

Optimizing Reachability Sets in Temporal Graphs by Delaying ^{*}

Argyrios Deligkas[†] Igor Potapov[‡]

Abstract

A temporal graph is a dynamic graph where every edge is assigned a set of integer time labels that indicate at which discrete time step the edge is available. In this paper, we study how changes of the time labels, corresponding to delays on the availability of the edges, affect the reachability sets from given sources. The questions about reachability sets are motivated by numerous applications of temporal graphs in network epidemiology, which aim to minimise the spread of infection, and scheduling problems in supply networks in manufacturing with the opposite objectives of maximising coverage and productivity. We introduce control mechanisms for reachability sets that are based on two natural operations of delaying time events which significantly affecting the chains of these events. The first operation, termed merging, is global and batches together consecutive time labels in the whole network simultaneously. This corresponds to postponing all events until a particular time. The second, imposes independent delays on the time labels of every edge of the graph. We provide a thorough investigation of the computational complexity of different objectives related to reachability sets when these operations are used. For the merging operation, i.e. global lockdown effect, we prove NP-hardness results for several minimization and maximization reachability objectives, even for very simple graph structures. For the second operation, independent delays, we prove that the minimization problems are NP-hard when the number of allowed delays is bounded. We complement this with a polynomial-time algorithm for minimising the reachability set in case of unbounded delays.

Keywords: Temporal Graphs; Reachability Sets; Optimisation; NP-hard

1 Introduction

A plethora of real life scenarios can be modelled as a dynamic network that changes over time. These scenarios range from train, bus, and distribution schedules, to livestock movements between farms and virus spreading. Many of these dynamic networks consist of a fixed set of nodes and what changes over time is the connectivity between pairs of them; the locations of the stations, distribution centers, and farms remain the same over time, while the connections between any two of them can change every few minutes, hours, or days. An equivalent way to see these networks is as a sequence of static networks that change according to a predetermined, known in advance, schedule. These type of networks, known as *temporal graphs*, were formalized in the seminal work of [19]. Since then, there is flourish of work on temporal graphs [8, 9, 15, 21, 22, 35]. Formally, in a temporal graph every edge is assigned a set of integer time labels that indicate at which discrete time steps the edge is available. In other words, a temporal graph is a schedule of edge sets E_1, E_2, \dots, E_t over a fixed set of vertices.

^{*}A preliminary conference version of this work appeared in Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI), New York, New York, USA, 2020.

[†]Department of Computer Science, Royal Holloway University of London, UK. Email: argyrios.deligkas@rhul.ac.uk. This work was done while the author was a postdoc at Materials Innovation Factory at the University of Liverpool, funded by LRC.

[‡]Department of Computer Science, University of Liverpool, UK. Email: potapov@liverpool.ac.uk. Research partially funded by the grant EP/R018472/1 “Application driven Topological Data Analysis”.

In this paper, we study questions related to *reachability sets* on temporal graphs. The reachability set of vertex v is the set of vertices that there exist *temporal paths* from v to them. Informally speaking, a temporal path is a path whose edges, additionally to the usual connectivity, use strictly increasing time labels [33]. However, in contrast to static graphs, temporal paths do not preserve transitivity. As a result, some well known concepts from graph theory, like Menger’s theorem, do not hold for temporal graphs and have to be redefined [2,16,20].

Reachability sets emerge naturally in a wide range of models and real life applications [3, 25,27]. One of them in relation to temporal graphs, is the minimization of spread of infectious diseases over contact networks. The recent case of coronavirus COVID-19 once again reveals the importance of effective prevention and control mechanisms [1]. Data provide significant evidence that commuter patterns and airline flights [6,10], and livestock movements between farms [23] could spread an infectious disease. In particular, the impact of animal movements has been extensively studied within the epidemics community. Recent studies have shown that restricting of animal movements is among the most effective methods for controlling the spread of an infection [7,18,28–30] and changes in the network of animal movements between farms (nodes in a graph) can significantly decrease the spread of an infection [17,24]. Contrary to the applications above, there are cases where we wish to maximize the reachability sets. Consider for example a distribution schedule, where the use of the underlying network for every day comes at a fixed cost. Then, the goal is to optimally utilize the infrastructure of the network by choosing which days to use the network in order to achieve maximum reachability while maintaining low cost.

The importance of these problems combined with their inherit temporal nature, made temporal graph theory a tool for analyzing epidemiology in (animal) networks [5,11–14,26,31,32]. [12] and [14] studied how reachability sets on temporal graphs change, when the schedule of the edges can change according to specific operations. More specifically, they studied MINMAXREACH and MAXMINREACH problems where the goal, respectively, was to minimize the maximum, or to maximize the minimum size of reachability sets between a given set of sources. [12] studied these objectives under the operation of *reordering* of the edge sets and it was proven that both problems are NP-hard. For MINMAXREACH, [14] studied the operations of deletion of whole edge sets, or individual edges within an edge set. It was proven that MINMAXREACH under both notions of deletion is NP-hard, but they admit an FPT algorithm when they are parameterized by the treewidth of the network and the number of non-reachable vertices.

Although optimization of reachability sets capture important real life problems, some of the proposed solutions are not completely satisfying due to big changes in infrastructure. For example, reordering of edge sets can be difficult, costly, or even impossible to perform due to the physical constraints of the network, or due to the number of changes required in the existing infrastructure. For edge-deletions, an upper bound on the number of allowed deletions is required. This is crucial, since the deletion of every edge trivially minimizes the reachability sets, but makes the existing network infrastructure useless. In addition, edge deletions can create a “bottleneck” problem in the network even if their number is bounded. The deletion of an edge can create a sink to the network, or completely isolate some of its parts. Instead, we wish to study the following problem.

Given a temporal graph and a set of sources, optimize the size of the reachability set of the sources using only “natural” and “infrastructure-friendly” operations.

Natural operations should be intuitive and should be easy to apply; deletion, or postponement, of a temporal edge can naturally happen. On the other hand, bringing forward a temporal edge cannot be always feasible; a bus cannot go faster than its top speed. Infrastructure-friendly

Problem	Graph Class	Sources	Labels/Edge	Edges/Step
MINREACH	Path	$O(n)$	1	3
MINREACH MINMAXREACH MINAVGREACH	Tree $\Delta = 3$	1	1	1
MAXREACH	Path	$O(n)$	1	4
MAXREACH MAXMINREACH MAXAVGREACH	Bipart. $\Delta = 3$	1	1	4
MAXREACH MAXMINREACH MAXAVGREACH	Tree $\Delta = 3$	1	1	10

Table 1: NP-hardness results for the λ -merge operation for any $\lambda \geq 2$. Δ denotes the maximum degree of the graph. The results are approximation preserving and they rule out PTASs and they hold even if the underlying graph is directed.

operations should not be too difficult to perform and that do not require many changes to the given network and temporal schedule.

Our contribution. Our contribution is twofold. Firstly, we introduce and study two operations, *merging* and *delaying*, that are natural and infrastructure-friendly and were not studied in the past. The idea behind both operations is the postponement of the edges of the graph. Merging operation is parameterized by λ and it batches together λ consecutive edge sets; a λ -merge on E_1, \dots, E_λ changes the first $\lambda - 1$ edge sets to the empty sets, and the λ -th edge set is the union of all λ edge sets. The delaying operation, independently delays a temporal edge; a δ -delay on the label i of edge uv changes it to $i + \delta$. In contrast to deletion of temporal edges that can directly isolate vertices, our operations isolate some vertices *only temporarily*. Our second contribution is a thorough investigation of the computational complexity of reachability objectives under these operations. In particular, we study the MINMAXREACH, MINREACH, and MINAVGREACH, where in the last two cases the goal is to minimize the number, or the average number respectively, of reachable vertices from a given set of sources. With respect to maximization objectives, we study MAXREACH, MAXMINREACH, and MAXAVGREACH. We proved that these problems are NP-hard under the merging operation even for very restricted classes of graphs. Our results are depicted in Table 1. For the delaying operation, we studied the minimization problems. They remain NP-hard when we *bound* the number of times we are allowed to use this operation. We complement these results with a polynomial time algorithm for the case of the unbounded number of delays that works for any δ .

