# Approximating Optimal Labelings for Temporal Connectivity (Extended Abstract)

Daniele Carnevale[1], Gianlorenzo D'Angelo[1] and Martin Olsen[2]

[1]Gran Sasso Science Institute, Viale Francesco Crispi, 7, L'Aquila, 67100, Italy
[2]Aarhus University, Birk Centerpark 15, 7400, Herning, Denmark

**Abstract**

Temporal graphs, in which edges are annotated with time-labels indicating their availability at specific time steps, provide a flexible model for dynamic networks. Two nodes are said to be *temporally connected* if there exists a path between them such that the edges are traversed in strictly increasing order of their time-labels. We study the *Minimum Aged Labeling* (MAL) problem, which consists of assigning time-labels so that all node pairs are temporally connected within a global deadline $a$, while minimizing the total number of labels used. MAL highlights the trade-off between ensuring temporal connectivity and minimizing the number of edge activations under strict resource constraints. It has significant implications for energy-aware scheduling in domains such as logistics, transportation, and communication, where reducing the number of active time slots can directly translate into lower energy use, reduced emissions, or infrastructure simplification.

Our results establish strong inapproximability bounds for MAL: we prove that no algorithm can approximate the minimum number of labels within a factor better than $O(\log n)$ for $a \geq 2$, unless P $=$ NP, and not within $2^{\log^{1-\epsilon} n}$ for $a \geq 3$, unless NP $\subseteq$ DTIME$(2^{\text{polylog}(n)})$. Subsequently, we develop approximation algorithms that, under certain assumptions, almost match these lower bounds. Notably, the approximation performance depends on the relationship between $a$ and the diameter of the input graph.

Moreover, we establish a connection with the classical *Diameter Constrained Spanning Subgraph* (DCSS) optimization problem on static graphs and prove that our hardness results apply to DCSS.

**Keywords**

Temporal graphs, Approximation algorithms, Scheduling, Routing

## 1. Introduction

We study a scheduling problem on dynamic networks, inspired by practical applications in logistics, distribution scheduling, and information diffusion within social networks. Consider a real-world scenario of parcel delivery where a warehouse $W$ serves three cities arranged in a star topology. Each city has parcels destined for the other cities, and for each pair of cities $(A, B)$, there is at least one parcel that must be sent from $A$ to $B$. A fleet of vehicles located at $W$ is responsible for transporting parcels between the warehouse and the cities. Each vehicle can depart from $W$ at any hour. Upon arrival at city $A$, a vehicle delivers parcels destined for $A$ that were previously stored at $W$, collects parcels originating from $A$, and returns to the warehouse. A round trip from $W$ to a city and back is referred to as a *trip*. For simplicity, we assume travel times are negligible. When a vehicle $V_1$ returns from city $A$ to the warehouse, it deposits all parcels originating from $A$ and destined for other cities. When another vehicle $V_2$ departs towards city $B$, it carries all parcels currently stored at $W$ with final destination $B$. If the trip of $V_1$ is scheduled before that of $V_2$, parcels from $A$ to $B$ are successfully delivered; otherwise, these parcels must wait for the next trip. Performing a thorough scheduling of all vehicle trips can simultaneously reduce delivery times and operational costs such as the number of trips, vehicle usage, fuel consumption, and greenhouse gas emissions.

For example, the 1st schedule in Figure 1 requires 6 trips to deliver all parcels, with the last deliveries completed at 8 a.m. The 2nd schedule minimizes the latest delivery time, ensuring that all parcels are

*2nd Workshop on Green-Aware Artificial Intelligence, 28th European Conference on Artificial Intelligence (ECAI 2025), October 25–30, 2025, Bologna, Italy*
daniele.carnevale@gssi.it (D. Carnevale); gianlorenzo.dangelo@gssi.it (G. D'Angelo); martino@btech.au.dk (M. Olsen)
0009-0006-3997-869X (D. Carnevale); 0000-0003-0377-7037 (G. D'Angelo); 0000-0002-2740-545X (M. Olsen)

