

# An Oblivious Spanning Tree for Buy-at-Bulk Network Design Problems

Srivathsan Srinivasagopalan, Costas Busch, S.S. Iyengar  
Computer Science Department  
Louisiana State University  
{ssrini1, busch, iyengar}@csc.lsu.edu

## Abstract

We consider the problem of constructing a single spanning tree for the single-source buy-at-bulk network design problem for doubling-dimension graphs. We compute a spanning tree to route a set of demands (or data) along a graph to or from a designated root node. The demands could be aggregated at (or symmetrically distributed to) intermediate nodes where the fusion-cost is specified by a non-negative concave function  $f$ . We describe a novel approach for developing an oblivious spanning tree in the sense that it is independent of the number of data sources (or demands) and cost function at intermediate nodes.

To our knowledge, this is the first paper to propose a single spanning tree solution to this problem (as opposed to multiple overlay trees). There has been no prior work where the tree is oblivious to both the fusion cost function and the set of sources (demands). We present a deterministic, polynomial-time algorithm for constructing a spanning tree in low doubling graphs that guarantees  $\log^3 D \cdot \log n$ -approximation over the optimal cost, where  $D$  is the diameter of the graph and  $n$  the total number of nodes. With constant fusion-cost function our spanning tree gives a  $O(\log^3 D)$ -approximation for every Steiner tree to the root.

## 1 Introduction

Buy-at-bulk network design problems arise in scenarios where economies of scale applies or when the availability of capacity in discrete units result in concave cost function on the edges. As observed in Chekuri et al. [2006], a commonly seen application is in telecommunication networks where bandwidth on a link can be purchased in some discrete units  $u_1 < u_2 < \dots < u_n$  with respective costs  $c_1 < c_2 < \dots < c_n$ . The economies of scale exhibits the property where the cost per bandwidth decreases as the number of units purchased gets larger:  $c_1/u_1 > c_2/u_2 > \dots c_n/u_n$ . This behavior justifies the sale of network capacity in “wholesale” (or “volume discount”) where more the capacity is bought, the cheaper is the price per unit of bandwidth.

We study the single-source buy-at-bulk (SSBB) network design problem with the following constraints: an unknown number of source (or demand) nodes and unknown concave transportation cost function  $f$ . An abstraction of this problem can be found in many applications, one of which is data fusion in wireless sensor networks where the given constraints are assumed unknown or vary over time. Others include design of VLSI power circuitry, Transportation & Logistics (railroad, water, oil, gas pipeline construction) etc. For simplicity, we consider data fusion problem in communication networks, though SSBB can also be applied to data distribution problems; our solution holds for both the cases.

As mentioned in Goel and Estrin [2003], if information flows from  $k$  different sources over a link, then, the total information that needs to be transmitted is  $f(k)$ , where the function  $f$  is

called a *canonical* fusion function where  $f$  is concave, non-decreasing and  $f(0) = 0$ . If  $f$  is known in advance, then, the problem of building an optimal fusion tree is well understood Bartal [1998]. However, here we consider the case where both  $f$  and the number of data sources are unknown. The focus of this paper is to construct an oblivious spanning tree according to the following definition:

**Definition 1.1 (*k*-Oblivious Spanning Tree (*k*-OST)).** *A spanning tree  $T$  with root node  $s$  in a graph  $G$  is  $k$ -oblivious if it provides a  $k$ -factor approximation to any data fusion problem toward sink node  $s$  with an arbitrary set of data sources and arbitrary canonical fusion function  $f$ .*

In this paper, we consider building an oblivious spanning tree for doubling dimension graphs. Doubling dimension graphs has been used in many different contexts including compact routing in wired networks [Abraham et al., 2006, Konjevod et al., 2008], traveling salesman, navigability and problems related to modeling the structural properties of the Internet distance matrix for distance estimation [Kleinberg et al., 2009, Fraigniaud, 2007]. As noted in Fraigniaud et al. [2006], it has become a key concept to measure the ability of network to support efficient algorithms or to realize specific tasks efficiently. For wireless networks, this concept has found many uses in solving many distributed communication problems [Kuhn et al., 2005], distributed resource-management [Gao et al., 2009], information exchange among producers and consumers [Funke et al., 2006], and for determining other performance qualities such as energy-conservation in wireless sensor networks [Pemmaraju and Pirwani, 2006].

SSBB is NP-Hard as the Steiner Tree problem is its special case (when  $f(x) = 1$ ) due to reduction from Steiner tree problem Salman et al. [2000].

## 1.1 Contribution

We build a *single* oblivious spanning tree for doubling-dimension graphs, such that the tree is independent of the data sources, and can accommodate *any* canonical fusion-cost function. Our approach gives a deterministic, polynomial-time algorithm that guarantees  $O(2^{10\rho} \log^3 D \log n)$ -OST (measured as the total involved communication cost), where  $\rho$  is the doubling dimension factor. For constant fusion cost functions, the  $k$ -OST behaves like a  $k$ -oblivious steiner tree. We define a *k-oblivious steiner tree* as a steiner tree that provides a  $k$ -approximation to any data fusion problem towards the sink with arbitrary number of data sources. For constant fusion-cost function ( $f(\cdot) = c$ ), we obtain a  $O(2^{10\rho} \log^3 D)$ -oblivious steiner tree to sink  $s$ .

Our spanning tree construction is based on the following techniques. We partition the nodes in a hierarchical fashion. The selection of nodes for a given ‘level’ of hierarchy is based on their mutual distances proportional to the level. Nodes of successive levels are connected by shortest paths. The intersecting paths are appropriately modified to result in a spanning tree. A modified tree is built from the spanning tree to ensure that all paths have appropriate end-nodes. Analysis is done on this modified tree.

## 1.2 Related Work

SSBB problems have been primarily considered in both Operations Research and Computer Science literatures in the context of flows with concave costs. SSBB problem was first introduced by Salman et al. [2000] and a  $O(\log^2 n)$ -approximation was given by Awerbuch and Azar [1997] using metric tree embedding. Bartal [1998] further improved this result to  $O(\log n)$ . Guha et al. [2001] provided the first constant-factor approximation to the problem.

Goel and Estrin [2003] build an overlay tree on graphs that satisfy triangle-inequality based on maximum matching algorithm that guarantees  $1 + \log k$  approximation, where  $k$  is the number of

Table 1: Our results and comparison with previous results for data-fusion schemes. IAF - Independent of fusion function, ISS - Independent of source-set,  $n$  is the total number of nodes in the topology,  $k$  is the total number of source nodes.