2 Preliminaries

Temporal Graphs. A *temporal graph* $\langle G, \mathcal{T} \rangle$ is a pair of a (directed) graph $G = (V, E)$ and a function \mathcal{T} that maps every edge of G to a list of time steps at which the edge is available. The maximum time step, t_{\max} , on the edges of G defines the *lifetime* of the temporal graph. Another interpretation of a temporal graph, which is useful in our case, is to see it as a schedule

of subgraphs, or edge-sets, $E_1, E_2, \dots, E_{t_{\max}}$ of G , known in advance and defined by the function \mathcal{T} ; at time step t function \mathcal{T} defines a set $E_t \subseteq E$ of edges available for this time step. We will say that an edge has the label i , if it is available at time step i . The size of a temporal graph $\langle G, \mathcal{T} \rangle$ is $|V| + \sum_{t \leq t_{\max}} |E_t|$. A *temporal path* in $\langle G, \mathcal{T} \rangle$ from v_1 to v_k is a sequence of edges $v_1v_2, v_2v_3, \dots, v_{k-1}v_k$ such that each edge $v_i v_{i+1}$ is available at time step t_i and $t_i < t_{i+1}$ for every $i \in [k-1]$.

Reachability Sets. A vertex u is *reachable* from vertex v if there exists a temporal path from v to u . We assume that a vertex is reachable from itself. It is possible that u is reachable from v , but v is not reachable from u . The reachability set of v , denoted $\text{reach}(v, \langle G, \mathcal{T} \rangle)$, is the set of vertices reachable from v in $\langle G, \mathcal{T} \rangle$. Given a temporal graph with lifetime t_{\max} and a time step $t < t_{\max}$, the set $\text{reach}_t(v, \langle G, \mathcal{T} \rangle)$ contains all the vertices reachable from v up to time t . The set $\text{reach}(v, \langle G, \mathcal{T} \rangle)$ can be computed in polynomial time with respect to the size of $\langle G, \mathcal{T} \rangle$; it suffices to check whether there exists a temporal path from v to every vertex in V , which can be done efficiently [34].

Merging. A *merging* operation on \mathcal{T} postpones some of the edge-sets in a particular way. Intuitively, a merging operation “batches together” a number of consecutive edge-sets.

Definition 1 (λ -merge) For every positive integer λ , a λ -merge of $E_i, E_{i+1}, \dots, E_{i+\lambda-1}$, replaces $E_j = \emptyset$ for every $i \leq j < i + \lambda - 1$ and $E_{i+\lambda-1} = \bigcup_{i \leq j \leq i+\lambda-1} E_j$.

So, every merge corresponds to the global delay of events from times $i, i+1, \dots, i+\lambda-1$ to the time $i+\lambda-1$. We say that two λ -merges are *independent*, if they merge $E_i, \dots, E_{i+\lambda-1}$ and $E_j, \dots, E_{j+\lambda-1}$ respectively, and $i+\lambda-1 < j$.¹

When it is clear from the context, instead of writing that we merge E_i with E_{i+1} we will write that we merge i with $i+1$.

Definition 2 ((λ, μ) -merging scheme) For positive integers λ and μ , a (λ, μ) -merging scheme applies μ independent λ' -merges on $E_1, E_2, \dots, E_{t_{\max}}$, where $\lambda' \leq \lambda$. A merging scheme is maximal if there is no other feasible λ -merge available.



Figure 1: Left: a temporal graph where $E_1 = \{xy\}$ and $E_2 = \{yz\}$; vertex z is reachable from x . Right: the resulted graph after merging E_1 with E_2 ; vertex z is no longer reachable from x .

A (λ, μ) -merging scheme for a temporal graph $\langle G, \mathcal{T} \rangle$ essentially produces a new temporal graph by modifying the schedule \mathcal{T} using μ independent λ' -merges, where $\lambda' \leq \lambda$. We will use $\mathcal{T}_{(\lambda, \mu)}^M$ to denote the modified schedule and $\langle G, \mathcal{T}_{(\lambda, \mu)}^M \rangle$ the corresponding modified temporal graph.

Our goal is to compute merging schemes that optimise some objectives regarding reachability sets from a given set of vertices. The input to the problems we study consists of a temporal

¹Independent merges (delays) are commutative in respect to their application. For example let us consider a temporal graph with 4 nodes $x_1 \xrightarrow{1} x_2 \xrightarrow{2} x_3 \xrightarrow{3} x_4$ and two non-independent merges E_1, E_2 and E_2, E_3 . Following the order of E_1, E_2 and then E_2, E_3 we will have a graph $x_1 \xrightarrow{3} x_2 \xrightarrow{3} x_3 \xrightarrow{3} x_4$ and the alternative order of merges E_2, E_3 and then E_1, E_2 will give $x_1 \xrightarrow{2} x_2 \xrightarrow{3} x_3 \xrightarrow{3} x_4$, so it creates some ambiguity for non-independent merges.

$(\lambda, \mu) - \text{MERGING}$	Objective
1. MINREACH	$\min \bigcup_{v \in S} \text{reach}(v, \langle G, \mathcal{T}_{(\lambda, \mu)}^M \rangle) $
2. MINMAXREACH	$\min \max_{v \in S} \text{reach}(v, \langle G, \mathcal{T}_{(\lambda, \mu)}^M \rangle) $
3. MINAVGREACH	$\min \sum_{v \in S} \text{reach}(v, \langle G, \mathcal{T}_{(\lambda, \mu)}^M \rangle) $
4. MAXREACH	$\max \bigcup_{v \in S} \text{reach}(v, \langle G, \mathcal{T}_{(\lambda, \mu)}^M \rangle) $
5. MAXMINREACH	$\max \min_{v \in S} \text{reach}(v, \langle G, \mathcal{T}_{(\lambda, \mu)}^M \rangle) $
6. MAXAVGREACH	$\max \sum_{v \in S} \text{reach}(v, \langle G, \mathcal{T}_{(\lambda, \mu)}^M \rangle) $

Table 2: Problems 1 - 3 are minimization problems, while Problems 4 - 6 are maximization problems. If $|S| = 1$, then the solution for all minimization problems is the same; similarly for the maximization problems.

graph $\langle G, \mathcal{T} \rangle$, two positive integers λ, μ , and a subset of vertices $S \subseteq V$ which will term *sources*. The objectives we study are formally defined in Table 2.

Delaying. While merging operations affect globally the whole graph; *edge delays* affect only one label of an edge. We parameterize the delay operation by δ ; the maximum delay we can impose on a label of an edge. Hence, a δ -delay on edge uv at time step i changes the label i to $i' \leq i + \delta$. We denote \mathcal{T}_δ^D the schedule produced after applying δ -delay operations on \mathcal{T} . When the number of allowed δ -delay operations is κ , we denote $\mathcal{T}_{(\delta, \kappa)}^D$ produced schedule. Under delaying operations we study MINREACH, MINMAXREACH, and MINAVGREACH.

