Minimum Latency Aggregation Scheduling for Arbitrary Tree Topologies under the SINR Model

Guanyu Wang, Qiang-Sheng Hua, and Yuexuan Wang

Institute for Interdisciplinary Information Science, Tsinghua University, China
{wang-gy10,qshua,wangyuexuan}@mail.tsinghua.edu.cn

Abstract. Almost all the existing wireless data aggregation approaches need a topology construction step before scheduling. These solutions assume the availability of flexible topology controls. However, in real scenarios, lots of factors (impenetrable obstacles, barriers, etc.) limit the topology construction for wireless networks. In this paper we study a new problem called Minimum-Latency Aggregation Scheduling for Arbitrary Tree Topologies (MLAT). We first provide an NP-hardness proof for MLAT. Second, we draw an important conclusion that two frequently used greedy scheduling algorithms result in a large overhead compared with the optimal solution: the scheduling latency generated by these two greedy solutions are $\sqrt{n}$ times the optimal result, where n is the total number of links. We finally present an approximation algorithm for MLAT which works well for the tree with a small depth. All the above results are based on the SINR (Signal-to-Interference-plus-Noise Ratio) model.

1 Introduction

Data aggregation is a fundamental operation in wireless sensor networks. Given a set of sensor nodes distributed on the Euclidean plane, the data aggregation problem is to compute an aggregate function (e.g. a maximum or average function) on the data from all nodes in the wireless sensor network, and let the final aggregated value to be sent to a sink node in the fewest timeslots. To solve the data aggregation problem, also called as the $MLAS$ (Minimum-Latency Aggregation Scheduling) problem in the literature, the interference models employed will play an important role. Compared with exceedingly simplified graph based models or the protocol models used in many previous studies of data aggregation [1,7,17,19], a more realistic SINR (Signal-to-Interference-plus-Noise-Ratio) interference model [3] has been widely adopted in the community [12,11,10,4]. The SINR model is also called the physical model since it reflects the physical reality more accurately. The advantages and robustness of the SINR model are analyzed in [14]. In this paper, we employ the SINR model to study the data aggregation problem for arbitrary tree topologies.

For the MLAS problem, the best result to date under the SINR model is $O(\log n)$ given by Halldórsson et al [4]. There is a hardness result for the MLAS

X.-Y. Li, S. Papavassiliou, S. Ruehrup (Eds.): ADHOC-NOW 2012, LNCS 7363, pp. 139–152, 2012.

with uniform power assignment in [10]. The data aggregation problem for arbitrary directed acyclic networks under the SINR model is also studied in [6]. In this paper, the authors first show that this problem is NP-hard and give both heuristic and approximation algorithms. Note that the NP-hardness result for the directed acyclic networks may not mean the same hardness result for the tree topologies and the latter is the most frequently used topology for data aggregation. In addition, compared with the directed acyclic networks [6], one may get better scheduling results for the restricted tree topologies.

Also by using the SINR model, some other related wireless scheduling problems have been studied in the literature. Moscibroda et al. in 2006 [14] first initiated the connectivity scheduling problem (to construct a spanning tree over a set of sensor nodes on the plane in the fewest number of timeslots). This kind of connectivity scheduling problem has been further studied and better results have been proposed in [15,13,16]. The NP-hardness of the One-Shot scheduling problem (to pick the maximum number of links to be scheduled in the same timeslot) with uniform power (all the nodes take the same power) was proposed by Goussevskaia et al.[16]. This result was extended to the non-uniform power version later [8]. Very recently, some further hardness results have been given: Halldórsson and Wattenhofer [5] proved that One-Shot scheduling with uniform power assignment is in APX (the set of NP optimization problems that allow constant-factor approximation algorithms) and Kesselheim [9] extended the result to the power control version.

Note that, the typical solution for MLAS involves the construction of an appropriate data aggregation tree, followed by scheduling its transmission links. For example, the nearest neighbor tree is one of the widely used topologies. However, in a dynamic physical environment, it can not be guaranteed that any two nearest neighbors can communicate with each other successfully. This problem arises if there are obstacles or barriers restricting the kinds of links that can be formed between nodes that could otherwise be within communication range. Based on this observation, we study the Minimum Latency Aggregation Scheduling for Arbitrary Tree Topologies ($MLAT$). The only difference between MLAT and MLAS is that the tree topology is given in advance for the MLAT problem instead of first constructing a tree in the MLAS problem.

1.1 Formal Description of the MLAT Problem

We are given a tree consisting of nodes $V = \{v_0, v_1, v_2, \ldots, v_n\}$ with root v_0. We divide time into timeslots, defined to be the unit of time required to transmit once for any link. All the nodes are arbitrarily distributed in the Euclidean plane and can be both a sender and receiver, but only in different timeslots. The distance between any two nodes v_i, v_j is denoted by $d(v_i, v_j)$. Each edge $l_{ij} = (v_i, v_j)$ represents a communication request from a sender v_i to a receiver v_j. The length of link l_{ij} is denoted by $d_{ij} = d(v_i, v_j)$, where $d_{gj} = d(v_g, v_j)$ denotes the distance between the sender of link l_{gh} and receiver of link l_{ij}.