Related Work	Algorithm Type	Graph Type	IAF	ISS	Approx Factor	Tree Type
Lujun Jia <i>et al.</i> Jia et al. [2006]	Deterministic	Random Deployment	×	✓	$O(\log n)$	One Overlay
Ashish Goel <i>et al.</i> Goel and Estrin [2003]	Randomized	General Graph $\Delta$ -inequality	✓	×	$O(\log k)$	One Overlay
Ashish Goel <i>et al.</i> Goel and Post [2009]	Probabilistic	General Graph	✓	×	$O(1)$	Multiple Overlay
This paper	Deterministic	Low Doubling Dimension	✓	✓	$O(\log^3 D \cdot \log n)$	One Spanning

sources. An overlay tree, if projected to a graph, may not be a tree (could have cycles). In a related paper by Goel and Post [2009], the authors construct (in polynomial time) a set of overlay trees from a given general graph such that the expected cost of a tree for any  $f$  is within an  $O(1)$ -factor of the optimum cost for that  $f$ .

Jia et al. [2006] build Group Independent Spanning Tree Algorithm (GIST) that constructs an overlay tree for randomly deployed nodes in 2D. The tree (that is oblivious to the number of data sources) simultaneously achieves  $O(\log n)$ -approximate fusion cost and  $O(1)$  approximate delay. However, their solution assumes a constant fusion cost function. We summarize and compare the related work in Table 1.

Jia et al. [2005] provide approximation algorithms for TSP, Steiner Tree and set cover problems. They present a polynomial-time ( $O(\log(n)), O(\log(n))$ )-partition scheme for general metric spaces. An improved partition scheme for doubling metric spaces is also presented that incorporates constant dimensional Euclidean spaces and growth-restricted metric spaces. The authors present a polynomial-time algorithm for Universal Steiner Tree (UST) that achieves polylogarithmic stretch with an approximation guarantee of  $O(\log^4 n / \log \log(n))$  for arbitrary metrics and derive a logarithmic stretch,  $O(\log(n))$  for any doubling, Euclidean, or growth-restricted metric space over  $n$  vertices.

An earlier version of this work with preliminary results in the context of data aggregation appeared as a brief announcement in Srinivasagopalan et al. [2009].

## 2 Definitions

Consider a weighted graph  $G = (V, E, w)$ ,  $w : E \rightarrow \mathbb{Z}^+$ . Let  $s \in V$  be the sink node. For any two nodes  $u, v \in V$  let  $\text{dist}_G(u, v)$  denote the *distance* between  $u, v$  (measured as the total weight of the shortest path that connects  $u$  and  $v$ ). Let  $D$  denote the *diameter* of  $G$ , that is,  $D = \max_{u, v \in V} \text{dist}_G(u, v)$ . Given a subset  $V' \subseteq V$ , we denote  $\text{dist}_G(u, V')$  the smallest distance between  $u$  and any node in  $V'$ . We also define  $\text{near}_G(u, V') = \{v \in V' : \text{dist}_G(u, v) = \text{dist}_G(u, V')\}$

to be the set of nodes in  $V'$  that  $u$  is closest to. Given a set of nodes  $H \subseteq V$  and parameter  $d$ , we define *Maximal Independent Set of  $G$  for distance  $d$*  as  $I = MIS(G, H, d)$  to be an arbitrary maximal set of nodes in  $G$  such that  $H \subseteq I$  and for any  $u, v \in G \setminus H$ ,  $\text{dist}_G(u, v) \geq d$  and  $\text{dist}_G(u, x) \geq d$  where  $x \in H$ . There is a simple greedy algorithm to compute it. A special case is the notion of  $d$ -independent nodes computed by  $I = MIS(G, \phi, d)$ .

We adapt the definition of doubling-dimension graph from Nieberg [2006], Gupta et al. [2003].

**Definition 2.1 (r-Neighborhood).** Given a graph  $G = (V, E)$ , the  $r$ -neighborhood of any vertex  $u \in V$ ,  $N_G(u, r)$ , is defined as the set of nodes whose distance (in hops) is at most  $r$  from  $u$ ;  $N_G(u, r) = \{v \mid \text{dist}_G(u, v) \leq r\}$ .

**Definition 2.2 (Doubling Dimension Graph).** A graph  $G$  with the smallest  $\rho$  such that every  $R$ -neighborhood is completely covered by at most  $2^\rho$   $R/2$ -neighborhood is said to doubling. If  $\rho$  is bounded by a constant and is small, we say that  $G$  is **doubling** and has a low dimension.

**Observation 2.3.** In any  $R$ -neighborhood of doubling dimension graphs, there are at most  $m$   $R/(2^j)$ -neighborhoods where  $m = 2^{\rho \log(\frac{R}{R/(2^j)})}$ ,  $j \geq 0$ .

### 3 Spanning Tree Construction

We start with an informal description of the construction of the spanning tree. We build the tree in a hierarchical manner that has  $\kappa = O(\log D)$  levels. At level  $\kappa = \lceil \log D \rceil$  is the sink ( $I_\kappa = \{s\}$ ) and at level  $\kappa = 0$  are the individual nodes ( $I_0 = V$ ). Each level  $i$  of the hierarchy is built by identifying a set of independent nodes,  $I_i = \{\ell_i^1, \ell_i^2, \dots, \ell_i^n\}$ , ( $n \geq 0$  and the subscript corresponds to level), that are  $2^i$  distance (in hops) apart. The hierarchy consists of  $\kappa$  levels of independent nodes  $I_0, \dots, I_\kappa$  where  $I_i \subseteq V$ ,  $0 \leq i \leq \kappa$ . Members of  $I_i$  are also called leaders of level  $i$ . Some leaders could belong to more than one level (eg., sink  $s$ ). Leaders of consecutive levels are connected by shortest paths to build a tree. A formal description appears in Algorithm 1.

The construction of hierarchical levels of independent nodes is top-down.  $I_i$  is computed by  $MIS(G, I_{i+1}, 2^i)$ , for  $0 \leq i \leq \kappa - 1$ .  $I_i$  will contain all the  $2^j$ -independent nodes of higher levels  $j$ ,  $i < j \leq \kappa$  as well as a  $2^i$ -independent set of nodes. We enforce the constraint that  $s \in I_i$  for every  $I_i$ . Note that each node  $v \in I_i \setminus I_{i+1}$  has to be within distance  $2^{i+1} - 1$  to at least one node in  $I_{i+1}$  (otherwise  $v$  must be a member of  $I_{i+1}$ ).

