

Distributed Symmetry Breaking in Hypergraphs

Shay Kutten^{1,*}, Danupon Nanongkai^{2,**}, Gopal Pandurangan^{3,***}, and Peter Robinson^{4†}

¹ Faculty of IE&M, Technion, Haifa, Israel

² Faculty of Computer Science, University of Vienna, Austria

³ Department of Computer Science, University of Houston, USA

⁴ Department of Computer Science, National University of Singapore

Abstract. Fundamental local symmetry breaking problems such as Maximal Independent Set (MIS) and coloring have been recognized as important by the community, and studied extensively in (standard) graphs. In particular, fast (i.e., logarithmic run time) randomized algorithms are well-established for MIS and $\Delta + 1$ -coloring in both the LOCAL and CONGEST distributed computing models. On the other hand, comparatively much less is known on the complexity of distributed symmetry breaking in *hypergraphs*. In particular, a key question is whether a fast (randomized) algorithm for MIS exists for hypergraphs.

In this paper, we study the distributed complexity of symmetry breaking in hypergraphs by presenting distributed randomized algorithms for a variety of fundamental problems under a natural distributed computing model for hypergraphs. We first show that MIS in hypergraphs (of arbitrary dimension) can be solved in $O(\log^2 n)$ rounds (n is the number of nodes of the hypergraph) in the LOCAL model. We then present a key result of this paper — an $O(\Delta^\epsilon \text{polylog } n)$ -round hypergraph MIS algorithm in the CONGEST model where Δ is the maximum node degree of the hypergraph and $\epsilon > 0$ is any arbitrarily small constant. We also present distributed algorithms for coloring, maximal matching, and maximal clique in hypergraphs.

To demonstrate the usefulness of hypergraph MIS, we present applications of our hypergraph algorithm to solving problems in (standard) graphs. In particular, the hypergraph MIS yields fast distributed algorithms for the *balanced minimal dominating set* problem (left open in Harris et al. [ICALP 2013]) and the *minimal connected dominating set problem*.

Our work shows that while some local symmetry breaking problems such as coloring can be solved in polylogarithmic rounds in both the LOCAL and CONGEST models, for many other hypergraph problems such as MIS, hitting set, and maximal clique, it remains challenging to obtain polylogarithmic time algorithms in the CONGEST model. This work is a step towards understanding this dichotomy in the complexity of hypergraph problems as well as using hypergraphs to design fast distributed algorithms for problems in (standard) graphs.

1 Introduction

The importance, as well as the difficulty, of solving problems on hypergraphs was pointed out recently by Linial, in his Dijkstra award talk [31]. While standard graphs⁵ model *pairwise* interactions well, hypergraphs can be used to model *multi-way* interactions. For example, social network interactions include several individuals as a group, biological interactions involve several entities (e.g., proteins) interacting at the same time,

* Research supported in part by the Israel Science Foundation and by the Technion TASP center.

** This work was partially done while at ICERM, Brown University USA and Nanyang Technological University, Singapore.

*** This work was done while at Nanyang Technological University and Brown University. Research supported in part by the following research grants: Nanyang Technological University grant M58110000, Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier 2 grant MOE2010-T2-2-082, Singapore MOE AcRF Tier 1 grant MOE2012-T1-001-094, and a grant from the US-Israel Binational Science Foundation (BSF).

† Research supported by the grant Fault-tolerant Communication Complexity in Wireless Networks from the Singapore MoE AcRF-2.

⁵ Henceforth, when we say a graph, we just mean a standard (simple) graph.

distributed systems can involve several agents working together, or multiple clients who share a server (e.g., a cellular base station), or multiple servers who share a client, or shared channels in a wireless network. In particular, hypergraphs are especially useful in modelling social networks (e.g., [44]) and wireless networks (e.g., [3]). Unfortunately, as pointed out by Linial, much less is known for hypergraphs than for graphs. The focus of this paper is studying the complexity of fundamental local symmetry breaking problems in *hypergraphs*⁶. A related goal is to utilize these hypergraph algorithms for solving (standard) graph problems.

In the area of distributed computing for (standard) graphs, fundamental local symmetry breaking problems such as Maximal Independent Set (MIS) and coloring have been studied extensively (see e.g., [33,30,6,41,26] and the references therein). Problems such as MIS and coloring are “local” in the sense that a solution can be *verified* easily by purely local means (e.g., each node communicating only with its neighbors), but the solution itself should satisfy a global property (e.g., in the case of coloring, every node in the graph should have a color different from its neighbors and the total number of colors is at most $\Delta + 1$, where Δ is the maximum node degree). Computing an MIS or coloring locally is non-trivial because of the difficulty of *symmetry breaking*: nodes have to decide on their choices (e.g., whether they belong to the MIS or not) by only looking at a *small* neighbourhood around it. (In particular, to get an algorithm running in k rounds, each node v has to make its decision by looking only at information on nodes within distance k from it.) Some of the most celebrated results in distributed algorithms are such fast localized algorithms. In particular, $O(\log n)$ -round (randomized) distributed algorithms are well-known for MIS [33] and $\Delta + 1$ -coloring [6] in both the LOCAL and CONGEST distributed computing models [41].

Besides the interest in understanding the complexity of fundamental problems, the solutions to such localizable symmetry breaking problems had many obvious applications. Examples are scheduling (such as avoiding the collision of radio transmissions, see e.g. [16], [11]), or matching nodes such that each pair can communicate in parallel to the other pairs, see e.g. [5]), resource management (such as assigning clients to servers, see, e.g. [4]), and even for obtaining $O(\text{Diameter})$ solutions to global problems that cannot be solved locally, such as MST computation [17,29].

In contrast to graphs which have been extensively studied in the context of distributed algorithms, many problems become much more challenging in the context of hypergraphs. An outstanding example is the MIS problem, whose local solutions for graphs were mentioned above. On the other hand, in hypergraphs (of arbitrary dimension) the complexity of MIS is wide open. (In a hypergraph, an MIS is a maximal subset I of hypernodes such that no subset of I forms an hyperedge.) Indeed, determining the parallel complexity (in the PRAM model) of the Maximal Independent Set (MIS) problem in hypergraphs (for arbitrary dimension) remains as one of the most important open problems in parallel computation; in particular, a key open problem is whether there exists a polylogarithmic time PRAM algorithm [20,8,22]. As discussed later, efficient CONGEST model distributed algorithms that uses simple local computations will also give efficient PRAM algorithms.

1.1 Main Results

We present distributed (randomized) algorithms for a variety of fundamental problems under a natural distributed computing model for hypergraphs (cf. Section 2).

Hypergraph MIS. A main focus is the hypergraph MIS problem which has been the subject of extensive research in the PRAM model (see e.g., [20,21,22,8,34]). We first show that MIS in hypergraphs (of arbitrary dimension) can be solved in $O(\log^2 n)$ distributed rounds (n is the number of nodes of the hypergraph) in the LOCAL model (cf. Theorem 1). We then present an $O(\Delta^\epsilon \text{polylog } n)$ round algorithm for finding a MIS in hypergraphs of arbitrary dimension in the CONGEST model, where Δ is the maximum degree of the hypergraph (we refer to Theorem 1 for a precise statement of the bound) and $\epsilon > 0$ is any small

⁶ Formally, a hypergraph (V, F) consists of a set of (hyper)nodes V and a collection F of subsets of V ; the sets that belong to F are called *hyperedges*. The *dimension* of a hypergraph is the maximum number of hypernodes that belong to a hyperedge. Throughout, we will use n for the number of nodes, m for the number of hyperedges, and Δ for the degree of the hypergraph which is the maximum node degree (i.e., the maximum number of edges a node is in). A standard graph is a hypergraph of dimension 2.

positive constant. In the distributed computing model (both LOCAL and CONGEST), computation within a node is free; in one round, each node is allowed to compute any function of its current data. However, in our CONGEST model algorithms, each processor will perform very simple computations (but this is not true in the LOCAL model). In particular, each step of any node v can be simulated in $O(d_v)$ time by a single processor or in $O(\log m)$ time with d_v processors. Here, d_v is the degree of the node in the *server-client* computation model — cf. Section 2; $d_v = O(m)$, where m is the number of hyperedges. From these remarks, it follows that our algorithms can be simulated on the PRAM model to within an $O(\log m)$ factor slowdown using $O(m + n)$ processors. Thus our CONGEST model algorithm also implies a PRAM algorithm for hypergraph MIS running in $O(\Delta^\epsilon \text{polylog } n \log m)$ rounds using a linear number of processors for a hypergraph of arbitrary dimension.

Algorithms for standard graph problems using hypergraph MIS. In addition to the importance of hypergraph MIS as a hypergraph problem, we outline its importance to solving several natural symmetry breaking problems in (standard) graphs too. For the results discussed below, we assume the CONGEST model.

Consider first the following graph problem called the *restricted minimal dominating set (RMDS)* problem which arises as a key subproblem in other problems that we discuss later. We are given a (standard) graph $G = (V, E)$ and a subset of nodes $R \subseteq V$, such that R forms a dominating set in G (i.e., every node $v \in V$ is either adjacent to R or belongs to R). It is required to find a *minimal* dominating set *in* R that dominates V . (The minimality means that no subset of the solution can dominate V ; it is easy to verify the minimality condition locally.) Note that if R is V itself, the problem can be solved by finding a MIS of G , since a MIS is also a minimal dominating set (MDS); hence an $O(\log n)$ algorithm exists. However, if R is some arbitrary proper subset of V (such that R dominates V), then no distributed algorithm running even in sublinear (in n) time (let alone polylogarithmic time) is known. Using our hypergraph MIS algorithm, we design a distributed algorithm for RMDS running in $O(\min\{\Delta^\epsilon \text{polylog } n, n^{o(1)}\})$ rounds in the CONGEST model (Δ is the maximum node degree of the graph) — cf., Section 4.1.

RMDS arises naturally as the key subproblem in the solution of other problems, in particular, the *balanced minimal dominating set (BMDS)* problem [19] and the *minimal connected dominating set (MCDS)* problem. Given a (standard) graph, the BMDS problem (defined formally in Section 4.2) asks for a minimal dominating set whose average degree is small with respect to the average degree of the graph; this has applications to load balancing and fault-tolerance [19]. It was shown that such a set exists and can be found using a *centralized* algorithm [19]. Finding a fast distributed algorithm was a key problem left open in [19]. In Section 4.3, we use our hypergraph MIS algorithm of Section 3 to present an $\tilde{O}(D + \min\{\Delta^\epsilon, n^{o(1)}\})$ round algorithm (the notation \tilde{O} hides a polylog n factor) for BMDS problem (in the CONGEST model), where D is the diameter (of the input standard graph) and Δ is the maximum node degree.

The MCDS problem is a variant (similar to variants studied in the context of wireless networks, e.g. [12]) of the well-studied *minimum* connected dominating set problem (which is NP-hard) [10,13]. In the MCDS problem, we require a dominating set that is connected and is *minimal* (i.e., no subset of the solution is a MCDS). In contrast to the approximate minimum connected dominating set problem (i.e., finding a connected dominating set that is not too large compared to the optimal) which admits efficient distributed algorithms [15,18] (polylogarithmic run time algorithms are known for both the LOCAL and CONGEST model for the unweighted case), we show that it is impossible to obtain an efficient distributed algorithm for MCDS. In Section 4.3, we use our hypergraph MIS algorithm of Section 3 as a subroutine to construct a distributed algorithm for MCDS that runs in time $\tilde{O}(D(D \min\{\Delta^\epsilon, n^{o(1)}\} + \sqrt{n}))$. We also show that $\tilde{\Omega}(D + \sqrt{n})$ is a lower bound on the run time for any distributed MCDS algorithm.