Formally, the SINR model is defined as follows. The signal power $P_i(j)$ received at v_j from sender v_i depends on the transmission power P_{ij} of v_i and the distance d_{ij}. The path loss radio propagation model for the reception of signals says the signal strength that v_j receives degrades at $d_{ij}^{-\alpha}$ (α denotes the path-loss exponent, and is usually a constant between 2 and 6), i.e. $P_i(j) = P_{ij}/d_{ij}^{\alpha}$. Every sender v_g (with corresponding receiver v_h) that sends concurrently with v_i causes an interference $I_g(j) = P_g(j) = P_{gh}/d_{gj}^{\alpha}$ at receiver v_j. All interferences accumulate. The total interference $I(v_j)$ experienced by receiver j is given as the sum of all interferences caused by other concurrently sending nodes, i.e. $I(v_j) = \sum_{l_{gh} \neq l_{ij}} I_g(j)$. A receiver v_j successfully receives a message from its sender v_i if and only if it obeys the *precedence* constraint (a node cannot send its data to the parent node until it has received data from *all* children nodes) and the following SINR threshold holds:

$$SINR_S(v_j) = \frac{P_i(j)}{\sum_{l_{gh} \in S \setminus l_{ij}} I_g(j) + N} \geq \beta$$

where N is ambient noise, $\beta \geq 1$ denotes the minimum SINR required for a message to be successfully received, and S is the set of concurrently transmitting links. We call the set SINR-feasible set. Denote all edges of the given tree as $E = \{l_1, l_2, \ldots, l_n\}$ (for notational simplicity, we omit the sender and receiver suffix here), we strive to find a sequence of t sets, i.e. a schedule: $S = \{L_1, L_2, \ldots, L_t\}$, $L_1 \cup L_2 \cup \cdots \cup L_t = E$ and $L_i \cap L_j = \emptyset$, $\forall i, j \in [t]$ $i \neq j$, such that:

$$S = \underset{S' = \{L_1, L_2, \ldots, L_t\}}{\text{argmin}} \quad t$$

$$h > g \qquad \forall i, j, k \quad l_{ij} \in L_g \ and \ l_{jk} \in L_h,$$

$$SINR_{L_m}(v_j) \geq \beta, \ \forall m \ \forall l_{ij} \in L_m$$

1.2 Results

In Section 2, we will present the first NP-hardness proof for MLAT under two real conditions: (i) ambient noise exists, and (ii) all nodes have limited power ranges. We then analyze the gap between local optimal solution and global optimal solution for the MLAT Problem. Section 3 provides the evidence of that most local greedy approaches of existing aggregation scheduling algorithms perform poorly compared with the optimal result. The timeslots (scheduling latency) needed by the local optimal methods could be $\sqrt{n}$ times larger than the global optimal solution, where n is the total number of wireless links. We then present an approximation algorithm for MLAT in Section 4, which adopts an existing strategy of iteratively maximizing concurrently transmitting links. We derive the exact approximation ratio bounded by $O(\min\{d \cdot \log n, n/d\})$, where d is the depth of the given tree. Even though the greedy approaches have been proved to perform poorly, this analysis shows that it still has guaranteed efficiency for the tree with a small depth.

2 NP-Hardness Proof for MLAT

In this section, we will show the following NP-hardness result.

Theorem 1. *MLAT with given power range and nonzero background noise is NP-hard.*

Proof. Let $P_{v_i} \in [P_{min}, P_{max}]$ be the transmission power assigned to every sender v_i, and let $N > 0$.

We will give a polynomial time reduction that expands on methods used in [16,8] from the Partition Problem to the decision version of MLAT, when one must decide whether there exists a schedule of a specified length for a given aggregation tree.

Partition Problem: Do there exist sets $\mathcal{I}_1$, $\mathcal{I}_2 \subset \mathcal{I}$ where $\mathcal{I} = \{i_1, i_2, \ldots, i_n\}$ is a set of integers s.t.

$$\mathcal{I}_1 \cup \mathcal{I}_2 = \mathcal{I}, \mathcal{I}_1 \cap \mathcal{I}_2 = \emptyset,$$

$$\sum_{i_j \in \mathcal{I}_1} i_j = \sum_{i_j \in \mathcal{I}_2} i_j = \frac{1}{2} \sum_{i_j \in \mathcal{I}} i_j = \frac{1}{2}\sigma.$$

This problem was proved to be NP-complete by Karp [2]. We construct a many-to-one reduction from an arbitrary Partition Problem instance to an instance of MLAT. We will argue that the instance of MLAT can be scheduled in $T \leq n+3$ timeslots if and only if the reduced Partition Problem instance can be solved.

Lemma 1. *The Partition Problem can be reduced to MLAT in polynomial time.*

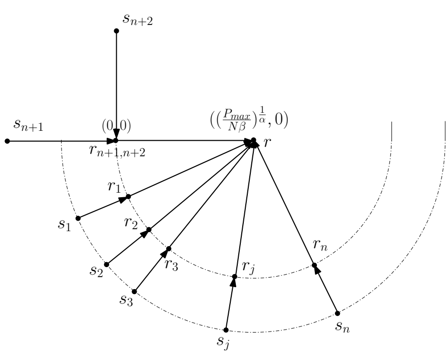

Fig. 1. Example of constructed instance of MLAT from Partition Problem

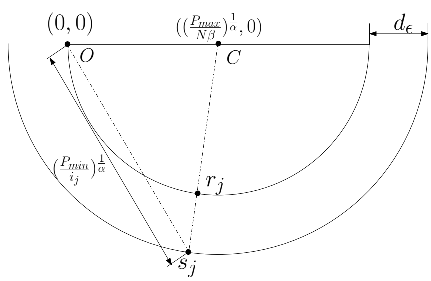

Fig. 2. The semicircle in the plane

Without loss of generality, we assume all elements in the Partition Problem instance $\mathcal{I} = \{i_1, i_2, \ldots, i_n\}$ to be distinct and positive. Next, we construct an instance of MLAT with $2n+3$ links $L = \{l_1, l_2, \ldots, l_{2n+3}\}$ (cf. Fig. 1). We define the sender and receiver of link l_i as s_i and r_i, respectively.