Paths are also constructed in a top-down fashion. A path at any level  $i$ ,  $p_i$  starts at some leader at level  $i$  and ends at a leader at level  $i + 1$ . A set of all paths at level  $i$  is denoted as  $P_i$  and the set of all paths of all levels is denoted by  $P = \{P_{\kappa-1}, P_{\kappa-2}, \dots, P_2, P_1\}$ . We begin constructing paths from level  $\kappa - 1$  to sink  $s$  by computing shortest paths from each node of  $I_{\kappa-1}$  to  $s$ . We continue constructing the spanning tree by computing the paths for the remaining levels. Suppose we have constructed paths for level  $i$ ; we will construct  $P_{i-1}$ . Each node  $v \in I_{i-1}$  prefers some arbitrary node in  $\text{near}_G(v, I_i)$ ; if  $s \in \text{near}_G(v, I_i)$ , then,  $s$  is preferred. Shortest paths from each node  $v$  are constructed to their respective leaders.

When paths for all levels are built, the resulting structure may not be a tree. It could result in a graph that might have intersecting paths. Define *regular* paths as paths that do not intersect any (higher-level) path on their way to their end-nodes. The paths of  $P_{\kappa-1}$ , are regular paths, since there were no higher-level paths to intersect and are included in  $P_{\kappa-1}^{reg}$ .

Define *pruned paths* as those paths that intersect paths of higher level. If a path  $p_i$  intersects a path  $p_j$  ( $j > i$ ) along its way to  $\ell_{i+1}$ ,  $p_i$  is pruned from the intersection point to its destination. Such paths are included in  $P_i^{pr}$ . This pruning of intersecting paths ensures the structural property of a spanning tree (see Figure 1).

---

**Algorithm 1:** Spanning Tree

---

**Input:** Graph  $G$  with sink  $s$ .**Output:** A spanning tree  $T_s$ .

```
 $P \leftarrow \phi; I_\kappa \leftarrow \{s\};$  //  $\kappa \leftarrow \lceil \log D \rceil$   
 $P^{reg} \leftarrow \phi; P^{pr} \leftarrow \phi;$  // List of regular and pruned paths  
foreach level  $i = \kappa - 1$  to 0 do  
   $I_i \leftarrow MIS(G, I_{i+1}, 2^i);$   
  foreach  $v \in I_i$  do  
    //  $v$  chooses nearest  $\ell \in I_{i+1}$  as parent.  
     $\ell \leftarrow near_G(v, I_{i+1});$  // If  $s \in near(v, I_{i+1})$ ,  $s$  is chosen.  
     $p_v \leftarrow$  Compute shortest path from  $v$  to  $\ell$ ;  
    if  $p_v$  intersects any path at level  $> i$  at point  $u$  then  
      // Prune path  $p_v$  by removing segment from  $u$  to  $\ell$   
       $p' \leftarrow$  path segment from  $v$  to  $u$ ;  
       $P_i^{pr} \leftarrow P_i^{pr} \cup p'$ ;  
    else  
       $P_i^{reg} \leftarrow P_i^{reg} \cup p_v$ ;  
 $P \leftarrow \bigcup_{i=0}^{i=\kappa-1} P_i^{reg} \cup \bigcup_{i=0}^{i=\kappa-1} P_i^{pr}$   
return  $T_s$  formed by the paths in  $P$ ;
```

---

Note that regular paths of the same level could intersect and continue on different directions to reach a common leader. In this case, one of the paths is modified to use the same segment as the other after the intersection point. Another scenario is when two paths (say from  $u$  and  $v$  of level  $i$ ) intersect at  $m$  and proceed to their respective endnodes  $x$  and  $y$ . In this case, either  $v$  or  $u$  will choose a common leader and appropriately modify its path. In both these scenarios, the resulting paths remain regular even if they overlap. Note that in both the cases, the path segments, after intersection, should have the same length.

The spanning tree algorithm executes in  $\kappa$  rounds and each round computes  $MIS$  (in  $O(|E|)$  time, assuming that the input is given by an adjacency list) and shortest paths (in  $|E| \log n$ ) time steps). This amounts to a total running time of  $O(\log D \cdot |E| \log n)$ .

## 4 Modified Tree Construction

The pruned paths in the spanning tree  $T$  will not have leaders as end-nodes. To ensure that end-nodes of all paths are leaders, we modify  $T$  to  $\bar{T}$ . We begin with an overview of the modified tree construction. We construct  $\bar{T}$  from  $T$  by assigning alternate leaders to those paths whose ‘upper’ sections have been pruned. We first begin by assigning *levels* to all the nodes of regular paths by `AssignLevels` (see appendix) and including those paths in  $\bar{T}$ . Then, we begin a top-down, level-by-level process where we ‘modify’ the pruned paths by extending the pruned paths to their newly assigned alternate leaders. Note that a modified path could be a concatenation of multiple pruned paths. Then, we assign levels to the nodes of the recently modified path as well and include this modified path in  $\bar{T}$ . The end of this process results in a tree  $\bar{T}$ . A more formal description appears in Algorithm 2.

Define `AssignLevels`( $p_i, H, i$ ), where  $H$  is a pair of end-nodes of  $p_i$ , to assign levels to all the nodes of  $p_i$  by identifying maximal independent nodes (excluding the end nodes of  $p_i$ ). Levels

---

**Algorithm 2:** Modified Tree

---

**Input:** Spanning Tree  $T$  rooted at  $s$ .

**Output:** A modified tree  $\bar{T}$ .

```
 $\bar{T} \leftarrow \phi$ ; //  $T = P = \{P_{\kappa-1}, P_{\kappa-2}, \dots, P_1, P_0\}$ 
// Assign Levels to all nodes in all regular paths in  $T$ .
 $i \leftarrow \kappa - 1$ ; // start from second level from top
while  $i \geq 0$  do
  foreach  $p_i \in P_i^{reg}$  do
    //  $v_a$  and  $v_b$  are the start and end nodes of path  $p_i$ 
     $H \leftarrow \{v_a, v_b\}$ ; //  $v_a$  is at same level as that of  $p_i$ .
    AssignLevels ( $p_i, H, i$ );
     $\bar{T} \leftarrow \bar{T} \cup p_i$ ;
   $i \leftarrow i - 1$ ;
// Pruned paths in  $\bar{T}$  - Modify paths and assign levels.
 $\omega \leftarrow \kappa - 2$ ;
while  $\omega > 0$  do
  foreach  $p_\omega \in P_\omega^{pr}$  do
     $\bar{p}_\omega \leftarrow \text{ModifyPath}(p_\omega, p_i)$ ; //  $p_\omega$  intersects  $p_i$ ,  $i > \omega$  and  $v'_b$  be the elected
    pseudo-leader.  $p_i$  may be a modified path itself.
     $\bar{T} \leftarrow \bar{T} \cup \bar{p}_\omega$ ;
     $H \leftarrow \{v_a, v'_b\}$ ; //  $v_a$  and  $v'_b$  are the start and end nodes of  $\bar{p}_\omega$ .
    AssignLevels ( $\bar{p}_\omega, H, \omega$ );
   $\omega \leftarrow \omega - 1$ ;
return  $\bar{T}$ ;
```

---

are assigned in the range  $(i - 1)$  to 0. A modified path is connected to an alternate leader called *pseudo-leader* by the function  $\text{ModifyPath}(p_\omega, p_i)$  which chooses the nearest level- $(\omega + 1)$  node on  $p_i$  from the intersection point.

