. Given an orientation of G' , we call a vertex v a **sink if it has exactly 3 incoming edges**. Clearly, a sinkless orientation (i.e. an orientation where there are no sinks) in G' corresponds to a weak $\lfloor \deg(v)/3 \rfloor$ -orientation of G . Moreover, one round in G' can be emulated in G by using one round.

Now we start with an arbitrary orientation of G' . Some vertices might be sinks. We will use an augmenting path approach to eliminate sinks.

Divide the vertices into the following three types. Type I vertices are the sinks. Type II vertices are those u such that $\deg(u) = 3$ and $\text{indeg}(u) = 2$. Type III vertices are those with $\text{indeg}(u) \leq 1$ or $\deg(u) < 3$.

An augmenting path is a path $P = (u_1, \dots, u_l)$ such that:

1. u_1 is a Type I vertex. u_l is a Type III vertex. u_i is a Type II vertex for $1 < i < l$.
2. u_{i+1} is oriented towards u_i for $1 \leq i < l$.

If P is an augmenting path then flipping P will make u_1 no longer a sink. Moreover, it will not create any new sink.

Consider a Type I vertex u . An augmenting path of length $O(\log n)$ starting from u can be found as follows. Let $L_0 = \{u\}$. Given L_{i-1} , let L_i be the set of incoming neighbors of every vertex in L_{i-1} . If L_i contains a Type III vertex, then an augmenting path is found. Otherwise, it must be the case that L_i contains all Type II vertices. Since there are at least $2 \cdot |L_{i-1}|$ incoming edges from L_i to L_{i-1} ($\text{indeg}(u) = 2$ for $u \in L_{i-1}$) and Type II vertices have out-degree 1, we have $|L_i| \geq 2 \cdot |L_{i-1}|$. Hence, this process can only continue for at most $O(\log n)$ times.

Every Type I vertex would be able to find an augmenting paths of length $O(\log n)$ this way. Moreover, note that these augmenting paths can only overlap at their ending vertex, since the intermediate Type II vertices have out-degree 1. Each ending vertex is a Type

III vertex, which can only be the ending vertex of at most 3 augmenting paths since it has at most 3 outgoing edges. It selects an arbitrary augmenting path to accept. Therefore, at least $1/3$ fraction of augmenting paths will be accepted. We flip along the accepted augmenting paths to fix Type I vertices so that we eliminate at least $1/3$ fraction of the sinks. Therefore, it takes $O(\log n)$ repetitions to eliminate all the sinks. The total number of rounds is $O(\log n \cdot \log n) = O(\log^2 n)$. The process can be easily implemented in the CONGEST model. \blacktriangleleft

Before we describe how to adapt the splitting algorithm of [22], we need to introduce the following definition. Given a function $\delta : V \rightarrow \mathbb{R}_{\geq 0}$ and $\lambda \in \mathbb{Z}_{\geq 0}$. A (δ, λ) -path decomposition \mathcal{P} is a partition of E into paths P_1, \dots, P_ρ such that

1. Each vertex v is an endpoint of at most $\delta(v)$ paths.
2. Each path P_i is of length at most λ .

Given a (δ, λ) -path decomposition \mathcal{P} . The virtual graph $G_{\mathcal{P}} = (V, E_{\mathcal{P}})$ consists of exactly ρ edges, where each path $P_i = (v_{i,\text{start}}, \dots, v_{i,\text{end}})$ corresponds to an edge $(v_{i,\text{start}}, v_{i,\text{end}})$ in $E_{\mathcal{P}}$. Lemma 15 is the adaption of [22, Lemma 2.11] to the CONGEST model.