Max2Sat(3). To produce many of our results we use the problem MAX2SAT(3). An instance of MAX2SAT(3) is a CNF formula ϕ with n boolean variables and m clauses, where each clause involves exactly two variables and every variable appears in at most three clauses. The goal is to maximize the number of satisfied clauses. Without loss of generality we will assume that every variable in ϕ appears exactly one time as a positive literal and at most two times as a negative literal. In [4] it was proven that MAX2SAT(3) is NP-hard and that it is even hard to approximate better than 1.0005.

3 Merging: Minimization problems

In this section we study minimization problems under merging operations. To begin with, we prove that MINREACH under $(2, \mu) - \text{MERGING}$ is NP-hard even when G is a path with many sources. Then, using this result, we explain how to get NP-hardness for any $\lambda \geq 2$. Next, we show how to extend our construction and get a bipartite, planar graph of maximum degree 3 and only one source, and thus we prove NP-hardness for all Problems 1 - 3. Our next result is NP-hardness for the same set of problems on trees with one source. For this result, we derive a new construction. Although all of our results are presented for undirected graphs, they can be trivially extended for directed graphs. In addition, all of our reductions are approximation preserving hence we get that all Problems 1 - 3 are NP-hard even to approximate.




Figure 2: The gadget for the clause (\bar{x}_i, x_j) . The labels on the edges denote the time steps these edges are available; $M = 4m + 4$. The black vertex, c_k , is a source, i.e., it belongs to S .

3.1 MinReach on paths

Construction. We will reduce from MAX2SAT(3). The k -th clause of ϕ will be associated with the time steps $4k, 4k + 1$ and $4k + 2$, while the i -th variable will be associated with the time steps $M + 4i, M + 4i + 1$ and $M + 4i + 2$, where $M = 4m + 4$. For every clause of ϕ we construct a path with nine vertices, where the middle vertex is a source. Consider the path for the k -th clause, that involves the variables x_i and x_j . An example of such a path can be found on Figure 2. The middle piece of the path consists of the vertices $c_k, y_k^l, z_k^l, y_k^r, z_k^r$; where $c_k \in S$. Edge $c_k y_k^l$ has the label $4k$, edges $y_k^l z_k^l$ and $c_k y_k^r$ have the label $4k + 1$, and edge $y_k^r z_k^r$ has the label $4k + 2$. The labels on the left and the right pieces of the path depend on the literals of the variables of the clause. If variable x_i appears in the clause with a positive literal, then we pick an arbitrary side of the path, say the left, and add the label $M + 4i + 1$ to the edge $z_k^l w_k^l$ and the label $M + 4i + 2$ to the edge $w_k^l v_k^l$. If x_i appears in the clause with a negative literal, then we add the label $M + 4i$ to the edge $z_k^l w_k^l$ and the label $M + 4i + 1$ to the edge $w_k^l v_k^l$. In order to create a single path, we arbitrarily connect the endpoints of the paths we created for the clauses; every edge that connects two such paths has label 1. Observe that for the constructed temporal graph, $\langle G, \mathcal{T} \rangle$, the following hold: $\langle G, \mathcal{T} \rangle$ has $9m$ vertices and lifetime $4m + 4n + 6$; every vertex is reachable from a one of sources; at every time step there are at most three edges available; every edge has only one label, i.e., it is available only at one time step. Clearly, the size of $\langle G, \mathcal{T} \rangle$ is polynomial to the size of ϕ and $|S| = m$. We will ask for a $(2, m + n)$ -merging scheme.

Intuition. The correctness of our reduction relies on two ideas. The first idea is that in every subpath, under any $(2, \mu)$ -merging scheme, *at most* four vertices are not reachable from S . In order to make four vertices not reachable within a gadget, the following synergy must happen. Vertex c_k should choose which side of the path to “save”; merging $4k$ with $4k + 1$ make the vertices z_k^l, w_k^l, v_k^l unreachable; merging $4k + 1$ with $4k + 2$ make the vertices z_k^r, w_k^r, v_k^r unreachable. Observe that only one of the two merges can happen, since the two merges together are not independent. Hence, such a merge makes three vertices of one side unreachable. In order to make the v -vertex of the other side unreachable, one extra merge has to happen. This merge will be translated to a truth assignment for a variable, which is the second idea of our reduction.

The merge of $M + 4i$ with $M + 4i + 1$ corresponds to setting x_i to False and the merge of $M + 4i + 1$ with $M + 4i + 2$ to True.

Lemma 1 *If there exists an assignment that satisfies l clauses of ϕ , then there exists a $(2, n+m)$ -merging scheme such that $3m + l$ vertices are not reachable from S .*

Proof. Given an assignment we produce a merging scheme as follows. If variable x_i is False, then merge $M + 4i$ with $M + 4i + 1$; if it is True, we merge $M + 4i + 1$ with $M + 4i + 2$. Then, we consider every path that corresponds to a clause that is satisfiable and we check the side that our previous merges created a “cut” with respect to reachability in the temporal path. For the k -th clause, if there is a cut from the side where the edge with label $4k$ lies, then we merge $4k$ with $4k + 1$; if the cut is from the side where the edge $4k + 2$ lies, then we merge $4k + 1$ with $4k + 2$; if there are cuts in both sides, or there are no cuts then we arbitrarily make one of the two merges. It is not hard to check that under this merging scheme: in every gadget that corresponds to a satisfied clause we have four vertices not reachable, and in every gadget that corresponds to a non satisfied clause only three vertices are not reachable. Hence, we have made $m + n$ 2-merges and there are $4m + l$ not reachable vertices. ■

Lemma 2 *If there exists an optimal $(2, n + m)$ -merging scheme such that $3m + l$ vertices are not reachable from S , then there exists an assignment for ϕ that satisfies l clauses.*

Proof. Firstly, observe that for the time steps $4k - 1, 4k, 4k + 1, 4k + 2$, and $4k + 3$, for every $k \in [m]$, no more than one 2-merge is necessary. This is because at time steps $4k - 1$ and $4k + 3$ there are no edges available, hence any merge that includes these time steps does not change the set of reachable vertices. In addition, this merge has to involve time step $4k + 1$, else it is meaningless as we argued above. The same holds for the time steps $M + 4i - 1, M + 4i, M + 4i + 1, M + 4i + 2$, and $M + 4i + 3$ for every $i \in [n]$. Hence, in any optimal merging scheme there is only one 2-merge for every triple $4k, 4k + 1, 4k + 2$, where $k \in [m]$, and only one 2-merge for every triple $M + 4i, M + 4i + 1, M + 4i + 2$, where $i \in [n]$. Under any such merging scheme, for every k , in the k -th gadget at least three vertices are not reachable (due to the merge that involves $4k + 1$). In addition, at most four vertices are not reachable due to a, potential, extra merge that makes a v -vertex not reachable.

In order to create a truth assignment for the variables of ϕ , we proceed as follows. We check the gadgets that have four vertices unreachable and we focus on the “side” of the gadget that has only one vertex unreachable; for the k -th gadget, it is the left side if $4k + 1$ is merged with $4k$, and it is the right side if $4k + 1$ is merged with $4k + 2$. Assume that it is the left side of the k -th gadget that has only one vertex, the vertex v_k^l unreachable. If edge $w_k^l v_k^l$ has the label $M + 4i + 1$, then we set x_i to False. This means that $M + 4i + 1$ is merged with $M + 4i$; else v_k^l would be reachable. If edge $w_k^l v_k^l$ has the label $M + 4i + 2$, then we set x_i to True. This means that $M + 4i + 1$ is merged with $M + 4i + 2$. By the construction of the gadget, x_i satisfies the k -th clause. Any vertices with undefined value, we set them to True. So, if there are l gadgets with four unreachable vertices, the constructed assignment satisfies l clauses of ϕ . To complete our proof, we need to argue that the truth assignment for the variables is well defined, i.e. we did not set x_i both to True and False. For contradiction assume that the value of x_i is not well defined. This means that $M + 4i + 1$ is merged with $M + 4i$, and that $M + 4i + 1$ is merged with $M + 4i + 2$ as well; a contradiction. ■

The combination of Lemmas 1 and 2 already yield NP-hardness for MINREACH under $(2, \mu)$ -MERGING, when μ is part of the input. Furthermore, our construction allows us to restrict the

search only to maximal merging schemes. In Lemma 1 we can arbitrarily extend the produced merging scheme to a maximal one without any issues in the correctness of the proof, while in Lemma 2, without loss of generality, we can assume that the merging scheme is maximal.