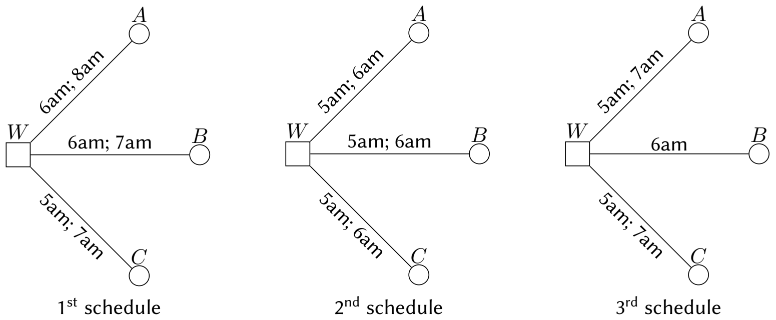


**Figure 1:** The scheduled time of trips are reported close to the edges. **1st schedule**: All 6 scheduled trips are necessary to deliver all parcels, with the last parcel delivered at 8 a.m. **2nd schedule**: All parcels are first deposited at $W$ in the 3 trips at 5 a.m., and then delivered from $W$ to their destinations in 3 trips at 6 a.m. While 6 trips are still required, the last parcel is delivered by 6 a.m. **3rd schedule**: Parcels from cities $A$ and $C$ are deposited at $W$ in the 5 a.m. trips. The 6 a.m. trip to $B$ delivers the parcels to $B$ that were deposited earlier at $W$, and collects the parcels from $B$, bringing them back to $W$. Finally, the 7 a.m. trips deliver the parcels from $B$ to $A$ and $C$. Only 5 trips are used, but the last parcel is delivered at 7 a.m.

delivered by 6 a.m., two hours earlier than in the first schedule, while still requiring 6 trips. The 3rd schedule, instead, minimizes the total number of trips, reducing them to 5, but delays the last delivery to 7 a.m., one hour later than in the second case.

The scheduling problem becomes significantly more complex when considering a general network in which each vertex can act both as a warehouse and as a city, and where connections between vertices are represented by an arbitrary graph. Moreover, similar challenges arise in other domains such as distribution scheduling [1] and information spreading, where the goal is to arrange a small number of meetings among employees of a company so that each individual can share information with any other within a given time frame (see [2]).

Motivated by these applications, we consider the following question: *What is the minimum number of trips required to deliver all parcels within a given time frame?*

We model a schedule of trips along the edges of a given network using a *temporal graph*, where the scheduling time of a vehicle is represented as an edge label. A path in this graph is considered valid (or *temporal*) only if its edges are traversed in strictly increasing order of their labels. We then study the optimization problem of assigning the minimum number of labels to the edges of the given network so that every pair of vertices is connected by a temporal path, and the largest label does not exceed a given integer $a$, referred to as the *maximum allowed age*. This problem, called *Minimum Aged Labeling* (MAL), was introduced by Mertzios et al. [3], who proved that it is APX-hard on *directed graphs*. Later, Klobas et al. [4] showed that MAL is NP-complete even on *undirected graphs*. To the best of our knowledge, no previous hardness or approximation algorithm results are known for MAL on undirected graphs. Furthermore, the reduction used to prove the APX-hardness on directed graphs [3] does not easily extend to the undirected setting, since it relies heavily on edge directions to constrain vertex reachability. This paper explores the approximability of MAL on undirected graphs, highlighting how the problem's approximation complexity varies with respect to the parameters $a$ and the diameter $D_G$ of the input graph. MAL also has a theoretical motivation, as it can be seen as a *dynamic version* of the classical *Diameter Constrained Spanning Subgraph* (DCSS) problem. The DCSS problem asks to find a spanning subgraph $H$ of a graph $G$ such that the diameter of $H$ does not exceed a given integer, while minimizing the number of edges in $H$.