► **Lemma 15.** *Assume that $T(n, \Delta) \geq \log n$ is the running time of an algorithm \mathcal{A} that finds a weak $\lfloor \deg(v)/3 \rfloor$ -orientation in the CONGEST model. Then for any positive integer i , there is a deterministic distributed algorithm \mathcal{A} that finds a $((\frac{2}{3})^i \cdot \deg(v) + 12, 2^i)$ -path decomposition \mathcal{P} in time $O(2^i \cdot T(n, \Delta))$ in the CONGEST model.*

The reason of why such an adaptation works is due to the fact that all the paths in \mathcal{P} are disjoint. Therefore, one round in $G_{\mathcal{P}}$ can be simulated using $O(\lambda)$ rounds in G in the CONGEST model, given \mathcal{P} is a (δ, λ) -path decomposition.

For completeness, we explain how the path-decomposition in Lemma 15 can be obtained. Let \mathcal{P}_0 denote the initial path decomposition where each path is a single edge in E . Given \mathcal{P}_{i-1} , \mathcal{P}_i can be built as follows: Obtain a $\lfloor \deg(v)/3 \rfloor$ -orientation on $G_{\mathcal{P}_{i-1}}$. For each vertex u , group the outgoing edges into at least $\lfloor \lfloor \deg(u)/3 \rfloor / 2 \rfloor$ pairs. For each such edge pair (u, x) and (u, w) , we reverse the path that corresponds to (u, x) and append it with (u, w) . The new degree of u becomes at most $\deg(u) - 2 \cdot \lfloor \lfloor \deg(u)/3 \rfloor / 2 \rfloor \leq \frac{2}{3} \deg(u) + 4$. Therefore, if \mathcal{P}_{i-1} is a $(\delta, \lambda(v))$ -decomposition then \mathcal{P}_i is a $((\frac{2}{3})^i \cdot \delta(v) + 4, 2\lambda(v))$ decomposition.

Since \mathcal{P}_0 is a $(\deg(v), 1)$ -path decomposition, \mathcal{P}_i a $(z_i(v), 2^i)$ -path decomposition, where

$$z_i(v) = \left(\frac{2}{3}\right)^i \cdot \deg(v) + 4 \sum_{k=0}^{i-1} \left(\frac{2}{3}\right)^k \leq \left(\frac{2}{3}\right)^i \cdot \deg(v) + 12.$$

The running time of the j 'th iteration is $O(2^j \cdot T(n, \Delta))$. Therefore, the total running time is $\sum_{j=1}^i O(2^j \cdot T(n, \Delta)) = O(2^i \cdot T(n, \Delta))$. By setting $i = \log_{3/2}(1/\epsilon)$ and $T(n, \Delta) = \log^2 n$ from Lemma 14, we get a $(\epsilon \cdot \deg(v) + 12, (1/\epsilon)^{\log_{3/2} 2})$ -path decomposition in $O((1/\epsilon)^{\log_{3/2} 2} \cdot \log^2 n)$ rounds.

Suppose \mathcal{P} is a (δ, λ) -path decomposition. If we orient the edges on each path in \mathcal{P} in consistent with the direction of the path, then we must have $|\text{outdeg}(v) - \text{indeg}(v)| \leq \delta(v)$. Therefore, we obtain the following theorem.

► **Theorem 16.** *For $\epsilon > 0$, there exists a $O((1/\epsilon)^{1.71} \cdot \log^2 n)$ rounds deterministic algorithm in the CONGEST model that computes an orientation such that for each vertex v , $|\text{outdeg}(v) - \text{indeg}(v)| \leq \epsilon \deg(v) + 12$.*

