

Fast Distributed Backup Placement in Sparse and Dense Graphs

Leonid Barenboim* Gal Oren**

February 26, 2019

Abstract

We consider the Backup Placement problem in networks in the *CONGEST* distributed setting. Given a network graph $G = (V, E)$, the goal of each vertex $v \in V$ is selecting a neighbor, such that the maximum number of vertices in V that select the same vertex is minimized. The backup placement problem was introduced by Halldorsson, Kohler, Patt-Shamir, and Rawitz [12], who obtained an $O(\log n / \log \log n)$ approximation with randomized polylogarithmic time. Their algorithm remained the state-of-the-art for general graphs, as well as specific graph topologies. In this paper we obtain significantly improved algorithms for various graph topologies. Specifically, we show that $O(1)$ -approximation to optimal backup placement can be computed *deterministically* in $O(1)$ rounds in wireless networks, certain social networks, claw-free graphs, and more generally, in any graph with *neighborhood independence bounded by a constant*. At the other end, we consider sparse graphs, such as trees, forests, planar graphs and graphs of constant arboricity, and obtain constant approximation to optimal backup placement in $O(\log n)$ deterministic rounds.

Clearly, our constant-time algorithms for graphs with constant neighborhood independence are asymptotically optimal. Moreover, we show that our algorithms for sparse graphs are not far from optimal as well, by proving several lower bounds. Specifically, optimal backup placement of unoriented trees requires $\Omega(\log n)$ time, and approximate backup placement with a polylogarithmic approximation factor requires $\Omega(\sqrt{\log n / \log \log n})$ time. Our results extend the knowledge regarding the question of "what can be computed locally?" [20], and reveal surprising gaps between complexities of distributed symmetry breaking problems.

1 Introduction

We consider the Backup Placement problem in networks. The problem is defined as follows. Given a network graph $G = (V, E)$, the goal of each vertex $v \in V$ is selecting a neighbor, such that the maximum number of vertices in V that select the same vertex is minimized. The number of neighbors that select a certain vertex $v \in V$ is called the *load* on v . The load of a backup placement solution for $G = (V, E)$ is the maximum load on a vertex $v \in V$. For a positive integer $t \geq 1$, *t-backup placement* for G is a solution in which each vertex $v \in V$ selects a neighbor, such that the maximum load in V is at most t times the load of an optimal solution for G . To illustrate this problem, consider the following two simple examples. First, suppose that G is a cycle graph. (See Figure 1). In this case, the optimal solution to the backup placement problem is when each node selects its succeeding neighbor to be its backup node (in blue in Figure 1). By that, the maximal load in this graph is of only one unit. Next, suppose that G is a rooted tree as in Figure 1 (right). The optimal solution in this case forces all the nodes to choose their parents

*Open University of Israel. E-mail: leonidb@openu.ac.il

** Ben-Gurion University of the Negev, Nuclear Research Center - Negev. E-mail: orenw@post.bgu.ac.il

This work was supported by the Lynn and William Frankel Center for Computer Science, the Open University of Israel's Research Fund, and ISF grant 724/15.

(besides the root) to be their backup nodes. By that, the maximal load is $\Delta = \Delta(G)$, the maximum degree in the graph. Similarly, in any regular tree, no solution with load smaller than $\Delta - 1$ exists. On the other hand, some non-regular trees admit very good solutions, of load as small as 2 units.




Figure 1: Optimal backup placement (in blue) in C_6 cycle graph (left) and a rooted tree (right)

Finding a backup placement in a network while minimizing the load on the network vertices is a very important goal [12, 13, 23]. First of all, it allows each vertex to perform a backup to a neighboring node, rather than a more distant destination, and thus improves network performance. In addition, nodes' memories are used to the minimum extent for the purpose of backups, which makes it possible to maximize the memory available for other purposes of the vertices. Moreover, if a certain node fails, the number of backups that become unavailable is minimized. For these and other reasons, the backup placement problem is considered as a central distributed problem by Halldorsson, Kohler, Patt-Shamir and Rawitz [12], who initiated the study of this problem in the distributed setting in 2015.

We would like to point out an additional important aspect. This problem is also a central symmetry-breaking problem, and is interesting to compare to such important problems as Maximal Matching and Maximal Independent Set. In a certain sense the backup placement problem generalizes Maximal Matching. Note that a perfect Maximal Matching is also a perfect backup placement. Indeed, all vertices can perform backups to their pairs in the matching. When a perfect matching is not possible, then an optimal backup placement assigns up to c selecting neighbors to each vertex, for a parameter $c > 0$. Therefore, this generalizes Maximal Matching, in which $c = 1$, but certain vertices may be selected by no vertices at all. Since Maximal Matching and other symmetry-breaking problems cannot be solved in a constant number of rounds [19, 17], even on paths and cycles, the question of whether (approximate) backup placement can be solved in various networks within this number of rounds is of great interest.

In this paper we answer this question in the affirmative. Specifically, we show that $O(1)$ -backup-placement can be computed in $O(1)$ rounds in wireless networks, certain social networks, claw-free graphs, and more generally, in any graph with *neighborhood independence bounded by a constant*. (That is, the number of independent vertices in any immediate neighborhood is bounded by a constant.) At the other end, we consider sparse graphs, such as trees, forests, planar graphs and graphs of constant arboricity, and obtain constant approximation to backup placement in $O(\log n)$ rounds. These results suggest that the problem is harder on sparse graphs than on dense ones. We prove this formally, by providing a lower bound for computing (approximate) backup placement on unoriented trees of $\Omega(\sqrt{\log n / \log \log n})$ rounds. Our lower bound holds for approximation ratios up to polylogarithmic in n . This is in contrast to oriented trees, where a constant approximation to backup placement is achieved within $O(1)$ rounds, and various dense graphs in which it is possible to compute a sparse infrastructure based on oriented trees within a constant number of rounds. All our algorithms work in the *CONGEST* model, where only $O(\log n)$ bits are sent per edge per round. See Table 1 below for a comparison of our results with various symmetry-breaking results in various graph families. This reveals interesting separations between complexities of several problems and several graph topologies.

Topology Problem	Very dense: Graphs of bounded growth	Dense: Graphs of constant neighborhood independence	Sparse: Graphs of constant arboricity	General Graphs
Maximal Independent Set	$\Theta(\log^* n)$ [19, 27]	$2^{O(\sqrt{\log n})}$ [25]	$O(\frac{\log n}{\log \log n})$ [5]	$2^{O(\sqrt{\log n})}$ [25] $\Omega(\sqrt{\frac{\log n}{\log \log n}})$ [17]
Maximal Matching	$\Theta(\log^* n)$ [19, 27]	$O(\log n)$ [9]	$O(\frac{\log n}{\log \log n})$ [5] $\Omega(\sqrt{\frac{\log n}{\log \log n}})$ [17]	$O(\log^3 n)$ [11] $\Omega(\sqrt{\frac{\log n}{\log \log n}})$ [17]
$O(\frac{\log n}{\log \log n})$ - Backup Placement	$O(\frac{\log^6 n}{\log^4 \log n})$ (randomized) [12]			
$O(1)$ -Backup Placement	$O(1)$ This paper	$O(1)$ This paper	$O(\log n)$ $\Omega(\sqrt{\frac{\log n}{\log \log n}})$ This paper	

Table 1: Comparison of running times of state-of-the-art symmetry-breaking and backup placement deterministic results. (With the exception of $O(\log n / \log \log n)$ -backup-placement, for which the only non-trivial previous result is randomized.)