To begin, we must scale all the integers in $\mathcal{I}$ by the same factor k such that $k \cdot i_{min} \geq \frac{N\beta P_{min}}{2^\alpha P_{max}}$, where the i_{min} and i_{max} (used in the definition of d_ϵ below) are the smallest and largest integers in $\mathcal{I}$, respectively. In the following we regard the set $\mathcal{I}$ to be properly scaled.

We assign every node a position in the Euclidean plane. First, we fix a semicircle of radius $R_1 = (P_{max}/(N\beta))^{1/\alpha}$ to the plane centered at some point C, followed by another semicircle with radius $R_2 = R_1 + d_\epsilon$ also centered at C, d_ϵ a small constant. Let I_{min} be defined as:

$$I_{min} = \min_{i_g, i_h \in \mathcal{I}, i_g \neq i_h} |(\frac{P_{min}}{i_h})^{\frac{1}{\alpha}} - (\frac{P_{min}}{i_g})^{\frac{1}{\alpha}}|.$$

We will show that for any small $\epsilon > 0$,

$$d_\epsilon = \min\{\frac{I_{min}}{(\frac{(1+\epsilon)n\beta P_{max}}{\epsilon P_{min}})^{\frac{1}{\alpha}} + 1}, (\frac{P_{min}}{N\beta(1+\epsilon)})^{\frac{1}{\alpha}}, (\frac{P_{min}}{i_{max}})^{\frac{1}{\alpha}}\}$$

is sufficiently small for our reduction.

For each integer $i_j \in \mathcal{I}$, we place sender s_j on the larger semicircle such that its distance from the leftmost point of the smaller semicircle (origin O) is $(P_{min}/i_j)^{1/\alpha}$ (cf. Fig.2). Because of our scaling of $\mathcal{I}$ and choice of d_ϵ, such a point will always exist.

$$d(s_j, O) = (\frac{P_{min}}{i_j})^{\frac{1}{\alpha}} \quad \forall 1 \leq j \leq n$$

Next, we designate the position for every receiver r_j, $1 \leq j \leq n$ to be the intersecting point on the smaller semicircle of the line which passes through both s_j and C. Note that the distance between any pair s_j, r_j, $1 \leq j \leq n$ is always d_ϵ.

Finally we place four nodes s_{n+1}, s_{n+2}, $r_{n+1,n+2}$ and r. Note that $r_{n+1,n+2}$ is the receiver corresponding to senders s_{n+1} and s_{n+2}. Node r is the parent node for all receivers $r_1, r_2, \ldots, r_n, r_{n+1,n+2}$ in the tree (cf. Fig. 1).

$$pos(s_{n+1}) = (-(\frac{P_{max}}{\beta(N+\frac{\sigma}{2})})^{\frac{1}{\alpha}}, 0), \quad pos(r_{n+1,n+2}) = (0,0)$$

$$pos(s_{n+2}) = (0, (\frac{P_{max}}{\beta(N+\frac{\sigma}{2})})^{\frac{1}{\alpha}}), \quad pos(r) = ((\frac{P_{max}}{N\beta})^{1/\alpha}, 0)$$

Next $n+3$ links of the tree are constructed as follows:

$$l_{n+1} = (s_{n+1}, r_{n+1,n+2}), \quad l_{n+2} = (s_{n+2}, r_{n+1,n+2})$$

$$l_{n+3} = (r_{n+1,n+2}, r), \quad l_{n+(i+3)} = (r_i, r) \;\; 1 \leq i \leq n$$

We then prove four properties of this tree:

1. l_{n+1}, l_{n+2} must transmit in different timeslots.
2. All l_i, $1 \leq i \leq n$, and one of l_{n+1}, l_{n+2} can transmit concurrently (i.e. in the same timeslot).
3. l_{n+1}, l_{n+2} can transmit successfully if and only if the total interference from other senders is not greater than $\sigma/2$.
4. l_j, $n+3 \leq j \leq 2n+3$, can only transmit alone.

The first property arises directly from the observation that l_{n+1} and l_{n+2} have a common receiver $r_{n+1,n+2}$. If they transmit in the same timeslot, there is no possibility that both of their SINR is greater than β.

