



Distributed Online Data Aggregation in Dynamic Graphs

Quentin Bramas¹(✉), Toshimitsu Masuzawa³, and Sébastien Tixeuil²

¹ University of Strasbourg, ICUBE, CNRS, Strasbourg, France
bramas@unistra.fr

² Sorbonne University, LIP6, CNRS, Paris, France

³ Osaka University, Suita, Japan

Abstract. We consider the problem of aggregating data in a dynamic graph, that is, aggregating the data that originates from all nodes in the graph to a specific node, the sink. We are interested in giving lower bounds for this problem, under different kinds of adversaries.

In our model, nodes are endowed with unlimited memory and unlimited computational power. Yet, we assume that communications between nodes are carried out with pairwise interactions, where nodes can exchange control information before deciding whether they transmit their data or not, given that each node is allowed to transmit its data at most once. When a node receives a data from a neighbor, the node may aggregate it with its own data.

We consider three possible adversaries: the online adaptive adversary, the oblivious adversary, and the randomized adversary that chooses the pairwise interactions uniformly at random. For the online adaptive and the oblivious adversaries, we give impossibility results when nodes have no knowledge about the graph and are not aware of the future. Also, we give several tight bounds depending on the knowledge (be it topology related or time related) of the nodes. For the randomized adversary, we show that the Gathering algorithm, which always commands a node to transmit, is optimal if nodes have no knowledge at all. Also, we propose an algorithm called Waiting Greedy, where a node either waits or transmits depending on some parameter, that is optimal when each node knows its future pairwise interactions with the sink.

1 Introduction

Dynamic graphs, that is, graphs that evolve over time, can conveniently model dynamic networks, which recently received a lot of interest from the academic community (*e.g.* mobile sensor networks, vehicular networks, disruption tolerant networks, interaction flows, etc.). Depending on the problem considered, various

This work was performed within the Labex SMART supported by French state funds managed by the ANR within the Investissements d'Avenir program under reference ANR-11-IDEX-0004-02. **A preliminary 2 pages poster of this work appeared in IEEE ICDCS 2016 [4].**

© Springer Nature Switzerland AG 2019

M. F. Atig and A. A. Schwarzmann (Eds.): NETYS 2019, LNCS 11704, pp. 365–380, 2019.

https://doi.org/10.1007/978-3-030-31277-0_24

models were used: among others, static graphs can be used to represent a snapshot in time of a dynamic graph, functions can be used to define continuously when an edge appears over time, and sequences of tuples can represent atomic interactions between nodes over time.

The problem we consider in this paper assumes an arbitrary dynamic network, such as sensors deployed on a human body, cars evolving in a city that communicate with each other in an ad hoc manner, etc. We suppose that initially, each node in the network originates some data (*e.g.* that originates from a sensor, or from computation), and that these data must be aggregated at some designated node, the *sink*. To this goal, a node may send its data to a communication neighbor at a given time (the duration of this communication is supposed to be one time unit). We assume that there exists an aggregation function that takes two data as input and gives one data as output (the function is aggregating in the sense that the size of the output is supposed to be the same as a single input, such functions include min, max, etc.).

The main constraint for communications between nodes is that a node is allowed to send its data (be it its original data, or aggregated data) exactly once (*e.g.* to keep energy consumption low). A direct consequence of this constraint is that a node must aggregate data anytime it receives some, provided it did not send its data previously. It also implies that a node cannot participate in the data aggregation protocol once it has transmitted its data. A nice property of any algorithm implementing this constraint is that the number of communications is minimum. The problem of aggregating all data at the sink with minimum duration is called the *minimum data aggregation time problem* [6]. The essence of such a data aggregation algorithm is to decide whether or not to send a node's data when encountering a given communication neighbor: by waiting, a node may be able to aggregate more data, while by sending a node disseminates data but excludes itself for the rest of the computation.

In this paper, we consider that nodes may base their decision on their initial knowledge and past experience (past interactions with other nodes) only. Then, an algorithm accommodating those constraints is called an *online distributed data aggregation* algorithm. The existence of such an algorithm is conditioned by the (dynamic) topology, initial knowledge of the nodes (*e.g.* about their future communication neighbors), etc.

For simplicity, we assume that interactions between the nodes are carried out through pairwise operations. Anytime two nodes a and b are communication neighbors (or, for short, are interacting), either no data transfer happens, or one of them sends its data to the other, that executes the aggregation function on both its previously stored data and the received data, the output is then stored in the (new) stored data of the receiver. In the sequel, we use the term *interaction* to refer to a pairwise interaction.

We assume that an adversary controls the dynamics of the network, that is, the adversary decides which are the interactions. As we consider atomic interactions, the adversary decides what sequence of interactions is to occur in a given execution. Then, the sequence of static graphs to form the evolving graph can

be seen as a sequence of single edge graphs, where the edge denotes the interaction that is chosen by the scheduler at this particular moment. Hence, the time when an interaction occurs is exactly its index in the sequence. Our model of dynamic graphs as a sequence of interactions differs from existing models on several points. First, general models like *Time-varying-graph* [8] make use of continuous time, which adds a lot of complexity. Also, discrete time general models such as *evolving graph* [8] capture the network evolution as a sequence of static graphs. Our model is a simplification of the evolving graph model where each static graph has a single edge. *Population protocols* [2] also consider pairwise interactions, but focus on finite state anonymous nodes with limited computational power and unlimited communication power (a given node can transmit its information many times), while we consider powerful nodes (that can record their past interactions) that are communication limited (they can send their data only once). Finally, *Dynamic edge-relabeling* [7] is similar to population protocols, but the sequence of pairwise interactions occurs inside an evolving graph. This model shares the same differences as population protocols with our model.