References

- 1 Mohamad Ahmadi, Fabian Kuhn, and Rotem Oshman. Distributed approximate maximum matching in the CONGEST model. In *Proc. 32nd International Symposium on Distributed Computing (DISC)*, pages 6:1–6:17, 2018.
- 2 Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(6):121–164, 2012. doi:10.4086/toc.2012.v008a006.
- 3 Baruch Awerbuch. Complexity of network synchronization. *J. ACM*, 32(4):804–823, 1985.
- 4 Bahman Bahmani, Ashish Goel, and Kamesh Munagala. Efficient primal-dual graph algorithms for mapreduce. In *WAW*, volume 8882 of *Lecture Notes in Computer Science*, pages 59–78. Springer, 2014.
- 5 Bahman Bahmani, Ravi Kumar, and Sergei Vassilvitskii. Densest subgraph in streaming and mapreduce. *PVLDB*, 5(5):454–465, 2012.
- 6 Leonid Barenboim and Michael Elkin. Sublogarithmic distributed MIS algorithm for sparse graphs using nash-williams decomposition. *Distributed Computing*, 22(5-6):363–379, 2010.
- 7 Edvin Berglin and Gerth Stølting Brodal. A simple greedy algorithm for dynamic graph orientation. *Algorithmica*, 82(2):245–259, 2020.
- 8 Alex Beutel, Wanhong Xu, Venkatesan Guruswami, Christopher Palow, and Christos Faloutsos. Copycatch: Stopping group attacks by spotting lockstep behavior in social networks. In *Proceedings of the 22Nd International Conference on World Wide Web (WWW)*, pages 119–130, 2013.
- 9 Sayan Bhattacharya, Monika Henzinger, Danupon Nanongkai, and Charalampis E. Tsourakakis. Space- and time-efficient algorithm for maintaining dense subgraphs on one-pass dynamic streams. In *Proc. 47th ACM Symposium on Theory of Computing (STOC)*, pages 173–182, 2015.
- 10 Gerth Stølting Brodal and Rolf Fagerberg. Dynamic representations of sparse graphs. In *Proc. 6th International Workshop on Algorithms and Data Structures (WADS)*, pages 342–351, 1999.
- 11 Yi-Jun Chang and Thatchaphol Saranurak. Improved distributed expander decomposition and nearly optimal triangle enumeration. *CoRR*, abs/1904.08037, 2019.
- 12 Moses Charikar. Greedy approximation algorithms for finding dense components in a graph. In *APPROX*, volume 1913 of *Lecture Notes in Computer Science*, pages 84–95. Springer, 2000.
- 13 Jie Chen and Yousef Saad. Dense subgraph extraction with application to community detection. *IEEE Trans. on Knowl. and Data Eng.*, 24(7):1216–1230, 2012.
- 14 Yon Dourisboure, Filippo Geraci, and Marco Pellegrini. Extraction and classification of dense communities in the web. In *Proc. 16th International Conference on World Wide Web (WWW)*, pages 461–470, 2007.
- 15 Hossein Esfandiari, MohammadTaghi Hajiaghayi, and David P. Woodruff. Applications of uniform sampling: Densest subgraph and beyond. *CoRR*, abs/1506.04505, 2015.
- 16 Manuela Fischer. Improved deterministic distributed matching via rounding. In *31st International Symposium on Distributed Computing (DISC)*, pages 17:1–17:15, 2017.
- 17 Manuela Fischer, Mohsen Ghaffari, and Fabian Kuhn. Deterministic distributed edge-coloring via hypergraph maximal matching. In *Proc. 58th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 180–191, 2017.
- 18 Eugene Fratkin, Brian T. Naughton, Douglas Brutlag, and Serafim Batzoglou. Motifcut: Regulatory motifs finding with maximum density subgraphs. *Bioinformatics (Oxford, England)*, 22:e150–7, August 2006.
- 19 Harold N. Gabow and Herbert H. Westermann. Forests, frames, and games: Algorithms for matroid sums and applications. *Algorithmica*, 7(1):465, June 1992.
- 20 G. Gallo, M. D. Grigoriadis, and R. E. Tarjan. A fast parametric maximum flow algorithm and applications. *SIAM J. Comput.*, 18(1):30–55, 1989.