Consider that we are at some level  $\omega$  where  $0 \leq \omega \leq \kappa - 1$  and suppose that there are several pruned paths in  $P_\omega$ . Let  $p_\omega \in P_\omega$  be one such path and let  $y \in p_i$  be the intersection point, where  $i > \omega$ . A *pseudo-leader*,  $v_{\omega+1}$ , is chosen on  $p_i$  using  $\text{ModifyPath}(p_\omega, p_i)$  (see appendix). Note that this may alter  $I_i$  to  $\bar{I}_i$  by replacing the original leader by the pseudo-leader. The path  $p_\omega$  is extended from  $y$  to  $v_{\omega+1}$  and this new extended path  $\bar{p}_\omega$  replaces  $p_\omega$  in the modified tree  $\bar{T}$ . Once a new path  $\bar{p}_\omega$  is established, all the nodes in it are assigned levels using  $(\text{AssignLevels}(\bar{p}_\omega, H, \omega))$ , where  $H$  is the set of end-nodes of  $\bar{p}_\omega$ . This procedure of modifying pruned paths, replacing the old pruned paths by new, extended, *modified* paths and assigning levels to all nodes in those paths is repeated for all levels down to 0. The resulting tree is a modified tree with normal leaders and pseudo-leaders for respective types of paths.

## 5 Analysis

We will analyze the algorithm based on the data fusion algorithm in the modified tree  $\bar{T}$ . The fusion algorithm works in phases, where each phase consists of  $\kappa$  rounds. In the beginning of each phase, several source nodes may have data to send. In the first round, each source node sends its

data to its respective parent at level 1. In the second round, the leaders at level 1 send the fused data to their respective parent at level 2. In general, in round  $i$ , where  $1 \leq i \leq \kappa$ , the leaders at level  $i - 1$  send their data to their respective parent at level  $i$  which then fuses the received data. At the end of round  $\kappa$ , the sink node  $s$  would have fused all the received data.

The modified tree  $\bar{T}$  naturally defines a hierarchical partition of  $G$ . For each node  $u \in \bar{T}_i$ , we define the respective cluster  $Z_i^u$  to contain all the nodes in  $V$  which appear in the subtree of  $\bar{T}$  rooted at  $u$  at level  $i$ . Node  $u$  is the *leader* of  $Z_i^u$ . We denote by  $Z_i = \bigcup_{u \in \bar{T}_i} Z_i^u$  the partition of  $G$  at level  $i$ . Let  $Z = \bigcup_{0 \leq i \leq \kappa} Z_i$  denote a *partition hierarchy* of  $V$  with  $\kappa$  levels. Note that  $Z_\kappa$  consists of one cluster containing all the nodes in  $V$ . Clearly, each cluster  $Z_i^u$  induces a connected subgraph in  $G$ .

Let  $A$  denote all the source nodes. We assume that each data item generated at a source node has size 1. Let  $A_i^u = A \cap Z_i^u$ . According to the data fusion algorithm, at the end of a round no later than  $i$ , the data items in  $A_i^u$  will appear fused at node  $u$  with size  $f(|A_i^u|)$ . Let  $B_i^j$  denote the set of nodes at level  $i$  which at the end of round  $i$  hold data with size  $f(|A_i^u|) \in [2^j, 2^{j+1} - 1]$ , where  $0 \leq j \leq \lambda$ , and  $\lambda = \lceil \log |A| \rceil = O(\log n)$ . Let  $B_0^0 = A$ . Lemmas (5.3 and 5.4) establishes the lower-bound on communication cost at each round.

A path  $p_i^j$  could be intersected by multiple lower-level paths. Even though the leaders at a level  $i$  are sufficiently far off, due to intersection by other paths, the leader at level  $i$  might be close to many leaders of lower level paths. However, the number of such leaders that are close is *limited*. Lemmas A.3, A.4 and 5.2 establishes the maximum number of pseudo-leaders in a given neighborhood. If all the nodes of a path  $p_i$  belong to a cluster  $Z_i^u$ , we call such a path to be a *total internal path*. Let  $\delta = 3 \cdot 2^i$ .

**Lemma 5.1.** *The maximum distance in  $G$  between any node  $v \in Z_i^u$  to  $u$  is  $\delta = 3 \cdot 2^i$  and there is a total internal path from any node  $v \in Z_i^u$  to  $u$  with respect to  $Z_i^u$ .*

*Proof.* Consider a path  $p_{v,u} \in Z_i^u$ . In the worst case, this path could be a concatenation of several modified paths, ranging from level 0 to  $i - 1$ . The total length of  $p_{v,u}$  would be equal to the sum of maximum lengths of each of those segments:  $\sum_{j=0}^{i-1} (3 \cdot 2^j - 2) = 3 \cdot 2^i - 2i - 1 < 3 \cdot 2^i$ .

By construction, any cluster  $Z_i^u$  will contain those nodes of  $V$  that appear in the subtree of  $\bar{T}$  rooted at  $u$  at level  $i$ . The path  $\bar{p}_i \in \bar{T}$  from any of the member nodes  $v \in Z_i^u$  to  $u$  contains *only* nodes that lie within the cluster  $Z_i^u$ . Also,  $\bar{p}_i$  translates to a path  $p \in T$ .  $\square$

**Lemma 5.2.** *The total number of pseudo-leaders at level  $i$ , where  $i \leq r$ , which are inside  $N_G(x, 2^r)$  is at most  $2^{2\rho(r-i+3)} \cdot (\kappa - i + 1)^2$ .*