Our results significantly improve upon previously known results. The best previous results are those for general graphs, due to Halldorsson et al. [13]. These are randomized algorithms with polylogarithmic running times and approximation ratio $O(\frac{\log n}{\log \log n})$. In the current paper, however, all our algorithms are *deterministic*. Our running time is $O(1)$ for graphs with constant neighborhood independence (i.e., dense graphs) and $O(\log n)$ for graphs with constant arboricity (i.e., sparse graphs). Our approximation ratio is $O(1)$ for graphs with constant neighborhood independence, as well as for graphs with constant arboricity. Thus we significantly outperform the previous results, both in terms of running time and approximation ratio, for various dense and sparse graph families we mentioned above.

Our results also extend the knowledge with respect to the question of "what can be computed locally?". This fundamental question, asking what can be computed in $O(1)$ rounds, was raised in the seminal paper of Naor and Stockmeyer [20]. The importance of algorithms with a constant number of rounds follows from the fact that they are not affected by network size, nor by other parameters, such as maximum degree. Consequently, such algorithms are extremely scalable. Since the publication of Naor and Stockmeyer, several distributed algorithm with a constant number of rounds were devised, including weak-coloring [16, 20], $O(\Delta)$ -forests-decomposition [24], minimum coloring approximation [4, 7], certain network decompositions [4, 7], approximate minimum dominating set [18, 7], and approximate minimum spanner [7]. However, several fundamental symmetry-breaking problems cannot be solved in $O(1)$ -rounds, including Maximal Matching and Maximal Independent Set. A lower bound of $\Omega(\log^* n)$ holds for these problems on paths, cycles, graphs of bounded growth, graphs of bounded independence and more [19]. Thus, we find it intriguing that a related problem of $O(1)$ -backup-placement can be solved in these and other graph topologies in a constant number of rounds.

1.1 Our Upper-Bound Techniques

Previous techniques for distributed backup placement and related problems were based on selfish improvement policies, generalizations of Maximal Matching, and Maximal Packing [13, 14]. In contrast, in this paper we employ an entirely different approach. Specifically, our main technical tool is a set of

spanning trees or forests that cover all vertices (but not necessarily all edges) of the input graph. In most of our constructions, the number of edges is in the order of the number of vertices. Interestingly, such spanning forests turn out to be very useful not only for graphs of small arboricity, but also for dense graphs in which the arboricity is very large, and the number of edges in the input graphs can be as large as $\Theta(n^2)$. Thus our approach is more general than the popular approach of [5] that places all edges in forests, and works especially well on low-arboricity graphs. The new structures we devise in the current paper are used to efficiently coordinate backups, and reduce redundant load. For example, each parent in a tree can select one of its children for a backup. Since each vertex has at most one parent, this causes a load of only 1 unit. But leaves have no children, and thus constitute a challenging instance. Selecting their parents for backups can cause excessive load. Thus, our algorithms select siblings for backups, and coordinate them in such a way that no sibling is selected by too many vertices.

Another challenging instance is bipartite graphs. Even if one side of the partition has a constant degree, but on the other side the degree is unbounded, it is challenging to avoid the selection of the same vertex by many neighbors. We overcome this difficulty by gradually selecting vertices of sufficiently low degrees, and temporarily removing them and their neighbors from the graph. This produces additional vertices of small degrees, which allows us to repeat the process, until all vertices succeed to make a selection for a backup. Since in each stage vertices select neighbors of sufficiently low degrees, the resulting load is bounded. Since each time a constant fraction of vertices is removed, the algorithm terminates within $O(\log n)$ rounds. This idea can be generalized to work for graphs with constant arboricity. These graphs can be partitioned into few forests. All non-leaf vertices can be easily handled, by selecting their children for backups. Since each vertex has a bounded number of parents, the load is bounded as well. It remains to handle the leaves. But the set of leaves and the set of non-leaves constitute a bipartite graph on which one side has a bounded degree. This can be handled as explained above.

1.2 Our Lower-Bound Techniques

We consider two trees that are identical, except the very last layer of leaves. While one of the trees is d -regular, for a parameter $d > 0$, in the other one each leaf is a single child, rather than having $d - 2$ siblings, as in the d -regular tree. We prove that an optimal backup placement of the former tree is such that the root is selected for backups by all its children. On the other hand, for the latter tree, in the optimal backup placement the children of the root must select their own children, rather than the root. Since the trees are indistinguishable (by the roots) within fewer rounds than the height of the trees, any deterministic algorithm with such a number of rounds will fail at least on one of them.

The above argument works for an optimal backup placement, but not for an approximate one. For the latter case, we provide a more sophisticated proof. We consider a graph in which all neighborhoods of sufficiently low radius look like one of the above-mentioned trees, possibly with single-child leaves in the last layer. Consequently, a backup placement algorithm executed on such a graph must assign a small number of backups to each vertex. (Each vertex can be seen as a root of a certain tree, in which most of its children do not select it.) But this turns out to be impossible, unless the trees are sufficiently high and the algorithm is invoked for sufficiently many rounds.

1.3 The Importance of Backup Placement in Wireless Networks

Internet of Things (IoT) devices involve very differentiated requirements in terms of communication fashion, computation speed, memory limit, transmission rate and data storage capacity [2]. A massive volume of heterogeneous multi-attributed data is created by those diverse devices periodically, and needs to be sent to other locations. As those devices may operate in a nonstop fashion, and may be placed in geographically-diverse locations, they introduce a challenge in terms of communication overhead and

storage capacity, since transfer and storage of data are resources-consuming activities within an IoT data management system [1]. Hence, these factors are crucial in data management techniques for IoT.

The challenging situation can be described as follows. As data is aggregated in a device or in the concentration storage points within the IoT, the amount of vacant storage becomes very limited. Thus the data usually needs to be sent further up the system, either to a storage relay node, or for storage end-node [10]. As wireless broadband communications is most likely to be used all along this process in the IoT world, there are also strict limitations to the amount of communications each power-limited device can perform. Moreover, since crashes of nodes, failures of communication links, and missing data are unavoidable in wireless networks, fault-tolerance becomes a key-issue. Among the causes of these constant failures are environmental factors, damaged communications links, data collision and overloaded nodes [15]. Thus, the problem of fast data transfer and storage is a major topic of interest in IoT, as it addresses factors that are the most challenging in terms of hardware and power resources limits [26]. A specific topic of interest is backup placement in wireless sensor networks, which was considered in [23].

2 Backup Placement in Trees and Spanning Forests of Dense Graphs

We begin with devising a procedure for computing $O(1)$ -backup placement in trees. We assume that each vertex knows its parent in the tree. Later we will show how to drop this assumption. The procedure receives a tree $T = (V, E)$ as input, and proceeds as follows. Each vertex that is not a leaf selects an arbitrary child for the backup placement. Each leaf selects its parent for the backup placement. All these selections are performed in parallel within a single round. This completes the description of the algorithm. Its pseudocode is provided in Algorithm 1. Its action is illustrated in Figure 2. Next, we prove correctness and analyze running time.