Related Work. The problem of data aggregation has been widely studied in the context of wireless sensor networks. The literature on this problem can be divided in two groups depending on the assumption made about the collisions being handled by an underlying MAC layer.

In the case when collisions are not handled by the MAC layer, the goal is to find a collision-free schedule that aggregates the data in minimum duration. The problem was first studied by Annamalai *et al.* [3], and formally defined by Chen *et al.* [9], which proved that the problem is NP-complete. Then, several papers [12–14, 16] proposed centralized and distributed approximation algorithms for this problem. The best known algorithm is due to Nguyen *et al.* [12]. More recently, Bramas *et al.* [6] considered the generalization of the problem to dynamic wireless sensor networks (modeled by evolving graphs). Bramas *et al.* [6] show that the problem remains NP-complete even when restricted to dynamic WSNs of degree at most 2 (compared to 3 in the static case).

When collisions are handled by the MAC layer, various problems related to data aggregation have been investigated. The general term *in-network aggregation* includes several problems such as gathering and routing information in WSNs, mostly in a practical way. For instance, a survey [11] relates aggregation functions, routing protocols, and MAC layers with the objective of reducing resource consumption. *Continuous aggregation* [1] assumes that data have to be aggregated, and that the result of the aggregation is then disseminated to all participating nodes. The main metric is then the delay before aggregated data is delivered to all nodes, as no particular node plays the role of a sink. Most related to our concern is the work by Cornejo *et al.* [10]. In their work, each node starts with a token, the time is finite and no particular node plays the role of a sink node. Then, the topology evolves with time, and at each time instant, a node has at most one neighbor with whom it can interact and send or not its token. The goal is to minimize the number of nodes that own at least one token. Assuming an algorithm does not know the future, Cornejo *et al.* [10] prove that

its competitive ratio is $\Omega(n)$ with high probability (w.r.t. the optimal offline algorithm) against an oblivious adversary.

Our Contributions. In this paper we define the problem of distributed online data aggregation in dynamic graphs, and study its complexity. It turns out that the problem difficulty strongly depends on the power of the adversary (that chooses which interactions occur in a given execution).

For the oblivious and the online adaptive adversaries, we give several impossibility results when nodes have no knowledge about the future evolution of the dynamic graph, nor about the topology. Also, when nodes are aware of the underlying graph (where an edge between two nodes exists if those nodes interact at least once in the execution), the data aggregation is impossible in general. To examine the possibility cases, we define a cost function whose purpose is to compare the performance of a distributed online algorithm to the optimal offline algorithm for the same sequence of interactions. Our results show that if all interactions in the sequence occur infinitely often, there exists a distributed online data aggregation algorithm whose cost is finite. Moreover, if the underlying graph is a tree, we present an optimal algorithm.

For the randomized adversary, we first present tight bounds when nodes have full knowledge about the future interactions in the whole graph. In this case, the best possible algorithm terminates in $\Theta(n \log(n))$ interactions, in expectation and with high probability. Then, we consider nodes with restricted knowledge, and we present two optimal distributed online data aggregation algorithms that differ in the knowledge that is available to nodes. The first algorithm, called *Gathering*, assumes nodes have no knowledge whatsoever, and terminates in $O(n^2)$ interactions on average, which we prove is optimal with no knowledge. The second one, called *Waiting Greedy*, terminates in $O\left(n^{3/2}\sqrt{\log(n)}\right)$ interactions with high probability, which we show is optimal when each node only knows the time of its next interaction with the sink (the knowledge assumed by *Waiting Greedy*).

We believe our research paves the way for stimulating future researches, as our proof arguments present techniques and analysis that can be of independent interest for studying dynamic networks.

2 Model

A dynamic graph is modeled as a couple (V, I) , where V is a set of nodes and $I = (I_t)_{t \in \mathbb{N}}$ is a sequence of pairwise interactions (or simply interactions). An interaction occurs when two nodes in the network can exchange information. A special node in V is the *sink* node, and is denoted by s in the sequel. In the sequence $(I_t)_{t \in \mathbb{N}}$, the index t of an interaction also refers to its *time of occurrence*. In the sequel V always denotes the set of nodes, $n \geq 3$ its size, and $s \in V$ the sink node.

In general, we consider that nodes in V have unique identifiers, unlimited memory and unlimited computational power. However, we sometimes consider

nodes with no persistent memory between interactions; those nodes are called *oblivious*.

Initially, each node in V receives a data. Nodes have different data and a node can transmit its data at most once. Formally, during an interaction $I_t = \{u, v\}$, if both nodes have not yet transmitted their data, then one of the node has the possibility to transmit its data to the other node. The receiver aggregates the received data with its own data. The transmission and the aggregation take exactly one time unit. If a node decides to transmit its data, then it does not own any data, and is not able to receive other's data anymore.

Problem Statement. The data aggregation problem consists in choosing at each interaction whether a node transmits (and which one) or not so that after a finite number of interactions, the sink is the only node that owns data. In this paper we study distributed and online algorithms that solve this problem. Such algorithms are called *distributed online data aggregation* (DODA) algorithms.