We furthermore explain how to get NP-hardness for any $\lambda \geq 2$. Although our construction does not work when $\lambda > 2$, we can modify it and get the result for any $\lambda \geq 2$, by “spreading” the edges of $\langle G, \mathcal{T} \rangle$ over time. More formally, between any two time steps of $\langle G, \mathcal{T} \rangle$ we add $\lambda - 2$ dummy time steps where no edge is available. The crucial observation is that, after the spreading, any λ' -merge with $\lambda' < \lambda$ does not affect the set reachable from S . Hence, after the spreading we can replace the 2-merges with λ -merges in Lemmas 1 and 2 and get NP-hardness for any $\lambda \geq 2$.

Theorem 1 *MINREACH under (λ, μ) – MERGING is NP-hard for any $\lambda \geq 2$, even when G is a path, and the following constraints hold:*

1. every edge is available only at one time step;
2. at any time step there are at most three edges available;
3. the merging scheme is maximal.

3.2 Graphs with a single source

Next, we explain how to get NP-hardness for the case where $|S| = 1$. This result immediately implies NP-hardness for MINMAXREACH and MINAVGREACH under (λ, μ) – MERGING, since when $|S| = 1$ these problems coincide with MINREACH. It is not hard to get this result, given the previous construction. We can simply add a vertex c_0 in the constructed graph that would be the only source and it is connected with every vertex that used to be a source, so c_0 is connected with all the c -vertices from the previous construction, with an edge available at time step 1. Since the next time step that an edge exists in is time step 4, this means that under any $(2, \mu)$ -merging scheme, at time step 2 every vertex of S has been reached by c_0 . Hence, The NP-hardness follows from the previous construction. However, this is not the strongest result we can get, since this graph has a vertex with non constant degree. Instead, we modify $\langle G, \mathcal{T} \rangle$ as we describe next in order to get a tree of maximum degree 3.




Figure 3: The construction for graphs with one source.

The construction is depicted in Figure 3. We create a path with $m + 1$ vertices where c_0 is the only source. In addition, every vertex k of the path, different than c_0 , is connected with vertex c_k from the construction of the previous section via an edge with label $2k + 2$. The labels in any two consecutive edges of this path differ by two. Finally, every edge of the construction from the previous section is shifted by $2m + 4$, i.e., we add the number $2m + 4$ to its label. The crucial observation is that under any merging scheme that uses 2-merges at time step $2m + 2$ every c_k

vertex for every $k \in [m]$ will be reached. Hence, at time step $2m + 2$, under any $(2, \mu)$ -merging scheme, we get an instance that it is equivalent, with respect to reachability from this time step and on, to the instance of the previous section. Hence, we can get NP-hardness for $\lambda = 2$. Using the trick of “spreading” described before, we can get the result for any $\lambda \geq 2$.

Theorem 2 *Problems 1 - 3 are NP-hard for any $\lambda \geq 2$ even when there exists only one source, G is a tree of maximum degree three, and the Constraints 1- 3 from Theorem 1 hold.*

4 Merging: Maximization problems

In this section we prove NP-hardness for maximization problems. Before we proceed with the exposition of the results though, we should discuss some issues about maximization reachability problems and merging. Any merge weakly decreases the reachability set from the sources. Hence, while in the minimization problems, in principle, we would like to perform as many merges as possible, for maximization problems we would like to perform the minimum number of merges that are allowed. In addition, the reachability set weakly decreases with λ , i.e., the size of the merge. Hence, for maximization problems it is better to do the smallest merge possible, i.e., perform only 2-merges. For this reason, if we get NP-hardness for $(2, \mu)$ -merging schemes, then we can immediately conclude that finding the optimal (λ, μ) -merging scheme is NP-hard, for any $\lambda \geq 2$. So, for maximization problems, we *require* that *at least* μ λ -merges need to happen. This is motivated by distribution networks; the use of the network comes at a cost, thus we would like to use the network as few times as possible.

Again, we prove our results by reducing to MAX2SAT(3). This time though we need to be more careful; in order to make our reduction valid, we should not allow for “dummy” merges, i.e., merges that do not change the reachable vertices from the sources. As before, we first prove NP-hardness for paths with multiple sources. Then, we modify our reduction to get NP-hardness for graphs with one source and we provide a reduction for trees with just one source.

4.1 MaxReach on paths

Construction. We will reduce from MAX2SAT(3). Every variable x_i of ϕ will be associated with the time steps $3i, 3i + 1$, and $3i + 2$. For every variable we create a path of length 5 with ends the vertices y_i and z_i ; this path is depicted at the bottom of Figure 4. In this path, the edges at the ends of the paths have labels $3i + 2$ and the rest of the edges have label $3i + 3$. The paths that are created from variables will be termed *variable-paths*. Both ends of every variable-path are sources. For every clause we create a path with five vertices. So, if the k -th clause involves the variables x_i and x_j , we construct a path with ends the vertices v_k^i and v_k^j , and middle the vertex c_k , which will be termed *c-vertex*. If x_i appears with a positive literal in the clause, then the labels on the two edges that connect v_k^i and c_k are $3i$ and $3i + 1$ respectively; else the labels are $3i + 1$ and $3i + 2$. The top side of Figure 4 depicts the path for the clause (\bar{x}_i, x_j) . Both ends of the path are sources. These paths will be termed *clause-paths*. To create one path, we arbitrarily connect the constructed paths. The temporal graph $\langle G, \mathcal{T} \rangle$ we constructed has $6n + 5m$ vertices, lifetime $3n + 2$, and at any time step at most four edges are available. Furthermore all the vertices are reachable from S .

Intuition. The labels on the variable-paths guarantee that there exists a $(2, n)$ -merging scheme that maximizes the number of reachable vertices such that it does not apply a 2-merge that merges time step $3i + 2$ with $3i + 3$ for any $i \in [n]$. This guarantees two things. Firstly,




Figure 4: Gadgets for MAXREACH on paths. Top: the gadget for the clause (\bar{x}_i, x_j) . Bottom: the gadget for variable x_i .

there exists an optimal $(2, n)$ -merging scheme where all the vertices, except the c -vertices, are reachable from S . Hence, any optimal merging scheme has to maximize the number of reachable c -vertices. Second, given the subset of c -vertices that are reachable under the produced merging scheme, we can easily deduce a truth assignment that satisfies the clauses that correspond to the reachable c -vertices.

In what follows, we will assume that we do not allow any merge that involves time steps 1 and 2. This is because such merges would be “dummy”, since there are no edges available at these time steps. However, this assumption is without loss of generality; we can simply decrease all the labels by 2, hence an edge with label i will get label $i - 2$.