*Proof.* From Lemma A.3, there are  $2^{\rho(r-i+3)} \cdot (\kappa - i + 1)$  path segments  $p_{i+j} \in T$ ,  $j \geq 0$ , crossing  $N(x, 2^r)$ . From Lemma A.4, each such path segment can have multiple modified path segments at level  $i$  or higher passing through it ( $\leq 2^{\rho(r-i+1)} \cdot (\kappa - i + 1)$ ), the total number of modified path segments that cross  $N(x, 2^r)$  would be atmost  $2^{2\rho(r-i+3)} \cdot (\kappa - i + 1)^2$ . This gives also an upper bound to the number of pseudo-leaders at level  $i$  or higher.  $\square$

**Lemma 5.3.** *The communication cost of the data fusion algorithm at round  $i$  is at most  $|B_{i-1}^j| 6 \cdot 2^{i+j}$ , for any  $1 \leq i \leq \kappa$ .*

*Proof.* Let  $v$  be a leader in  $\bar{T}_{i-1}$  and suppose that  $v \in B_{i-1}^j$ . The cost of sending data from  $v$  to its parent at level  $i$  is at most  $\delta \cdot (2^{j+1} - 1) \leq 3 \cdot 2^{i+j+1}$ , which is obtained by the product of the total path length connecting  $v$  to its parent (bounded by  $\delta$ ), from Lemma 5.1, to the size of the data (bounded by  $2^{j+1} - 1$ ). Therefore, the total cost at round  $r$  for communicating data in set  $B_{i-1}^j$  is at most  $|B_{i-1}^j| 6 \cdot 2^{i+j}$ .  $\square$

We proceed with estimating a lower bound on the communication cost at each round  $k$ . Let  $C^*(A, G)$  denote the optimal communication cost for  $A$  data sources in graph  $G$ . We start with a simple lower bound.

**Lemma 5.4.** *If a message is sent in round  $i$  from a node  $v \in B_{i-1}^j$  then  $C^*(A, G) > \max(2^{i+j-1}, 1)$ .*

*Proof.* Since  $v \in B_{i-1}^j$ , the message sent by  $v$  during round  $i$  has size at least  $2^j$ . If  $i = 1$  then the lower bound 1 of the claim follows trivially. Suppose now that a message is sent during round  $i > 1$ . Let  $\Gamma \subseteq A$  be the set of source nodes such that for every  $x \in \Gamma$ ,  $\text{dist}_G(x, s) > 2^{i-1}$ . We will show that  $f(|\Gamma|) \geq 2^j$ . Assume for contradiction that  $f(|\Gamma|) < 2^j$ . From the construction of  $\bar{T}$  it can be verified that every node within distance  $2^{i-1}$  from  $s$  is a member of the subtree of  $\bar{T}$  rooted at  $s$  at level  $i - 1$ , since  $s$  is always preferred as a leader choice; these are the nodes of set  $A \setminus \Gamma$ . Thus, the last round of the algorithm in which a message is sent carrying information from data sources in  $A \setminus \Gamma$  is  $i - 1$ . Therefore, the message which is sent at round  $i$  must carry information only from the data sources in  $\Gamma$ . Therefore,  $f(|\Gamma|) \geq 2^j$ . The smallest cost to transfer the data from the  $\Gamma$  data sources to the sink  $s$  is more than  $f(|\Gamma|)2^{i-1} \geq 2^j 2^{i-1} = 2^{i+j-1}$ .  $\square$

We continue with an alternative lower bound. We construct a new family of multi-graphs by contracting the clusters of the partition hierarchy  $Z$ . Let  $\hat{G}_i = (\hat{V}_i, \hat{E}_i, \hat{w})$  be complete graph that we obtain when we contract the edges in the clusters of  $Z_i$  and the whole cluster  $Z_i^u$  is replaced by  $u$ . Each edge  $e = (x, y) \in E$  is replaced by edge  $e' = (x', y') \in \hat{E}_i$  where  $x'$  and  $y'$  are the respective leaders of  $x$  and  $y$  at partition  $Z_i$  and  $\hat{w}(x', y') = \text{dist}_G(x', y')$ .

**Lemma 5.5.** *Given an arbitrary set of nodes  $X \subseteq \hat{V}_i$ , where  $0 \leq i \leq \kappa$ , there is a subset  $X' \subseteq X$  such that  $|X'| \geq |X|/(2^{10\rho} \cdot (\kappa - i + 1)^2 + 1)$ , and for each pair of distinct  $x, y \in X'$ ,  $\text{dist}_{\hat{G}_i}(x, y) > 3 \cdot \delta$ .*

*Proof.* Construct a new unweighted graph  $H = (X, E_H)$  such that  $(x, y) \in E_H$  if  $\text{dist}_{\hat{G}_i}(x, y) \leq 3 \cdot \delta$ . Let  $x \in X$ . From Lemma 5.2, the number of pseudo-leaders of level  $i$  which are in  $N_G(x, 3\delta)$  is bounded by  $m = 2^{2\rho(i+2-i+3)} \cdot (k - i + 1)^2 = 2^{2\rho(5)} \cdot (k - i + 1)^2 = 2^{10\rho} \cdot (\kappa - i + 1)^2$ . By construction of  $\hat{G}_i$ ,  $\text{dist}_{\hat{G}_i}(x, y) \leq 3\delta$  only if  $y \in N_G(x, 3\delta)$ . Therefore, the degree of  $H$  is bounded by  $m$ . Hence,  $H$  accepts a  $m + 1$  coloring which implies that there is an independent set  $X' \subseteq X$  in  $H$  of size  $|X'| \geq |X|/(2^{10\rho} \cdot (\kappa - i + 1)^2 + 1)$ . Clearly, for each  $x, y \in X'$ ,  $\text{dist}_{\hat{G}_i}(x, y) > 3 \cdot \delta$ .  $\square$

**Lemma 5.6.**  $C^*(A, G) \geq 2^{i+j-1} \{ |B_i^j| / [2^{10\rho} \cdot (\kappa - i + 1)^2 + 1] - 1 \}$ , for every  $i, 0 \leq i \leq \kappa, 0 \leq j \leq \lambda$ .

*Proof.* Consider a particular  $B_i^j$ . Let  $\mu = |B_i^j|$  and  $B_i^j = \{u_1, \dots, u_\mu\}$ . We have that  $A_i^{u_q}$  data sources are in  $Z_i^{u_q}$  and  $f(|A_i^{u_q}|) \in [2^j, 2^{j+1} - 1]$ . Let  $X = \bigcup_{q=1}^\mu A_i^{u_q}$ . Let  $C^*(X, G)$  denote the optimal cost of fusing the data from the  $X$  data sources to the sink node  $s$  in graph  $G$ .