## 2. Related Work

Due to their versatility, temporal graphs have been studied from various perspectives and under different names, such as *dynamic*, *evolving*, or *time-varying* graphs or networks (see [5]). One area that has attracted significant attention, mainly motivated by virus-spread minimization (e.g., [6, 7]), concerns

**Table 1**
Overview of approximation hardness for the MAL problem, classified by parameter $a$.

| | value of $a$ | Approximation hardness | Complexity assumption |
|---|---|---|---|
| $a = D_G$ | $a = 2$ | $\Theta(\log n)$ | $\mathsf{P} \neq \mathsf{NP}$ |
| | $a \geq 3$ | $\Omega(\log n)$ | $\mathsf{P} \neq \mathsf{NP}$ |
| | | $\Omega(2^{\log^{1-\epsilon} n})$ | $\mathsf{NP} \not\subseteq \mathsf{DTIME}(2^{\mathsf{polylog}(n)})$ |

the modification of temporal networks to optimize certain objectives; a recent survey is in [8]. Several types of operations have been considered, including delaying labels and merging consecutive time steps [1], edge- and label-deletion [9], and altering the relative order of time labels [10]. Furthermore, Molter et al. [11] analyzed how the choice between edge-deletion and delaying affects the parameterized complexity of the reachability-minimization objective. In [2], the authors studied a problem related yet orthogonal to MAL, where the goal is to minimize the maximum time needed for a subset of vertices to reach every other vertex by shifting labels. Klobas et al. [4] considered a generalization of MAL in which only a subset of *terminal* vertices must be temporally connected within the given maximum allowed age $a$, and proved its $W[1]$-hardness when parameterized by the number of labels.

# 3. Our results

This extended abstract outlines our main results on the MAL problem, covering both hardness of approximation and the design of efficient approximation algorithms. Further details can be found in the conference version of this work [12], while a full version is available on ArXiv [13]. We denote by $D_G$ the diameter of graph $G$.

## 3.1. Hardness of Approximation

We begin by showing that, for any fixed maximum labeling age $a \geq 2$, the MAL problem cannot be approximated within a factor better than $O(\log n)$, unless P = NP. Furthermore, under a different complexity assumption, we prove a stronger inapproximability bound: for any $\epsilon \in (0, 1)$, there exists no $2^{\log^{1-\epsilon} n}$-approximation algorithm for MAL, unless NP $\subseteq$ DTIME($2^{\text{polylog}(n)}$), even when $a$ is fixed to a value at least 3.

These results advance our understanding of the computational complexity of MAL in two directions: from the exact computation perspective and from the approximation point of view.

(1) From the perspective of exact computation, the NP-hardness result in [4] applies only when $a = D_G = 10$. In contrast, we prove that MAL remains NP-hard for every fixed $a \geq 2$ (still requiring $a = D_G$). This completes the complexity classification of MAL with respect to the parameter $a$, since the case $a = 1$ is straightforward. Additionally, our reduction is considerably simpler than that of [4].

(2) From an approximation standpoint, the reduction in [3] establishes that MAL is APX-hard when $a = D_G = 9$ in directed graphs. We strengthen this result by proving two stronger inapproximability bounds: first, MAL cannot be approximated within a logarithmic factor unless P = NP; second, under a stronger complexity assumption, no $2^{\log^{1-\epsilon} n}$-approximation algorithm exists for any $\epsilon \in (0, 1)$. The latter bound indicates that achieving better than a polynomial-factor approximation is unlikely. Furthermore, these hardness results are established for undirected graphs (with the condition $a = D_G$), and hold for all fixed values of $a \geq 2$ and $a \geq 3$, respectively. Finally, we note that the same inapproximability bounds also apply to DCSS.

A summary of our hardness of approximation results is presented in Table 1.

## 3.2. Approximation Algorithms

As in [4] and [3], all our reductions require the condition $a = D_G$. Thus, we study the approximability of MAL when $a \geq D_G$, addressing an open question posed in [4]. We present three sets of results that

**Table 2**
Summary of the approximation results for MAL as a function of the age bound $a$ and the graph diameter $D_G$. In the first four rows, the condition on $D_G$ indicates the range within which the approximation ratio improves upon the trivial $O(n)$ bound.