Algorithms for other hypergraph problems. Besides MIS (and the above related standard graph problems), we also study distributed algorithms for coloring, maximal matching, and maximal clique in hypergraphs. We show that a $\Delta + 1$ -coloring of a hypergraph (of any arbitrary dimension) can be computed in $O(\log n)$ rounds (this generalizes the result for standard graphs). We also show that maximal matching in hypergraphs can be solved in $O(\log m)$ rounds. Maximal clique is a less-studied problem, even in the case of graphs, but nevertheless interesting. Given a (standard) graph $G = (V, E)$, a maximal clique (MC) L is subset of V such that L is a clique in G and is maximal (i.e., it is not contained in a bigger clique).

MC is related to MIS since any MIS in the complement graph G^c is an MC in G . For a hypergraph, one can define an MC with respect to the server graph (cf. Section 2). Finding MC has applications in finding a *non-dominated coterie* in quorum systems [35]. We show that an MC in a hypergraph can be found in $O(\text{DIM} \log n)$ rounds, where DIM is the dimension of the hypergraph and n is the number of nodes. All the above results hold in the CONGEST model as well.

1.2 Technical Overview and Other Related Work

We study two natural network models for computing with hypergraphs — the *server-client* model and the *vertex-centric* models (cf. Section 2). The server-client model is commonly used in packing and covering problems such as set cover and packing LPs (e.g., [42,2,40,25,7,24]). It is also a natural model for the facility location problem (e.g., [38,39]). The vertex-centric model was considered in, e.g., [23]. Our algorithmic results apply to both models (except the one on maximal matching).

The distributed MIS problem on hypergraphs is significantly more challenging than that on (standard) graphs. Simple variants/modifications of the distributed algorithms on graphs (e.g., Luby’s algorithm and its variants [33,36,41]) do not seem to work for higher dimensions, even for hypergraphs of dimension 3. For example, running Luby’s algorithm or its permutation variant [33] on a (standard) graph by replacing each hyperedge with a clique does not work — in the graph there can be only one node in the MIS, whereas in the hypergraph all nodes of the clique, except one, can be in the MIS. It has been conjectured by Beame and Luby [8] that a generalisation of the permutation variant of an algorithm due to Luby [33] can give a $\text{polylog}(m+n)$ run time in the PRAM model, but this has not been proven so far (note that this bound itself can be large, since m can be exponential in n).

Our distributed hypergraph MIS algorithm (Section 3) consists of several ingredients. A key ingredient is the *decomposition lemma* (cf. Lemma 3) that shows that the problem can be reduced to solving a MIS problem in a low diameter network. The lemma is essentially an application of the *network decomposition* algorithm of Linial and Saks [32]. This applies to the CONGEST model as well — the main task in the proof is to show that the Linial-Saks decomposition works for (both) the hypergraph models in the CONGEST setting. The polylogarithmic run time bound for the LOCAL model follows easily from the decomposition lemma. However, this approach fails in the CONGEST model, since it involves collecting a lot of information at some nodes. The next ingredient is to show how the PRAM algorithm of Beame and Luby [8] can be simulated efficiently in the distributed setting; this we show is possible in a low diameter graph. Kelsen’s analysis [22] of Beame-Luby’s algorithm (which shows a polylogarithmic time bound in the PRAM model for *constant* dimension hypergraphs) immediately gives a polylogarithmic round algorithm in the CONGEST model for a hypergraph of *constant* dimension. To obtain the $\tilde{O}(\Delta^\epsilon)$ algorithm (for any constant $\epsilon > 0$) for a hypergraph of *arbitrary* dimension in the CONGEST model, we use another ingredient: we generalize a theorem of Turan (cf. Theorem 6) for hypergraphs — this shows that a hypergraph of low average degree has a *large* independent set. We show further that such a large independent set can be found when the network diameter is $O(\log n)$. Combining this theorem with the analysis of Beame and Luby’s algorithm gives the result for the CONGEST model for any dimension. Our CONGEST model algorithm, as pointed out earlier, also implies a $\tilde{O}(\Delta^\epsilon)$ round algorithm for the PRAM model. Recently, independently of our result, Bercea et al.[9] use a similar approach to obtain an improved algorithm for the PRAM model. In particular, they improve Kelsen’s analysis of Beame-Luby algorithm to apply also for slightly *super-constant* dimension. This improved analysis of Kelsen also helps us in obtaining a slightly better bound (cf. Theorem 1).

We apply our hypergraph MIS algorithm to solve two key problems — BMDS and MCDS. The BMDS problem was posed in Harris et al. [19], but no efficient distributed algorithm was known. A key bottleneck was solving the RMDS problem which appears as a subroutine in solving BMDS. In the current paper, we circumvent this bottleneck by treating the RMDS problem as a problem on hypergraphs.

The MCDS problem, to the best of our knowledge, has not been considered before and seems significantly harder to solve in the distributed setting compared to the more well-studied approximate version of the connected dominating set problem [15,18]. The key difficulty is being *minimal* with respect to *both* connectivity and domination. We use a layered approach to the problem, by first constructing a breadth-first tree (BFS) and then adding nodes to the MCDS, level by level of the tree (starting with the leaves). We make sure that

nodes added to the MCDS in level i dominates the nodes in level $i + 1$ and is also minimal. To be minimal with respect to connectivity we cluster nodes that are in MCDS at level $i + 1$ by connected components and treat these as super-nodes. To minimally dominate these super nodes we use the hypergraph MIS algorithm; however there is a technical difficulty of simulating the hypergraph algorithm on super-nodes. We show that such a simulation can be done efficiently by reducing the dimension of the constructed hypergraph (cf. Lemma 10) which show that hypergraph MIS on a hypergraph of arbitrary dimension can be reduced to solving a equivalent problem in a hypergraph of $\text{polylog}(m + n)$ dimension with only $O(\log n)$ factor slow down. We also show a lower bound of $\Omega(D + \sqrt{n})$ for the MCDS problem by using the techniques of Das Sarma et al. [14]. This lower bound holds even when $D = \text{polylog } n$. In this case, our upper bound is tight up to a $\text{polylog } n$ factor. We also show that $\Omega(D)$ is a *universal* lower bound for MCDS as well as for maximal clique and spanning tree problems, i.e., it applies essentially to all graphs.

2 Preliminaries

A hypergraph \mathcal{H} consists of a set $V(\mathcal{H})$ of n (hyper)nodes and a set family $E(\mathcal{H})$ of m hyperedges, each of which is a subset of $V(\mathcal{H})$. We define the *degree of node u* to be the total number of hyperedges that u is contained in. Furthermore, we define the *degree of the hypergraph*, denoted by Δ , as the maximum over all hypernode degrees. The size of each hyperedge is bounded by the *dimension* DIM of \mathcal{H} ; note that a hypergraph of dimension 2 is a graph.

We now introduce our main model of computation. In our distributed model, \mathcal{H} is realized as a (standard) undirected bipartite graph G with vertex sets S and C where $|S| = n$ and $|C| = m$. We call S the set of *servers* and C the set of *clients* and denote this realization of a hypergraph as the *server-client model*. That is, every vertex in S corresponds to a vertex in \mathcal{H} and every vertex in C corresponds to a hyperedge of \mathcal{H} . For simplicity, we use the same identifiers for vertices in C as for the hyperedges in \mathcal{H} . There exists a (2-dimensional) edge in G from a server $u \in S$ to a client $e \in C$ if and only if $u \in e$. See Figure 1a for an example. Thus, the degree of \mathcal{H} is precisely the maximum degree of the servers and the dimension of \mathcal{H} is given by the maximum degree of the clients.

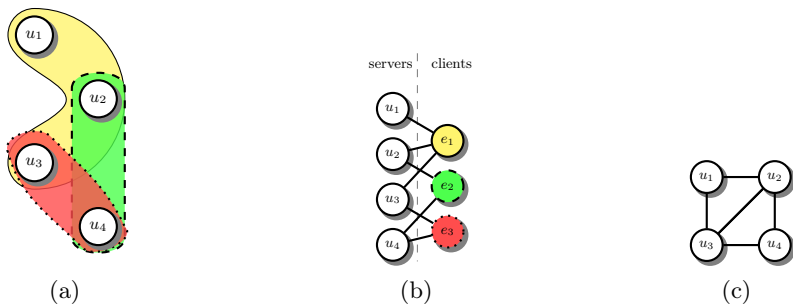


Fig. 1: Figure (a) depicts a hypergraph consisting of vertices u_1, \dots, u_4 and edges $e_1 = \{u_1, u_2, u_3\}$, $e_2 = \{u_2, u_4\}$, and $e_3 = \{u_3, u_4\}$. Figures (b) and (c) respectively show this hypergraph in the bipartite server-client model and the vertex-centric model.

An alternative way to model a hypergraph \mathcal{H} as a distributed network is the *vertex-centric* model (cf. Figure 1c). Here, the nodes are exactly the nodes of \mathcal{H} and there exists a communication link between nodes u and v if and only if there exists a hyperedge $e \in E(\mathcal{H})$ such that $u, v \in e$. Note that in this model, we assume that every node locally knows all hyperedges in which it is contained. For any hypergraph \mathcal{H} , we call the above underlying communication graph in the vertex-centric model (which is a standard graph) the *server graph*, denoted by $G(\mathcal{H})$.

We consider the standard synchronous round model (cf. [41]) of communication. That is, each node has a unique id (arbitrarily assigned from some set of size polynomial in n) and executes an instance of a distributed

algorithm that advances in discrete *rounds*. To correctly model the computation in a hypergraph, we assume that each node knows whether it is a server or a client. In each round every node can communicate with its neighbors (according to the edges in the server-client graph) and perform some local computation. We do not assume shared memory and nodes do not have any a priori knowledge about the network at large.

We will consider two types of models — CONGEST and LOCAL [41]. In the CONGEST model, only a $O(\log n)$ -sized message can be sent across a communication edge per round. In the LOCAL model, there is no such restriction. Unless otherwise stated, we use the CONGEST model in our algorithms.

3 Distributed Algorithms for Hypergraph MIS Problem

We present randomized distributed algorithms and prove the following for the hypergraph MIS problem:

Theorem 1. *The hypergraph MIS problem can be solved in the following expected time⁷ in both vertex-centric and server-client representations.*

1. $O(\log^2 n)$ time in the LOCAL model.
2. $O(\log^{(d+4)!+4} n)$ time⁸ in the CONGEST model when the input hypergraph has constant dimension d .
3. $O(\min\{\Delta^\epsilon \log^{(1/\epsilon)^{O(1/\epsilon)}} n, \sqrt{n}\})$ time in the CONGEST model for any dimension, where ϵ is such that $1 \geq \epsilon \geq \frac{1}{\frac{\log \log n}{c \log \log \log n} - 1}$ from some (large) constant c . (In particular, $\Delta^\epsilon \log^{(1/\epsilon)^{O(1/\epsilon)}} n$ becomes $\Delta^{o(1)} n^{o(1)}$ when we use $\epsilon = \frac{1}{\frac{\log \log n}{c \log \log \log n} - 1}$.)

In Section 3.1, we prove a *decomposition lemma* which plays an important role in achieving all the above results.

3.1 Low-Diameter Decomposition

First, we note that, for solving MIS, it is sufficient to construct an algorithm that solves the following *subgraph-MIS* problem on low-diameter networks.