A DODA is an algorithm that takes as input an interaction $I_t = \{u, v\}$, and its time of occurrence $t \in \mathbb{N}$, and outputs either u , v or \perp . If a DODA outputs a node, this node is the receiver of the other node's data. In more details, if u is the output, this means that before the interaction both u and v own data, and the algorithm orders v to transmit its data to u . The algorithm is able to change the memory of the interacting nodes, for instance to store information that can be used in future interactions. In the sequel, \mathcal{D}_{ODA} denotes the set of all DODA algorithms. And $\mathcal{D}_{\text{ODA}}^{\emptyset}$ denotes the set of DODA algorithms that only require oblivious nodes. Executing a DODA algorithm A on a sequence of interactions I correspond to executing A on each interaction $I_t \in I$, and the duration of the execution is the first time of occurrence of the interaction when the sink becomes the only node that owns data.

A DODA can require some knowledge to work. A knowledge is a function (or just an attribute) given to every node that gives some information about the future, the topology or anything else. By default, a node $u \in V$ has two pieces of information: its identifier $u.ID$ and a boolean $u.isSink$ that is true if u is the sink, and false otherwise. A DODA algorithm may use additional functions associated with different knowledge. $\mathcal{D}_{\text{ODA}}(i_1, i_2, \dots)$ denotes the set of DODA algorithms that use the functions i_1, i_2, \dots , representing the *knowledge* of the nodes. The first function we define for a node $u \in V$ in a dynamic graph (V, I) , is the function $u.meetTime$ that maps a time $t \in \mathbb{N}$ with the smallest time $t' > t$ such that $I_{t'} = \{u, s\}$ i.e., the time of the next interaction with the sink (for $u = s$, we define $s.meetTime$ as the identity, $t \mapsto t$). Then $\mathcal{D}_{\text{ODA}}(meetTime)$ refers to the set of DODA algorithms that use the information *meetTime*.

Adversary Models. In this paper we consider three models of adversaries:

- *The oblivious adversary.* This adversary knows the algorithm's code, and must construct the sequence of interactions before the execution starts.
- *The adaptive online adversary.* This adversary knows the algorithm's code and can use the past execution of the algorithm to construct the next interaction. However, it must make its own decision as it does not know in advance

the decision of the algorithm. In the case of deterministic algorithms, this adversary is equivalent to the oblivious adversary.

- *The randomized adversary.* This adversary constructs the sequence of interactions by picking pairwise interactions uniformly at random.

Section 3 presents our results with the oblivious and the adaptive online adversaries. The results with the randomized adversary are given in Sect. 4.

Definition of Cost. To study and compare different DODA algorithms, we use a tool slightly different from the competitive analysis that is generally used to study online algorithms. The competitive ratio of an algorithm is the ratio between its performance and the optimal offline algorithm's performance. However, one can hardly define objectively the performance of an algorithm. For instance, if we just consider the number of interactions before termination, then an oblivious adversary can construct a sequence of interactions starting with the same interaction repeated an arbitrary number of time. In this case, even the optimal algorithm has infinite duration. Moreover, the adversary can choose the same interaction repeatedly after that the optimal offline algorithm terminates. This can prevent any non optimal algorithm from terminating and make it have an infinite competitive-ratio.

To prevent this we define the cost of an algorithm. Our cost is a way to define the performance of an algorithm, depending on the performance of the optimal offline algorithm. We believe our definition of cost is well-suited for a lots of problems where the adversary has a strong power, especially in dynamic networks. One of its main advantages is that it is invariant by trivial transformation of the sequence of interactions, like inserting or deleting duplicate interactions.

For the sake of simplicity, the execution of an offline optimal data aggregation algorithm, having minimum duration, is called a *convergecast*. Consider a sequence of interactions I . Let $opt(t)$ be the ending time of a convergecast on I , starting at time $t \in \mathbb{N}$. If the ending time is infinite (if the optimal offline algorithm does not terminate) we write $opt(t) = \infty$. Let $T : \mathbb{N}_{\geq 1} \mapsto \mathbb{N} \cup \{\infty\}$ be the function defined as follows:

$$T(1) = opt(0), \quad \forall i \geq 1 \quad T(i+1) = opt(T(i) + 1)$$

$T(i)$ is the duration of i successive convergecasts (two convergecasts are consecutive if the second one starts just after the first one completes).

Let $duration(A, I)$ be the termination time of algorithm A executed on the sequence of interactions I . Now, we define the cost $cost_A(I)$ of an algorithm A on the sequence I , as the smallest integer i such that $duration(A, I) \leq T(i)$:

$$cost_A(I) = \min\{i \mid duration(A, I) \leq T(i)\}$$

This means that $cost_A(I)$ is an upper bound on the number of successive convergecasts we can perform during the execution of A , on the sequence I . It follows from the definition that an algorithm performs an optimal data aggregation if and only if $cost_A(I) = 1$.

Also, if $\text{duration}(A, I) = \infty$, then it is possible that $\text{cost}_A(I) < \infty$. Indeed, if $i_{\max} = \min_i \{i \mid T(i) = \infty\}$ is well-defined, then $\text{cost}_A(I) = i_{\max}$, otherwise $\text{cost}_A(I) = \infty$.

3 Oblivious and Online Adaptive Adversaries

In this section we give several impossibility results when nodes have no knowledge, and then show several results depending on the amount of knowledge. We choose to limit our study to some specific knowledge, but one can be interested in studying the possible solutions for different kind of knowledge.