The second property can be derived from the following lemma:

Lemma 2. *Every transmission $l_i \in L' = \{l_1, l_2, \ldots, l_n\}$ is successful using transmission power P_{min}, no matter how many other links $l_j \in L'$ along with either l_{n+1} or l_{n+2} transmit concurrently, even if all transmitting links, except for l_i, use power P_{max}.*

Proof. For links in L', the worst case scenario is that all senders s_i, $1 \leq i \leq n$, and either s_{n+1} or s_{n+2} transmit concurrently with P_{max}. Recall that we have chosen a very small d_ϵ (i.e. the distance between s_i and r_i, $1 \leq i \leq n$). It is easy to see that $d(s_{n+1}, r_{n+1,n+2}) = d(s_{n+2}, r_{n+1,n+2}) \geq d_\epsilon$. We can bound the distance between r_i and s_j, $1 \leq i \leq n, 1 \leq j \leq n+2, i \neq j$.

$$d(r_i, s_j) \geq d(s_j, s_i) - d(s_i, r_i) \geq |d(s_j, O) - d(s_i, O)| - d_\epsilon \tag{1}$$

$$\geq I_{min} - d_\epsilon \geq ((\frac{(1+\epsilon)n\beta P_{max}}{\epsilon P_{min}})^{\frac{1}{\alpha}} + 1 - 1)d_\epsilon \tag{2}$$

$$= (\frac{(1+\epsilon)n\beta P_{max}}{\epsilon P_{min}})^{\frac{1}{\alpha}} d_\epsilon \tag{3}$$

The first inequality follows from two triangle inequalities (cf. Fig.3). The second inequality follows from the definition of I_{min} and d_ϵ. Thus, we can derive an SINR lower bound for all receivers r_i, $1 \leq i \leq n$:

$$SINR(r_i) = \frac{\frac{P_{s_i}}{d(s_i, r_i)^\alpha}}{N + \sum_{l_j \in L'/l_i \cup \{l_{n+1} \text{ or } l_{n+2}\}} \frac{P_{s_j}}{d(r_i, s_j)^\alpha}}$$

$$\geq \frac{\frac{P_{min}}{d_\epsilon^\alpha}}{N + \frac{nP_{max}}{d(s_j, r_i)^\alpha}} = \frac{\frac{P_{min}}{d_\epsilon^\alpha}}{N + \frac{\epsilon P_{min}}{(1+\epsilon)d_\epsilon^\alpha \beta}}.$$

On the other hand, according to the value of d_ϵ, we know $d_\epsilon \leq (P_{min}/(N\beta(1+\epsilon)))^{1/\alpha}$. So we obtain $N \leq \frac{P_{min}}{(1+\epsilon)d_\epsilon^\alpha \beta}$.

Combining these, we get

$$SINR(r_i) \geq \frac{\frac{P_{min}}{d_\epsilon^\alpha}}{(\frac{P_{min}}{(1+\epsilon)d_\epsilon^\alpha \beta}) + \frac{\epsilon P_{min}}{(1+\epsilon)d_\epsilon^\alpha \beta}} = \beta.$$

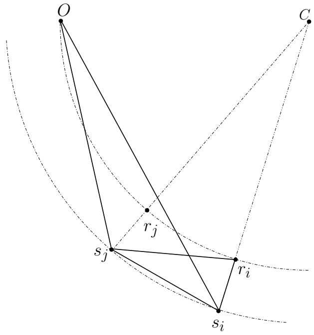

Fig. 3. The distance between r_i and s_j

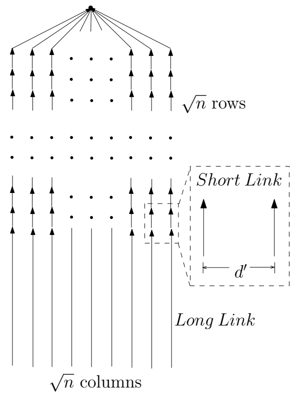

Fig. 4. Counterexample for the Leaf-First approach

The third property can be derived from the total interference suffered by $r_{n+1,n+2}$ from s_k, $1 \leq k \leq n$. When these senders transmit with minimum power P_{min},

$$I_{r_{n+1,n+2}}(s_k) = \frac{P_{min}}{((\frac{P_{min}}{i_k})^{\frac{1}{\alpha}})^{\alpha}} = i_k.$$

However, even if s_{n+1} uses transmission power P_{max}, we have:

$$P_{s_{n+1}}(r_{n+1,n+2}) = \frac{P_{max}}{((\frac{P_{max}}{\beta(N+\frac{\sigma}{2})})^{\frac{1}{\alpha}})^{\alpha}} = \beta(N + \frac{\sigma}{2}).$$

If we want s_{n+1} to transmit successfully, the following inequality must hold:

$$\frac{P_{s_{n+1}}(r_{n+1,n+2})}{N+I} = \frac{\beta(N+\frac{\sigma}{2})}{N+I} \geq \beta.$$

It is easy to see that if s_{n+1} transmits using a smaller power or if other senders transmit using larger powers, then the SINR balance will be destroyed. The same analysis also holds for s_{n+2}. Thus the following lemma can be derived from these three properties:

Lemma 3. *There exists a 2-slot schedule for all links in $L'' = \{l_1, l_2, \ldots, l_{n+2}\}$ if and only if there is a solution to instance $\mathcal{I}$ of the Partition Problem.*

Proof. By the second property, we only need to consider l_{n+1}, l_{n+2}.

If $\{\mathcal{I}_1, \mathcal{I}_2\}$ is a solution to $\mathcal{I}$, then $\sum_{i_j \in \mathcal{I}_1} i_j = \sum_{i_k \in \mathcal{I}_2} i_k = \sigma/2$. This means we can let l_{n+1} and all l_j, $\forall i_j \in \mathcal{I}_1$ transmit concurrently in the first timeslot, and l_{n+2}, l_k, $\forall i_k \in \mathcal{I}_2$ transmit in the second timeslot. The correctness of this schedule is guaranteed by the third property.