Definition 2 (Subgraph-MIS Problem). *In the Subgraph-MIS problem, we are given an n -node network G . This network is either in a vertex-centric or server-client representation of some hypergraph \mathcal{H} . Additionally, we are given a subnetwork G' of G representing a sub-hypergraph⁹ \mathcal{H}' of \mathcal{H} . The goal is to find an MIS of \mathcal{H}' .*

Lemma 3 (Decomposition Lemma). *For any function T , if there is an algorithm \mathcal{A} that solves subgraph-MIS on CONGEST server-client (respectively vertex-centric) networks G of $O(\log n)$ diameter in $T(n)$ time (where n is the number of nodes in G), then there is an algorithm for MIS on CONGEST server-client (respectively vertex-centric) networks of any diameter that takes $O(T(n) \log^4 n)$ time.*

The main idea of the lemma is to run the *network decomposition* algorithm of Linial and Saks [32] and simulate \mathcal{A} on each cluster resulting from the decomposition. The only part that we have to be careful is that running \mathcal{A} simultaneously on many clusters could cause a congestion. We show that this can be avoided by a careful scheduling. The details are as follows.

The network decomposition algorithm of [32] produces an $O(\log n)$ -decomposition with weak-diameter $O(\log n)$. That is, given a (two-dimensional) graph G , it partitions nodes into sets S_1, S_2, \dots, S_k and assigns color $c_i \in \{1, 2, \dots, O(\log n)\}$ to each set S_i with the following properties:

- the distance between any two nodes in the same set S_i is $O(\log n)$, and

⁷ Our time bounds can also be easily shown to hold with high probability, i.e., with probability $1 - 1/n$.

⁸ As is common, we use the notation $\log^f n$ which is the same as $(\log n)^f$.

⁹ Given a subset $V' \subseteq V$, a sub-hypergraph of \mathcal{H} is simply a hypergraph induced by V' — except hyperedges that contain vertices that do not belong to V' , all other hyperedges of \mathcal{H} (which intersect with V') are present in the sub-hypergraph.

- any two neighboring nodes of the same color must be in the same set (in other words, any two “neighboring” sets must be assigned different colors).

This algorithm takes $O(\log^2 n)$ time even in the CONGEST model [32]. We use the above decomposition algorithm to decompose the server graph $G(\mathcal{H})$ (cf. Section 2) of the input hypergraph. The result is the partition of hypernodes (servers) into colored sets satisfying the above conditions (in particular, two nodes sharing the same hyperedge must be in the same partition or have different colors). In addition, we modify the Linial-Saks (LS) algorithm to produce low-diameter subgraphs that contain these sets with the property that subgraphs of the same color have “small overlap”.

Lemma 4. *Let G be the input network (server-client or vertex-centric model) representing hypergraph \mathcal{H} . In $O(\log^3 n)$ time and for some integer k , we can partition hypernodes into k sets S_1, \dots, S_k , produce k subgraphs of G denoted by G_1, G_2, \dots, G_k , and assign color $c_i \in \{1, 2, \dots, O(\log n)\}$ to each subgraph G_i , with the following properties:*

1. For all i , G_i has diameter $O(\log n)$ and $S_i \subseteq V(G_i)$.
2. For any S_i and S_j that are assigned the same color (i.e. $c_i = c_j$), there is no hyperedge in \mathcal{H} that contains hypernodes (servers) in both S_i and S_j .
3. Every edge in G is contained in $O(\log^3 n)$ graphs G_{i_1}, G_{i_2}, \dots

Observe that the first two properties in Lemma 4 are similar to the guarantees of the Linial-Saks algorithm, except that Lemma 4 explicitly gives low-diameter graphs that contain the sets S_1, \dots, S_k . The third property guarantees that such graphs have “small congestion”.

Proof. Note that the Linial-Saks algorithm works as follows. The algorithm runs in iterations where in the i^{th} iteration it will output sets of color i . In the i^{th} iteration, each vertex y selects an integer radius $r_y \in \{1, \dots, O(\log n)\}$ at random (according to some distribution). Then it broadcasts its ID and the value r_y to all nodes within distance r_y of it. For every node v , after receiving all such messages from other nodes, selects the node with highest ID from among nodes y that sends their IDs to v ; denote such node by $C(v)$. For any node y , define set S_y as the set that contains every node v that has $C(v) = y$ and its distance to y is strictly less than r_y . We call S_y the set centered at y (note that y might not be in S_y). All sets in this iteration receives color i . The distance between every pair of nodes u and v in any set S_y is $O(\log n)$ since their distance to y is $O(\log n)$. We can guarantee that there are no two neighboring nodes u and v in different sets because otherwise $C(u) = C(v)$ (this crucially uses the fact that sets are formed by nodes v whose distance to $C(v)$ is strictly less than $r_{C(v)}$). By carefully picking the distribution of r_y , [32] shows that the number of iterations is $O(\log n)$.

The following is one simple (although not the most efficient) way to simulate the above algorithm in the server-client CONGEST model to compute S_1, \dots, S_k . We implement each iteration of the above algorithm in *sub-iterations*. In the beginning of the j^{th} sub-iteration, every server y with $r_y = j$ sends its ID to its neighboring clients. We then repeat the following for $2j - 1$ steps: every node (client or server) sends the maximum ID that it receives to its neighbors. It is easy to see that after all sub-iterations every server v receives the maximum ID among the IDs of servers y such that $r_y = j$ and the distance between y and v in the server graph is at most j . Since $r_y = O(\log n)$ for every y , there are $O(\log n)$ sub-iterations and each sub-iteration takes $O(\log n)$ time. After all sub-iterations, every server v can select $C(v)$. Thus, we can simulate the Linial-Saks algorithm in $O(\log^3 n)$ time. (Simulating Linial-Saks algorithm on the vertex-centric model can be done similarly except that we will have $j - 1$ sub-iterations instead of $2j - 1$.)

We now construct G_1, \dots, G_k . At any sub-iteration above, if a node v sends the ID of some node y to its neighbors, we add its neighbors and all edges incident to v to G_y (corresponding to set S_y). Clearly, S_y is contained in $V(G_y)$ since G_y contains all nodes that receive the ID of y . This process also guarantees that G_y has $O(\log n)$ diameter since every node in G_y can reach G_y in $O(\log n)$ hops by following the path that the ID of y was sent to it. Additionally, since the simulation of the Linial-Saks algorithm finishes in $O(\log^3 n)$ rounds, and in each round we add an edge (u, v) to at most two subgraphs, we have that every edge is in $O(\log^3 n)$ subgraphs.

Proof (Proof of Lemma 3). We decompose the network as in Lemma 4. Then, we use \mathcal{A} to compute MIS iteratively in $O(\log n)$ iterations as follows. At the i^{th} iteration, we consider each set S_t and graph G_t of color i . We will decide whether each node in S_t will be in the final solution of MIS or not. We assume that we already did so for sets of colors $1, 2, \dots, i - 1$.

Let \mathcal{H}_t be the following sub-hypergraph. \mathcal{H}_t consists of all hypernodes in S_t . For each hyperedge e that contains a node in S_t , we add an edge $e' = e \cap S_t$ to \mathcal{H}_t if e contains *none* of the following hypernodes: (1) a hypernode in set S' of color $j > i$, and (2) a node in set S'' of color $j < i$ that is already decided to be *not* in the MIS. We can construct \mathcal{H}_t quickly since each server (hypernode) can decide locally whether each client (hyperedge) adjacent to it satisfies the above property or not.

Now we compute MIS of \mathcal{H}_t by simulating \mathcal{A} to solve the subgraph-MIS problem on G_t where the subgraph we want to solve is the subgraph G'_t of G_t representing \mathcal{H}_t . Note that since G_t has diameter $O(\log n)$, \mathcal{A} will finish in $T(n)$ time if we simulate \mathcal{A} on only G_t . However, we will actually simulate \mathcal{A} on *all* graphs G_{t_1}, G_{t_2}, \dots of color i *simultaneously*. Since each edge is contained in $O(\log^3 n)$ such graphs, we can finish simulating \mathcal{A} on all graphs in $O(T(n) \log^3 n)$ time.

After we finish simulating \mathcal{A} on \mathcal{H}_t , we use the solution as a solution of MIS of the original graph \mathcal{H} ; that is, we say that a hypernode is in the MIS of \mathcal{H} if and only if it is in the MIS of \mathcal{H}_t . We now prove the correctness. Let M_t be the MIS of \mathcal{H}_t . First, observe that any hypernode in M_t can be added to the MIS solution of \mathcal{H} without violating the independent constraint since \mathcal{H}_t contains all hyperedges of \mathcal{H} except those that contain some hypernode of higher color (which is not yet added to the MIS of \mathcal{H}) and hypernode of lower color that is already decided not to be in the MIS of \mathcal{H} . Secondly, the fact that any hypernode v in S_t that is not in M_t implies that there is a hyperedge e' in \mathcal{H}_t that contains all hypernodes in \mathcal{H}_t except v . Let e be a hyperedge in \mathcal{H} such that $e' \subseteq e$. Note that e does not contain any hypernode in other set $S_{t'}$ of the same color as S_t . Also observe that every hypernode in $e \setminus S_t$ must be already decided to be in the MIS of \mathcal{H} (otherwise, we will not have $e' = e \cap S_t$ in \mathcal{H}_t). Thus, every hypernode in e' except v is already in the MIS of \mathcal{H} as well; in other words, v cannot be in the MIS of \mathcal{H} . This completes the correctness of the algorithm. Thus, after we finish simulating \mathcal{A} on graphs of all colors, we obtain the MIS of \mathcal{H} . Since we need $O(T(n) \log^3 n)$ time for each color, we need $O(T(n) \log^4 n)$ time in total.

Lemma 5. *MIS can be solved in $O(\log^2 n)$ rounds in the LOCAL models (both vertex-centric and server-client representations).*

Proof. Using Lemma 4, we partition the hypernodes of the input network into subgraphs each of which have $O(\log n)$ diameter and no two subgraphs assigned the same colour share a hyper edge. Our algorithm proceeds in the same way as in the proof of Lemma 3, except that there is no congestion in the LOCAL model when we simulate \mathcal{A} (as specified in Lemma 3) on all graphs of color i . Thus, we need $O(T(n))$ time per color instead of $O(T(n) \log^3 n)$. Moreover, we can solve the subgraph-MIS problem on a network of $O(\log n)$ diameter in $O(\log n)$ time by collecting the information about the subgraph to one node, locally compute the MIS on such node, and send the solution back to every node. Thus, $T(n) = O(\log n)$. It follows that we can solve MIS on networks of any diameter in $O(\log^2 n)$ time.

3.2 $O(\log^{(d+4)!+4} n)$ time in the CONGEST model assuming constant dimension d

Let $(\mathcal{H}, \mathcal{H}')$ be an instance of the subgraph-MIS problem such that the network G representing \mathcal{H} has $O(\log n)$ diameter. We now show that we can solve this problem in $O(\log^{(d+4)!} n)$ time when \mathcal{H}' has a constant dimension d , i.e. $|e| \leq d$ for every hyperedge e in \mathcal{H}' . By Lemma 3, we will get a $O(\log^{(d+4)!+4} n)$ -time algorithm for the MIS problem in the case of constant-dimensional hypergraphs (of any diameter) which works in both vertex-centric and server-client representations and even in the CONGEST model. This algorithm is also an important building block for the algorithm in the next section.