3.1 Impossibility Results when Nodes Have No Knowledge

Theorem 1. *For every algorithm $A \in \mathcal{D}_{\text{ODA}}$, there exists an adaptive online adversary generating a sequence of interactions I such that $\text{cost}_A(I) = \infty$.*

Proof. Let I be the sequence of interactions among 3 nodes a , b , and the sink s , defined as follows. $I_0 = \{a, b\}$. If a transmits, then $\forall i > 0$, $I_{2i+1} = \{a, s\}$ and $I_{2i+2} = \{a, b\}$ so that b is never able to transmit. Symmetrically if b transmits the same thing happens. If no node transmits, then $I_1 = \{b, s\}$. If b transmits, then $\forall i > 0$, $I_{2i+2} = \{a, b\}$ and $I_{2i+3} = \{b, s\}$ so that a is never able to transmit. Otherwise $I_2 = \{a, b\}$ and we continue as in the first time. A never terminates, and a convergecast is always possible for the offline optimal algorithm, so that $\text{cost}_A(I) = \infty$. \square

In the case of deterministic algorithms, the previous theorem is true even with an oblivious adversary. However, for a randomized algorithm, the problem is more complex. The following theorem states that the impossibility results for oblivious randomized algorithm, leaving the case of general randomized algorithms against oblivious adversary as an open question.

Theorem 2. *For every randomized algorithm $A \in \mathcal{D}_{\text{ODA}}^\theta$, there exists an oblivious adversary generating a sequence of interactions I such that $\text{cost}_A(I) = \infty$ with high probability¹.*

Proof. Let $V = \{s, u_0, \dots, u_{n-2}\}$. In the sequel, indexes are modulo $n - 1$ i.e., $\forall i, j \geq 0$, $u_i = u_j$ with $i \equiv j \pmod{n - 1}$. Let I^∞ defined by, for all $i \in \mathbb{N}$, $I_i^\infty = \{u_i, s\}$. Let I^l be the finite sequence, prefix of length $l > 0$ of I^∞ . For every $l > 0$, the adversary can compute the probability P_l that no node transmits its data when executing A on I^l . $(P_l)_{l>0}$ is a non-increasing sequence, it converges to a limit $\mathcal{P} \geq 0$. For a given l , if $P_l \geq 1/n$, there is at least two nodes whose probability not to transmit when executing A on I^l is at least $n^{-\frac{1}{n-2}} = 1 - O\left(\frac{1}{\sqrt{n}}\right)$. To prove this, we can see the probability P_l as the product of $n - 1$ probabilities p_0, p_1, \dots, p_{n-2} where p_i is the probability that node u_i

¹ An event A occurs with high probability if $P(A) > 1 - O(1/\log(n))$.

does not transmit during I^l . Those events are independent since the algorithm is oblivious. Let $p_d \geq p_{d'}$ be the two greatest probabilities in $\{p_i\}_{0 \leq i \leq n-2}$, we have:

$$\begin{aligned} \left(\prod_{i=0}^{n-2} p_i \geq \frac{1}{n}\right) &\Rightarrow \left(\sum_{i=0}^{n-2} \log(p_i) \geq \log\left(\frac{1}{n}\right)\right) \\ &\Rightarrow \left((n-2) \log(p_{d'}) \geq \log\left(\frac{1}{n}\right)\right) \Rightarrow \left(p_{d'} \geq n^{-\frac{1}{n-2}}\right) \end{aligned}$$

This implies that, if $\mathcal{P} \geq 1/n$, then A does not terminate on the sequence I^∞ with high probability.

Otherwise, let l_0 be the smallest index such that $P_{l_0} < 1/n$. So that with high probability, at least one node transmits when executing A on I^{l_0} . Also, $P_{l_0-1} \geq 1/n$ so that the previous argument implies that there is at least two nodes u_d and $u_{d'}$ whose probability to still have a data (after executing A on I^{l_0-1}) is at least $n^{-\frac{1}{n-2}}$. If $l_0 = 0$ we can choose $\{u_d, u_{d'}\} = \{u_1, u_2\}$. We have $u_d \neq u_{l_0}$ or $u_{d'} \neq u_{l_0}$. Without loss of generality, we can suppose $u_d \neq u_{l_0}$, so that the probability that u_d transmits is the same in I^{l_0-1} and in I^{l_0} .

Now, u_d is a node whose probability not to transmit when executing A on I^{l_0} is at least $n^{-\frac{1}{n-2}} = 1 - O\left(\frac{1}{\sqrt{n}}\right)$. Let I' be the sequence of interactions defined as follows:

$$\forall i \in [0, n-2] \setminus \{d-1\}, I'_i = \{u_i, u_{i+1}\}, I'_{d-1} = \{u_{d-1}, s\}$$

I' is constructed such that u_d (the node that has data with high probability) must send its data along a path that contains all the other nodes in order to reach the sink. But this path contains a node that does not have a data.

Let I be the sequence of interaction starting with I^{l_0} and followed by I' infinitely often. We have shown that with high probability, after l_0 interactions, at least one node transmits its data and the node u_d still has a data. The node that does not have data prevents the data owned by u_d from reaching s . So that A does not terminate, and since a convergecast is always possible for the offline optimal algorithm, then $cost_A(I) = \infty$. □

3.2 When Nodes Know the Underlying Graph

Let \bar{G} be the underlying graph *i.e.*, $\bar{G} = (V, E)$ with $E = \{(u, v) \mid \exists t \in \mathbb{N}, I_t = \{u, v\}\}$. The following results assume that the underlying graph is given initially to every node.