From the first property, if there is a 2-slot schedule for L'', l_{n+1}, l_{n+2} must transmit in different timeslots. Without loss of generality, we assume l_{n+1} transmits in the first timeslot, l_{n+2} in the second. Let L_1 and L_2 be the sets of links transmitting in first and second timeslot, respectively. According to the third property, $\sum_{l_j \in L_1} i_j \leq \sigma/2$, and $\sum_{l_k \in L_2} i_k \leq \sigma/2$. However, we already know $\sum_{l_j \in L_1 \cup L_2} i_j = \sigma$. So the following equation holds:

$$\sum_{l_j \in L_1} i_j = \sum_{l_k \in L_2} i_k = \sigma/2$$

which means we have a solution for the Partition Problem instance $\mathcal{I}$.

The fourth property follows naturally. Since the lengths of all l_j, $n+3 \leq j \leq 2n+3$ are $(P_{max}/(N\beta))^{1/\alpha}$, i.e., the radius of the smaller semicircle, receivers $r_j \in \{r_1, r_2, \ldots, r_n, r_{n+1,n+2}\}$ become senders with transmission power P_{max}. Then $P_{r_j}(r) = (P_{max})/((P_{max}/N\beta)^{1/\alpha})^{\alpha} = N\beta$, which means any other additional interference will make $SINR_{r_j}(r)$ fall below the threshold β, i.e., l_j, $n+3 \leq j \leq 2n+3$, can only transmit alone.

Combining all four properties, we conclude that the constructed instance of MLAT can be scheduled in $T \leq n+3$ timeslots if and only if the reduced Partition Problem instance can be solved. Therefore, if we have a polynomial time algorithm A for MLAT, then we may also solve the Partition Problem using A as a subroutine in polynomial time.

3 Gap between the Local and Global Optimization

In this section, we will show that two greedy approaches (layer-first and leaf-first) result in very poor schedules: the scheduling latencies generated by greedy solutions could be $\sqrt{n}$ times the optimal result, where n is the total number of links. Note that most existing data aggregation scheduling algorithms use the greedy ideas after the topology construction step, even the best $O(\log n)$ result for the MLAS problem [4]. This may give some hint that using an appropriate topology construction algorithm could help reducing the scheduling latency for the data aggregation problem.

3.1 Leaf-First Method

Assume we have a black box which can find the maximum size set of concurrently transmitting links, from all the given links. This black box is used for *greedily select* operation in this section. We want to show that even we can find the local optimal concurrent transmissions using this black box, it still leads to a very poor performance compared with the global optimal solution in the worst case.

Definition 1. *For any given data aggregation tree defined in Section 1.1, in any round, greedily select the leaves of the tree (i.e. choose the maximum number of links that can transmit simultaneously) at the beginning of that round to transmit, without violating the SINR threshold. This approach is called Leaf-First.*

Theorem 2. *Given a data aggregation tree with n links, assume the output of the Leaf-First approach is the schedule S_{leaf} and the minimum-latency schedule of this tree is S_{opt}. In the worst case, $|S_{leaf}| = \Omega(\sqrt{n})|S_{opt}|$.*

Proof. Construct the data aggregation tree with $\sqrt{n}$ layers and every layer consists of $\sqrt{n}$ links, as shown in Fig.4. Except for the links in the top layer, this tree has only two types of links: *long link* and *short link*. The long links only appear in the deepest layer and their lengths are all $d_{long} = (\frac{P_{max}}{N\beta})^{1/\alpha}$. It is easy to show that every long link can only transmit alone since $\frac{P_{max}/d_{long}^{\alpha}}{N} = \beta$, i.e. any additional interference from other senders let the transmission of a long link fail. The remaining links are short links which is very short compared with distance (denoted as d' in the figure) between any two short links such that all the short links in the same layer can transmit concurrently. Simply set their length as $d_{short} = (\frac{P_{min}}{\beta(\sqrt{n}P_{max}/d'^{\alpha}+N)})^{\frac{1}{\alpha}}$, then the required property above for short link holds.

After the construction, we apply the Leaf-First approach on this data aggregation tee. Its performance is shown in Fig.5. Obviously, the Leaf-First approach needs $\Theta(n)$ timeslots in total to finish the aggregation. However, a much better aggregation should be finishing all the long links one by one and then short links layer by layer, which results in a $2\sqrt{n} - 2$ time-slot schedule, and all the links in the top layer still needs extra $\sqrt{n}$ timeslots. Therefore, $3\sqrt{n} - 2$ timeslots is enough to finish the aggregation. So we get $|S_{leaf}| = \Omega(\sqrt{n})|S_{opt}|$.

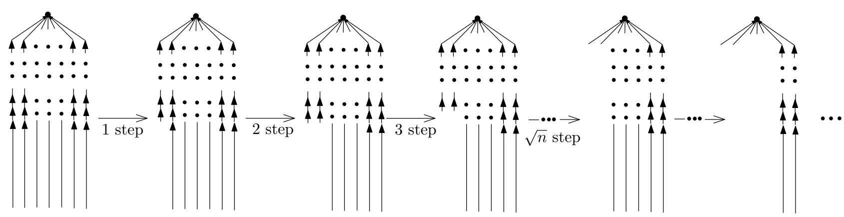

Fig. 5. Performance of the Leaf-First approach on counterexample

3.2 Layer-First Approach

For simplicity, we assume $P_{max} = P_{min} = P$ and $\beta = 1$ in this part, i.e., adopting the special case –uniform power assignment– in the following analysis. In addition, the ambient noise is so small compared with P that it can be ignored. We give the formal definition of the Layer-First approach like above.