Our algorithm simulates the PRAM algorithm of Beame and Luby [8] which was proved by Kelsen [22] to finish in $O(\log^{(d+4)!} n)$ time when the input hypergraph has a constant dimension d and this running time

was recently extended to any $d \leq \frac{\log \log n}{4 \log \log \log n}$ by Bercea et al. [9]¹⁰. The crucial part in the simulation is to compute a number $\zeta(\mathcal{H}')$ defined as follows. For $\emptyset \neq x \subseteq V(\mathcal{H}')$ and an integer j with $1 \leq j \leq d - |x|$ we define: $N_j(x, \mathcal{H}') = \{y \subseteq V(\mathcal{H}') \mid x \cup y \in E(\mathcal{H}') \wedge x \cap y = \emptyset \wedge |y| = j\}$, and $d_j(x, \mathcal{H}') = (|N_j(x, \mathcal{H}')|)^{1/j}$. Also, for $2 \leq i \leq d$, let¹¹ $\zeta_i(\mathcal{H}') = \max\{d_{i-|x|}(x, \mathcal{H}') \mid x \subseteq V(\mathcal{H}') \wedge 0 < |x| < i\}$ and $\zeta(\mathcal{H}') = \max\{\zeta_i(\mathcal{H}') \mid 2 \leq i \leq d\}$. We now explain how to compute $\zeta(\mathcal{H}')$ in $O(\log^{(d+4)!} n)$ time. First, note that we can assume that every node knows the list of members in each hyperedge that contains it: this information is already available in the vertex-centric representation; and in the server-client representation, every hyperedge can send this list to all nodes that it contains in $O(d)$ time in the CONGEST model. Every node v can now compute, for every i , $\zeta_i(v, \mathcal{H}') = \max\{d_{i-|x|}(x, \mathcal{H}') \mid x \subseteq V(\mathcal{H}') \wedge 0 < |x| < i \wedge v \in x\}$. This does not require any communication since for any x such that $v \in x$, node v already knows all hyperedges that contain x (they must be hyperedges that contain v). Now, we compute $\zeta(\mathcal{H}') = \max\{\zeta_i(v, \mathcal{H}') \mid 2 \leq i \leq d \wedge v \in V(\mathcal{H}')\}$ by computing through the breadth-first search tree of the network representing \mathcal{H} (this is where we need the fact that the network has $O(\log n)$ diameter).

Once we get $\zeta(\mathcal{H}')$, the rest of the simulation is trivial; we refer to the full paper for details. We provide some detail here for completeness. We mark each hypernode in \mathcal{H}' with probability $p = \frac{1}{2^{d+1}\zeta(\mathcal{H}'')}$. If a hyperedge has all of its nodes marked, unmark all of its nodes. Remove the hypernodes that are still marked from \mathcal{H}' and add them to the independent set. We also remove these hypernodes from \mathcal{H}' , thus reducing the size of some hyperedges in \mathcal{H}' . In the remaining hypergraph do the following: eliminate any edges properly containing another edge; remove any hypernodes that form a 1-dimension edge (i.e. remove every hypernode v such that there is a hyperedge $\{v\}$); finally, remove isolated vertices (i.e., those not contained in any edge) and add them to the independent set. Let \mathcal{H}'' be the resulting hypergraph. Repeat this procedure until there is no hypernodes left. It is easy to see that all steps (before we repeat the procedure) takes $O(1)$ rounds. Kelsen [9] and Bercea et al. [9] showed that we have to repeat this procedure only $O(\log^{(d+4)!} n)$ time (in expectation and with high probability) when $d \leq \frac{\log \log n}{4 \log \log \log n}$ (there is no guarantee for any other values of d); so, our simulation finishes in $O(\log^{(d+4)!} n)$ rounds.

3.3 $\Delta^\epsilon \log^{(1/\epsilon)^{O(1/\epsilon)}} n$ and $\Delta^{o(1)} n^{o(1)}$ Time in the CONGEST model

We rely on a modification of Turán's theorem, which states that a (two-dimensional) graph of *low* average degree has a *large* independent set (see e.g. Alon and Spencer [1]). We show that this theorem also holds for high-dimensional hypergraphs, and show further that such a large independent set can be found w.h.p when the network diameter is $O(\log n)$.

Lemma 6 (A simple extension of Turán's theorem). *Let $d \geq 2$ and $\delta \geq 2$ be any integers. Let \mathcal{H} be any hypergraph such that every hyperedge in \mathcal{H} has dimension at least d , there are n hypernodes, and the average hypernode degree is δ . (Note that the diameter of the network representing \mathcal{H} can be arbitrary.) If every node knows δ and d , then we can find an independent set M whose size in expectation is at least $\frac{n}{\delta^{1/(d-1)}}(1 - \frac{1}{d})$ in $O(1)$ time.*

Proof. We modify the proof of Theorem 3.2.1 in [1, pp.29]. Let $p = (1/\delta)^{1/(d-1)}$ (note that $p < 1$) and S be a random set of hypernodes in \mathcal{H} defined by $\Pr[v \in S] = p$ for every hypernode v . Let $X = |S|$, and let Y be the number of hyperedges in \mathcal{H} contained in S (i.e. hyperedge $e \in E(\mathcal{H})$ such that $e \subseteq S$). For each hyperedge e , let Y_e be the indicator random variable for the event $e \subseteq S$; so, $Y = \sum_{e \in E(\mathcal{H})} Y_e$. Observe that for any hyperedge e , $E[Y_e] = p^{|e|} \leq p^d$ since e contains at most d hypernodes. So, $E[Y] = \sum_{e \in E(\mathcal{H})} E[Y_e] \leq \frac{n\delta}{d} p^d$

¹⁰ The original running time of Kelsen [22] is in fact $O((\log n)^{f(d)})$ where $f(d)$ is defined as $f(2) = 7$ and $f(i) = (i-1) \sum_{j=2}^{i-1} f(j) + 7$ for $i > 2$. The $O(\log^{(d+4)!} n)$ time (which is essentially the same as Kelsen's time) was shown in [9]. We will use the latter running time for simplicity. Also note that the result in this section holds for all $d \leq \frac{\log \log n}{4 \log \log \log n}$ due to [9].

¹¹ A note on the notation: [8,22] use Δ to denote what we use ζ to denote here. We use a different notation since we use Δ for another purpose.

(the inequality is because the number of hyperedges in \mathcal{H} is at most $\frac{n\delta}{d}$). Clearly, $E[X] = np$; so,

$$E[X - Y] \geq np - \frac{n\delta}{d}p^d = np(1 - \frac{\delta}{d}p^{d-1}) = n(\frac{1}{\delta})^{\frac{1}{d-1}}(1 - 1/d)$$

where the last equality is because $p = (\frac{1}{\delta})^{\frac{1}{d-1}}$. Our algorithm will pick such a random set S . (Every node can decide whether it will be in S locally.) Then it selects one vertex from each edge of S and deletes it. (This can be done in $O(1)$ time.) This leaves a set S^* with at least $n(\frac{1}{\delta})^{\frac{1}{d-1}}(1 - \frac{1}{d})$ hypernodes in expectation. All edges having been destroyed, S^* is an independent set.

Algorithm. We use the following algorithm to solve the subgraph-MIS problem on a sub-hypergraph \mathcal{H}' of \mathcal{H} , assuming that the network representing \mathcal{H} has $O(\log n)$ diameter. Let $n' = |V(\mathcal{H}')|$. Let d be an arbitrarily large constant. Let \mathcal{H}'_d be the sub-hypergraph of \mathcal{H}' where $V(\mathcal{H}'_d) = V(\mathcal{H}')$ and we only keep hyperedges of dimension (i.e. size) at least d in \mathcal{H}'_d . (It is possible that \mathcal{H}'_d contains no edge.) We then find an independent set of expected size at least $\frac{n'}{\Delta^{1/(d-1)}}(1 - 1/d)$ in \mathcal{H}'_d , denoted by S ; this can be done in $O(1)$ time by Lemma 6 (note that we use the fact that $\delta \leq \Delta$ here). Let \mathcal{H}'_S be the sub-hypergraph of \mathcal{H}' induced by nodes in S . Note that \mathcal{H}'_S does not contain any hyperedge in \mathcal{H}'_d and thus has dimension at most d , which is a constant. So, we can run the $O(\log^{(d+4)!} n)$ -time algorithm from Section 3.2 to find an MIS of \mathcal{H}'_S . We let M'_S be such a MIS of \mathcal{H}'_S .

Our intention is to use M'_S as part of some MIS M' of \mathcal{H}' . Of course, any hypernode v in $V(\mathcal{H}'_S) \setminus M'_S$ cannot be in such M' since $M' \cup \{v\}$ will contain some hyperedge e in \mathcal{H}'_S which is also a hyperedge in \mathcal{H}' . It is thus left to find which hypernodes in $V(\mathcal{H}') \setminus S$ should be added to M'_S to construct an MIS M' of \mathcal{H}' . To do this, we use the following hypergraph. Let \mathcal{H}'' be the sub-hypergraph of \mathcal{H}' such that $V(\mathcal{H}'') = V(\mathcal{H}') \setminus S$ and for every hyperedge $e \in E(\mathcal{H}')$, we add a hyperedge $e \cap V(\mathcal{H}'')$ to \mathcal{H}'' if and only if $e \subseteq M'_S \cup V(\mathcal{H}'')$; in other words, we keep edge e that would be “violated” if we add every hypernode in \mathcal{H}'' to M' . We now find an MIS M'' of \mathcal{H}'' by recursively running the same algorithm with \mathcal{H}'' , instead of \mathcal{H}' , as a subgraph of \mathcal{H} . The correctness follows from the following claim.

Claim. $M' = M'_S \cup M''$ is a MIS of \mathcal{H}' .

Proof. First, we show that M' is an independent set of \mathcal{H}' . Assume for a contradiction that there is a hyperedge e in \mathcal{H}' such that $e \subseteq M'$. This means that $e \subseteq M'_S \cup V(\mathcal{H}'')$ since $M'_S \cup M'' \subseteq M'_S \cup V(\mathcal{H}'')$. It follows from the construction of \mathcal{H}'' that there is an edge $e' = e \cap V(\mathcal{H}'')$ in \mathcal{H}'' . Note that $e \cap V(\mathcal{H}'') \subseteq M''$; in other words $e' \subseteq M''$. This, however, contradicts the fact that M'' is an MIS in \mathcal{H}'' .

Now we show that M' is maximal. Assume for a contradiction that there is a hypernode v in $V(\mathcal{H}') \setminus M'$ such that $M' \cup \{v\}$ is an independent set. If v is in S , then $M'_S \cup \{v\}$ is an independent set in \mathcal{H}'_S (since it is a subset of $M' \cup \{v\}$), contradicting the fact that M'_S is an MIS in \mathcal{H}'_S . So, v must be in $V(\mathcal{H}'')$. This, however, implies that $M'' \cup \{v\}$ is an independent set in \mathcal{H}'' (again, since it is a subset of $M' \cup \{v\}$), contradicting the fact that M'' is an MIS in \mathcal{H}'' .