Algorithm 1 Tree Backup Placement

```

1: procedure TREE-BP(TREE  $T$ )
2:   foreach node  $v \in T$  in parallel do:
3:     if  $v$  is not a leaf then
4:        $v.BP \leftarrow \text{Arbitrary}(v.children)$ .
5:     else
6:        $v.BP \leftarrow v.parent$ 

```

Lemma 2.1. *The algorithm computes an $O(1)$ -backup placement of T .*

Proof. Let S be an optimal solution, and S' is our solution for T . Each leaf must have selected its parent because this is its only neighbor. Let $f' : V \rightarrow N$ denote the number of backups in each vertex $v \in V$ in our solution. Let $f : V \rightarrow N$ denote the number of backups in each vertex $v \in V$ in the optimal solution. Next, we show that $f'(v) \leq f(v) + 1$, for each $v \in V$. It holds that $f(v)$ is at least the number of the leaf children of v . On the other hand, $f'(v)$ is the number of the leaf children of v plus at most 1, since v 's non-leaf children do not select v , but v 's parent may select v . \square

The above algorithm can be directly extended to a forest of trees in which vertices know their parents, by executing the algorithm in parallel on all trees. More interestingly, we can extend this technique to graphs that are very dense, as opposed to trees and forests. As a first step, we need to compute a spanning forest of any input graph G , such that the trees of the forest are *vertex-disjoint*. This can be done by extending the algorithm of Panconesi and Rizzi for Δ -forests-decomposition [24]. (Where Δ is the maximum degree in the input graph.) The latter algorithm produces edge-disjoint trees, that are not necessarily vertex-disjoint. Specifically, it partitions the edges of any graph into Δ acyclic subgraph, such

that each vertex has at most one parent in each subgraph. In other words, each of the Δ subgraphs is a forest in which vertices know their parents. The algorithm of Panconesi and Rizzi requires just a single round on any graph. An important property of the algorithm of [24] we employ is that for any pair of neighbors in the input graph G , the higher-ID neighbor is the parent of the lower-ID neighbor in some forest of the output.

Consider a subgraph $G = (V, E'), E' \subseteq E$, obtained as follows. First, invoke the Δ -forest-decomposition algorithm of [24], and initialize $E' = \emptyset$. Next, each vertex $v \in V$ that has at least one parent in at least one forest, selects an arbitrary parent u , and adds (u, v) to E' . Note that the number of edges added this way is bounded by $|V| = n$, and the resulting graph G' is acyclic. Indeed, a cycle would imply that either a vertex has selected two parents, or the cycle is consistently oriented, i.e., a vertex is an ancestor of itself. But each vertex selects at most one parent, and each parent has a higher ID than that of its child, thus a vertex cannot be its own ancestor. Thus there are no cycles in G' . Note that the forest G' is not necessarily a spanning forest in this stage. This is because some vertices that are roots in all forests of Panconesi-Rizzi's execution may belong to none of the edges in E' . This can be corrected as follows. Each such root vertex v , that has not been selected by any of its children in the previous stage, selects an arbitrary child u and adds (u, v) to E' . Note that such root vertices form an independent set, i.e., they cannot be neighbors in G . (Otherwise, one of them would be the parent of the other in some forest of Panconesi-Rizzi's execution.) Consequently, all edges added in this stage do not close cycles, since each such edge has one endpoint of degree 1 in the resulting graph $G' = (V, E')$. Moreover, in this stage each vertex of V belongs to at least one edge of E' , thus G' is a spanning forests. In addition, each vertex belongs to exactly one tree, thus the trees are vertex-disjoint.

Since each stage can be done in parallel by all vertices within a single round, the entire spanning forest construction requires just 3 rounds. Specifically, Δ -forest-decomposition stage, followed by parent selection stage, followed by root correction stage. We summarize this discussion below.

Corollary 2.1.1. *A spanning forest of any input graph G can be constructed within $O(1)$ rounds.*




Figure 2: Tree Backup Placement




Figure 3: General Backup Placement

Next, we devise an algorithm that computes a backup placement for any graph G in which a spanning forest has been computed. For simplicity, we describe our algorithm for a tree T in the forest. If there are more trees, the algorithm should be executed on all trees of the forest in parallel. (However, each vertex is aware only of the execution in the tree T it belongs to.) We compute backup placement for T as follows. Each non-leaf vertex in T selects one of its children arbitrarily. Each leaf $v \in T$ selects a sibling w of v in T with the property that $(v, w) \in E(G), ID(w) < ID(v)$, and there is no other sibling z of v such that $(v, z) \in E(G)$, and $ID(w) < ID(z) < ID(v)$. If there is no such sibling w , then v selects its parent. This completes the description of the algorithm. Note that if there is such a sibling w , it is unique. The pseudocode is provided in Algorithm 2. Its action is illustrated in Figure 3. We analyze the algorithm below.

Algorithm 2 General Backup Placement

```
1: procedure GENERAL-BP(GRAPH  $G = (V, E)$ , TREE  $T$ )
2:   foreach node  $v \in T$  in parallel do:
3:     if  $v$  is not a leaf then
4:        $v.BP \leftarrow \text{Arbitrary}(v.children)$ .
5:     else if  $\exists w$  sibling of  $v$ , such that  $(v, w) \in E$  and  $(ID(w) < ID(v)) \wedge \neg \exists z$  sibling of  $v$ , such that
       $(v, z) \in E$  and  $(ID(w) < ID(z) < ID(v))$  then
6:        $v.BP \leftarrow w$ .
7:     else
8:        $v.BP \leftarrow v.Parent$ .
```

First, we prove that for graphs with bounded neighborhood independence, the algorithm provides a good backup placement.

Lemma 2.2. *Let G be a graph with neighborhood independence at most c , for $c > 0$, and T be a tree in a spanning forest of G . Then each vertex of G is selected by at most c of its children in T in Algorithm 2.*

Proof. Assume for contradiction that $v \in V$ of G has been selected by $c + 1$ of its children in T . Let V_{c+1} denote the set of those children. Since G is a graph with neighborhood independence at most c , it implies that there must be at least two vertices $u, w \in V_{c+1}$ which are connected by an edge in G . Assume without loss of generality that $ID(u) < ID(w)$. In this case, according to our algorithm, w must have selected either u or some other sibling, but not its parent v . This is a contradiction. Hence, v has been selected by at most c of its children in T . \square

Lemma 2.3. *Let G be a graph with neighborhood independence at most c , for $c > 0$, and T be a tree in a spanning forest of G . Then each vertex of G is selected by at most c of its siblings in T in Algorithm 2.*

Proof. If more than c siblings select v , two of them are connected by an edge in G , since neighborhood independence is bounded by c . Denote these vertices by u, w . (Note that $(u, v) \in E(G), (u, w) \in E(G), (v, w) \in E(G)$). Assume without loss of generality that $ID(u) < ID(w)$. Since u has selected v , it means that $ID(v) < ID(u)$. Therefore, $ID(v) < ID(u) < ID(w)$, and w could not have selected v , since it should have selected a sibling with the closest ID to its own. This is a contradiction. \square

Since each vertex can be selected only by its children, by its siblings that are connected to it in G , and by its parent, and cannot be selected by vertices outside T , we obtain the next corollary.