Definition 2. *For any given data aggregation tree defined in Section 1.1, in any round, greedily select the wireless links belong to the deeper layers (i.e. transmit the deepest links as many as possible, then the second deepest ones, and so on) at the beginning of that round to transmit, without violating the SINR threshold. This approach is called Layer-First.*

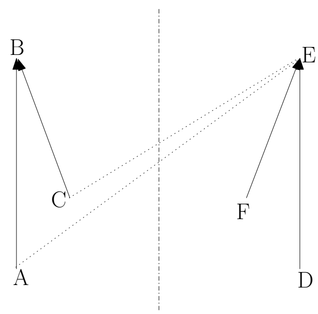

Fig. 6. Example of counter-block links

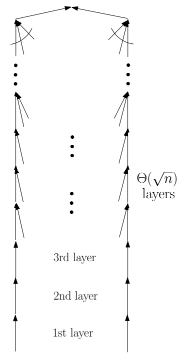

Fig. 7. Counterexample for the Layer-First approach

Theorem 3. *Given a data aggregation tree with n links, assume the output of the Layer-First approach is the schedule S_{layer} and the minimum-latency schedule of this tree is S_{opt}. In the worst case, $|S_{layer}| = \Omega(\sqrt{n})|S_{opt}|$.*

Before we start the proof, we need to define a special pair of links called *counter-block links.*

Definition 3. *Two links are called counter-block links if and only if they are axially symmetric (as l_{AB}, l_{DE} and l_{CB}, l_{FE} shown in Fig.6) and can transmit concurrently without any additional interference (i.e. the existence of any other transmission makes their SINRs lower than the threshold).*

Proof. As shown in Fig.6, we just need to choose appropriate $||AB||$ and $||AE||$ such that $\frac{P/||AB||^\alpha}{P/||AE||^\alpha} = \beta$ then l_{AB} and l_{DE} are counter-block links. It is easy to show that if $||CB||$ is a little shorter than $||AB||$, we can also set $||CE||$ to make sure that l_{CB} and l_{FE} are also counter-block links. Obviously, this method allows us to have infinitely many pairs of counter-block links with two common receivers. Using this technique, we construct the data aggregation tree which is shown in Fig.7, whose ith layer has $\lceil i/3 \rceil$ pairs of counter-block links (except for the top layer). Assume this tree has n links in total and m layers, then $3(m-1)m/2 = (n-2)/2$, i.e., the number of layers $m = \Theta(\sqrt{n})$.

Still, we need to apply the Layer-First approach on this data aggregation tree. For every layer, the maximum number of concurrently transmitting links is two due to the fact that there are only two receivers. One pair of counter-block links can be selected to transmit any timeslot by the Layer-First approach, so the whole aggregation needs $(n-2)/2+2$ timeslots. Next, we will show a more efficient scheduling which needs just $m = \Theta(\sqrt{n})$ timeslots.

Except for the top layer, we separate the tree into left part and right part. For the left (or the right) part, in the ith timeslot, one link in the $(3k+i-1)$th layer, $k = 1, 2, \cdots$, can transmit. In other words, for one single side, links are chosen at intervals of three layers in one timeslot (no vertical links can be chosen except for the deepest layer). So after two timeslots (one for each side), the 1st layer of the tree finishes transmission. At the same time, the 4th layer leaves one vertical link for each side, which means after eight timeslots, all the deepest four layer are removed from the tree after transmission, and so on. Finally, this scheduling can finish aggregation of this tree with $2m$ timeslots. We show $|S_{leaf}| = \Omega(\sqrt{n})|S_{opt}|$ by plugging in the definition of m. Next lemma explains that this schedule satisfies the SINR constraint.

Lemma 4. *For a set of links in one side of the tree constructed above, it is SINR-feasible if any two links in this set are at least three layers far away from each other.*

Proof. Assume this SINR-feasible set is L. For any link l_0 in L, there are at most two links (denoted as l_3 from higher layer and l_{-3} from lower layer) which are three layers far away from it. The distance between the sender of l_3 (or l_{-3}) and the receiver of l_0 is at least $2||l_0||$ (or $4||l_0||$). Similarly, there are two links l_6, l_{-6} which are six layers far away from l_0. Corresponding distances between senders and l_0's receiver are at least $5||l_0||$ and $7||l_0||$, and so on. Since the total number of links in one side of the tree is bounded by $n/2$, the SINR of l_0 can be easily derived:

$$SINR(l_0) \geq \frac{\frac{P}{||l_0||^\alpha}}{\frac{P}{2||l_0||)^\alpha} + \frac{P}{4||l_0||)^\alpha} + \cdots + \frac{P}{(2+3(\frac{n}{2}-1))||l_0||)^\alpha} + \frac{P}{(4+3(\frac{n}{2}-1))||l_0||)^\alpha}}$$

$$> \frac{\frac{P}{||l_0||^\alpha}}{2\sum_{k=1}^{n/2} \frac{P}{(k||l_0||)^\alpha}} > \frac{1}{2\sum_{k=1}^{\infty} \frac{1}{k^\alpha}} > \frac{1}{2\frac{\alpha}{\alpha-1}} > 1$$

where the second-to-last inequality follows the Riemann's zeta-function and the last one based on the fact that $\alpha \in [2, 6]$.