We now analyze the running time of this algorithm. Recall that $E[|S|] \geq \frac{n'}{\delta^{1/(d-1)}}(1 - 1/d)$. In other words, the expected value of $|V(\mathcal{H}'')| \leq (1 - \frac{c(d)}{\Delta^{1/(d-1)}})|V(\mathcal{H}')|$ where $c(d) = \frac{1}{2}(1 - 1/d)$ is a constant which is strictly less than one (recall that d is a constant). It follows that the expected number of recursion calls is $O(\Delta^{\frac{1}{d-1}})$. Since we need $O(\log^{(d+4)!} n)$ time to compute M'_S and to construct \mathcal{H}'' , the total running time is $O(\Delta^{\frac{1}{d-1}} \log^{(d+4)!} n)$. By Lemma 3, we can compute MIS on any hypergraph \mathcal{H} (of any diameter) in

$$O(\Delta^{\frac{1}{d-1}} \log^{(d+4)!+4} n)$$

time. For any constant $\epsilon > 0$, we set $d = 1 + 1/\epsilon$ to get the claimed running time of

$$O(\Delta^\epsilon \log^{(5+1/\epsilon)!+4} n) = \Delta^\epsilon \log^{(1/\epsilon)^{O(1/\epsilon)}} n. \quad (1)$$

Moreover, by the recent result of Bercea et al. [9], we can in fact set d as large as $\frac{\log \log n}{4 \log \log \log n}$. In this case, note that for some constant c' ,

$$(d+4)! = d^{c'd} = e^{c'd \log d} = e^{c' \cdot \frac{\log \log n}{c \log \log \log n} \cdot \log \log \log n} = \log^{1/10} n$$

where the last equality holds when we set $c = 10c'$. Thus,

$$\log^{(d+4)!} n = \log^{\log^{1/10} n} n = 2^{(\log^{1/10} n) \log \log n} = n^{o(1)}.$$

The running time thus becomes $\Delta^{o(1)} n^{o(1)}$.

3.4 $O(\sqrt{n})$ Time in the CONGEST model

We obtain the $O(\sqrt{n})$ time by modifying the PRAM algorithm of Karp, Upfal, and Wigderson [21, Section 4.1]. (Note that we do not need the fact that the network diameter is $O(\log n)$ for this algorithm.) Their algorithm is as follows. Let v_1, v_2, \dots, v_n be a random permutation of hypernodes. The algorithm gradually adds a hypernode to the independent set one by one, starting from v_1 . It stops at some hypernode v_k when v_k cannot be added to the independent set. Thus, v_1, \dots, v_{k-1} are added to the independent set; the algorithm removes these hypernodes from the graph. It also removes *all* hypernodes that cannot be added to the independent set (i.e. any v such that $\{v_1, \dots, v_{k-1}, v\}$ contains some hyperedge) and all hyperedges that contain them. It repeats the same process to find a MIS of the remaining graph. It is easy to show (see [21] for detail) that the union of a MIS of the remaining graph and $\{v_1, \dots, v_{k-1}\}$ is a MIS of the input graph. The key to proving the efficiency of this algorithm is the following.

Claim ([21]). The expected number of removed hypernodes (v_1, \dots, v_{k-1} and hypernodes that cannot be added to the independent set) in the above process is $\Omega(\sqrt{n})$.

It follows almost immediately that we have to repeat the process only $O(\sqrt{n})$ times in expectation (see [21, Appendix] for detail). We now show how to modify this algorithm to our setting. Every hypernode v picks a random integer $r(v)$ between 1 and n^2 . It can be guaranteed that hypernodes pick different numbers with high probability. Then every hypernode v marks itself to the independent set if for any hyperedge e that contains v , $r(v) < \max_{u \in e} r(u)$, i.e., its number is not the maximum in any hyperedge. We add all marked hypernodes to the independent set, remove them from the graph, and eliminate hypernodes that cannot be added to the independent set (i.e. a hypernode v marks itself as “eliminated” if there is a hyperedge e such that $e \setminus \{v\}$ is a subset of marked hypernodes). We then repeat this process until there is no hypernode left.

Using Section 3.4, we show that our algorithm has to repeat only $O(\sqrt{n})$ times, as follows. Consider an ordering v_1, \dots, v_n where $r(v_i) < r(v_{i+1})$. This is a random permutation. Let k be such that v_1, \dots, v_k are added to the independent set by Karp et al.’s algorithm and v_{k+1}, \dots, v_n are not. Observe that for every $1 \leq i \leq k$ and every hyperedge e that contains v_i , $r(v_i) < \max_{u \in e} r(u)$ (otherwise edge e will be violated when we add v_1, \dots, v_k to the independent set). In other words, our algorithm will also add v_1, \dots, v_k to the independent set (but it may add other hypernodes as well). It follows that our algorithm will eliminate every hypernode that is eliminated by Karp et al.’s algorithm. In other words, the set of hypernodes removed by our algorithm is a superset of the set of hypernodes removed by Karp et al.’s algorithm. Thus, by Section 3.4, the expected number of hypernodes removed in each iteration of our algorithm is $\Omega(\sqrt{n})$. By the same analysis as Karp et al., our algorithm will need only $O(\sqrt{n})$ iterations in expectation. Each iteration can be easily implemented in $O(1)$ rounds, so our algorithm takes $O(\sqrt{n})$ time in expectation.

4 Applications of Hypergraph MIS algorithms to standard graph problems

In this section we show that our distributed hypergraph algorithms have direct applications in the standard graph setting. As a first application of our MIS algorithm, we show how to solve the restricted minimal dominated set (RMDS) problem in Section 4.1. We will use this RMDS-algorithm to obtain a distributed algorithm for solving the balanced minimal dominating set (BMDS) problem, thereby resolving an open problem of [19].

Let R be the set of restricted nodes (which are part of the MDS).
 Simulate a server client network H . Every node (locally) adds vertices to the clients C resp. servers S , and simulates the edges in H .

- 1: **for** every node u **do**
- 2: Node u adds a client e_u to C .
- 3: **if** $u \in R$ **then**
- 4: Node u adds a server s_u to S , and an edge (s_u, e_u) to $E(H)$.
- 5: **for** all nodes u, v where $(u, v) \in E(G)$ **do**
- 6: If server s_u exists in H , add edge (s_u, e_u) to H .

7: Find an MIS on H and let $O_{MIS} \subseteq S$ be the servers that are in the output set.
 8: **for** every node u where s_u exists **do**
 9: If $s_u \notin O_{MIS}$, then node u adds itself to the RMDS.

Algorithm 4.1: An RMDS-algorithm: Finding a minimal dominating set on a graph G that is a subset of a given dominating set R .

4.1 Restricted Minimal Dominating Set (RMDS)

We are given a (standard) graph $G = (V, E)$ and a subset of nodes $R \subseteq V$, such that R forms a dominating set in G (i.e., every node $v \in V$ is either adjacent to R or belongs to R). We are required to find a *minimal* dominating set that is a subset of R and dominates V . Since a minimal vertex cover is the complement of a maximal independent set, we can leverage our MIS algorithm (cf. Section 3). To this end, we show that the RMDS problem can be solved by finding a minimal hitting set (or minimal vertex cover) on a specific hypergraph H . The server client representation of H is determined by G and R as follows: For every vertex in V we add a client (i.e. hyperedge) and, for every vertex in R , we also add a server. Thus, for every vertex $u \in V$, we have a client e_u and, if $u \in R$, we also have a server s_u . We then connect a server s_u to a client e_v , iff either u and v are adjacent in G , or $u = v$. Algorithm 4.1 contains the complete pseudo code of this construction. Note that we can simulate this server client network on the given graph with constant overhead in the CONGEST model. We have the following result by virtue of Theorem 1:

Theorem 7. RMDS can be solved in expected time $\tilde{O}(\min\{\Delta^\epsilon, n^{o(1)}\})$ (for any const. $\epsilon > 0$) on graph G in the CONGEST model and in time $O(\log^2 n)$ in the LOCAL model where Δ is the maximum degree of G .

4.2 Balanced Minimal Dominating Set

We define the *average degree* of a (standard) graph G , denoted by δ , as the total degrees of its vertices (degree of a vertex is its degree in G) divided by the number of vertices in G . A *balanced minimal dominating set (BMDS)* (cf. [19]) is a minimal dominating set D in G that minimizes the ratio of the average degree of D to that of the graph itself (the average degree of the set of nodes D is defined as the average degree of the subgraph induced by D). The BMDS problem is motivated by applications in fault-tolerance and load balancing (see [19] and the references therein). For example, in a typical application, an MDS can be used to form clusters with low diameter, with the nodes in the MDS being the “clusterheads” [37]. Each clusterhead is responsible for monitoring the nodes that are adjacent to it. Having an MDS with low degree is useful in a resource/energy-constrained setting since the number of nodes monitored *per* node in the MDS will be low (on average). This can lead to better load balancing, and consequently less resource or energy consumption per node, which is crucial for ad hoc and sensor networks, and help in extending the lifetime of such networks while also leading to better fault-tolerance. For example, in an n -node star graph, the above requirements imply that it is better for the leaf nodes to form the MDS rather than the central node alone. In fact, the average degree of the MDS formed by the leaf nodes — which is 1 — is within a constant factor of the average degree of a star (which is close to 2), whereas the average degree, $n - 1$, of the MDS consisting of the central node alone is much larger. A *centralized* polynomial time algorithm for computing a BMDS with (the

Nodes compute the average network degree δ .

Every node u of degree $> 2\delta$ marks itself with probability $\frac{\log t}{t}$ where $t = \frac{2\delta \log \delta}{\log \log \delta}$.

Every node of degree $\leq 2\delta$ marks itself.

If a node v is not marked, and none of the neighbors of v are marked, then v marks itself.

Let MARKED be the set of nodes that are marked. Invoke the RMDS algorithm (cf. Section 4.1) on G where the restricted set is given by MARKED.

Every node that is in the solution set of the RMDS algorithm remains in the final output set.

Algorithm 4.2: A distributed BMDS-algorithm.

best possible in general ¹²) average degree $O(\frac{\delta \log \delta}{\log \log \delta})$ was given in [19]. A distributed algorithm that gives the same bounds was left a key open problem. We now present a distributed variant of this algorithm (cf. Algorithm 4.2) that uses our hypergraph MIS-algorithm as a subroutine. Note that since the BMDS problem is defined on standard graphs, we assume that Algorithm 4.2 executes on a standard synchronous network adhering to the CONGEST model of communication.

Theorem 8. *Let δ be the average degree of a graph G . There is a CONGEST model algorithm that finds a BMDS with average degree $O(\frac{\delta \log \delta}{\log \log \delta})$ in expected $\tilde{O}(D + \min\{\Delta^\epsilon, n^{o(1)}\})$ rounds, where D is the diameter, Δ is the maximum node degree of G , and $\epsilon > 0$ is any constant.*

Proof. Computing the average degree in Step 1 of Algorithm 4.2 can be done by first electing a leader, then building a BFS-tree rooted at the leader, and finally computing the average degree by convergecast.

It was shown in [19] that marking the nodes according to Algorithm 4.2 yields an average degree of $O(\frac{\delta \log \delta}{\log \log \delta})$. The runtime bound follows since the first part of the algorithm can be done in $O(D)$ rounds and the running time of the RMDS-algorithm (cf. Theorem 7).

4.3 Minimal Connected Dominating Sets (MCDS)

Given a graph G , the MCDS problem requires us to find a minimal dominating set M that is connected in G . We now describe our distributed algorithm for solving MCDS in the CONGEST model (see Algorithm 4.3 for the complete pseudo code) and argue its correctness. We first elect a node u as the leader using a $O(D)$ time algorithm of [27]. Node u initiates the construction of a BFS tree B , which has $k \leq D$ levels, after which every node knows its level (i.e. distance from the leader u) in the tree B . Starting at the leaf nodes (at level k), we convergecast the maximum level to the root u , which then broadcasts the overall maximum tree level to all nodes in B along the edges of B .