Corollary 2.3.1. *Let G be a graph with neighborhood independence at most c , for $c > 0$, and T be a tree in a spanning forest G . Then each vertex of T is selected by at most $2c + 1$ of its neighbors in G in Algorithm 2.*

In the next lemma we analyze the running time of the algorithm.

Lemma 2.4. *The Running time of the algorithm is $O(1)$.*

Proof. The algorithm consists of two parts, each of which has a constant running time: (1) distributed selection of a backup placement node, in which each non-leaf selects one of its children, and (2) a distributed selection of a backup placement node, in which each leaf selects one of its neighbors. Therefore, the total running time is $O(1)$. \square

By starting with a computation of a spanning forest of the input graph G , and then invoking our algorithm on all trees of the spanning forest of G in parallel, we obtain the following result.

Corollary 2.4.1. *For an input graph G with neighborhood independence c , we compute backup placement with load at most $2c + 1$ in $O(1)$ rounds.*

Next, we analyze the message complexity of our algorithm.

Lemma 2.5. *The algorithm can be implemented with messages of size $O(\log n)$ per link per round.*

Proof. Initially, all vertices have unique IDs of size $O(\log n)$ bits each. Computing a spanning forest requires each vertex to know who of its neighbors have greater IDs than its own, and whether some neighbor has selected the vertex. This requires $O(\log n)$ bits per edge. Once the spanning forest has formed, each vertex can send its own ID and its parent ID to all of its neighbors. By that, each vertex learns if a neighbor is also one of its siblings in a tree T . (They are siblings if they have the same parent.) Also, using this technique each vertex can identify the sibling of a smaller ID that is closest to its own.

This information is sufficient to select a neighbor for a backup placement by our algorithm. Sending two IDs over an edge requires $O(\log n)$ bits. Thus, the algorithm can be implemented with messages of size $O(\log n)$ per link per round. \square

We note that for a variety of real-life network topologies, the neighborhood independence c is bounded by a small constant. This includes Unit Disk Graphs that model wireless networks with the same transmission range for all nodes, and Bounded Disk Graphs in which transmission ranges may differ [3]. We elaborate on this in Appendix A.1. Our algorithm is applicable not only to wireless networks in which the number of independent nodes is bounded in each r -hop neighborhood, for a constant r (as UDG and BDG), but also to more general networks. Specifically, it is sufficient that the number of neighbors is bounded only in the 1-hop neighborhood. A notable example is line graphs, in which the number of independent nodes in 1-hop neighborhoods is bounded by 2, but it is unbounded in r -hop neighborhoods for $r \geq 2$. The proof of this is provided in Appendix A.1, Lemma A.3. Line graphs are an example of a family with *constant diversity* [8]. (Their diversity is bounded by 2.) Graphs with diversity c are those in which all vertices belong to at most c cliques. Hence, each vertex has at most c independent neighbors. Thus, our algorithms are applicable also to this more general family of graphs with constant diversity.

3 Backup Placement in Bipartite Graphs

In this section we devise a backup placement algorithm for bipartite graphs $G = (U, V, E)$, in which the maximum degree of vertices in U is bounded by a parameter a , and the maximum degree of vertices in V is unbounded. Such graphs are motivated by client-server systems in which each client may connect to a bounded number of servers, but each server has an unrestricted number of clients. For now, we assume that all vertices know the load value t of the optimum backup placement, i.e., the maximum number of selected vertices in an optimum solution. Later, we will drop this assumption. The goal of our algorithm is obtaining a maximum load of $O(at)$. To this end, each vertex of V may select for a backup an arbitrary neighbor in U . Since each vertex in U has at most a neighbors, the maximum load on vertices of U is going to be at most a as well. It remains to define an algorithm for the vertices of U . The algorithm is the following. It proceeds in phases. In each phase, each vertex in U that has a neighbor in V with degree at most $2at$, selects such a vertex. (The selection is arbitrary, if more than one such a neighbor exists.) Next, we remove from V all vertices of degree at most $2at$, and also remove from U all neighbors of vertices that have been removed from V . (Note that all these vertices of U have already made selections in V . Note also that once some vertices are removed, certain vertex degrees become smaller.) Next, we proceed to the next phase, that is performed in the same way. Specifically, vertices in U with neighbors of degree at most $2at$ select such neighbors, one neighbor each. Then, all vertices in V of degree bounded by $2at$, and their neighbors, are removed. We repeat this until no vertices remain in G .




Figure 4: Round 1.




Figure 5: Round 2.




Figure 6: Round 3.

Figure 7: The Bipartite Graph Distributed Backup Placement Execution: An example of a bipartite graph, with $a = 2$ and an estimation $t' = 1$ for t , and we allow up to $2 \cdot at' = 4$ backups per $v \in V$. Black arrows represent edges of the input graph, while blue ones represent selections of backup placement.

The pseudocode of the algorithm is provided in Algorithm 3. Its action is illustrated in Figure 7. Next we prove its correctness and analyze running time.

Algorithm 3 The Bipartite Graph Distributed Backup Placement Algorithm

- 1: for each $v \in V$, $v.allow_backups = True$.
 - 2: **procedure** BIPARTITE-BP($G = (U, V, E)$, a, t)
 - 3: In each round each vertex $v \in V$ in parallel does:
 - 4: **if** $deg(v) < 2at$ and $deg(v) \neq 0$ **then**
 - 5: v sends a message "allow backup" to all of its neighbors.
 - 6: Each vertex in U that receives an "allow backup" message, selects for a backup an arbitrary neighbor that sent this message.
 - 7: $v.allow_backups \leftarrow False$
 - 8: $V = V \setminus v$, $U = U \setminus v.neighbors$
-

Lemma 3.1. *In each phase, at most $1/2$ of remaining vertices in V have more than $2at$ remaining neighbors each.*

Proof. We denote the number of remaining vertices in V in the beginning of each phase i by n_i . We denote this set of remaining vertices of V (respectively, U) by V_i (resp., U_i). Denote also the number of remaining neighbors of a vertex v in round i by $deg_i(v)$. We prove our claim by contradiction. Assume that in a certain phase i the number of vertices $v \in V_i$ with $deg_i(v) > 2at$ is more than $\frac{1}{2}n_i$. Thus, the degrees sum of vertices in V_i in this stage is at least $\sum deg_i(v), v \in V_i > \frac{n_i}{2} \cdot 2at = n_i \cdot at$. Recall that t is a parameter that quantifies the optimal solution. It is important to emphasize that the meaning of t as the optimal backup is the following. For each $u \in U$, and in particular $u \in U_i$, there exists a selection of some $v \in V$ for backup, such that after those selections have been made, the degree of each $v \in V$ in the subgraph induced by the selected edges is $deg(v) \leq t$. Note also that all neighbors of vertices of U_i are in V_i . Indeed, any vertex in U with a neighbor outside of V_i has been removed in an earlier stage. Hence, in

the optimal solution, all selection of vertices of U_i are vertices of V_i , and the maximum load is t . Thus, the size $|U_i|$ of U_i , in the optimal solution, is $|U_i| \leq n_i \cdot t$. (*)