Lemma 3 *If there exists an assignment that satisfies l clauses of ϕ , then there exists a $(2, n)$ -merging scheme such that $6n + 4m + l$ vertices are reachable from S .*

Proof. Given an assignment for the variables of ϕ we produce a $(2, n)$ -merging scheme as follows. If variable x_i is false, then we merge $3i$ with $3i + 1$; else we merge $3i + 1$ with $3i + 2$. Observe, all the vertices in the variable-paths are reachable under this merging scheme. Consider now the clause-paths. Observe that all the vertices except the middle ones, the c -vertices, in these paths are reachable. Hence, we have to argue only about the c -vertices of $\langle G, \mathcal{T} \rangle$. Take a clause-path that corresponds to a satisfiable clause; assume that this is the k -th clause that is satisfied by the variable x_i . Then, observe that the c -vertex of this path is reachable by the vertex v_k^i . If x_i appears as a positive literal, then the path from v_k^i to c_k uses the labels $3i$ and $3i + 1$, while the merging scheme merges $3i + 1$ with $3i + 2$. If x_i appears as a negative literal, then the path from v_k^i to c_k uses the labels $3i + 1$ and $3i + 2$, while the merging scheme merges $3i + 1$ with $3i$. Hence, we can conclude that $6n + 4m + l$ vertices are reachable under the produced merging scheme. ■

Lemma 4 *If there exists an optimal $(2, n)$ -merging scheme such that $6n + 4m + l$ vertices are reachable from S , then there exists an assignment for ϕ that satisfies l clauses.*

Proof. To begin with, observe that it is without loss of generality to assume that the $(2, n)$ -merging scheme that maximizes the reachable set from S does not merge $3i + 2$ with $3i + 3$ for any $i \in [n]$. If $3i + 1$ is not merged, then the merging scheme is not optimal since we can merge $3i + 1$ with $3i + 2$ and increase the reachable set by at least one vertex. If $3i + 1$ is merged, then there exists a j such that $3j + 1$ that is not merged and at least one of $3j$ and $3j + 2$ is not merged. So we replace the merge of $3i + 2$ and $3i + 3$ with a merge that uses $3j + 1$. This way, the number of reachable vertices did not decrease; we gained two vertices by “unmerging” $3i + 2$ with $3i + 3$ and we lost at most one vertex by merging $3j + 1$ with $3j + 2$, or at most two vertices by merging $3j + 1$ with $3j$ (this is due to the fact that every variable appears one time

with a positive literal and at most two times with a negative literal). Thus, since there are n 2-merges, $3i + 1$ is merged with one of $3i$ or $3i + 1$ for every $i \in [n]$.

So, given a $(2, n)$ -merging scheme that satisfies the constraints above, we construct the truth assignment for the variables of ϕ as follows. If $3i + 1$ is merged with $3i$, then we set x_i to False; else we set it to True. Observe that for every clause-path where vertex c_k is reachable under the merging scheme, we get that the corresponding clause is satisfied by the produced assignment; this is due to our construction. It remains to show that there exist l clause paths where the c -vertex is reachable. This is not hard to see. Firstly, all the vertices of the variable-paths are reachable under the merging scheme. Then, for all clause-paths observe that under any $(2, n)$ -merging scheme all the vertices except the c -vertices are reachable; these are $4m$ in total. Hence, it must be true that under the assumed merging scheme l c -vertices are reachable from S . ■

Lemmas 3 and 4 imply that MAXREACH under $(2, \mu)$ - MERGING is NP-hard. As we have already explained, this means that the problem is NP-hard for any $\lambda \geq 2$.

Theorem 3 MAXREACH under (λ, μ) - MERGING is NP-hard for any $\lambda \geq 2$, even when the underlying graph is a path, every edge has one label, and at any time step there exist at most four edges available.

4.2 Graphs with a single source

In this section we modify the construction of the previous section and prove NP-hardness for MAXREACH with only one source. Thus, we get as a corollary NP-hardness for MAXMINREACH and MAXAVGREACH.

Construction. The idea is similar to the one we used before: we construct a tree whose leaves are the sources of the previous construction, so we can use the result of the previous section. To achieve this, we will use the gadget depicted in Figure 5. This gadget consists of a line of length $2m + n$. The left end of the line has degree 1 and it is the only source, while the rest of the vertices of the line have degree 2 or 3. On the degree-3 vertices we add the clause-paths and the variable-paths, depicted as c -paths and v -paths in the figure, by adding $m + n + 1$ to their labels. The end points of c -paths are connected with the j -th and the $j + 2$ -th vertices of the gadget where the $j + 1$ -th vertex has degree 2; this is done in order to get a bipartite graph. The end points of v -paths are connected with the j -th and the $j + 1$ -th vertices. The crucial observation is that in an optimal $(2, n)$ -merging scheme we do not merge any time step lower than $2m + n + 1$. Any such merge will make unreachable at least 5 or 6 vertices (the vertices of a clause-path or a variable-path) while a merge at a time step higher than $2m + n + 1$ will make unreachable at most 2 vertices. Hence, at time step $2m + n + 1$ all the vertices that correspond to the endpoints of the paths from the previous section have been reached by the source, so we can use the proof from the previous section and get our result.

Theorem 4 Problems 4 - 6 are NP-hard even when there is one source, the underlying graph is bipartite, it has maximum degree 3, every edge has one label, and at any time step there are at most four edges available.

4.3 Trees with one source

In this section we prove that MAXREACH is NP-hard even on trees with only one source. Again, our reduction is from MAX2SAT(3). First, we explain how to get NP-hardness for forests where




Figure 5: The construction used to prove NP-hardness for MAXREACH with one source. c-path and v-path stand for a clause-path and variable-path respectively from the previous section where we add $n + m + 1$ to the labels of these paths. The filled vertex is the only source of the graph.




Figure 6: The gadgets we use in Section 4.3, where $M = 3m + 2$. The left gadget corresponds to the k -th clause of ϕ equal to (x_i, \bar{x}_j) ; the middle gadget is for the k -th clause; the right gadget is for the i -th variable

every connected component of the forest has only one source. Then, using the idea from the previous section, we connect the components of the forest and create a single tree with only one source.

Construction. We will use a similar approach as before and we associate each variable with three consecutive time steps. This time though we associate every clause with three consecutive time steps as well. For every clause we create a tree with 9 vertices where only one of them is a source; an example of this tree is depicted in the left part of Figure 6. Each such tree consists of three pieces: the middle piece, the left piece, and the right piece. For the k -th clause, the middle piece consists of the vertices $c_k, y_k^l, z_k^l, y_k^r, z_k^r$; c_k is a source. Edge $c_k y_k^l$ has the label $3k$, edges $y_k^l z_k^l$ and $c_k y_k^r$ have the label $3k + 1$, and edge $y_k^r z_k^r$ has the label $3k + 2$. The labels on the left and the right pieces of the path depend on the literals of the variables of the clause. If variable x_i appears in the clause with a positive literal, then we pick an arbitrary side of the path, say the left, and add the label $M + 3i$ to the edge $z_k^l w_k^l$ and the label $M + 3i + 1$ to the edge $w_k^l v_k^l$, where $M = 3m + 2$. If x_i appears in the clause with a negative literal, then we add the label $M + 3i + 1$ to the edge $z_k^l w_k^l$ and the label $M + 3i + 2$ to the edge $w_k^r v_k^r$. In addition, we create the following. For every $k \in [m]$ we create eight paths of length 2; the middle part of Figure 6. The one end point of every such path is a source. The labels of the edges for every path, from the source to the other end, are $3k + 2$ and $3k + 3$. For every $i \in [n - 1]$ we create eight paths of length 2; the right part of Figure 6. The one end point of every path is a source. The labels of

the edges for every path, from the source to the other end, are $M + 3i + 2$ and $M + 3i + 3$. The constructed temporal graph $\langle G, \mathcal{T} \rangle$ has $33m + 24n$ vertices and lifetime $3m + 3n + 4$. All the vertices of the graph are reachable from S . We ask for a $(2, m + n)$ -merging scheme. Observe that at any time step there exist at most 10 edges available. Again, we assume, without loss of generality, that no merging that involves labels 1 and 2 is allowed.