We then proceed in iterations processing two adjacent tree levels at a time, starting with nodes at the maximum level k . Note that since every node knows k and its own level, it knows after how many iterations it needs to become active. Therefore, we assume for simplicity that all leafs of B are on level k . We now describe a single iteration concerning levels i and $i - 1$: First, consider the set L_i of level i nodes that have already been added to the output set M in some previous iteration; initially, for $i = k$, set L_i will be empty. We run the $O(D + \sqrt{n})$ time algorithm of [43] to find maximal connected components among the nodes in L_i in the graph G ; let $\mathcal{C} = \{C_1, \dots, C_\alpha\}$ be the set of these components and let ℓ_j be the designated component leader of component $C_j \in \mathcal{C}$.

We now simulate a hypergraph that is defined as the following bipartite server client graph H : Consider each component in \mathcal{C} as a *super-node*; we call the other nodes on level i *non-super-nodes*. The set \mathcal{C} of clients contains all super-nodes in \mathcal{C} and all nodes on level i that are neither adjacent to any super-node nor have been added to the output set O . The set S of servers contains all nodes on level $i - 1$. The edges of H are the induced inter-level edges of G between servers and non-super-node clients. In addition, we add an edge between a server $s \in S$ and a (super-node) client $C_j \in \mathcal{C}$, iff there exists a $v \in C_j$ such that $(v, s) \in E(G)$. Conceptually, we can think of the edges incident to C_j as pointing to the component leader node ℓ_j . Next,

¹² That is, there exists graphs with average degree δ , where this bound is essentially the optimal.

we find a MIS (cf. Section 3) on the (virtual) hypergraph H . We sketch how we simulate the run of the MIS algorithm on H in G : If a node $v \in C_j$ receives a message from a node in S , then v forwards this message to the component leader ℓ_j . (If a node receives multiple messages at the same time, it simply forwards all messages sequentially by pipelining.) After waiting for $\tilde{O}(D)$ rounds, the component leader ℓ_j locally simulates the execution of ℓ_j according to the MIS algorithm by using the received (forwarded) messages. Any messages produced by the simulation at ℓ_j are then sent back through the same paths to the neighbors of C_j . Let O_i be the set of nodes (on level $i - 1$) that are not in the MIS; note that O_i forms a minimal vertex cover on the hypergraph given by H . At the end of this iteration, we add O_i to the output set M and then proceed to process levels $i - 1$ and $i - 2$.

Theorem 9. MCDS can be solved in the CONGEST model in expected time $\tilde{O}(D(D \min\{\Delta^{o(1)}, n^{o(1)}\} + \sqrt{n}))$.

Proof. We first argue the correctness of the algorithm. It is straightforward to see that, after k iterations, the solution set $M = \bigcup_{i=1}^k O_i$ forms a dominating set of G . For connectivity, note that since O_i is a minimal vertex cover on the induced subgraph H , it follows that every super-node in the client set has a neighboring node in O_i . This guarantees that M remains connected (in G) after adding O_i .

Next, we consider minimality. Suppose that there exists a node w in the solution set M that is *redundant* in the sense that it can be removed from M such that $M \setminus \{w\}$ is a MCDS of G . Assume that w became redundant in the iteration when processing levels j and $j - 1$. Note that by the properties of the BFS tree, w must be either on levels j or $j - 1$, since, in this iteration, we only add new nodes to M that are themselves on level $j - 1$. By the correctness of the MIS algorithm, w does not become redundant in the same iteration that it is added to M , thus w can only be on level j . Moreover, observing that w can only have been added to M in the preceding iteration to dominate some node x on level $j + 1$, it follows that w cannot be made redundant by adding some node z on level $j - 1$, since z cannot dominate x . This shows that the set M is minimal as required.

We now argue the running time bound. The pre-processing steps of electing a leader and constructing a BFS tree can be completed in $O(D)$ rounds. The for-loop of the MCDS algorithm has $O(D)$ iterations, thus it is sufficient if we can show that we can simulate a single iteration (including finding a MIS on the constructed hypergraph H) in $\tilde{O}(\sqrt{n} + D \min\{\Delta^\epsilon, n^{o(1)}\})$ rounds. Consider the iteration that determines the status of nodes in level i , i.e., the nodes on level i form the set of servers as defined in MCDS algorithm. First, we run the algorithm of [43], which, given a graph G and a subgraph G' , yields maximal connected components (w.r.t. G') in time $\tilde{O}(D + \sqrt{n})$, where D is the diameter of G . Then, we simulate the MIS algorithm Section 3 on the hypergraph H given by the set of servers and the clients (some of which are super-nodes). Consider a super-node C_j . We can simulate a step of the MIS algorithm by forwarding all messages that nodes in C_j receive (from servers on level $i - 1$) to the component leader node ℓ_j by sequentially pipelining simultaneously received messages. The following lemma shows that we can assume that each client has at most $O(\log n)$ incident servers, i.e., the dimension of the hypergraph H is bounded by $O(\log n)$:

Lemma 10. If there is an algorithm \mathcal{A} that solves MIS on n -node m -edge hypergraphs of dimension up to $3 \log(m + n)$ in $T(n)$ rounds for some function T , then there is an algorithm \mathcal{A}' that solves hypergraph MIS on any n -node m -edge hypergraph with any dimension in $\tilde{O}(T(n))$ rounds.

Proof (Proof of Lemma 10). \mathcal{A}' works as follows. Let \mathcal{H} be the input graph. We will use M as a final MIS solution for \mathcal{A}' ; initially, $M = \emptyset$. First, we mark every hypernode with probability $1/2$. Let \mathcal{H}' be the subgraph of \mathcal{H} induced by marked nodes (i.e. \mathcal{H}' consist of every edge such that every node it contains is marked). Observe that, with probability at least $1 - 1/m^2$, every hyperedge in \mathcal{H}' has dimension at most $3 \log m$ because every hyperedge that contains more than $3 \log m$ nodes will have all its nodes marked with probability at most m^{-3} . We now run \mathcal{A} to solve hypergraph MIS on \mathcal{H}' . We add all nodes in the resulting MIS to M and remove them from \mathcal{H} . Additionally, we remove from \mathcal{H} all other nodes in \mathcal{H}' (that are not in the MIS of \mathcal{H}') and edges containing them. (These nodes cannot be added to the MIS of \mathcal{H} so they are removed.) We then repeat this procedure to find the MIS of the remaining graph. Observe that this procedure removes $n/2$ nodes in expectation. So, we have to repeat it only $O(\log n)$ times in expectation. Each of this procedure takes $O(T(n))$ time, so we have the running time of $\tilde{O}(T(n))$ in total.

- 1: Let M be the final output set; initially $M = \emptyset$.
- 2: We perform leader election using an $O(D)$ time algorithm of [27], yielding some leader ℓ .
- 3: Node ℓ initiates the construction of a breadth-first-search tree B of $k \leq D$ levels.
- 4: The leafs of B report their level (i.e. distance from the root) to ℓ by convergecast and the leader ℓ then rebroadcasts the maximum level to all children along the tree edges. At the end of this step, every node knows its level in B and the maximum tree level.
- 5: **for** tree level $i = k, \dots, 1$ **do**
- 6: Let $L_i \subseteq M$ denote the nodes on level i that have been added to M . (Note that L_k is empty initially.) Find a set of maximal connected components $\mathcal{C} = \{C_1, \dots, C_\alpha\}$ of the nodes in L_i using the $O(D + \sqrt{n})$ time algorithm of [43]; let $\ell_1, \dots, \ell_\alpha$ denote the roots of the respective components.

Solving MIS on the hypergraph induced by levels i and $i - 1$:
- 7: We construct the following bipartite server client graph H . Consider each component in \mathcal{C} as a “super-node”. The set \mathcal{C} of clients contains all super-nodes in \mathcal{C} and all nodes on level i that are neither adjacent to any super-node nor have been added to the output set O . The set S of servers contains all nodes on level $i - 1$. The edges of H are the induced inter-level edges of G between servers and clients that do not form a component. In addition, add an edge between $s \in S$ and $C_j \in \mathcal{C}$, iff there exists an $v \in C_j$ such that $(v, s) \in E(G)$. Conceptually, we can think of the edges incident to C_j to point to the component leader node ℓ_j .
- 8: Find a MIS (cf. Section 3) on the virtual hypergraph H :
We sketch how we simulate the run of the MIS algorithm on H in G : If a node $v \in C_j$ receives a message from a node in S , v forwards this message to the component leader ℓ_j . (If a node receives multiple messages at the same time, it simply forwards all messages sequentially by pipelining.) After waiting for $\tilde{O}(D)$ rounds, the component leader ℓ_j locally simulates the execution of the MIS algorithm by using the received (forwarded) messages. Any messages produced by the simulation at ℓ_j are then sent back through the same paths to the neighbors of C_j .
- 9: Add every node on level $i - 1$ that is not in the MIS to the output set M .

Algorithm 4.3: A distributed MCDS-algorithm.

It follows from Lemma 10 that forwarding messages towards the component leader can incur a delay of at most $O(\log n)$ additional rounds due to congestion. This means that one step of the MIS algorithm can be implemented in $\tilde{O}(D)$ rounds, and thus the total time complexity of a single iteration of the for-loop takes time $\tilde{O}(D \min\{\Delta^\epsilon, n^{o(1)}\} + \sqrt{n})$, as required.

4.4 Lower Bounds

In this section we show lower bounds on the number of rounds for computing a minimal connected dominating set (MCDS). First, we show that $\tilde{\Omega}(D + \sqrt{n})$ rounds are necessary in the worst case for solving MCDS in the CONGEST model showing a reduction to the spanning connected subgraph problem (SCS).

We then consider the LOCAL model where nodes can send messages of unbounded size. Here we present a lower bound of $\Omega(D)$ rounds for computing a minimal connected dominating set (MCDS). While it is easy to see that this lower bound holds on a cycle of n nodes, we show that $\Omega(D)$ is a *universal* bound in the sense that, for any given diameter $D = D(n)$ as a function of n , we can construct a graph where the algorithm takes $\Omega(D)$ time. As a byproduct of our proof, we obtain the same lower bound for spanning tree computation and maximal clique.

Theorem 11. *There exists an $\epsilon > 0$ and a graph G of n nodes and diameter D , such that any ϵ -error MCDS algorithm R takes $\tilde{\Omega}(D + \sqrt{n})$ rounds in the CONGEST model.*

Proof. We will show the lower bound via reduction from the spanning connected subgraph (SCS) problem, for which an $\tilde{\Omega}(D + \sqrt{n})$ lower bound is already known (cf. [14]). Consider an instance of the SCS problem: we are given a set of edges defining a subgraph H of graph G and every node must output “yes” if H spans G and is connected; otherwise at least one node must output “no”.

Suppose that we are given the MCDS algorithm R as stated in the theorem. We will first show that, as long as graph G has $O(n)$ edges, we can instantiate R to yield a solution for SCS without significant overhead. Since the lower bound graph for the SCS problem in [14] has $O(n)$ edges, this will yield the result. For an instance of SCS given by G and $H \subseteq G$, we will construct a graph $G' = G'(G, H)$ of $\Theta(|V(G)|)$ vertices and $\Theta(|E(G)|)$ edges: We initialize $V(G')$ to $V(G)$ and subdivide each edge $(u, v) \in E(G)$ by adding a *subdividing vertex* $b_{u,v}$ to $V(G')$ and edges $(u, b_{u,v}), (b_{u,v}, v)$ to $E(G')$. Let $B[H]$ be the set of all vertices that subdivide an edge in H and let $B[G \setminus H]$ be the set of vertices subdividing other edges in $E(G) \setminus E(H)$. Then, for every vertex $u \in V(G) \cup B[H]$ we add an *outer vertex* g_u to $V(G')$ and attach it to $u \in V(G')$ by adding the edge (g_u, u) to $E(G')$. In other words, we attach outer vertices to all nodes that were part of the original graph G and to all nodes that subdivide an edge in H .