But each vertex in U_i has degree at most a , and $\sum \deg_i(v), v \in V_i$ is greater than $n_i \cdot at$, by the above assumption. Since in phase i all remaining neighbors of vertices in V_i are in U_i , it follows that the number of vertices in U_i is greater than $n_i \cdot at/a = n_i \cdot t$. This is a contradiction to (*). \square

Lemma 3.2. *The algorithm terminates within $O(\log n)$ rounds.*

Proof. As proven in the lemma above, at any phase i the number of vertices $v \in V$ with $\deg_i(v) > 2at$ is at most $\frac{1}{2}$ of remaining vertices in that phase. Thus, in the first round of the algorithm, the size of the subset $\{v \in V \mid \deg_1(v) > 2at\}$ is at most $\frac{1}{2}n_1 = \frac{1}{2}|V|$. Similarly, in the next round, the size of the subset $\{v \in V \mid \deg_2(v) > 2at\}$ is at most $\frac{1}{2}n_2 = \frac{1}{4}n_1$, and so forth. After R rounds, for a positive integer R , the number of vertices in V with degree greater than $2at$ is bounded by $n_1/2^R$. Since the algorithm completes all backup placements in the round when no vertices in V with degree greater than $2at$ remain, it terminates within $O(\log n_1) = O(\log n)$ rounds. \square

Lemma 3.3. *Each vertex in V is selected by at most $2at$ vertices in U .*

Proof. In each round vertices in U select neighbors in V only if these neighbors have degree at most $2at$. Moreover, each vertex in V is selected in the same round by all neighbors that choose it during the algorithm. (This is because in the end of the round when a vertex is selected for the first time, it is removed from V .) But a vertex cannot be selected by more than $2at$ neighbors of U , since in that round the vertex degree is bounded by $2at$. \square

Next, we extend our algorithm to the scenario that the optimum value t is not known to the vertices. (But vertices know a . In Section 4 we elaborate on the reason of this assumption regarding knowledge of a .) In this case when t is not known, we invoke our algorithm several times, starting with an estimation $t' = 1$ for t , and doubling t' after each invocation. As long as $t' < t$, an invocation may remove only some of the vertices, in contrast to the case $t' \geq t$, in which all vertices are removed, as shown in the proof of Lemma 3.1. Thus, each invocation is performed on the set of remaining vertices from the previous invocation. In the first time when t' becomes at least t , all remaining vertices are removed, and we are done. Note that in this stage it holds that $t' < 2t$. In each invocation a vertex is selected by at most $2at'$ neighbors, and so the maximum number of selections of the same vertex in all invocations is $O(at)$. The number of rounds required for this computation is $O(\log n \log t)$. We summarize this in the next theorem.

Theorem 3.4. *Suppose that the optimum backup placement value t is not known to the vertices before execution. In this case we compute a backup placement with a load of $O(at)$ within $O(\log n \log t) = O(\log^2 n)$ rounds.*

We note that the $O(\log t)$ executions can be performed in parallel. To this end, we assign arrays A_u of size $\log n$ to each vertex $u \in U$, and execute the algorithm independently for $\log n$ times in parallel. (We do not know t , but we know that $t < n$.) Each such execution requires $O(\log n)$ time, but their parallel invocations result in $O(\log n)$ rounds overall, rather than $O(\log^2 n)$. Now, we have $\log n$ results for each vertex in U . We note that $A_u[i]$ may have not reached an answer, if $2^i < t$. Otherwise, $A_u[i]$ must contain a selection for u , by the correctness of the original algorithm, that works with any value that is at least the optimum t . For each vertex $u \in U$, we take the result for the smallest i for which a backup placement $A_u[i]$ has been found. We know that $2^i < 2t$, where t is the optimum.

Next, we analyze how many vertices in U may select the same vertex in V in the worst case. Consider such a vertex $v \in V$, and denote u_1, u_2, \dots, u_q the set of vertices that selected v . Our goal is to bound q . Recall that each selection of a vertex $u_j, j \in [q]$ is stored in $A_{u_j}[i]$, where $i < \log(2t)$. Moreover, there

are at most $2a$ vertices that selected v in $\{A_{u_1}[1], A_{u_2}[1], \dots, A_{u_q}[1]\}$, at most $4a$ vertices that selected v in $\{A_{u_1}[2], A_{u_2}[2], \dots, A_{u_q}[2]\}$, at most $8a$ vertices that selected v in $\{A_{u_1}[3], A_{u_2}[3], \dots, A_{u_q}[3]\}$, etc. In general, for $i > 0$, there at most $2^i \cdot a$ vertices that selected v in $\{A_{u_1}[i], A_{u_2}[i], \dots, A_{u_q}[i]\}$. Thus, the overall number of vertices that selected v in all these arrays cells with indices $1, 2, \dots, i$ is bounded by $2a \cdot \sum_{j=1}^i 2^j = 2a \cdot (2^{i+1} - 2) < 8a \cdot t$.

Corollary 3.4.1. *The number of selections of the same vertex in V by vertices in U is bounded by $8at$.*

The last result gives rise to the next theorem that summarizes this section.

Theorem 3.5. *Given a bipartite graph $G = (U, V, E)$ with maximum degree a in U , one can compute a backup placement with a load of $O(at)$ in $O(\log n)$ rounds, even if the optimum load t is not known to the vertices.*

4 Backup Placement in Graphs of Bounded Arboricity

Following our method for backup placement in bipartite graphs, we can also prove $O(\log n)$ time complexity for backup placement in graphs of constant arboricity, and in particular, in planar graphs. We still assume that all vertices know the load value t of the optimum backup placement, i.e., the maximum number of selected vertices in an optimum solution. The goal of our algorithm is obtaining a maximum load of $O(a \cdot t)$, where a is the arboricity of the input graph. We note that our algorithm can be extended to the scenario where the optimum load t is not known, similarly to Section 3. Nevertheless, vertices still need to know a . We note, however, that this information is often available to the vertices. For example, in planar graphs, vertices know that the arboricity a is bounded by 3 [21, 22].

In order to fulfill our goal, we use the Procedure Partition algorithm [5, 6], which receives as input a graph G with arboricity a , and partitions it into $\ell = O(\log n)$ sets, H_1, H_2, \dots, H_ℓ , such that the number of neighbors of a vertex $v \in H_i$, $i \in [\ell]$, in the set $H_i \cup H_{i+1} \cup \dots \cup H_\ell$ is at most $(2 + \epsilon)a$, for an arbitrarily small constant $\epsilon > 0$. The time complexity of this algorithm is $O(\log n)$. For the sake of simplicity, we set $\epsilon = 1$. Consequently, the algorithm partitions the vertices of G into $\lceil \frac{2}{\epsilon} \log n \rceil = 2 \log n$ sets, such that each vertex v has at most $3a$ neighbors in the union of sets with greater or equal index to the H -set of v .