| Value of $a$ | Approximation ratio and conditions on $D_G$ |
|---|---|
| $a \leq D_G + 1$ | $O(D_G \cdot n^{3/5+\epsilon})$, for $D_G \in o(n^{2/5})$ |
| $D_G + 2 \leq a \leq D_G + 3$ | $O(D_G \cdot n^{1/2})$, for $D_G \in o(n^{1/2})$ |
| $D_G + 4 \leq a \leq D_G + 5$ | $O(D_G \cdot n^{2/5})$, for $D_G \in o(n^{3/5})$ |
| $D_G + 6 \leq a < \lceil 3/2 D_G \rceil$ | $O(D_G \cdot n^{1/3})$, for $D_G \in o(n^{2/3})$ |
| $\lceil 3/2 D_G \rceil \leq a < \lceil 5/3 D_G \rceil$ | $O(\sqrt{n \cdot \log n})$, for $D_G \in \Omega(\sqrt{n^{1/3} \cdot \log n})$ |
| | $O(D_G \cdot n^{1/3})$, otherwise |
| $a \geq \lceil 5/3 D_G \rceil$ | $O(\sqrt{n \cdot \log n})$, for $D_G \in \Omega(\sqrt{n / \log n})$ |
| | $O\big((D_G \cdot n \cdot \log^2 n)^{1/3}\big)$, for $D_G \in \Omega(\log n) \cap O(\sqrt{n/\log n})$ |
| | $O(D_G \cdot n^{1/3})$, otherwise |

highlight how the approximation guarantees for MAL depend on the relationship between $a$ and $D_G$.

(1) We begin by analyzing the case where $a$ is sufficiently larger than $R_G$, the radius of the graph $G$. Specifically, if $a \geq 2R_G$ (resp., $a \geq 2R_G + 1$), we provide a polynomial-time algorithm that computes a solution requiring at most 2 (resp., 1) more labels than the optimum. These additive approximation guarantees translate into asymptotic multiplicative bounds, with approximation factors arbitrarily close to 1 as the input size increases. Furthermore, if $a \geq 2D_G + 2$, an optimal solution can be computed in polynomial time. Since MAL does not admit any feasible solution when $a < D_G$, this first set of results leaves open the intermediate regime where $D_G \leq a < 2R_G$.

(2) We next consider the case where $a$ is only slightly larger than $D_G$, leveraging a connection between MAL and the DCSS problem. Specifically, we show that the approximability of MAL is within a factor $a$ of that of DCSS. We begin by showing that when $a = D_G = 2$, MAL admits a logarithmic-factor approximation, which asymptotically matches our first inapproximability bound. Note that if $a = 2$ and $D_G = 1$, the graph must be a clique. In this case, $R_G = 1$ and $a = 2R_G$, so MAL can be solved using at most 2 labels more than the optimum. For $a \geq D_G + 2$, we obtain an $O(D_G \cdot n^{1/2})$-approximation, which is sublinear in $n$ when $D_G$ is sufficiently small. This bound improves as $a$ increases: we achieve approximation factors of $O(D_G \cdot n^{2/5})$ and $O(D_G \cdot n^{1/3})$ for $a \geq D_G + 4$ and $a \geq D_G + 6$, respectively. Finally, for any $a \geq D_G$, we obtain an approximation factor of $O(D_G \cdot n^{3/5+\epsilon})$ for any constant $\epsilon > 0$. All these approximation factors depend linearly on $D_G$, as our approach approximates DCSS via generalizations of it, and then transforms the resulting solution into one for MAL.

(3) Our main algorithmic contribution addresses the case $D_G \leq a < 2R_G$, where we approximate MAL without relying on DCSS, thereby avoiding the linear dependence on $D_G$ in the approximation ratio. We show that when $a \geq \lceil \frac{3}{2} \cdot D_G \rceil$ (resp., $a \geq \lceil \frac{5}{3} \cdot D_G \rceil$), MAL can be approximated within a factor of $O(\sqrt{n \log n})$ (resp., $O((D_G n \log^2 n)^{1/3})$). Both bounds are sublinear, and the second algorithm yields better performance when $D_G = o(\sqrt{n / \log n})$, although it requires larger values of $a$.

Table 2 summarizes the approximation guarantees we obtain for MAL.

## 4. Conclusion and Open Questions

In this work, we investigated the complexity of approximating the Minimum Aged Labeling problem (MAL). We established strong hardness of approximation results for the case $a = D_G$, summarized in Table 1. When relaxing the parameter to $a > D_G$, we provided approximation algorithms with guarantees that depend on the relationship between $a$ and $D_G$, as detailed in Table 2. Additionally, we highlighted a connection between MAL and the Diameter Constrained Spanning Subgraph problem, showing that comparable hardness and approximation results hold for DCSS.