4 Approximation Algorithm for MLAT

In this section, we describe a greedy algorithm that solves MLAT, using existing techniques for the Wireless Capacity Maximization Problem (the same as the One-Shot Scheduling problem) [9,18] ([9] needs much larger P_{max}). Our algorithm is performed in a layer-by-layer style. Note that even though the greedy approaches have been proved to perform poorly without appropriate topology control, we show that we can still acccomplish an acceptable approximation ratio when data aggregation trees have small but reasonable depths.

The basic assumption for this section:

– The maximum power P_{max} for all senders is large enough that every link can transmit successfully in some case: $P_{max}/(Nd_{max}^\alpha) \geq \beta$, where d_{max} is the length of the longest edge of the aggregation tree.

– All the edges in the aggregation tree have lengths greater than 1. Any instance can be transformed into this case by scaling.

According to the analysis in [9,18], the correctness of the SINR-feasible sets generated in line 18 of Algorithm 1 is guaranteed. The approximation ratio for maximizing concurrently transmitting links is constant, which leads to an approximation algorithm for minimum latency scheduling with an approximation ratio bounded by $O(\log n)$. Note that the algorithm in [9] needs a very large maximum transmitting power.

Algorithm 1 labels all the edges of the given tree by first using a depth first search approach and then finds the depth of this tree in lines 3-7. Simply, the currently deepest edges of the tree are selected as scheduling candidates in lines 11-15. Then the existing scheduling algorithm for maximizing concurrent transmissions is used to select an approximated maximum SINR-feasbile link set. It repeats this process until all the links have been scheduled.

Algorithm 1. The Layer-by-Layer Algorithm for MLAT

Input: An arbitrary aggregation tree $T = \{V, E\}$ and N, α, β, P_{min}, P_{max};
Output: A schedule S in which every edge can transmit successfully under SINR;
1: $S := \emptyset$, $depth := 0$, $t := 1$;
2: Use Depth First Search(DFS) to label every edge e with its layer, to be stored in $layer(e)$
3: **for** every edge e in T **do**
4:   **if** $depth < layer(e)$ **then**
5:     $depth := layer(e)$;
6:   **end if**
7: **end for**
8: $L := E$;
9: **while** $|L| > 0$ **do**
10:   $L' := \emptyset$;
11:   **for** every edge l_i in L **do**
12:     **if** $layer(l_i) = depth$ **then**
13:       $L' := L' \cup l_i$;
14:     **end if**
15:   **end for**
16:   $L := L \setminus L'$;
17:   **while** $|L'| > 0$ **do**
18:     Given N, α, β, P_{min}, P_{max}, use the constant-approximation algorithm for maximizing concurrent transmissions (please refer to [9] or [18]), to compute an approximate maximum SINR-feasible link set L'' in L';
19:     $S_t := L''$; $S := S \cup \{S_t\}$; $L' := L' \setminus S_t$; $t := t + 1$;
20:   **end while**
21:   $depth := depth - 1$;
22: **end while**
23: Return S;

In Algorithm 1, we divide the schedule generated by the algorithm into sub-schedules: $S = \{S_1, S_2, \ldots, S_d\}$. d is the depth of the input aggregation tree, and S_i is the sub-schedule in which only links in the ith layer can appear: $S_1 \cup S_2 \cup \cdots \cup S_d = S_{alg}$, and $S_i \cap S_j = \emptyset$, $\forall i, j \in [d], i \neq j$.

Define $T_{opt} = |S_{opt}|$, where S_{opt} is the optimal solution of MLAT. If we only schedule the links in the ith layer, the minimum number of timeslots T_i needed must not exceed T_{opt}. This is because if T_i is greater than T_{opt}, we can find a subschedule $S_{i,opt}$ in S_{opt} which schedules all links in the ith layer, i.e., $|S_{i,opt}| \leq T_{opt} < T_i$, which contradicts the fact that T_i is the minimum number of timeslots required to schedule these links. Thus, we have $|S_{1,opt}| + |S_{2,opt}| + \cdots + |S_{d,opt}| \leq d\dot{T}_{opt}$. We already know that every sub-schedule from our algorithm adheres to $|S_i| \leq \log n |S_{i,opt}|$, which implies that $T_{alg} \leq d \log n \cdot T_{opt}$.

On the other hand, it is obvious that T_{alg}, T_{opt}, n, d are all greater than 0 and $T_{opt} \geq d$, $T_{alg} \leq n$. Therefore, we derive another bound: $T_{alg} \leq (n/d) T_{opt}$.

From these, we bound the approximation ratio by $O(\min\{d \cdot \log n, n/d\})$. This gives an approximation ratio of $O(\log^2 n)$ when d is either very small ($d \leq O(\log n)$), or very large ($d \geq \Omega(n/\log^2 n)$). Fortunately, common aggregation trees often fall within these depth ranges. For example, applying our algorithm on a nearest neighbor tree (which has depth $O(\log n)$) leads to an $O(\log^2 n)$ approximation ratio.