Theorem 3. *If $n \geq 4$, then, for every algorithm $A \in \mathcal{D}_{ODA}(\bar{G})$, there exists an online adaptive adversary generating a sequence of interactions I such that $cost_A(I) = \infty$.*

Proof. $V = \{s, u_1, u_2, u_3\}$. We create a sequence of interactions with the underlying graph $\bar{G} = (V, \{(s, u_1), (u_1, u_2), (u_2, u_3), (u_3, s)\})$. We start with the following interactions:

$$(\{u_1, s\}, \{u_3, s\}, \{u_2, u_1\}, \{u_2, u_3\}). \tag{1}$$

If u_2 transmits to u_1 in I_2 , then we repeat infinitely often the three following interactions:

$$(\{u_1, u_2\}, \{u_2, u_3\}, \{u_3, s\}, \dots).$$

Else, if u_2 transmits to u_3 in I_3 , then we repeat infinitely often the three following interactions:

$$(\{u_3, u_2\}, \{u_2, u_1\}, \{u_1, s\}, \dots).$$

Otherwise, we repeat the four interactions (1), and apply the previous reasoning. Then, A never terminates, and a convergecast is always possible, so that $cost_A(I) = \infty$. □

Theorem 4. *If the interactions occurring at least once, occur infinity often, then there exists $A \in \mathcal{D}_{ODA}^\emptyset(\bar{G})$ such that $cost_A(I) < \infty$ for every sequence of interactions I . However, $cost_A(I)$ is unbounded.*

Theorem 5. *If \bar{G} is a tree, there exists $A \in \mathcal{D}_{ODA}^\emptyset(\bar{G})$ that is optimal.*

3.3 If Nodes Know Their Own Future

For a node $u \in V$, $u.future$ denotes the future of u *i.e.*, the sequence of interactions involving u , with their times of occurrences. In this case, according to the model, two interacting nodes exchange their future and non-oblivious nodes can store it. This may seem contradictory with the motivation of the problem (that aims to reduce the number of transmissions). However, it is possible that the data must be sent only once for reasons not related to energy (such as data that cannot be duplicated, tokens, etc.). So, even if, in general, oblivious algorithms should be favored, we still investigated this case for the sake of completeness.

Theorem 6. *There exists $A \in \mathcal{D}_{ODA}(future)$ such that $cost_A(I) \leq n$ for every sequence of interactions I .*

Proof. One can show that the duration of $n - 1$ successive convergecasts is sufficient to perform a broadcast from any source. So every node broadcast its future to the other nodes. After that, all the nodes are aware of the future of every nodes and can compute an optimal data aggregation schedule. So that it takes only one convergecast to aggregate the data of the whole network. In total, n successive convergecasts are sufficient. □

4 Randomized Adversary

The randomized adversary constructs the sequence of interactions by picking a couple of nodes among all possible couples, uniformly at random. Thus, the underlying graph is a complete graph of n nodes (including the sink) and every interaction occurs with the same probability $p = \frac{2}{n(n-1)}$.

In this section, the complexity is computed on expectation (because the adversary is randomized) and no more “in the worst case” as previously. In this case, considering the number of interactions is sufficient to represent the complexity of an algorithm. We see in Theorem 8 that an offline algorithm terminates in $\Theta(n \log(n))$ interactions w.h.p. This bound gives a way to convert the complexity in term of number of interaction to a cost. Indeed, if an algorithm \mathcal{A} terminates in $O(n^2)$ interactions, then its performance is $O(n/\log(n))$ times worse than the offline algorithm and $cost_{\mathcal{A}}(I) = O(n/\log(n))$ for a randomly generated sequence of interactions I . For the sake of simplicity, in the remaining of the section, we give the complexity in terms of number of interactions.

Since an interaction does not depend on previous interactions, the algorithms we propose here are oblivious i.e., they do not modify the memory of the nodes. In more details, the output of our algorithms depends only on the current interaction and on the information available in the node.

First, we introduce three oblivious DODA algorithms. For the sake of simplicity, we assume that the output is ignored if the interacting nodes do not both have a data. Also, to break symmetry, we suppose the nodes that interact are given as input ordered by their identifiers.

- Waiting ($\mathcal{W} \in \mathcal{D}_{\text{ODA}}^\emptyset$): A node transmits only when it is connected to the sink s : $\mathcal{W}(u_1, u_2, t)$ equals u_i if $u_i = s$, and \perp otherwise.
- Gathering ($\mathcal{GA} \in \mathcal{D}_{\text{ODA}}^\emptyset$): A node transmits its data when it is connected to the sink s or to a node having data: $\mathcal{GA}(u_1, u_2, t)$ equals u_2 if $u_2 = s$, otherwise it equals u_1
- Waiting Greedy with parameter $\tau \in \mathbb{N}$ ($\mathcal{WG}_\tau \in \mathcal{D}_{\text{ODA}}^\emptyset(\text{meetTime})$): The node with the greatest meet time transmits, if its meet time is greater than τ :

$$\mathcal{WG}_\tau : (u_1, u_2, t) = \begin{cases} u_1 & \text{if } m_1 \leq m_2 \wedge \tau < m_2 \\ u_2 & \text{if } m_1 > m_2 \wedge \tau < m_1 \\ \perp & \text{otherwise} \end{cases} \quad \text{with} \quad \begin{cases} m_1 = u_1.\text{meetTime}(t) \\ m_2 = u_2.\text{meetTime}(t) \end{cases}$$

One can observe that after time τ , the algorithm acts as the Gathering algorithm.

4.1 Lower Bounds

We show a lower bound $\Omega(n^2)$ on the number of interactions required for DODA against the randomized adversary. The lower bound holds for all algorithms (including randomized ones) that do not have knowledge about future of the evolving network. The lower bound matches the upper bound of the *Gathering* algorithm given in the next subsection. This implies that this bound is tight.

Theorem 7. *The expected number of interactions required for DODA is $\Omega(n^2)$.*

Proof. We show that the last data transmission requires $\Omega(n^2)$ interactions in expectation.

We consider any (randomized) algorithm A and its execution for DODA. Before the last transmission (from some node, say v , to the sink s), only v has data except for s .