We transform the data fusion problem with the  $X$  sources to a new data fusion problem in  $\hat{G}_i$  where all data sources in  $A_i^{u_q}$  are relocated to the respective leader node  $u_q$ . The size of the data held at the leader node is  $f(|A_i^{u_q}|)$ . Let  $\hat{X} = B_i^j$ . We have that the optimal cost of fusing the data from  $\hat{X}$  in  $\hat{G}_i$  is no larger than the cost of fusing the data from  $X$  in  $G$ , namely,  $C^*(\hat{X}, \hat{G}_i) \leq C^*(X, G)$ , since graph  $\hat{G}_i$  is obtained from  $G$  by contracting edges, and thus the optimal communication cost of  $G$  cannot get worse in  $\hat{G}_i$ .

From Lemma 5.5, there is a subset  $\hat{X}' \subseteq \hat{X}$ , such that  $|\hat{X}'| \geq |\hat{X}| / [2^{10\rho} \cdot (\kappa - i + 1)^2 + 1]$ , and for each distinct  $x, y \in \hat{X}'$ ,  $\text{dist}_{\hat{G}_i}(x, y) > 3\delta$ . Let  $x, y \in \hat{X}'$  be any two nodes  $u \in Z_i^x$ ,  $v \in Z_i^y$  and  $\text{dist}_G(x, y) \geq \delta$ . The cost of fusing the data stored in  $\hat{X}'$  to the sink  $s$  is at least the weight of the Steiner tree  $T_S$  in  $G_i$  that connects the set of nodes  $\hat{X}'$  (including  $s$  in the Steiner tree would only increase its cost) multiplied by  $2^j$  which is a lower bound on the size of the data from each source



in  $\widehat{X}'$ . The total weight of  $T_S$  is  $w(T_S) > \delta(|\widehat{X}'| - 1) \geq \delta\{|B_i^j|/[2^{10\rho} \cdot (\kappa - i + 1)^2 + 1] - 1\}$ . The claim follows, since  $w(T_S) \cdot 2^j \leq C^*(\widehat{X}', \widehat{G}_i) \leq C^*(\widehat{X}, \widehat{G}_i) \leq C^*(X, G) \leq C^*(A, G)$ .  $\square$

**Theorem 5.7** (Approximation of Spanning Tree). *Tree  $\overline{T}$  is  $\sigma$ -OST where  $\sigma = O(2^{10\rho} \log^3 D \cdot \log n)$ .*

*Proof.* Suppose that there is at least one data source different than the sink  $s$ . The algorithm runs in  $\kappa = O(\log n)$  rounds. Let  $Q_{i-1}^j$  be the cost of the algorithm when transferring data from the level  $i - 1$  nodes  $B_{i-1}^j$  to the respective parent leader in level  $i$ . We have that the total cost of the algorithm at round  $i$  is  $Q_i = \sum_{j=0}^{\lambda} Q_{i-1}^j$ . From Lemma 5.3,

$$Q_{i-1}^j \leq |B_{i-1}^j| 6 \cdot 2^{i+j}. \quad (1)$$

Let  $Q_{i'-1}^{j'} = \max_{i,j} Q_{i-1}^j$ . Since at least one message is sent in the algorithm,  $Q_{i'-1}^{j'} \geq 1$ . From Lemma 5.4,

$$C^*(A, G) > \max(2^{i'+j'-1}, 1). \quad (2)$$

From Lemma 5.6, we also have

$$C^*(A, G) \geq 2^{i'+j'-1} \left( \frac{|B_{i'-1}^{j'}|}{2^{10\rho} \cdot (\kappa - i + 1)^2 + 1} - 1 \right). \quad (3)$$

If  $|B_{i'-1}^{j'}| < 2(2^{10\rho} \cdot (\kappa - i + 1)^2 + 1)$ , then, from Equations 1 and 2 we have:

$$\frac{Q_{i'-1}^{j'}}{C^*(A, G)} < \frac{2(2^{10\rho} \cdot (\kappa - i + 1)^2 + 1) \cdot 6 \cdot 2^{i'+j'}}{2^{i'+j'-1}} \leq 24 \cdot (2^{10\rho} \cdot (\kappa + 1)^2 + 1) \quad (4)$$

If  $|B_{i'-1}^{j'}| \geq 2(2^{10\rho} \cdot (\kappa - i + 1)^2 + 1)$ , then, from (1) and (3) we also obtain (4). Let  $Q$  be the total cost of the algorithm,  $Q = \sum_i Q_i$ . Since there are  $\kappa$  levels and  $\lambda$  clusters, from Equation 4 we have  $Q \leq \kappa(\lambda + 1)Q_{i'-1}^{j'} \leq \kappa(\lambda + 1) \cdot C^*(A, G) \cdot 24 \cdot (2^{10\rho} \cdot (\kappa + 1)^2 + 1)$ . Thus, since  $\kappa = O(\log D)$  and  $\lambda = O(\log n)$ , we have  $\frac{Q}{C^*(A, G)} \leq \kappa(\lambda + 1) \cdot 24 \cdot (2^{10\rho} \cdot (\kappa + 1)^2 + 1) = O(2^{10\rho} \log^3 D \cdot \log n)$ .  $\square$

If  $f$  is the constant cost function, then,  $\lambda = 0$  and we obtain the following corollary.

**Corollary 5.8.** *The constant cost function  $f(x) = c$ , gives  $O(2^{10\rho} \log^3 D)$ -oblivious Steiner tree from Theorem 5.7.*

## 6 Simulation Results

We simulated the Oblivious Spanning Tree and compared its performance (fusion cost) with GRID\_GIST [Jia et al., 2006], Maximum Matching Algorithm [Goel and Estrin, 2003] and other common trees such as MST (Minimum Spanning Tree) and SPT (Shortest-Paths Tree). We used a 2-D grid topology for our simulation using NetworkX [Hagberg et al., 2008]. 2-D grids are a special case of doubling dimension graphs and they fall under a variation of the Steiner tree problem called ‘Rectilinear Steiner Problem’ (RSP) where the tree structure has only vertical and horizontal lines that interconnects all points and is proved to be NP-Complete [Garey and Johnson, 1977]. Since calculating a minimum weight tree structure in a 2-D grid topology (a doubling-dimension graph) is essentially an RSP, the problem we are addressing is NP-Hard.

We build a single spanning tree in a grid with 1600 nodes. We simulate it for random sets of data sources, upto 1445, that are randomly placed. Note that GRID\_GIST is a special algorithm designed for grids and ours is a generalized algorithm. Hence, GRID\_GIST performs slightly better than OST (in Fig 2 in appendix).