5 Conclusion

In this paper, based on the fact that lots of factors (impenetrable obstacles, barriers, etc.) limit the topology construction for wireless networks in real scenarios, we introduce the Minimum Latency Aggregation Scheduling for Arbitrary Tree Topologies (MLAT) problem. We give the first NP-hardness proof for MLAT. In addition, we prove that the scheduling latencies generated by the two frequently used greedy algorithms could be $\sqrt{n}$ times the optimal result, where n is the total number of links. Given the fact that the MLAS problem could be solved in $O(\log n)$ timeslots [4] using the greedy scheduling approaches and topology control, the gaps we find for the MLAT problem show that giving another freedom of topology control could help reduce the aggregation latency. Finally, we propose an approximation algorithm for MLAT with approximation ratio bounded by $O(\min\{d \cdot \log n, n/d\})$ which is acceptable when the data aggregation trees have small depths. One of our future work is to give an approximation algorithm with a much better approximation ratio. Second, based on the preliminary observation that a clever mixture of different greedy approaches could lead to a better performance, we hope to devise a new heuristic algorithm combining both layer-first and leaf-first approaches. Another interesting extension to our work would be to design a distributed solution for MLAT.

Acknowledgment. The authors thank the anonymous reviewers for their helpful comments. This work was supported in part by the National Basic Research Program of China Grant 2011CBA00300, 2011CBA00302, the National Natural Science Foundation of China Grant 61103186, 61073174, 61033001, 61061130540, the Hi-Tech research and Development Program of China Grant 2006AA10Z216, and Hong Kong RGC-GRF grants 714009E and 714311.

References

1. Chen, X., Hu, X., Zhu, J.: Minimum Data Aggregation Time Problem in Wireless Sensor Networks. In: Jia, X., Wu, J., He, Y. (eds.) MSN 2005. LNCS, vol. 3794, pp. 133–142. Springer, Heidelberg (2005)
2. Fanghänel, A., Keßelheim, T., Vöcking, B.: Improved Algorithms for Latency Minimization in Wireless Networks. In: Albers, S., Marchetti-Spaccamela, A., Matias, Y., Nikoletseas, S., Thomas, W. (eds.) ICALP 2009. LNCS, vol. 5556, pp. 447–458. Springer, Heidelberg (2009)
3. Gupta, P., Kumar, P.R.: The capacity of wireless networks. IEEE Transactions on Information Theory 46(2), 388–404 (2000)
4. Halldórsson, M.M., Mitra, P.: Wireless connectivity and capacity. In: SODA (2012)
5. Halldórsson, M.M., Wattenhofer, R.: Wireless Communication Is in APX. In: Albers, S., Marchetti-Spaccamela, A., Matias, Y., Nikoletseas, S., Thomas, W. (eds.) ICALP 2009. LNCS, vol. 5555, pp. 525–536. Springer, Heidelberg (2009)
6. Hua, Q.-S., Wang, Y., Yu, D., Tan, H.: Minimum latency link scheduling for arbitrary directed acyclic networks under precedence and sinr constraints. Journal of Interconnection Networks 12(1-2), 85–107 (2011)
7. Huang, S.C.H., Wan, P.J., Vu, C.T., Li, Y., Yao, F.: Nearly constant approximation for data aggregation scheduling in wireless sensor networks. In: INFOCOM (2007)
8. Katz, B., Völker, M., Wagner, D.: Energy efficient scheduling with power control for wireless networks. In: WiOpt (2010)
9. Kesselheim, T.: A constant-factor approximation for wireless capacity maximization with power control in the sinr model. In: SODA (2011)
10. Lam, N.X., An, M.K., Huynh, D.T., Nguyen, T.N.: Minimum latency data aggregation in the physical interference model. In: MSWiM (2011)
11. Li, H., Hua, Q.-S., Wu, C., Lau, F.C.M.: Minimum-latency aggregation scheduling in wireless sensor networks under physical interference model. In: MSWiM (2010)
12. Li, X.Y., Xu, X.H., Wang, S.G., Tang, S.J., Dai, G.J., Zhao, J.Z., Qi, Y.: Efficient data aggregation in multi-hop wireless sensor networks under physical interference model. In: MASS (2009)
13. Moscibroda, T.: The worst-case capacity of wireless sensor networks. In: IPSN (2007)
14. Moscibroda, T., Wattenhofer, R.: The complexity of connectivity in wireless networks. In: INFOCOM (2006)
15. Moscibroda, T., Wattenhofer, R., Zollinger, A.: Topology control meets sinr: the scheduling complexity of arbitrary topologies. In: MOBIHOC (2006)
16. Anne, Y., Goussevskaia, O.O., Wattenhofer, R.: Complexity in geometric sinr. In: MOBIHOC (2007)
17. Wan, P.-J., Huang, S.C.-H., Wang, L.X., Wan, Z.Y., Jia, X.H.: Minimum-latency aggregation scheduling in multihop wireless networks. In: MOBIHOC (2009)
18. Wan, P.-J., Ma, C., Tang, S., Xu, B.: Maximizing Capacity with Power Control Under Physical Interference Model in Simplex Mode. In: Cheng, Y., Do Eun, Y., Qin, Z., Song, M., Xing, K. (eds.) WASA 2011. LNCS, vol. 6843, pp. 84–95. Springer, Heidelberg (2011)
19. Yu, B., Li, J., Li, Y.: Distributed data aggregation scheduling in wireless sensor networks. In: INFOCOM (2009)