The probability that v and s interacts in the next interaction is $\frac{2}{n(n-1)}$. Thus, the expected number EI of interactions required for v to transmit to s is:

$$EI = \frac{n(n-1)}{2}$$

So that the whole aggregation requires at least $EI = \Omega(n^2)$. □

We also give a tight bound for algorithms that know the full sequence of interactions.

Theorem 8. *The best algorithm in $\mathcal{D}_{\text{ODA}}^0$ (full knowledge) terminates in $\Theta(n \log(n))$ interactions, in expectation and with high probability.*

Proof. First, we show that the expected number of interactions of a broadcast algorithm is $\Theta(n \log n)$. The first data transmission occurs when the source node (say v_0) interacts with another node. The probability of occurrence of the first data transmission is $\frac{2(n-1)}{n(n-1)}$. After the $(i-1)$ -th data transmission, i nodes (say $V_{i-1} = \{v_0, v_1, \dots, v_{i-1}\}$) have the data and the i -th data transmission occurs when a node in V_{i-1} interacts with a node not in V_{i-1} . This happens with probability $\frac{2i(n-i)}{n(n-1)}$.

Thus, if X is the number of interactions required to perform a broadcast, then we have:

$$\begin{aligned} E(X) &= \sum_{i=1}^{n-1} \frac{n(n-1)}{2i(n-i)} = \frac{n(n-1)}{2} \sum_{i=1}^{n-1} \frac{1}{i(n-i)} = \frac{n(n-1)}{2n} \sum_{i=1}^{n-1} \left(\frac{1}{i} + \frac{1}{n-i} \right) \\ &= (n-1) \sum_{i=1}^{n-1} \frac{1}{i} \in \Theta(n \log n). \end{aligned}$$

And the variance is

$$\begin{aligned} \text{Var}(X) &= \sum_{i=1}^{n-1} \left(1 - \frac{2i(n-i)}{n(n-1)} \right) / \left(\frac{2i(n-i)}{n(n-1)} \right)^2 = n(n-1) \sum_{i=1}^{n-1} \frac{n(n-1) - 2i(n-i)}{(2i(n-i))^2} \\ &= O \left(n^4 \sum_{i=1}^{\lfloor n/2 \rfloor - 1} \left(\frac{1}{i(n-i)} \right)^2 \right) \end{aligned}$$

The last sum is obtained from the previous one by observing that it is symmetric with respect to the index $i = n/2$, and the removed elements ($i = \lfloor n/2 \rfloor$ and

possibly $i = \lceil n/2 \rceil$) are negligible. We define $f : x \mapsto \frac{1}{x^2(n-x)^2}$. Since f is increasing between 1 and $n/2$, we have

$$\sum_{i=1}^{\lfloor n/2 \rfloor - 1} f(i) \leq \int_1^{n/2} f(x)dx = \frac{\frac{(n-2)n}{n-1} + 2 \log(n-1)}{n^3} = O\left(\frac{1}{n^2}\right)$$

So that the variance is in $O(n^2)$. Using the Chebyshev's inequality, we have

$$P(|X - E(X)| > n \log(n)) = O\left(\frac{1}{\log^2(n)}\right)$$

Therefore, a sequence of $\Theta(n \log(n))$ interactions is sufficient to perform a broadcast with high probability. By reversing the order of the interactions in the sequence of interactions, this implies that a sequence of $\Theta(n \log(n))$ interactions is also sufficient to perform a convergecast with the same probability. Aggregating data along the convergecast tree gives a valid data aggregation schedule. \square

Corollary 1. *The best algorithm in $\mathcal{D}_{\text{ODA}}(\text{future})$ terminates in $\Theta(n \log(n))$ interactions, in expectation and with high probability.*

Proof. If each node starts with its own future, $O(n \log(n))$ interactions are sufficient to retrieve with high probability the future of the whole network. Then $O(n \log(n))$ interactions are sufficient to aggregate all the data with the full knowledge. \square

4.2 Algorithm Performance Without Knowledge

Without any knowledge, we show that the Gathering algorithm is optimal.

Theorem 9. *The expected number of interactions the Waiting requires to terminate is $O(n^2 \log(n))$. The expected number of interactions the Gathering requires to terminate is $O(n^2)$.*

4.3 Algorithm Performance with *meetTime*

In this subsection we study the performance of our algorithm *Waiting Greedy*, find the optimal value of the parameter τ and prove that this is the best possible algorithm with only the *meetTime* information (even if nodes have unbounded memory). We begin by a lemma to find how many interactions are needed to have a given number of nodes interacting with the sink.

Lemma 1. *If f is a function such that $f(n) = o(n)$ and $f(n) = \omega(\log(n))$ then, in $nf(n)$ interactions, $\Theta(f(n))$ nodes interact with the sink w.h.p.*

Now we can state our theorem about the performance of Waiting Greedy depending on the parameter τ . To prove this Theorem, we partition the set of nodes in two important subsets, L that contains the nodes that interact with the sink between time $\tau/2$ and τ and L^c its complementary. We show that the duration of our algorithm comes from two phases, one before time $\tau/2$ when the nodes in L^c have a high probability to meet another node in L , and one after $\tau/2$ when the nodes in L meet the sink directly.