- 21 Mohsen Ghaffari, David G. Harris, and Fabian Kuhn. On derandomizing local distributed algorithms. In *Proc. 59th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 662–673, 2018.
- 22 Mohsen Ghaffari, Juho Hirvonen, Fabian Kuhn, Yannic Maus, Jukka Suomela, and Jara Uitto. Improved distributed degree splitting and edge coloring. In *DISC*, pages 19:1–19:15, 2017.
- 23 Mohsen Ghaffari, Silvio Lattanzi, and Slobodan Mitrović. Improved parallel algorithms for density-based network clustering. In *Proc. 36th International Conference on Machine Learning (ICML)*, volume 97, pages 2201–2210, 2019.
- 24 Mohsen Ghaffari and Hsin-Hao Su. Distributed degree splitting, edge coloring, and orientations. In *Proc. of 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2505–2523, 2017.
- 25 David Gibson, Ravi Kumar, and Andrew Tomkins. Discovering large dense subgraphs in massive graphs. In *Proceedings of the 31st International Conference on Very Large Data Bases (VLDB)*, pages 721–732, 2005.
- 26 A. V. Goldberg. Finding a maximum density subgraph. Technical Report UCB/CSD-84-171, EECS Department, University of California, Berkeley, 1984.
- 27 Anupam Gupta, Amit Kumar, and Cliff Stein. Maintaining assignments online: Matching, scheduling, and flows. In *Proc. 25th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, page 468–479, 2014.
- 28 David G. Harris. Distributed approximation algorithms for maximum matching in graphs and hypergraphs. In *Proc. 60th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 700–724, 2019.
- 29 R. Kannan and V. Vinay. Analyzing the structure of large graphs. *Technical report*, 1999.
- 30 Samir Khuller and Barna Saha. On finding dense subgraphs. In *Proc. 36th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 597–608, 2009.
- 31 Tsvi Kopelowitz, Robert Krauthgamer, Ely Porat, and Shay Solomon. Orienting fully dynamic graphs with worst-case time bounds. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *Proc. 41st International Colloquium on Automata, Languages and Programming (ICALP)*, pages 532–543, 2014.
- 32 Łukasz Kowalik. Approximation scheme for lowest outdegree orientation and graph density measures. In *Proc. of the 17th International Conference on Algorithms and Computation (ISAAC)*, pages 557–566, 2006.
- 33 Nathan Linial and Michael E. Saks. Low diameter graph decompositions. *Combinatorica*, 13(4):441–454, 1993.
- 34 Andrew McGregor, David Tench, Sofya Vorotnikova, and Hoa T. Vu. Densest subgraph in dynamic graph streams. In *MFCS (2)*, volume 9235 of *Lecture Notes in Computer Science*, pages 472–482. Springer, 2015.
- 35 Gary L. Miller, Richard Peng, and Shen Chen Xu. Parallel graph decompositions using random shifts. In *SPAA*, pages 196–203. ACM, 2013.
- 36 C. St.J. A. Nash-Williams. Decomposition of finite graphs into forests. *Journal of the London Mathematical Society*, s1-39(1):12–12, 1964.
- 37 Jean-Claude Picard and Maurice Queyranne. A network flow solution to some nonlinear 0-1 programming problems, with applications to graph theory. *Networks*, 12(2):141–159, 1982.
- 38 Barna Saha, Allison Hoch, Samir Khuller, Louiqa Raschid, and Xiao-Ning Zhang. Dense subgraphs with restrictions and applications to gene annotation graphs. In *Research in Computational Molecular Biology*, pages 456–472, 2010.
- 39 Atish Das Sarma, Ashwin Lall, Danupon Nanongkai, and Amitabh Trehan. Dense subgraphs on dynamic networks. In *DISC*, volume 7611 of *Lecture Notes in Computer Science*, pages 151–165. Springer, 2012.
- 40 Saurabh Sawlani and Junxing Wang. Near-optimal fully dynamic densest subgraph. In *Proc. 52th ACM Symposium on Theory of Computing (STOC)*, 2020. To appear.

15:18 Distributed Dense Subgraph Detection and Low Outdegree Orientation

- 41 Hsin-Hao Su and Hoa T. Vu. Distributed dense subgraph detection and low outdegree orientation. *CoRR*, abs/1907.12443, 2019.
- 42 M. Thorup. Fully-dynamic min-cut. *Combinatorica*, 27(1):91–127, 2007. doi:10.1007/s00493-007-0045-2.
- 43 Neal E. Young. Randomized rounding without solving the linear program. In *SODA*, pages 170–178. ACM/SIAM, 1995.