We note that the sets H_1, H_2, \dots, H_ℓ obtained using Procedure Partition have the following property. For each vertex in H_i with $i > 1$, there must be a neighbor in a set with a smaller H -index. For example, for a vertex in H_2 there must be a neighbor in H_1 , and for a vertex in H_3 there must be a neighbor in H_2 or H_1 , and so forth. The reason, according to Procedure Partition [5], is that all vertices that did not join H_1 have an initial degree that is greater than $3a$, and once they join an H -set of a greater index, their degree becomes at most $3a$. This means that each of them has a neighbor that has been removed earlier, i.e., moved to an H -set with a smaller index. For the backup placement problem, this implies that it is possible to create a backup placement from any $H_i, i > 1$, to a vertex in $H_j, j < i$. To this end, all vertices in $H_i, i > 1$, select in parallel neighbors in sets with smaller indices than their own, one neighbor each. Since the number of neighbors in sets of greater or equal index is bounded by $3a$, each vertex is selected by at most $3a$ neighbors in this stage. In the case of planar graphs, $a \leq 3$ [21, 22]. Therefore, it is possible to perform backup placements for all $H_2, \dots, H_\ell, \ell = O(\log n)$ sets, such that any vertex performs a backup placement to some neighbor, and since the degree with respect to greater or equal indices is at most $3a = 9$, there is no vertex with more than that number of backup placements.

It remains to select placements for vertices in H_1 . For each vertex in H_1 that has a neighbor in H_1 , we select such a neighbor arbitrarily. Since the maximum degree in H_1 is bounded by $3a$, each vertex can be selected by at most $3a$ neighbors. However, some vertices in H_1 may have no neighbors in this set, but only neighbors in sets of greater indices. Next, we describe a solution for these vertices, and thus, complete the description of placements selection for all vertices in H_1, H_2, \dots, H_ℓ . We denote

$U = \{\text{all vertices in } H_1 \text{ that have no neighbors in } H_1\}$, $V = H_2 \cup H_3 \cup \dots \cup H_\ell$, $E' = \{(u, v), u \in U, v \in V, (u, v) \in E\}$, and execute our bipartite algorithm from Section 3 on $G' = (U, V, E')$. Consequently, each vertex in U finds a backup placement in $V = H_2 \cup H_3 \cup \dots \cup H_\ell$. In addition, according to Procedure Partition, each vertex $v \in V$ has a neighbor in a smaller-index H -set, and thus each $v \in V$ can perform backup placement to such a neighbor in a smaller-index set. Finally, the vertices that are neither in U nor in V belong to H_1 , and have neighbors in H_1 . They select arbitrary neighbors in H_1 for backup placements. This completes the description of placements selection for all vertices in the input graph. The pseudocode of the algorithm is provided in Algorithm 4.

Algorithm 4 The Bounded Arboricity Graph Distributed Backup Placement Algorithm

- 1: **procedure** BOUNDED-ARBORICITY-BP($G = (V, E), a$)
 - 2: for each $v \in V$, $v.allow_backups = True$.
 - 3: $H_1, H_2, \dots, H_\ell = \text{Procedure-Partition}(G, a)$
 - 4: each vertex $v \in H_i, \ell \geq i > 1$ in parallel does:
 - 5: select one arbitrary neighbor in $H_j, j < i$
 - 6: each vertex $v \in H_1$ in parallel does:
 - 7: **if** v has a neighbor in H_1 **then**
 - 8: v selects a neighbor in H_1 arbitrarily
 - 9: BIPARTITE-BP($U = \{\text{all vertices in } H_1 \text{ that have no neighbors in } H_1\}$, $V = H_2 \cup H_3 \cup \dots \cup H_\ell$, $E' = \{(u, v), u \in U, v \in V, (u, v) \in E\}$, a)
-

The running time of this algorithm is dominated by the execution of Procedure Partition, followed by our Bipartite algorithm of Section 3. Both these algorithms require $O(\log n)$ time. It remains to analyze the load of the algorithm. This is done in the next theorem.

Theorem 4.1. *For any positive parameter a , we compute Backup Placement in graphs of arboricity a within $O(\log n)$ time. The resulting load is $O(at)$, where t is the optimum load of the input graph.*

Proof. The running time follows from Theorem 3.5, and from the running time $O(\log n)$ of Procedure Partition of [5]. Next, we analyze the load. The maximum number of backups assigned to a vertex in each step, except the step of executing the Bipartite algorithm, is $3a$. The maximum number of backups of the Bipartite algorithm is $24ta$, where t is the load of the optimal backup placement of G' . (See Corollary 3.4.1. Indeed, each vertex in U has at most $3a$ neighbors in V .) Our goal, however, is to analyze the optimal load of G , rather than G' . Nevertheless, we next argue that the optimal load of a backup placement for G cannot be significantly smaller than that for G' . Thus, the above solution with load at most $3a + 24ta \leq 27ta$ is an $O(a)$ -approximation to the optimal backup placement of G as well.

Observe that the optimal load on vertices of V in G' is not greater than the load on these vertices in the optimal solution for G . Otherwise, the selections of vertices of U in G constitute a solution for G' with a smaller load on V than the optimal one. This is a contradiction. Thus, the execution of our Bipartite algorithm on (U, V, E') results in an $O(a)$ -approximation for the optimal load on V in G . In addition to the selections made during this execution, our algorithm performs additional selections, made by vertices that do not belong to U . However, all these additional selections cause an additional load of at most $3a$ in each vertex. Therefore, the complete solution of the procedure of Algorithm 4 is an $O(a)$ -approximation to the optimal solution for G . □

5 Lower Bounds

In this section we demonstrate that our upper bounds are not far from optimal, by proving lower bounds for exact and approximate backup placement. We begin with a lower bound for an exact (optimal load) solution.

Lemma 5.1. *Optimal backup placement in unoriented trees of height h requires $\Omega(h)$ rounds, for any deterministic algorithm. Moreover, in terms of the number of vertices n , the running time is $\Omega(\log n)$.*

Proof. In order to prove our claim we employ two distinct trees. (1) A fully balanced tree, in which each vertex (except the leafs) has a constant number of children d , and the height h is odd. (This makes the number of levels in the tree even.) (2) A "single leafs" tree, in which each non-leaf vertex has $d - 1$ siblings, but all leafs are single children, e.g., they have no siblings at all. The only difference between the fully balanced tree and the single leafs tree is in the h level, the furthest level from the root. (See Figure 8).




Figure 8: Optimal backup placement in fully balanced (left) and "single leafs" (right) trees.

The optimal solution for each of the above trees is the following. Tree (1): all vertices in even layers perform backups to their parents, while all vertices in odd layers perform backups to their children. In particular, all leafs perform backups to their parents. Tree (2): each vertex which is not a leaf selects one of its children for backup, and only the leafs select their parents for backups. Consequently, in the optimal backup placement in the fully balanced tree the load is d , while in the "single leafs" tree the load is 1 in all levels except $h - 1$, and the load is 2 in level $h - 1$. The reason for the load in level $h - 1$ is that both the single leaf and its grand-parent send their backups to the parent of the leaf. Next, we consider the root vertices in the two trees. Notice that in the fully balanced tree all the backups in level 1 must be sent to the root, and thus its load is d . (Otherwise, some vertex is going to have a load greater than d .) On the other hand, in the "single leafs" tree there are no backups at all of the root's children at the root, since they select their own children. Thus the load on the root is 0.