Intuition. The high level idea is that in any optimal $(2, m + n)$ -merging scheme all the vertices of the paths of length 2 are reachable; this is guaranteed by the number of copies of the paths and the choice of the labels in these paths. Hence, an optimal merging scheme maximizes the number of reachable vertices in the gadgets for the clauses. In every gadget though, at most 6 vertices are reachable under any optimal merging scheme. If 6 vertices are reachable, then we can easily deduce an assignment that satisfies this clause.

Lemma 5 *If there exists an assignment for the variables of ϕ that satisfies l clauses, then there exists a $(2, m + n)$ -merging scheme for $\langle G, \mathcal{T} \rangle$ such that $29m + 24n + l$ vertices are reachable from S .*

Proof. Given an assignment that satisfies l clauses we create the a merging scheme as follows. If x_i is true, then we merge $M + 3i + 1$ with $M + 3i + 2$, else we merge $M + 3i + 1$ with $M + 3i$. In addition, we consider every clause separately. If the k -th clause is satisfied by the variable that corresponds to the left piece of the gadget we created for the clause, then we merge $3k$ with $3k + 1$; else we merge $3k + 1$ with $3k + 2$. If the clause is not satisfied, then we merge $3k$ with $3k + 1$. Observe that under this merging scheme all the vertices of the length-two paths are reachable. In addition, in every tree that corresponds to a clause at least three vertices of y_k^l, y_k^r, z_k^l , and z_k^r are reachable from c_k . In addition, for every satisfiable clause the vertices w_k^l, v_k^l , and u_k^l are reachable if the clause is satisfied by the “left” variable, or the vertices w_k^r, v_k^r , and u_k^r are reachable if the clause is satisfied by the “right” variable. It is not hard to see that under this merging scheme $29m + 24n + l$ vertices are reachable from S . ■

Lemma 6 *If there exists a reachability maximizing $(2, m + n)$ -merging scheme such that $29m + 24n + l$ vertices are reachable from S in $\langle G, \mathcal{T} \rangle$, then there exists a truth assignment that satisfies l clauses of ϕ .*

Proof. To begin with observe that there does not exist an optimal $(2, m + n)$ -merging scheme that merges $3k + 2$ with $3k + 3$ for any $k \in [m]$. Any such merge makes eight vertices of $\langle G, \mathcal{T} \rangle$ unreachable; these vertices belong to a subset of paths of length 2 we constructed. Observe that even if such a merge was made, then we could find a different merge and contradict the optimality of the merging scheme. For example, if we had merged $3k' + 1$ with $3k'$, or with $3k' + 2$, for some $k' \in [m]$, then we would make at most 4 vertices unreachable. Alternatively, if we had merged $M + 3i + 1$ with $M + 3i$ or with $M + 3i + 2$, then we would make at most 2 vertices unreachable. At least one of these two merges can be made without changing the rest of the merges of the merging scheme due to the fact that the lifetime of the graph is $3m + 3n + 4$ and we ask for only $m + n$ 2-merges. For the same reason there does not exist a $(2, m + n)$ -merging scheme that merges $M + 3i + 2$ with $M + 3i + 3$ for any $i \in [n]$. Hence, we can conclude that $3k + 1$ is merged with one of the $3k$ and $3k + 2$ for every $k \in [m]$. In addition, we conclude that $M + 3i + 1$ is merged with either $M + 3i$, or $M + 3i + 2$, for every $i \in [n]$. Observe however, that independently of how $3k + 1$ is merged, 3 vertices of the gadget for the k -th clause become unreachable. More specifically, if $3k + 1$ is merged with $3i$, then the vertices z_k^l, w_k^l , and v_k^l

become unreachable; if $3k + 1$ is merged with $3i + 2$, then the vertices z_k^r, w_k^r , and v_k^r become unreachable. In addition, observe that the merging scheme described above guarantees that in the k -th gadget exactly one of the w_k^l and w_k^r will be reachable independently from the rest of the merges. Hence, from every gadget at least 5 vertices are reachable under the optimal merging scheme. With a simple counting argument we can see that from the paths of length 2 we have $24m + 24n$ reachable vertices and $5m$ vertices reachable in the gadgets.

Hence, there exist l v -vertices that are reachable under an optimal merging scheme. We will use the edges adjacent to these vertices to create a truth assignment for the variables of ϕ . If the vertex v_k^l is reachable and the edge $w_k^l v_k^l$ has the label $M + 3i + 1$, then we set x_i to True; else we set x_i to False. The value of x_i is well defined. If the edge $w_k^l v_k^l$ has the label $M + 3i + 1$ and v_k^l is reachable, then it means that $M + 3i + 1$ is merged by $M + 3i + 2$, hence there cannot be a v -vertex that it is reachable and it is adjacent to an edge with label $M + 3i + 2$; the path will be “cut” due to the aforementioned merge. Similarly, if the edge $w_k^l v_k^l$ has the label $M + 3i + 2$ and v_k^l is reachable, then it means that $M + 3i + 1$ is merged with $M + 3i$ and there does not exist a reachable v -vertex adjacent to an edge with label $M + 3i + 1$. If there exist variables with undefined value, then we set their value to True. The constructed truth assignment satisfies l clauses of ϕ ; if v_k^l is reachable and the edge $w_k^l v_k^l$ has the label $M + 3i + 1$, then the k -th clause is satisfied by variable x_i . Since there are l v -vertices that are reachable, the assignment satisfies l clauses of ϕ . ■

The combination of the two lemmas above already yield NP-hardness for MAXREACH on forests. However, we can use the gadget from Figure 5 to get a tree of maximum degree 3, with only one source. Hence, we immediately get NP-hardness for MAXMINREACH and MAXAVGREACH on trees with a single source.

Theorem 5 *Problems 4 - 6 are NP-hard even when there exists only one source, G is a tree, has maximum degree three, and every edge has only one label.*

5 Delaying

In this section we study *edge-independent* delays. Firstly, we show that when the number of allowed delays is bounded, then the minimization problems are NP-hard. Then, we study the case where the number of δ -delaying operations is unbounded. In contrary to unbounded edge deletions where the solution to the problems become trivial by essentially isolating every source, an unbounded number of δ -delays does not trivialize the problems and most importantly does not destroy the underlying network.

Theorem 6 *MINREACH, MINMAXREACH, and MINAVGREACH are NP-hard under δ - DELAYING, for any $\delta \geq 1$, when the number of operations is bounded by κ . In addition, they are $W[1]$ -hard, when parameterized by κ .*

Proof. To begin with, we make the following observation that will be helpful later in the proof. For any temporal graph delaying any edge with label t_{\max} does not decrease the reachability set of any source. This is because if vertex u from edge uv is already reached by a source before t_{\max} , then v will be reached from uv independently from any delay on this edge. Hence, if we have a temporal graph with $t_{\max} = 2$, we can assume without loss of generality that in any optimal solution we delay edges with label 1 *only*.

Firstly, we study MINREACH δ, κ - DELAYING. We will prove that the problem is hard even when $|S| = 1$, i.e., there is only one source. As before, this result implies that Problems and are hard as well.