## 7 Conclusions and Future Work

We studied the problem of computing near-optimal fusion trees when the number of sources and the fusion cost function are unknown in the context of *Doubling Dimension Graphs*. We have demonstrated that a simple, deterministic, polynomial-time algorithm based on independent sets can provide a near-optimal data structure for data-fusion (under the assumptions of concave cost function and doubling dimension graphs). The algorithm constructs a spanning tree which is then altered to produce a modified tree. We have shown that this algorithm guarantees  $\log^4 n$ -approximation over the optimal cost. As part of our future work, we are looking into the same problem on planar graphs, extending to multiple sinks and incorporating fault-tolerance and load-balancing.

## References

- Ittai Abraham, Cyril Gavoille, Andrew V. Goldberg, and Dahlia Malkhi. Routing in networks with low doubling dimension. In *ICDCS '06*, page 75, Washington, DC, USA, 2006. ISBN 0-7695-2540-7.
- B. Awerbuch and Y. Azar. Buy-at-bulk network design. In *FOCS '97*, page 542, Washington, DC, USA, 1997. IEEE Computer Society.
- Yair Bartal. On approximating arbitrary metrics by tree metrics. In *STOC '98*, pages 161–168, New York, NY, USA, 1998. ACM.
- C. Chekuri, M. T. Hajiaghayi, G. Kortsarz, and M. R. Salavatipour. Approximation algorithms for non-uniform buy-at-bulk network design. In *FOCS '06*, pages 677–686, Washington, DC, USA, 2006.
- Pierre Fraigniaud. The inframetric model for the internet. Technical report, 2007.
- Pierre Fraigniaud, Emmanuelle Lebhar, and Zvi Lotker. A doubling dimension threshold  $\theta(\log \log n)$  for augmented graph navigability. In *ESA, LNCS 4168*, pages 376–386. Springer, 2006.
- Stefan Funke, Leonidas Guibas, An Nguyen, and Yusu Wang. Distance-sensitive information brokering in sensor networks. In *DCOSS 2006*, volume 4026 of *LNCS*, pages 234–251, San Francisco, USA, 2006. IEEE, Springer.
- Jie Gao, Leonidas J. Guibas, Nikola Milosavljevic, and Dengpan Zhou. Distributed resource management and matching in sensor networks. In *Proc. of (IPSN'09)*, pages 97–108, April 2009.
- M. R. Garey and D. S. Johnson. The rectilinear steiner tree problem is np-complete. *SIAM Journal on Applied Mathematics*, 32(4):826–834, 1977. URL <http://www.jstor.org/stable/2100192>.
- Ashish Goel and Deborah Estrin. Simultaneous optimization for concave costs: single sink aggregation or single source buy-at-bulk. In *SODA '03*, pages 499–505, Philadelphia, PA, USA, 2003. SIAM.
- Ashish Goel and Ian Post. An oblivious  $o(1)$ -approximation for single source buy-at-bulk. *CoRR*, abs/0908.3740, 2009.

- Sudipto Guha, Adam Meyerson, and Kamesh Munagala. A constant factor approximation for the single sink edge installation problems. In *STOC '01*, pages 383–388, New York, NY, USA, 2001. ACM. ISBN 1-58113-349-9.
- Anupam Gupta, Robert Krauthgamer, and James R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *FOCS '03*, page 534, Washington, DC, USA, 2003. IEEE Computer Society.
- Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. Exploring network structure, dynamics, and function using NetworkX. In *Proceedings of the 7th Python in Science Conference (SciPy2008)*, pages 11–15, Pasadena, CA USA, August 2008.
- Lujun Jia, Guolong Lin, Guevara Noubir, Rajmohan Rajaraman, and Ravi Sundaram. Universal approximations for tsp, steiner tree, and set cover. In *STOC '05: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 386–395, New York, NY, USA, 2005. ACM.
- Lujun Jia, Guevara Noubir, Rajmohan Rajaraman, and Ravi Sundaram. Gist: Group-independent spanning tree for data aggregation in dense sensor networks. In *DCOSS*, pages 282–304, 2006.
- Jon Kleinberg, Aleksandrs Slivkins, and Tom Wexler. Triangulation and embedding using small sets of beacons. *J. ACM*, 56(6):1–37, 2009. ISSN 0004-5411. doi: <http://doi.acm.org/10.1145/1568318.1568322>.
- Goran Konjevod, Andréa W. Richa, and Donglin Xia. Dynamic routing and location services in metrics of low doubling dimension. In *DISC '08*, pages 379–393. Springer-Verlag, 2008.
- Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. On the locality of bounded growth. In *PODC '05*, pages 60–68, New York, NY, USA, 2005. ACM.
- T. Nieberg. *Independent and Dominating Sets in Wireless Communication Graphs*. PhD thesis, University of Twente, Zwolle, April 2006.
- Sriram V. Pemmaraju and Imran A. Pirwani. Energy conservation via domatic partitions. In *MobiHoc 2006*, pages 143–154, New York, NY, USA, 2006. ACM. doi: <http://doi.acm.org/10.1145/1132905.1132922>.
- F. S. Salman, J. Cheriyan, R. Ravi, and S. Subramanian. Approximating the single-sink link-installation problem in network design. *SIAM J. on Optimization*, 11(3):595–610, 2000. ISSN 1052-6234.
- Srivathsan Srinivasagopalan, Costas Busch, and S. Sitharama Iyengar. Brief announcement: Universal data aggregation trees for sensor networks in low doubling metrics. In *Algorithmic Aspects of Wireless Sensor Networks: 5th International Workshop, ALGOSENSORS 2009, Rhodes, Greece, July 10-11, 2009.*, pages 151–152, Berlin, Heidelberg, 2009. Springer-Verlag. ISBN 978-3-642-05433-4.

## A Observations and Proofs

**Lemma A.1 (Presence of a Pseudo-Leader).** *The  $\text{ModifyPath}(p_\omega, p_i)$  function guarantees selection of a  $(\omega + 1)$ -level pseudo-leader.*

*Proof.* Suppose path  $p_\omega$  intersects a higher-level path  $p_i$ . Let the start-node of  $p_\omega$  be  $u$  and let the end-node of  $p_i$  be  $v$ . Note that a path  $p_i$  goes from level  $i$  to level  $i + 1$ . There could be two cases for the presence of a pseudo-leader in  $p_i$ . If level of  $v$  is  $\omega + 1$ , then,  $v$  itself acts as a pseudo-leader for  $u$ . If level of  $v$  is greater than  $\omega + 1$ , then,  $p_i$  must have some nodes (within its end-nodes) that have been assigned to level  $\omega + 1$  (by the `AssignLevels` function). Hence, in either case, a pseudo-leader is guaranteed to be found in  $p_i$  for  $u$ .  $\square$




Figure 1: Pruning and Tree Modification.