Lemma 12. *If M is an minimal connected dominating set of G' , then the following holds: $\forall u \in M: u \notin B[G \setminus H]$ if and only if H is a spanning connected subgraph of G'*

Proof. First, observe that to dominate an outer vertex g_u , it is necessary that either $g_u \in M$ or its (only) neighbor v is in M . In the former case, it follows that v must also be in M to satisfy connectivity. But then we could remove g_u from M and still guarantee domination; thus it follows that no outer vertex is in M and, every neighbor of an outer vertex is in M . In particular, this means that all vertices of G and all vertices that subdivide edges of H must be in M (since for each of these we added an outer vertex). Finally, we observe that H is not a connected subgraph if and only if M needs to contain vertices that subdivide edges *not* in H . This completes the proof of Lemma 12.

Armed with Lemma 12, we can simply invoke algorithm R to test whether H is a spanning connected subgraph of G . Assuming that $|E(G)| \in O(n)$, it follows that asymptotically G' and G have the same number of vertices and edges and thus it is straightforward to simulate the run of the MCDS algorithm R on the (virtual) graph G' on top of the actual network G . From [14] we know that there exists a graph of n nodes and $O(n)$ edges where SCS takes time $\Omega(D + \sqrt{n})$. This completes the proof of Theorem 11.

Theorem 13 (Universal Lower Bound). *Let R be an algorithm that computes a Spanning Tree (resp. maximal clique and MCDS) in the LOCAL model with probability at least $15/16 + \epsilon$, for any constant $\epsilon > 0$. Then, for every sufficiently large n and every function $D(n)$ with $2 \leq D(n) < n/4$, there exists a graph G of $n' \in \Theta(n)$ nodes and diameter $D' \in \Theta(D(n))$ where R takes $\Omega(D)$ rounds with constant probability.*

Proof. For a given n and a function $D(n)$, we construct the following lower bound graph G of $\Theta(n)$ nodes and diameter $\Theta(D(n))$. Let $d \geq 1$ be the largest integer such that $n = 4dD(n) + \ell$, for $0 \leq \ell < 4D$; we will construct a graph G of $n' = n - \ell \in \Theta(n)$ vertices and diameter $D' = \lfloor (n - \ell)/8d \rfloor \in \Theta(D)$. Let $u_0, \dots, u_{n'-1}$ be the vertices of G . We will consider the set of *bridge vertices* defined as $\{u_i \in V(G) \mid \exists k \geq 0: i = kd < n'\}$ to describe the edges of G ; vertices not in this set are the *non-bridge vertices* of G . That is, for every bridge vertex u_i , we add the arc edges $(u_i, u_{i+1}), (u_i, u_{i-1}), \dots, (u_i, u_{i+d}), (u_i, u_{i-d})$ (indices are modulo n'). See Figure 2 for a concrete instance of this graph.

We first observe that solving maximal clique in this graph provides a leader node by simply running the $O(1)$ time leader election algorithm of [28] on the clique, thus showing that maximal clique takes $\Omega(D)$ time.

We now describe how to solve the leader election problem on G given an MCDS-algorithm or an ST-algorithm. Let $b_0, \dots, b_{n'/4d-1}$ be an ordering of the bridge vertices according to their adjacencies in G . As there are no edges between non-bridge vertices, any MCDS M must contain all except possibly 1 bridge vertex to guarantee connectivity. Moreover, the fact that every bridge vertex b_i dominates b_{i-1} and b_{i+1} (modulo $n'/4d$) implies that M must omit a bridge vertex to be minimal.

Observation 1. *If M is an MCDS of G , then there is exactly one bridge vertex $b_i \in G$ such that $b_i \notin M$.*

We call the subgraph that consists of a bridge vertex b_i and its adjacent vertices a *bridge cluster*. Analogously to Observation 1, we have the following:

Observation 2. *Let B be a partitioning of G into edge-disjoint bridge clusters and let S be a spanning tree of G . Then, there is exactly one bridge cluster $b \in B$ such that the subgraph $b \cap S$ is disconnected.*




Fig. 2: The Lower Bound Graph of Theorem 13 for $n = 12$ and diameter 2. Bridge vertices are marked by a double frame. The shaded regions represent two bridge clusters that partition the vertices into edge-disjoint sets.

Suppose that R is an algorithm that solves ST (the argument is analogous for MCDS) with probability p in time T . We first run R to obtain a spanning tree of G and then instruct every bridge vertex to check whether its cluster is connected. By construction, every vertex locally knows if it is a bridge vertex since non-bridge vertices have exactly 2 edges while bridge vertices have degree > 2 . By Observation 2, exactly 2 bridge vertices b_i and b_{i+1} will determine that their (overlapping) clusters are disconnected. The nodes b_i and b_{i+1} determine which of them has the greater id; this node then elects itself as the leader, while all other nodes enter the non-elected state. Thus there is an algorithm that elects a leader in $O(T)$ rounds with probability p .

It was shown in Theorem 3.13 of [27] that there is a class of graphs G_n with diameter $D(n)$ such that leader election takes $\Omega(D(n))$ rounds with constant probability. The proof of this result relies on the fact that the vertices of G_n can be partitioned into 4 disjoint but symmetric sets C_1, \dots, C_4 such that the distance between C_1 and C_3 (resp. C_2 and C_4) is $\Omega(D)$. It is straightforward to check that these properties also hold true in our graph class G . In particular, all bridge vertices observe the same round r -neighborhood of G , for all $r \geq 1$. Thus the proof of Theorem 3.13 in [27] can be adapted to our graph G . (We defer the details of this adaptation to the full version of the paper.) Together with the above reduction from leader election, this implies the sought time bound of $\Omega(D)$ rounds (with constant probability) for computing a minimal connected dominating set and finding a spanning tree on G .

5 Distributed Algorithms for Other Hypergraph Problems

Many algorithms in this section will simulate an algorithm for finding a MIS on a (standard) graph developed by Luby [33] as a subroutine. One version of this algorithm is this: (1) Randomly assign unique priorities to nodes in G (which can be achieved w.h.p. by having each node in G randomly pick an integer between 1 and n^4). (2) We mark and add all nodes that has higher priority than all its neighbors to the independent set. (3) We remove these marked nodes and their neighbors from the graph and repeat the procedure. Luby [33] shows that this procedure will repeat only $O(\log n)$ times in expectation. So, it is sufficient to get $\tilde{O}(1)$ time if our algorithms can simulate the three steps above in $\tilde{O}(1)$ time.

5.1 Maximal Clique

Theorem 14. *Maximal clique can be computed in $\tilde{O}(D)$ time in the CONGEST vertex-centric model and $\tilde{O}(D + \text{DIM})$ -time in the CONGEST server-client model, where D is the network (i.e., server graph or the server-client bipartite graph) diameter and DIM is the hypergraph dimension.*

Proof. Recall that in this problem, we want a maximal set S of hypernodes such that every two hypernodes u and v in S are contained in some common hyperedge. This is equivalent to finding a maximal clique in the server graph (defined in Section 2).

Since the underlying network of the vertex-centric model is exactly the server graph, we can easily find a maximal clique in this model, as follows. Pick any node s . (This can be done in $O(D)$ time by, e.g. picking a node with smallest ID or using a leader election algorithm.) Let S be the set of all neighbors of s . Let G_S be the subgraph of the server graph induced by nodes in S . Observe that if M is a maximal clique in G_S then $\{s\} \cup M$ is a maximal clique in G . So, it is sufficient to find a maximal clique in G_S . Observe further that if \bar{G}_S is the complement graph of G_S (i.e. an edge (u, v) is in \bar{G}_S if and only if it is not in G_S), then finding a maximal clique in G_S is equivalent to finding a MIS in \bar{G}_S .

We now simulate Luby's algorithm to find a MIS in \bar{G}_S . We simulate the first step by letting node s generate a random permutation of nodes in \bar{G}_S , say $v_1, v_2, \dots, v_{|S|}$, and send a priority i to node v_i . This can be done in one round since all nodes in \bar{G}_S are neighbors of s . Now, for every node v in \bar{G}_S of priority, say i , checks whether its priority is higher than all its neighbors in \bar{G}_S (as required by the second step of Luby's algorithm). Observe that this is the case if and only if the priorities $i+1, i+2, \dots, |S|$ are given to v 's neighbors in G_S . Node v can check this in one round by receiving the priorities of all its neighbors in G_S . For simulating the third step, each node v has to know whether it has a neighbor in \bar{G}_S that is marked. We do this by counting the number of marked nodes (every node tells s whether it is marked or not). Let c be such number. Then, every node v counts how many of its neighbors in G_S are marked. If this is less than c , then v has a neighbor in \bar{G}_S that is marked. This takes $O(1)$ rounds.

The above simulation of Luby's algorithm can be extended to the server-client model with an extra $O(\text{DIM})$ factor cost: For s to distribute the priorities in the first step, it has to send up to DIM priorities to the same hyperedge. For the second step, where each node v has to check whether its priority is higher than all its neighbors in \bar{G}_S , v has to receive the priorities of all its neighbors in G_S , and it might have to receive up to DIM priorities from the same hyperedge. Finally, for the third step where every node has to know the number of neighbors in G_S that are marked, it has to receive the list of IDs of its marked neighbors, and it might have to receive up to DIM IDs from the same hyperedge.

Note that the dependence on the diameter in the running time is necessary, as shown in Theorem 13.

5.2 $(\Delta + 1)$ -Coloring

Theorem 15. *The $(\Delta + 1)$ -coloring problem on hypergraphs has the same complexity as the $(\Delta + 1)$ -coloring problem on standard (two-dimensional) graphs; in particular, it can be solved in $O(\log n)$ time. This holds in both vertex-centric and server-client representations and even in the CONGEST model.*

Proof. Recall that in the $(\Delta + 1)$ -coloring problem we want to color hypernodes so that there is no monochromatic hyperedge, i.e. all hypernodes it contains have the same color. We solve this problem by converting a hypergraph \mathcal{H} to a two-dimensional graph G on the same set of nodes as follows. For every hyperedge e in \mathcal{H} , pick arbitrary two distinct hypernodes it contains, say u and v , and create an edge $e' = (u, v)$ in G . Observe that G has maximum degree at most Δ and any valid coloring in G will be a valid coloring in \mathcal{H} (since if an edge e in \mathcal{H} is monochromatic, then the corresponding edge e' in G will also be monochromatic). Thus, it is sufficient to find a $(\Delta + 1)$ coloring in G . We can do this by simulating any $(\Delta + 1)$ -coloring algorithm for G on \mathcal{H} . This shows that $(\Delta + 1)$ -coloring on hypergraphs is *as easy as* $(\Delta + 1)$ -coloring on standard graphs.