To prove our result we follow a similar technique as the authors of [14] and we reduce from CLIQUE. An instance of CLIQUE consists of a graph $G' = (V', E')$ and an integer k . We want to decide if G' contains a clique of size k , i.e., a complete subgraph with k vertices. We construct a temporal graph $\langle G, \mathcal{T} \rangle$, where $G = (V, E)$, as follows. For every vertex and every edge of G' we create a vertex in G ; we term the former ones v -vertices and the latter e -vertices. In addition, we create the vertex v^* which will be the the only source. Hence, $|V| = |V'| + |E'| + 1$. The edge-set E consists of:

- v^*v , where v is a v -vertex;
- vu , where v is a v -vertex, u is a e -vertex and the corresponding edge from E' contains the vertex v .

Every edge in G that contains v^* has label 1 and the rest of the labels have label 2. Finally, we set $\kappa = k$ and $\delta \geq 1$. We claim that G' contains a clique of size k if and only if $|\text{reach}(v^*, \langle G, \mathcal{T}_{(\delta, \kappa)}^D \rangle)| = |V| - k^2$.

Assume now that G' contains a clique k and let $X = \{v_1, \dots, v_k\}$ denote the set of its vertices and let Y denote the edges of the clique. Then, we delay by δ all the edges v^*v , where $v \in X$. We claim that no vertex in Y is reachable from v^* . To see why this is the case, let $y \in Y$ that correspond to the edge $v_i v_j$ where $v_i \in X$ and $v_j \in Y$. Observe that by the construction of $\langle G, \mathcal{T} \rangle$ there are only two temporal paths from v^* to y : $v^* \xrightarrow{1} v_i \xrightarrow{2} y$ and $v^* \xrightarrow{1} v_j \xrightarrow{2} y$. Hence, after delaying the edges v^*v_i and v^*v_j , we see that none of these paths is valid any more. Since $|Y| = k^2$, we get that v^* reaches $|V| - k^2$ vertices.

To prove the other direction, recall that since the $\langle G, \mathcal{T} \rangle$ has lifetime 2, it suffices to delay only edges with label 1. So assume that we can delay $\kappa = k$ edges with label 1 by δ such that $|\text{reach}(v^*, \langle G, \mathcal{T}_{(\delta, \kappa)}^D \rangle)| = |V| - k^2$. Let v^*v_1, \dots, v^*v_k be the edges we delayed. Observe that all the v -vertices are reachable by v^* . Hence, after the delays, k^2 e -vertices are not reachable by v^* . So, by the construction of $\langle G, \mathcal{T} \rangle$, in G' there should be k^2 edges between the vertices v_1, \dots, v_k . Hence, G' contains a clique of size k . This establishes the NP-hardness of MINREACH δ, κ - DELAYING. Since $|S| = 1$, the hardness for the other two problems follows. Finally, since CLIQUE is $W[1]$ -complete and we have a parameterized m -reduction, we get the $W[1]$ -hardness for the problems when they are parameterized by the number of allowed delays κ . ■

Next, we provide a polynomial-time algorithm for the minimization problems under δ - DELAYING, for any $\delta \geq 1$. Let us define the *reachability network* which will be used by our algorithm. Given a temporal graph $\langle G, \mathcal{T} \rangle$ with lifetime t_{\max} and a set of sources S , for every $t \leq t_{\max}$, we define $RV_t(\langle G, \mathcal{T} \rangle, S)$ to be the set of vertices that are reached at time t for *first time* from a vertex of S . Put formally, $v \in RV_t(\langle G, \mathcal{T} \rangle, S)$, if there exists $s \in S$ such that the *earliest arrival* path from s to v arrives at time t and moreover for every $t' < t$ there is no path from any $s \in S$ to v that arrives at time t' . Additionally, we define $RE_t(\langle G, \mathcal{T} \rangle, S)$ to be the set of temporal edges with label t that are adjacent to vertices in $RV_t(\langle G, \mathcal{T} \rangle, S)$. Observe that we can decide if $v \in RV_t(\langle G, \mathcal{T} \rangle, S)$ can be computed via computing the earliest arrival paths between every $s \in S$ and v , which can be done in polynomial time with respect to the size of $\langle G, \mathcal{T} \rangle$ [34]. Similarly, we can efficiently compute $RE_t(\langle G, \mathcal{T} \rangle, S)$. Finally, we say that it is δ -possible to change a label of an edge from t to $t + 1$ if the difference between $t + 1$ and the original label of the edge is at most δ .

The next observation follows from the fact that a delay on any edge in $RE_t(\langle G, \mathcal{T}_\delta^D \rangle, S)$ weakly decreases $RV_{t+1}(\langle G, \mathcal{T}_\delta^D \rangle, S)$. Hence, delaying as many as possible from these edges minimizes the number of vertices that are reached by S at time step $t + 1$.

```

 $\mathcal{T}_\delta^D \leftarrow \mathcal{T}$ ;
for  $1 \leq t < t_{\max}$  do
    Compute  $RE_{t+1}(\langle G, \mathcal{T}_\delta^D \rangle, S)$ 
    foreach edge  $e \in RE_{t+1}(\langle G, \mathcal{T}_\delta^D \rangle, S)$  do
        Change the label  $i$  of  $e$  to  $t + 1$ , if this is  $\delta$ -possible;
        Update  $\mathcal{T}_\delta^D$ ;
    end
end

```

Algorithm 1: Algorithm for δ – DELAYING.




Figure 7: In the temporal graph $\langle G, \mathcal{T} \rangle$ above with $S = s$, we have $RV_1(\langle G, \mathcal{T} \rangle, s) = \{x, y\}$, $RV_2(\langle G, \mathcal{T} \rangle, s) = \{z\}$, $RE_1(\langle G, \mathcal{T} \rangle, s) = \{sx, sy\}$, and $RE_2(\langle G, \mathcal{T} \rangle, s) = \{yz\}$. Any other set is empty.

Observation 1 For any $1 \leq t < t_{\max}$, if we can delay only temporal edges with label t , Algorithm 1 minimizes $RV_{t+1}(\langle G, \mathcal{T}_\delta^D \rangle, S)$.

Observation 1 can be used as an intermediate step in an induction argument to prove optimality of Algorithm 1.

Theorem 7 Algorithm 1 is optimal for MINREACH, MINMAXREACH, MINAVGREACH under δ – DELAYING.

Proof. To prove our theorem, we will prove by induction that at any time step $1 \leq t < t_{\max}$ Algorithm 1 minimizes $\sum_{i=1}^t RV_{i+1}(G, \mathcal{T}_\delta^D, S)$, i.e., it minimizes the number of reachable vertices until time $t + 1$. For $t = 1$ the claim holds from Observation 1. So, assume that it holds up to time step t . This means that $\sum_{i=1}^t RV_{i+1}(G, \mathcal{T}_\delta^D, S) \leq \sum_{i=1}^t RV_{i+1}(G, \mathcal{T}_\delta^t, S)$ where \mathcal{T}_δ^t is any other δ -delaying scheme. In addition, from observation 1 we know that $RV_{t+2}(G, \mathcal{T}_\delta^D, S) \leq RV_{t+1}(G, \mathcal{T}_\delta^t, S)$. Hence, if we add these two inequalities, we get that $\sum_{i=1}^{t+1} RV_{i+1}(G, \mathcal{T}_\delta^D, S) \leq \sum_{i=1}^{t+1} RV_{i+1}(G, \mathcal{T}_\delta^t, S)$. This means that Algorithm 1 is optimal up to time step $t + 1$ and the theorem follows. ■

6 Discussion

Our hardness results immediately imply, or can be easily extended to prove, several other interesting results. Firstly, we observe that all of our reductions are approximation preserving, thus since we use MAX2SAT(3), we get that there are no approximation schemes for these problems, unless $P = NP$. In addition, we can add directions to the edges of the gadgets without breaking the correctness of our proofs. Another, less trivial, observation is that in our reductions we create unit disk graphs; a type of graphs prominently used in epidemics.