Theorem 10. *Let f be a function such that $f(n) = o(n)$ and $f(n) = \omega(\log(n))$. The algorithm Waiting Greedy with $\tau = \Theta(\max(nf(n), n^2 \log(n)/f(n)))$ terminates in τ interactions w.h.p.*

Proof. To have an upper bound on the number of interactions needed by Waiting Greedy to terminate, we decompose the execution in two phases, one between time 0 and a time t_1 and the other between time t_1 and a time $t_2 = \tau$ (for simplicity, one can take $t_1 = \tau/2$). In the last phase, a set of nodes $L \subset V$ interacts at least once directly with the sink. Nodes in L do not transmit to anyone in the first phase by definition of the algorithm (they have a meetTime smaller than τ). Nodes in L help the other nodes (in $L^c = V \setminus L$) to transmit their data in the first phase. Maybe nodes in L^c can transmit to L in the second phase, but we do not take this into account, that is why it is an upper bound.

If a node u in L^c interacts with a node in L in the first phase, either it transmits its data, otherwise (by definition of the algorithm) it has a meetTime smaller than τ (and smaller than t_1 because it is not in L). In every case, a node in L^c that meets a node in L in the first phase, transmits its data. To prove the theorem i.e., in order for the algorithm to terminates before τ with high probability, we prove two claims: (a) the number of nodes in L is $f(n)$ with high probability if $t_2 - t_1 = nf(n)$ and (b) all nodes in L^c interact with a node in L before t_1 with high probability if $t_1 = \Theta(n^2 \log(n)/f(n))$. The first claim is implied by Lemma 1. Now we prove the second claim.

Let X be the number of interactions required for the nodes in L^c to meet a node in L . The probability of the i -th interaction between a node in L^c (with a data) and a node in L , after $i - 1$ such interactions already occurred, is $2f(n)(n - f(n) - i)/n(n - 1)$. Then we have:

$$\begin{aligned}
 E(X) &= \sum_{i=1}^{n-f(n)-1} \frac{n(n-1)}{2f(n)(n-f(n)-i)} = \frac{n(n-1)}{2f(n)} \sum_{i=1}^{n-f(n)-1} \frac{1}{n-f(n)-i} \\
 &\sim_{+\infty} \frac{n^2}{2f(n)} \log(n-f(n)) = \frac{n^2}{2f(n)} \log(n(1-f(n)/n)) \sim_{+\infty} \frac{n^2 \log(n)}{2f(n)} \\
 \text{Var}(X) &= \sum_{i=1}^{n-f(n)-1} \frac{\left(1 - \frac{2f(n)(n-f(n)-i)}{n(n-1)}\right)}{\left(\frac{2f(n)(n-f(n)-i)}{n(n-1)}\right)^2} \sim \sum_{i=1}^{n-f(n)-1} \frac{n^4}{4f(n)^2 n^2} \sim \frac{n^3}{4f(n)^2}
 \end{aligned}$$

Like previously, using the Chebyshev's inequality, this implies that $X = O\left(\frac{n^2 \log(n)}{f(n)}\right)$ w.h.p. □

Corollary 2. *The algorithm Waiting Greedy, with $\tau = \Theta(n^{3/2}\sqrt{\log(n)})$ terminates in τ interactions with high probability.*

Proof. In the last theorem, the bound $O(\max(nf(n), n^2 \log(n)/f(n)))$ is minimized by the function $f: n \mapsto \sqrt{n \log(n)}$. □

From the previous corollary, we maximize the performance of our algorithm by choosing $\tau = \Theta(n^{3/2}\sqrt{\log(n)})$. To prove that no other algorithm has better performance, we show that, if an algorithm terminates before $\Theta(n^{3/2}\sqrt{\log(n)})$ interactions, then the number of nodes that do not meet the sink is so big that they cannot aggregate their data quickly enough. Indeed, when two nodes do not meet the sink before $\Theta(n^{3/2}\sqrt{\log(n)})$, then their *meetTime* information is useless for deciding which one should transmit, so we can analyze the aggregation speed of a large subset of the node.

Theorem 11. *Waiting Greedy with $\tau = \Theta(n^{3/2}\sqrt{\log(n)})$ is optimal in $\mathcal{D}_{\text{ODA}}(\text{meetTime})$.*

Proof. For the sake of contradiction, we suppose the existence of an algorithm $A \in \mathcal{D}_{\text{ODA}}(\text{meetTime})$ that terminates in $T(n)$ interactions with high probability, with $T(n) = o(n^{3/2}\sqrt{\log(n)})$. Without loss of generality we can suppose that A does nothing after $T(n)$ interactions. Indeed, the algorithm A' that executes A up to $T(n)$ and does nothing afterward has the same upper bound (since the bound holds with high probability).

Let L be the set of nodes that interact directly with the sink during the first $T(n)$ interactions. Let L^c be its complementary in $V \setminus \{s\}$. We know from Lemma 1 that $\#L = O(T(n)/n) = o(\sqrt{n \log(n)})$ w.h.p.

We can show that $T(n)$ interactions are not sufficient for all the nodes in L^c to interact with nodes in L . If nodes in L^c want to send their data to the sink, some data must be aggregated among nodes in L^c , then the remaining nodes in L^c that still own data must interact with a node in L before $T(n)$ interactions (this is not even sufficient to perform the DODA, but is enough to reach a contradiction).

When two nodes in L^c interact, their *meetTime* (that are greater than $T(n)$) and the previous interactions are independent with the future interactions occurring before $T(n)$. This implies that when two nodes in L^c interact, using those information to decide which node transmits is the same as choosing the sender randomly. From Theorem 9, this implies that the optimal algorithm to aggregate data in L^c is the Gathering algorithm.

Now, we show that, even after the nodes in L^c use the Gathering algorithm, there is with high probability at least one node in L^c that still owns data and that does not interact with any node in L . This node prevents the termination of the algorithm before $T(n)$ interactions with high probability, which is a contradiction.