It follows that the optimal solution in both of the trees depends solely on the h level. In order to prove that the lower bound for optimal solution of the backup placement problem in this case is $\Omega(h)$, we prove by contradiction that it is impossible to achieve the optimal solution while ignoring the h level. Assume that there is a deterministic optimal solution with time complexity smaller than $h - 1$. When this algorithm is executed on the two trees, the roots of both trees perform exactly the same computations, since the subtrees of height $h - 1$ are the same in both trees. In other words, the algorithm that is executed in the roots ignores the h level, i.e., the leafs. However, except the h level, both trees are identical. Therefore, the roots compute the same solutions in both trees (1) and (2). However, the correct solutions of the roots are distinct in each of these trees. Therefore, this execution of the algorithm will be correct only in at most one of the trees, while wrong in at least one of them necessarily. Therefore, there is no optimal solution with time complexity lower than $h - 1$. For d -regular trees with a constant d , the height is $\Omega(\log n)$, which completes the proof. \square

Next, we provide a lower bound for approximate backup placement.

Theorem 5.2. *$O(1)$ -Backup placement requires $\Omega(\sqrt{\frac{\log n}{\log \log n}})$ rounds, for any deterministic algorithm.*

Proof. The proof is by contradiction. In order to prove our lower bound we invoke any distributed algorithm A for $O(1)$ -backup placement with running time $k = o(\sqrt{\log n / \log \log n})$ on the Kuhn-Moscibroda-Wattenhofer [17] graph G . This graph has various helpful properties, but we are interested in two specific properties that are crucial to our proof. (1) The girth of the graph is at least $2k + 2$, and (2) the minimum degree in the graph is at least 2. The first property means that the view of a vertex v with respect to radius k (i.e., the subgraph of G induced by all vertices at distance at most k from v) is a tree. The second property implies that in such a view, all leafs are at distance k from v , but not closer to v . This is because the degrees are preserved in the tree, except the very last layer at distance k from v , in which the degrees become equal to 1. We denote this tree by T . The vertex v is considered to be the root of T .

Next, consider a tree T' obtained by adding a single child to each vertex at distance k from the root in T . These new vertices constitute the set of leafs of T' , and they all are at distance $k + 1$ from the root of T' . Within distance k from the root the trees T and T' are indistinguishable. Since the optimal backup placement on T' has load 2, the Algorithm A invoked on either T' , T or G must assign at most $O(1)$ backups to v . This is because T' , T and G are indistinguishable to A that is executed for $k = o(\sqrt{\log n / \log \log n})$ rounds. Observe that this is true for all vertices v in G , once A terminates on G . In other words, A computes an $O(1)$ -backup placement of G .

Now consider the subgraph $G' = (V, E')$ of G , where E' is the set of all edges (u, w) , such that either u selected w for a backup, w selected u , or both. Since each vertex selects just a single neighbor for a backup, and each vertex is selected by $O(1)$ neighbors, the maximum degree of G' is $O(1)$. However, it contains all vertices of G . Thus G' is a constant approximation to a minimum vertex cover of G . (Take a minimum vertex cover of G and add all neighbors in G' of its vertices. The result is $V(G') = V(G)$, and the size is greater by a multiplicative factor of $O(1)$.) But computing such an approximation requires $\Omega(\sqrt{\frac{\log n}{\log \log n}})$ time on G [17].

Note: it may be possible that such a graph G' that spans G and has maximum degree $O(1)$ does not exist at all. In this case assuming the existence of an algorithm A with running time $o(\sqrt{\log n / \log \log n})$ leads to a contradiction (the existence of such a graph). Anyway, we obtain that $\Omega(\sqrt{\log n / \log \log n})$ rounds are required for $O(1)$ -backup-placement. \square

We note that since the lower bound of [17] applies not only to a constant-factor approximation, but also to polylogarithmic approximation factors, our arguments are directly extended to the following result.

Corollary 5.2.1. *Backup placement with polylogarithmic approximation factor requires $\Omega(\sqrt{\frac{\log n}{\log \log n}})$ rounds, for any deterministic algorithm.*

References

- [1] Mervat Abu-Elkheir, Mohammad Hayajneh, and Najah Ali. Data management for the internet of things: Design primitives and solution. *Sensors*, 13(11):15582–15612, 2013.
- [2] Charu C Aggarwal, Naveen Ashish, and Amit Sheth. The internet of things: A survey from the data-centric perspective. In *Managing and mining sensor data*, pages 383–428. Springer, 2013.
- [3] Suman Banerjee and Samir Khuller. A clustering scheme for hierarchical control in multi-hop wireless networks. In *Proceedings IEEE INFOCOM 2001. Conference on Computer Communications. Twentieth Annual Joint Conference of the IEEE Computer and Communications Society*, pages 1028–1037. IEEE, 2001.
- [4] Leonid Barenboim. On the locality of some np-complete problems. In *International Colloquium on Automata, Languages, and Programming*, pages 403–415. Springer, 2012.
- [5] Leonid Barenboim and Michael Elkin. Sublogarithmic distributed mis algorithm for sparse graphs using nash-williams decomposition. *Distributed Computing*, 22(5-6):363–379, 2010.
- [6] Leonid Barenboim and Michael Elkin. Distributed graph coloring: Fundamentals and recent developments. *Synthesis Lectures on Distributed Computing Theory*, 4(1):1–171, 2013.
- [7] Leonid Barenboim, Michael Elkin, and Cyril Gavoille. A fast network-decomposition algorithm and its applications to constant-time distributed computation. *Theoretical Computer Science*, 751:2–23, 2018.
- [8] Leonid Barenboim, Michael Elkin, and Tzalik Maimon. Deterministic distributed $(\Delta + o(\Delta))$ -edge-coloring, and vertex-coloring of graphs with bounded diversity. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*, pages 175–184. ACM, 2017.
- [9] Leonid Barenboim and Tzalik Maimon. Distributed symmetry breaking in graphs with bounded diversity. In *2018 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 723–732. IEEE, 2018.
- [10] Tongrang Fan and Yanzhao Chen. A scheme of data management in the internet of things. In *Network Infrastructure and Digital Content, 2010 2nd IEEE International Conference on*, pages 110–114. IEEE, 2010.
- [11] Manuela Fischer. Improved deterministic distributed matching via rounding. In *31st International Symposium on Distributed Computing (DISC 2017)*, volume 91, page 17. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.
- [12] Magnús M Halldórsson, Sven Köhler, Boaz Patt-Shamir, and Dror Rawitz. Distributed backup placement in networks. In *Proceedings of the 27th ACM symposium on Parallelism in Algorithms and Architectures*, pages 274–283. ACM, 2015.
- [13] Magnús M Halldórsson, Sven Köhler, Boaz Patt-Shamir, and Dror Rawitz. Distributed backup placement in networks. *Distributed Computing*, 31(2):83–98, 2018.
- [14] Magnús M Halldórsson, Sven Köhler, and Dror Rawitz. Distributed approximation of k-service assignment. *Distributed Computing*, 32(1):27–40, 2019.