Our work creates many interesting directions for future research. The most obvious is to get approximation algorithms for the problems we study. In another, broader front, our paper

introduces a new conceptual direction in temporal graphs. Given a temporal network with an existing solution for a problem, can we utilize the current infrastructure in a better way and improve the solution without significantly changing the network?

Acknowledgements

The authors would like to thank Valentin Bura and Vladimir Gusev for several interesting discussions on reachability sets of temporal graphs at the beginning of the project.

References

- [1] Infection prevention and control and preparedness for covid-19 in healthcare settings. Technical report, European Centre for Disease Prevention and Control, 2020.
- [2] E. C. Akrida, J. Czyzowicz, L. Gasieniec, L. Kuszner, and P. G. Spirakis. Temporal flows in temporal networks. *Journal of Computer and System Sciences*, 103:46 – 60, 2019.
- [3] R. E. Allen, A. A. Clark, J. A. Starek, and M. Pavone. A machine learning approach for real-time reachability analysis. In *Proc. of IROS*, pages 2202–2208, 2014.
- [4] P. Berman and M. Karpinski. On some tighter inapproximability results. In *Proc. of ICALP*, pages 200–209. Springer, 1999.
- [5] A. Braunstein and A. Ingrosso. Inference of causality in epidemics on temporal contact networks. *Scientific reports*, 6:27538, 2016.
- [6] D. Brockmann and D. Helbing. The hidden geometry of complex, network-driven contagion phenomena. *Science*, 342(6164):1337–1342, 2013.
- [7] M. G. Buhnerkempe, M. Tildesley, et al. The impact of movements and animal density on continental scale cattle disease outbreaks in the United States. *PLoS One*, 9(3):e91724, 2014.
- [8] A. Casteigts, R. Klasing, Y. M. Neggaz, and J. G. Peters. Computing parameters of sequence-based dynamic graphs. *Theory of Computing Systems*, 63(3):394–417, Apr 2019.
- [9] J. Chen, H. Molter, M. Sorge, and O. Suchý. Cluster Editing in Multi-Layer and Temporal Graphs. In *Proc. of ISAAC*, volume 123, pages 24:1–24:13, 2018.
- [10] V. Colizza, A. Barrat, M. Barthélemy, and A. Vespignani. The role of the airline transportation network in the prediction and predictability of global epidemics. *Proceedings of the National Academy of Sciences*, 103(7):2015–2020, 2006.
- [11] J. Enright and R. R. Kao. Epidemics on dynamic networks. *Epidemics*, 24:88 – 97, 2018.
- [12] J. Enright and K. Meeks. Changing times to optimise reachability in temporal graphs. *CoRR*, abs/1802.05905, 2018.
- [13] J. Enright and K. Meeks. Deleting edges to restrict the size of an epidemic: a new application for treewidth. *Algorithmica*, 80(6):1857–1889, 2018.

- [14] J. Enright, K. Meeks, G. B. Mertzios, and V. Zamaraev. Deleting edges to restrict the size of an epidemic in temporal networks. In *Proc. of MFCS*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
- [15] T. Erlebach, M. Hoffmann, and F. Kammer. On temporal graph exploration. In *Proc. of ICALP*, pages 444–455, 2015.
- [16] T. Erlebach and J. T. Spooner. Faster Exploration of Degree-Bounded Temporal Graphs. In *Proc. of MFCS*, volume 117, pages 36:1–36:13, 2018.
- [17] M. C. Gates and M. E. Woolhouse. Controlling infectious disease through the targeted manipulation of contact network structure. *Epidemics*, 12:11–19, 2015.
- [18] A. E. Jones, J. Turner, C. Caminade, A. E. Heath, M. Wardeh, G. Kluiters, P. J. Diggle, A. P. Morse, and M. Baylis. Bluetongue risk under future climates. *Nature Climate Change*, 9(2):153, 2019.
- [19] D. Kempe, J. Kleinberg, and A. Kumar. Connectivity and inference problems for temporal networks. *Journal of Computer and System Sciences*, 64(4):820–842, 2002.
- [20] G. B. Mertzios, O. Michail, I. Chatzigiannakis, and P. G. Spirakis. Temporal network optimization subject to connectivity constraints. In *Proc. of ICALP*, pages 657–668, 2013.
- [21] O. Michail. An introduction to temporal graphs: An algorithmic perspective. *Internet Mathematics*, 12(4):239–280, 2016.
- [22] O. Michail and P. G. Spirakis. Traveling salesman problems in temporal graphs. *Theoretical Computer Science*, 634:1 – 23, 2016.
- [23] A. Mitchell, D. Bourn, J. Mawdsley, W. Wint, R. Clifton-Hadley, and M. Gilbert. Characteristics of cattle movements in britain—an analysis of records from the cattle tracing system. *Animal Science*, 80(3):265–273, 2005.
- [24] S. Mohr, M. Deason, M. Churakov, T. Doherty, and R. Kao. Manipulation of contact network structure and the impact on foot-and-mouth disease transmission. *Preventive veterinary medicine*, 157:8–18, 2018.
- [25] R. Niskanen, I. Potapov, and J. Reichert. On decidability and complexity of low-dimensional robot games. *Journal of Computer and System Sciences*, 107:124 – 141, 2020.
- [26] M. Nöremark and S. Widgren. Epicontacttrace: an r-package for contact tracing during livestock disease outbreaks and for risk-based surveillance. *BMC veterinary research*, 10(1):71, 2014.
- [27] I. Potapov. From post systems to the reachability problems for matrix semigroups and multicounter automata. In *Proc. of DLT*, pages 345–356, 2004.
- [28] H.-H. Thulke, D. Eisinger, and M. Beer. The role of movement restrictions and pre-emptive destruction in the emergency control strategy against CSF outbreaks in domestic pigs. *Preventive veterinary medicine*, 99(1):28–37, 2011.
- [29] J. Turner, R. G. Bowers, and M. Baylis. Modelling bluetongue virus transmission between farms using animal and vector movements. *Scientific reports*, 2:319, 2012.

- [30] J. Turner, A. Jones, A. Heath, M. Wardeh, C. Caminade, G. Kluiters, R. Bowers, A. Morse, and M. Baylis. The effect of temperature, farm density and foot-and-mouth disease restrictions on the 2007 uk bluetongue outbreak. *Scientific reports*, 9(1):112, 2019.
- [31] E. Valdano, L. Ferreri, C. Poletto, and V. Colizza. Analytical computation of the epidemic threshold on temporal networks. *Physical Review X*, 5(2):021005, 2015.
- [32] E. Valdano, C. Poletto, A. Giovannini, D. Palma, L. Savini, and V. Colizza. Predicting epidemic risk from past temporal contact data. *PLoS computational biology*, 11(3):e1004152, 2015.
- [33] J. Whitbeck, M. Dias de Amorim, V. Conan, and J.-L. Guillaume. Temporal reachability graphs. In *Mobicom*, pages 377–388, 2012.
- [34] H. Wu, J. Cheng, S. Huang, Y. Ke, Y. Lu, and Y. Xu. Path problems in temporal graphs. *Proc. of VLDB Endowment*, 7(9):721–732, 2014.
- [35] P. Zschoche, T. Fluschnik, H. Molter, and R. Niedermeier. The Complexity of Finding Small Separators in Temporal Graphs. In *Proc. of MFCS*, volume 117, pages 45:1–45:17, 2018.