5.3 Maximal Matching

Theorem 16. *The maximal matching problem on hypergraphs can be solved in $O(\log n)$ time in the CONGEST server-client model.*

Proof. Recall that this problem on a hypergraph \mathcal{H} asks for a maximal set $S \subseteq E(\mathcal{H})$ of *disjoint* hyperedges, i.e. $e \cap e' = \emptyset$ for all $e \neq e'$ in S . Consider the following *line graph* G : nodes of G is the hyperedges in \mathcal{H} ,

i.e. $V(G) = E(\mathcal{H})$, and there is an edge between two nodes $e, e' \in V(G)$ if and only if their corresponding hyperedges overlap, i.e. $e \cap e' \neq \emptyset$. Clearly, a set S is a maximal matching in \mathcal{H} if and only if it is a MIS in G . Thus, it is left to find a MIS in G . This can be done by simulating Luby's algorithm [33]. Observe that the first and second steps need no communication. For the third step, every node in G (hyperedges in \mathcal{H}) only needs to know the highest priority among its neighbors. This can be done in $O(1)$ rounds by having each hyperedge (client in the server-client representation) in \mathcal{H} send its priority to all hypernodes (server) that it contains, then these hypernodes send the maximum priority that it receives to all hyperedges that contain it. So, we can implement the three steps of Luby's algorithm in $O(1)$ rounds.

6 Concluding Remarks and Open Problems

Our work shows that while some local symmetry breaking problems such as coloring and maximal matching can be solved in polylogarithmic rounds in both the LOCAL and CONGEST models, for many others such as MIS, hitting set, and maximal clique it remains a challenge to obtain polylogarithmic time algorithms in the CONGEST model. This dichotomy manifests in hypergraphs of higher dimension. Understanding this dichotomy can be helpful to make further progress in improving the bounds or showing lower bounds, especially in the CONGEST model. In particular, an important open question is whether we can show super-polylogarithmic lower bounds for MIS for hypergraphs of high dimension in the CONGEST model?

Our results also have implications to solving hypergraph problems in the classical PRAM model. Our CONGEST model algorithms can be translated into PRAM algorithms running in (essentially) the same number of rounds (up to polylogarithmic factors). In particular, improving over the $\tilde{O}(\Delta^{o(1)})$ round algorithm for MIS in the CONGEST model can point to better PRAM algorithms for MIS which has been eluding researchers till now. A major question is whether $O(\text{polylog } n)$ or even $O(\text{polylog } m)$ round algorithms are possible in the CONGEST model for MIS (as shown here, the answer is “yes” in the LOCAL model).

Another aspect of this work, which was one of our main motivations, is using hypergraph algorithms for solving problems in graphs efficiently. In particular, our hypergraph MIS algorithm leads to fast distributed algorithms for the BMDS and the MCDS problems. In particular, it will be interesting to see if one can give an algorithm for MCDS that essentially matches the lower bound of $\tilde{\Omega}(D + \sqrt{n})$ (when D is large).

References

1. Alon, N., Spencer, J.: The Probabilistic Method. Wiley Series in Discrete Mathematics and Optimization, Wiley (2008)
2. Åstrand, M., Suomela, J.: Fast distributed approximation algorithms for vertex cover and set cover in anonymous networks. In: SPAA 2010: Proceedings of the 22nd Annual ACM Symposium on Parallelism in Algorithms and Architectures, Thira, Santorini, Greece, June 13-15, 2010. pp. 294–302 (2010), <http://doi.acm.org/10.1145/1810479.1810533>
3. Avin, C., Lando, Y., Lotker, Z.: Radio cover time in hyper-graphs. Ad Hoc Networks 12, 278–290 (2014)
4. Azar, Y., Naor, J., Rom, R.: The competitiveness of on-line assignments. In: Frederickson, G.N. (ed.) SODA. pp. 203–210. ACM/SIAM (1992), <http://dblp.uni-trier.de/db/conf/soda/soda92.html#AzarNR92>
5. Balakrishnan, H., Barrett, C.L., Kumar, V.S.A., Marathe, M.V., Thite, S.: The distance-2 matching problem and its relationship to the mac-layer capacity of ad hoc wireless networks. IEEE Journal on Selected Areas in Communications 22(6), 1069–1079 (2004)
6. Barenboim, L., Elkin, M.: Distributed Graph Coloring: Fundamentals and Recent Developments. Synthesis Lectures on Distributed Computing Theory, Morgan & Claypool Publishers (2013)
7. Bartal, Y., Byers, J.W., Raz, D.: Global optimization using local information with applications to flow control. In: 38th Annual Symposium on Foundations of Computer Science, FOCS '97, Miami Beach, Florida, USA, October 19-22, 1997. pp. 303–312 (1997), <http://doi.ieeecomputersociety.org/10.1109/SFCS.1997.646119>
8. Beame, P., Luby, M.: Parallel search for maximal independence given minimal dependence. In: SODA. pp. 212–218 (1990)
9. Bercea, I., Goyal, N., Harris, D., Srinivasan, A.: On computing maximal independent sets of hypergraphs in parallel. In: SPAA (June, 2014)
10. Chen, Y.P., Liestman, A.L.: Approximating minimum size weakly-connected dominating sets for clustering mobile ad hoc networks. In: MobiHoc. pp. 165–172 (2002)

11. Chlamtac, I., Kutten, S.: Tree-based broadcasting in multihop radio networks. *IEEE Trans. Computers* 36(10), 1209–1223 (1987)
12. Dai, F., Wu, J.: An extended localized algorithm for connected dominating set formation in ad hoc wireless networks. *IEEE Trans. Parallel Distrib. Syst.* 15(10), 908–920 (Oct 2004), <http://dx.doi.org/10.1109/TPDS.2004.48>
13. Das, B., Bharghavan, V.: Routing in ad-hoc networks using minimum connected dominating sets. In: *ICC* (1). pp. 376–380 (1997)
14. Das Sarma, A., Holzer, S., Kor, L., Korman, A., Nanongkai, D., Pandurangan, G., Peleg, D., Wattenhofer, R.: Distributed verification and hardness of distributed approximation. *SIAM J. Comput.* 41(5), 1235–1265 (2012)
15. Dubhashi, D.P., Mei, A., Panconesi, A., Radhakrishnan, J., Srinivasan, A.: Fast distributed algorithms for (weakly) connected dominating sets and linear-size skeletons. *J. Comput. Syst. Sci.* 71(4), 467–479 (2005)
16. Ephremides, A., Truong, T.V.: Scheduling broadcasts in multihop radio networks. *IEEE Transactions on Communications* 38(4), 456–460 (1990)
17. Garay, J.A., Kutten, S., Peleg, D.: A sublinear time distributed algorithm for minimum-weight spanning trees. *SIAM J. Comput.* 27(1), 302–316 (1998)
18. Ghaffari, M.: Near-optimal distributed approximation of minimum-weight connected dominating set. In: *ICALP* (2014)
19. Harris, D.G., Morsy, E., Pandurangan, G., Robinson, P., Srinivasan, A.: Efficient computation of balanced structures. In: *ICALP* (2). pp. 581–593 (2013)
20. Karp, R.M., Ramachandran, V.: Parallel algorithms for shared-memory machines. In *Handbook of Theoretical Computer Science (Vol. A)* pp. 869–942 (1990)
21. Karp, R.M., Upfal, E., Wigderson, A.: The complexity of parallel search. *J. Comput. Syst. Sci.* 36(2), 225–253 (1988), announced at *STOC 1985* and *FOCS 1985*
22. Kelsen, P.: On the parallel complexity of computing a maximal independent set in a hypergraph. In: *STOC*. pp. 339–350 (1992)
23. Koufogiannakis, C., Young, N.E.: Distributed algorithms for covering, packing and maximum weighted matching. *Distributed Computing* 24(1), 45–63 (2011), <http://dx.doi.org/10.1007/s00446-011-0127-7>, announced at *PODC’09* and *DISC’09*
24. Kuhn, F.: The price of locality: exploring the complexity of distributed coordination primitives. Ph.D. thesis, *ETH Zurich* (2005)
25. Kuhn, F., Moscibroda, T., Wattenhofer, R.: The price of being near-sighted. In: *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006, Miami, Florida, USA, January 22-26, 2006*. pp. 980–989 (2006), <http://dl.acm.org/citation.cfm?id=1109557.1109666>
26. Kuhn, F., Moscibroda, T., Wattenhofer, R.: Local computation: Lower and upper bounds. *CoRR* abs/1011.5470 (2010)
27. Kutten, S., Pandurangan, G., Peleg, D., Robinson, P., Trehan, A.: On the complexity of universal leader election. In: *PODC*. pp. 100–109 (2013)
28. Kutten, S., Pandurangan, G., Peleg, D., Robinson, P., Trehan, A.: Sublinear bounds for randomized leader election. In: *ICDCN’13*. pp. 348–362 (2013)
29. Kutten, S., Peleg, D.: Fast distributed construction of small k -dominating sets and applications. *J. Algorithms* 28(1), 40–66 (1998)
30. Linial, N.: Locality in distributed graph algorithms. *SIAM J. Comput.* 21(1), 193–201 (1992)
31. Linial, N.: Dijkstra award talk (Jerusalem, 2013), http://www.cs.huji.ac.il/~nati/PAPERS/disc_2013.pdf
32. Linial, N., Saks, M.E.: Low diameter graph decompositions. *Combinatorica* 13(4), 441–454 (1993)
33. Luby, M.: A simple parallel algorithm for the maximal independent set problem. *SIAM J. Comput.* 15(4), 1036–1053 (1986), announced at *STOC 1985*
34. Luczak, T., Szymanska, E.: A parallel randomized algorithm for finding a maximal independent set in a linear hypergraph. *J. Algorithms* 25(2), 311–320 (1997)
35. Makino, K., Kameda, T.: Efficient generation of all regular non-dominated coteries. In: *PODC*. pp. 279–288 (2000)
36. Métivier, Y., Robson, J.M., Saheb-Djahromi, N., Zemmari, A.: An optimal bit complexity randomized distributed mis algorithm. *Distributed Computing* 23(5-6), 331–340 (2011)
37. Moscibroda, T.: Clustering. Book Chapter in *Algorithms for Sensor and Ad Hoc Networks* pp. 37–60 (2007)
38. Moscibroda, T., Wattenhofer, R.: Facility location: distributed approximation. In: *Proceedings of the Twenty-Fourth Annual ACM Symposium on Principles of Distributed Computing, PODC 2005, Las Vegas, NV, USA, July 17-20, 2005*. pp. 108–117 (2005)

39. Pandit, S., Pemmaraju, S.V.: Return of the primal-dual: distributed metric facilitylocation. In: Tirthapura, S., Alvisi, L. (eds.) Proceedings of the 28th Annual ACM Symposium on Principles of Distributed Computing, PODC 2009, Calgary, Alberta, Canada, August 10-12, 2009. pp. 180–189. ACM (2009)
40. Papadimitriou, C.H., Yannakakis, M.: Linear programming without the matrix. In: Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, May 16-18, 1993, San Diego, CA, USA. pp. 121–129 (1993), <http://doi.acm.org/10.1145/167088.167127>
41. Peleg, D.: Distributed Computing: A Locality-Sensitive Approach. SIAM, Philadelphia (2000)
42. Suomela, J.: Survey of local algorithms. ACM Comput. Surv. 45(2), 24 (2013), <http://doi.acm.org/10.1145/2431211.2431223>
43. Thurimella, R.: Sub-linear distributed algorithms for sparse certificates and biconnected components. J. Algorithms 23(1), 160–179 (1997)
44. Wasserman, S., Faust, K.: Social network analysis: Methods and Applications. Cambridge University Press (1994)