Several questions remain open. In particular, the computational complexity of MAL when $D_G < a < 2D_G + 2$ is still unresolved, as our reductions do not seem to extend to cases where $a \geq D_G + 1$. Furthermore, it is unclear whether the inapproximability results for MAL hold when $D_G < a < 2R_G$.

## Declaration on Generative AI

The authors utilized ChatGPT and Grammarly to enhance language clarity and readability. The authors, who take full responsibility for the final version of the manuscript, carefully reviewed and refined all content generated by these tools.

## Acknowledgments

This work was partially supported by: the European Union - NextGenerationEU under the Italian Ministry of University and Research National Innovation Ecosystem grant ECS00000041 - VITALITY - CUP: D13C21000430001; PNRR MUR project GAMING "Graph Algorithms and MinINg for Green agents" (PE0000013, CUP D13C24000430001); the European Union - Next Generation EU under the Italian National Recovery and Resilience Plan (NRRP), Mission 4, Component 2, Investment 1.3, CUP J33C22002880001, partnership on "Telecommunications of the Future"(PE00000001 - program "RESTART"), project MoVeOver/SCHEDULE ("Smart interseCtions witH connEcteD and aUtonomous vehicLEs", CUP J33C22002880001); ARS01_00540 - RASTA project, funded by the Italian Ministry of Research PNR 2015-2020. The first and second authors acknowledge the support of the MUR (Italy) Department of Excellence 2023–2027.

## References

[1] A. Deligkas, I. Potapov, Optimizing reachability sets in temporal graphs by delaying, Proceedings of the AAAI Conference (2020) 9810–9817. doi:10.1609/AAAI.V34I06.6533.
[2] A. Deligkas, E. Eiben, G. Skretas, Minimizing reachability times on temporal graphs via shifting labels, in: Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23, 2023, pp. 5333–5340. doi:10.24963/ijcai.2023/592.
[3] G. Mertzios, O. Michail, P. Spirakis, Temporal network optimization subject to connectivity constraints, Algorithmica 81 (2019) 1416–1449. doi:10.1007/s00453-018-0478-6.
[4] N. Klobas, G. B. Mertzios, H. Molter, P. G. Spirakis, The complexity of computing optimum labelings for temporal connectivity, Journal of Computer and System Sciences 146 (2024) 103564. doi:10.1016/j.jcss.2024.103564.
[5] O. Michail, An introduction to temporal graphs: An algorithmic perspective*, Internet Mathematics 12 (2016). doi:10.1080/15427951.2016.1177801.
[6] A. Braunstein, A. Ingrosso, Inference of causality in epidemics on temporal contact networks, Scientific Reports 6 (2016) 27538. doi:10.1038/srep27538.
[7] J. Enright, R. R. Kao, Epidemics on dynamic networks, Epidemics 24 (2018) 88–97. doi:10.1016/j.epidem.2018.04.003.
[8] K. Meeks, Reducing reachability in temporal graphs: Towards a more realistic model of real-world spreading processes, in: Revolutions and Revelations in Computability, Springer International Publishing, Cham, 2022, pp. 186–195. doi:10.1007/978-3-031-08740-0_16.
[9] J. Enright, K. Meeks, G. B. Mertzios, V. Zamaraev, Deleting edges to restrict the size of an epidemic in temporal networks, Journal of Computer and System Sciences 119 (2021) 60–77. doi:10.1016/j.jcss.2021.01.007.
[10] J. A. Enright, K. Meeks, F. Skerman, Assigning times to minimise reachability in temporal graphs, Journal of Computer and System Sciences 115 (2021) 169–186. doi:10.1016/J.JCSS.2020.08.001.

[11] H. Molter, M. Renken, P. Zschoche, Temporal reachability minimization: Delaying vs. deleting, Journal of Computer and System Sciences 144 (2024) 103549. doi:10.1016/j.jcss.2024.103549.
[12] D. Carnevale, G. D'Angelo, M. Olsen, Approximating optimal labelings for temporal connectivity, Proceedings of the AAAI Conference on Artificial Intelligence 39 (2025) 26490–26497. doi:10.1609/aaai.v39i25.34849.
[13] D. Carnevale, G. D'Angelo, M. Olsen, Approximating optimal labelings for temporal connectivity, 2025. URL: https://arxiv.org/abs/2504.16837. arXiv:2504.16837.