**Lemma A.2 (Upper Bound on  $|p'_\omega|$ ).** *The upper bound on the length of  $p'_\omega$  is  $(3 \cdot 2^\omega - 2)$ .*

*Proof.* Consider a path  $p_\omega$  that starts at  $x \notin p_i$  and intersects another path  $p_i$  at  $y \in p_i$ . Since  $p_\omega$  is a pruned path, its distance from  $x$  to the intersection point  $y$  is at most  $2^\omega - 1$  (if it was  $2^\omega$  or more, point  $y$  would have been its original leader). `ModifyPath` will attempt to seek the nearest  $(\omega + 1)$ -level node (pseudo-leader) on  $p_i$  from  $y$  (Lemma A.1). Note that  $y$  itself cannot be the pseudo-leader for  $x$  because, if it was, then,  $p_\omega$  would not have been a pruned path. The distance from  $y$  to a pseudo-leader  $v$  on  $p_i$  would be at most  $2^{\omega+1} - 1$  because if this distance was more than  $2^{\omega+1} - 1$ , we would have found another pseudo-leader  $v'$  that is  $2^{\omega+1}$  distance away from  $v$  and closer to  $y$ . This is due to the presence of  $(2^{\omega+1})$ -independent set nodes on this path  $p_i$  computed by `AssignLevels`. Note that  $y$  cannot be an end-node of  $p_i$  and  $v$  could be one of the end-nodes of  $p_i$ . Hence, the length of  $p'_\omega$  could be at most  $2^\omega - 1 + 2^{\omega+1} - 1 = 3 \cdot 2^\omega - 2$ . Note that  $p_i$  itself could be a stretched pruned path and the upper bound holds irrespective of the length of  $p_i$ .  $\square$

Figure 1 gives an example of intersecting path and its modification to reach a pseudo-leader and form a modified path.

Consider an arbitrary node  $x \in G$  with its neighborhood  $N_G(x, 2^r)$ , where  $r \geq 0$ .

**Lemma A.3 (Max path segments).** *The total number of path segments  $p \in T$  at level  $i$  or higher that cross  $N_G(x, 2^r)$  is at most  $2^{\rho(r-i+3)} \cdot (\kappa - i + 1)$ .*

---

**Function** ModifyPath( $p_m, p_i$ )

---

**Input:** Paths  $p_i$  and  $p_m$  where  $p_m$  intersects  $p_i$  and  $i > m$

**Output:** A modified path  $\bar{p}_m$ .

// Let  $p_m$  start from  $x \notin p_i$  and intersect at  $y \in p_i$  along its path to leader

$\ell_{m+1}$ .

$v_{m+1} \leftarrow$  From  $y$ , identify the nearest level- $(m+1)$  node  $v \in p_i$ ;

$p_m^a \leftarrow$  subpath from  $x$  to  $y$  in  $p_m$ ;

$p_m^b \leftarrow$  subpath from  $y$  to  $v_{m+1}$  in  $p_i$ ;

$\bar{p}_m \leftarrow p_m^a + p_m^b$ ;

// Concatenate  $p_m^a$  and  $p_m^b$ .

**return**  $\bar{p}_m$ ;

---

---

**Function** AssignLevels( $p_i, H, i$ )

---

**Input:** Path  $p_i$ , set of end-nodes  $H$  of  $p_i$ , level  $i$ .

**Output:** Assignment of levels to all nodes in  $p_i$ .

$L_\lambda \leftarrow \phi$ ;

// Set of  $2^\lambda$ -independent nodes

**for**  $\lambda \leftarrow (i-1)$  **to** 0 **do**

    // Find  $2^\lambda$ -independent nodes at levels  $\lambda = (i-1), (i-2), \dots, 1, 0$ .

$L_\lambda \leftarrow MIS(p_i, H, 2^\lambda)$ ;

    Assign level  $\lambda$  to nodes in  $L_\lambda$ .

---

*Proof.* We know, by construction, that the length of a path  $p_{i+j} \in T$  is at most  $2^{i+j}$  where  $0 \leq j \leq (\kappa - i)$  and that there is at most one leader  $\ell_{i+j} \in I_i$  within  $N(x, 2^{i+j}/2)$ . Since we are looking at the number of path segments  $p_{i+j}$  that go through  $N(x, 2^r)$ , consider a large neighborhood  $N(x, (2^{i+j} + 2^r))$  and determine the number of neighborhoods of radius  $2^{i+j}/2$ ;  $N(x, 2^{i+j}/2)$ . This is equivalent to at most  $2^{\rho \log \frac{(2^{i+j} + 2^r)}{2^{i+j}/2}} \leq 2^{\rho \log(r-i+3)} = 2^{\rho(r-i+3)}$  because of doubling dimension. For all paths that span the levels from  $i$  to  $\kappa$ , the total number of path segments that cross  $N(x, 2^{i+j}/2)$  is equal to  $2^{\rho(r-i+3)} \cdot (\kappa - i + 1)$ .  $\square$

**Lemma A.4 (Max modified paths in a path segment).** *Consider a path segment  $p \in T$  that crosses  $N_G(x, 2^r)$ . The total number of modified paths  $\bar{p} \in \bar{T}$  at level  $i$  or higher that use nodes in  $p \cap N_G(x, 2^r)$  is at most  $2^{\rho(r-i+1)} \cdot (\kappa - i + 1)$ .*

*Proof.* Let  $Q = p \cap N_G(x, 2^r)$ . From Lemma A.2, we know that the maximum length of any modified path  $\bar{p}_{i+j}$  would be  $3 \cdot 2^{i+j} - 2 < 4 \cdot 2^{i+j} = 2^{i+j+2}$ . To find the total number of modified paths  $\bar{p}_{i+j}$  that passes through  $Q$ , we consider a larger neighborhood  $N(x, 2^{i+j+2} + 2^r)$  and find the number of  $N(y, 2^{\frac{i+j+2}{2}})$  that would cover the larger neighborhood. Since each  $\bar{p}_{i+j}$  has start node in  $I_{i+j}$  and by doubling property of the graph, it can be computed to  $2^{\rho \log \frac{2^{i+j+2} + 2^r}{2^{\frac{i+j+2}{2}}}} \leq 2^{\rho(r-i+1)}$ . Since  $j \in [0, (\kappa - i)]$ , the total number of paths would be  $2^{\rho(r-i+1)} \cdot (\kappa - i + 1)$ .  $\square$




Figure 2: Fusion Cost for varying set of source nodes in a 1600-node grid.