- [15] Gholamreza Kakamanshadi, Savita Gupta, and Sukhwinder Singh. A survey on fault tolerance techniques in wireless sensor networks. In *Green Computing and Internet of Things (ICGCIoT), 2015 International Conference on*, pages 168–173. IEEE, 2015.
- [16] Fabian Kuhn. Weak graph colorings: distributed algorithms and applications. In *Proceedings of the twenty-first annual symposium on Parallelism in algorithms and architectures*, pages 138–144. ACM, 2009.
- [17] Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. What cannot be computed locally! In *Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing*, pages 300–309. ACM, 2004.
- [18] Christoph Lenzen, Yvonne Anne Oswald, and Roger Wattenhofer. What can be approximated locally?: case study: dominating sets in planar graphs. In *Proceedings of the twentieth annual symposium on Parallelism in algorithms and architectures*, pages 46–54. ACM, 2008.
- [19] Nathan Linial. Locality in distributed graph algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992.
- [20] Moni Naor and Larry Stockmeyer. What can be computed locally? *SIAM Journal on Computing*, 24(6):1259–1277, 1995.
- [21] C St JA Nash-Williams. Edge-disjoint spanning trees of finite graphs. *Journal of the London Mathematical Society*, 1(1):445–450, 1961.
- [22] C St JA Nash-Williams. Decomposition of finite graphs into forests. *Journal of the London Mathematical Society*, 1(1):12–12, 1964.
- [23] Gal Oren, Leonid Barenboim, and Harel Levin. Distributed fault-tolerant backup-placement in overloaded wireless sensor networks. In *International Conference on Broadband Communications, Networks and Systems*, pages 212–224. Springer, 2018.
- [24] Alessandro Panconesi and Romeo Rizzi. Some simple distributed algorithms for sparse networks. *Distributed computing*, 14(2):97–100, 2001.
- [25] Alessandro Panconesi and Aravind Srinivasan. On the complexity of distributed network decomposition. *Journal of Algorithms*, 20(2):356–374, 1996.
- [26] Yongrui Qin, Quan Z Sheng, Nickolas JG Falkner, Schahram Dustdar, Hua Wang, and Athanasios V Vasilakos. When things matter: A survey on data-centric internet of things. *Journal of Network and Computer Applications*, 64:137–153, 2016.
- [27] Johannes Schneider and Roger Wattenhofer. A log-star distributed maximal independent set algorithm for growth-bounded graphs. In *Proceedings of the twenty-seventh ACM symposium on Principles of distributed computing*, pages 35–44. ACM, 2008.

Appendices

A Proofs

A.1 UDG, BDG, and line graphs neighborhood independence

Under the assumption regarding WSNs, as stated at [3], we only consider bi-directional communication links in the following two lemmas. Therefore, only when two nodes are in each other's transmission range, there is an edge which connect the two nodes.

Lemma A.1. *Let G be a unit disk graph, the neighborhood independence of a vertex u is at most $c = 5$. i.e. Every six neighbors of a vertex have at least one edge.*

Proof. Let R be the radius of of the UDG (Figure 9). Let assume the neighborhood independence of a vertex k is $c = 6$, with v_1, \dots, v_6 vertices set in a cyclic fashion around k . Because of the known UDG's features, $distance(k, v_i) \leq R$. If $(v_i, v_j) \notin E$, then $distance(v_i, v_j) > R$. Thus, $distance(v_i, v_j) > distance(k, v_i) \& distance(k, v_j)$. Hence, in $\Delta kv_i v_j$, (v_i, v_j) is the largest side, and so $\angle f$ is the largest angle, which must be $> \frac{\pi}{3}$. Therefore, $\sum_{i=1}^6 \angle v_i k v_j > 6 \times \frac{\pi}{3} = 2\pi$, a contradiction. \square




Figure 9: Unit Disk Graph model: A root node v_0 , which forms 5 different cliques, and $\{v_1, \dots, v_5\}$ nodes which must choose v_0 as their backup placement.




Figure 10: Bounded Disk Graph model: A root node v_0 , which forms $11 \times (\log \frac{R_{max}}{R_{min}})$ different cliques, such that $\{v_1, \dots, v_{k+1}\} \dots \{v_{k+1}, \dots, v_{k+11}\}$, $k = \log \frac{R_{max}}{R_{min}}$ nodes must choose v_0 as their backup placement.

Lemma A.2. *Let G be a bounded disk graph, the neighborhood independence of a vertex u is at most $c = 11 \times (\log \frac{R_{max}}{R_{min}})$. i.e. Every $12 \times (\log \frac{R_{max}}{R_{min}})$ neighbors of a vertex have at least one edge.*

Proof. Let $I(u)$ be the largest independent set in the subgraph induced by $u \cup N(u)$. We define a “moat” $b(i)$ for $i \geq 0$, which is the annulus defined by circles of radii $c^i R_{min}$ and $c^{i+1} R_{min}$, centered at u . Therefore we have $n = \log_c \frac{R_{max}}{R_{min}}$ moats at $\max (c^n R_{min} = R_{max})$. $c = \sqrt{3}$ for ease of proof using

geometric properties. Let $N_i(u)$ be the neighbors of u that are in moat $b(i)$. Let p, q be two vertices that are in $N_i(u)$, and without loss of generality let $R_p \leq R_q$.

The distance between p and q is at least $\min(R_p, R_q) = R_p$, since there is no edge between p and q . We can "shrink" the circle centered at p with radius R_p until u is *on the boundary* of the circle, $R'_p \leq R_p$ and $R'_p \geq c^i R_{min}$. The distance between q and u is at most $c^{i+1} R_{min}$, which is $\leq c R'_p$. q is *inside* the circle centered at u with radius $c R'_p$, but *outside* the circle centered at p with radius R'_p - this implies that q is in the crescent shaped shaded region.

In order to compute the angle α between R'_p and $c^{i+1} R_{min}$, we assume a triangle $\Delta R_p, c R'_p, R'_p$, as $c^{i+1} R_{min} \leq c R'_p$ and $R_p < R'_p$. Thus, $\arccos(\alpha) > (c^2 R_p'^2 + R_p'^2 - R_p'^2) / (2c R_p'^2) = c^2 / 2c = \frac{\pi}{6}^\circ$.

Under these circumstances, the angle between p and q at u (α) is $> \frac{\pi}{6}$, and therefore there cannot be more than 11 vertices in the moat $b(i) \rightarrow 11 \times (\log \frac{R_{max}}{R_{min}})$ (Figure 10). We proclaim $\log \frac{R_{max}}{R_{min}}$ is assumed in this article to be a constant as wireless networks tend to work in a relatively small $\log \frac{R_{max}}{R_{min}}$ proportions. \square

Lemma A.3. *Let $G = (V, E)$ be a graph, and let $L(G)$ be a line graph of G . the neighborhood independence of a vertex u in $L(G)$ is at most $c = 2$.*

Proof. A vertex v in $L(G)$ corresponds to an edge $e_v \in G$. An independent set I of neighbors of v corresponds to a set of edges $J_I \subseteq E$, such that each $e \in J_I$ share a common endpoint with e_v , but each pair of edges $e, e' \in J_I$ do not intersect. Therefore, $|J_I| \leq 2$, and consequently, $|I| \leq 2$. Thus, any line graph $L(G)$ has neighborhood independence bounded by 2. \square