Formally, we have the following lemmas. Due to space constraints, the proofs can be found in a technical report [5].

Lemma 2. *Let $g(n)$ be the number of nodes in L^c . After using the Gathering algorithm during $T(n)$ interactions, the number of nodes in L^c that still own data is in $\omega(\sqrt{n/\log(n)})$ w.h.p.*

Lemma 3. *Let $H \subset L^c$ be the nodes in L^c that still own data after the gathering. Then, $T(n)$ interactions are not sufficient for all the nodes in H to interact with nodes in L , w.h.p.*

End of the Proof of Theorem 11. We have shown that $T(n)$ interactions are not sufficient for the nodes in L^c to transmit their data (directly or indirectly) to the nodes in L . Indeed, we have shown that the nodes in L^c can apply the gathering algorithm so that $\omega(\sqrt{n \log(n)})$ nodes in L^c still own data with high probability. But, with high probability, one of the $\omega(\sqrt{n \log(n)})$ remaining nodes does not interact with a node in L in $T(n)$ interactions. This implies that, with high probability, at least one node cannot send its data to the sink in $T(n)$ interactions and an algorithm A with such a bound T does not exist. \square

5 Concluding Remarks

We defined and investigated the complexity of the distributed online data aggregation problem in dynamic graphs where interactions are controlled by an adversary. We obtained various tight complexity results for different adversaries and node knowledge, that open several scientific challenges:

1. What knowledge has a real impact on the lower bounds or algorithm efficiency?
2. What results can be generalized to a model where nodes can transmit a constant number of times instead of only once?
3. Can randomized adversaries that use a non-uniform probabilistic distribution alter significantly the bounds presented here in the same way as in the work by Yamauchi *et al.* [15]?

References

1. Abshoff, S., Meyer auf der Heide, F.: Continuous aggregation in dynamic ad-hoc networks. In: Halldórsson, M.M. (ed.) SIROCCO 2014. LNCS, vol. 8576, pp. 194–209. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-09620-9_16
2. Angluin, D., Aspnes, J., Eisenstat, D., Ruppert, E.: The computational power of population protocols. *Distrib. Comput.* **20**(4), 279–304 (2007)
3. Annamalai, V., Gupta, S.K.S., Schwiebert, L.: On tree-based convergecasting in wireless sensor networks. In: 2003 IEEE Wireless Communications and Networking, WCNC 2003, vol. 3, pp. 1942–1947. IEEE (2003)
4. Bramas, Q., Masuzawa, T., Tixeuil, S.: Distributed online data aggregation in dynamic graphs. In: 36th IEEE International Conference on Distributed Computing Systems, ICDCS 2016, Nara, Japan, 27–30 June 2016, pp. 747–748. IEEE Computer Society (2016)

5. Bramas, Q., Masuzawa, T., Tixeuil, S.: Distributed online data aggregation in dynamic graphs. arXiv preprint [arXiv:1602.01065](https://arxiv.org/abs/1602.01065) (2016)
6. Bramas, Q., Tixeuil, S.: The complexity of data aggregation in static and dynamic wireless sensor networks. In: Pelc, A., Schwarzmann, A.A. (eds.) SSS 2015. LNCS, vol. 9212, pp. 36–50. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-21741-3_3
7. Casteigts, A., Chaumette, S., Ferreira, A.: Characterizing topological assumptions of distributed algorithms in dynamic networks. In: Kutten, S., Žerovnik, J. (eds.) SIROCCO 2009. LNCS, vol. 5869, pp. 126–140. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-11476-2_11
8. Casteigts, A., Flocchini, P., Quattrociocchi, W., Santoro, N.: Time-varying graphs and dynamic networks. In: Frey, H., Li, X., Ruehrup, S. (eds.) ADHOC-NOW 2011. LNCS, vol. 6811, pp. 346–359. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22450-8_27
9. 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). https://doi.org/10.1007/11599463_14
10. Cornejo, A., Gilbert, S., Newport, C.: Aggregation in dynamic networks. In: Proceedings of the 2012 ACM Symposium on Principles of Distributed Computing, pp. 195–204. ACM (2012)
11. Fasolo, E., Rossi, M., Widmer, J., Zorzi, M.: In-network aggregation techniques for wireless sensor networks: a survey. *IEEE Wirel. Commun.* **14**(2), 70–87 (2007)
12. Nguyen, T.D., Zalyubovskiy, V., Choo, H.: Efficient time latency of data aggregation based on neighboring dominators in WSNs. In: 2011 IEEE Global Telecommunications Conference (GLOBECOM 2011), pp. 1–6. IEEE (2011)
13. Ren, M., Guo, L., Li, J.: A new scheduling algorithm for reducing data aggregation latency in wireless sensor networks. *Int. J. Commun. Netw. Syst. Sci.* **3**(8), 679 (2010)
14. XiaoHua, X., Li, M., Mao, X.F., Tang, S., Wang, S.G.: A delay-efficient algorithm for data aggregation in multihop wireless sensor networks. *IEEE Trans. Parallel Distrib. Syst.* **22**(1), 163–175 (2011)
15. Yamauchi, Y., Tixeuil, S., Kijima, S., Yamashita, M.: Brief announcement: probabilistic stabilization under probabilistic schedulers. In: Aguilera, M.K. (ed.) DISC 2012. LNCS, vol. 7611, pp. 413–414. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33651-5_34
16. Yu, B., Li, J., Li, Y.: Distributed data aggregation scheduling in wireless sensor networks. In: IEEE INFOCOM 2009, pp. 2159–2167. IEEE (2009)