Temporal Multi-Broadcast Optimization

Extended Abstract

Daniele Carnevale*
University of Geneva
Geneva, Switzerland
daniele.carnevale@unige.ch

Gianlorenzo D'Angelo
Gran Sasso Science Institute
L'Aquila, Italy
gianlorenzo.dangelo@gssi.it

ABSTRACT

We introduce the $\mathcal{D}$-*Temporal Multi-Broadcast* ($\mathcal{D}$-TMB) problem, which asks for scheduling the availability of edges so that a predetermined subset of sources temporally reach all other vertices while optimizing the worst-case temporal distance $\mathcal{D}$ from any source. We characterize the computational complexity and approximability of $\mathcal{D}$-TMB under six different definitions of temporal distance.

KEYWORDS

Temporal graphs; Approximation algorithms; Scheduling; Combinatorial optimization; Routing

ACM Reference Format:
Daniele Carnevale and Gianlorenzo D'Angelo. 2026. Temporal Multi-Broadcast Optimization: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/ITHO9185

1 INTRODUCTION

Temporal graphs provide a natural model for dynamic systems where interactions occur at specific times, such as communication networks, transportation systems, or epidemiological contact networks [8, 11]. In a temporal graph, each edge is associated with a set of timestamps indicating the times at which each edge is available. A *temporal path* is a path that respects the flow of time, i.e., the timestamps of the edges are traversed in non-decreasing order. The classical notion of shortest-path fails to capture temporal constraints, thus motivating the study of other distance measures for temporal paths. For example, one may seek a temporal path that arrives as early as possible (Earliest-Arrival Path, EA), starts as late as possible (Latest-Departure Path, LD), minimizes the overall duration (Fastest-Time Path, FT), minimizes total traveling time (Shortest-Traveling Path, ST), minimizes number of used edges (Minimum-Hop Path, MH), or minimizes the total waiting time (Minimum-Waiting Path, MW). Note that these temporal measures are not proper distances; for example, they might not satisfy the triangle inequality.

Motivated by applications in robotics, information spreading, and logistics, considerable research attention has recently been devoted to temporal network optimization problems, where the aim is to schedule the availability of edges in order to optimize some network property, see e.g. [3–7, 12, 13]. In this work, we focus on the problem of scheduling the availability of the edges of a graph in such a way that the worst-case temporal distance from a set of source vertices to all other vertices is optimal, according to a temporal distance measure $\mathcal{D}$. We introduce the $\mathcal{D}$-Temporal Multi-Broadcast problem, where we are given an underlying static graph, a *traversal* function that associates each edge-time pair with the time needed to traverse it, a *multiplicity* function indicating how many times an edge can appear, and a set of sources, and we must schedule the availability of edges over time in such a way that the multiplicity constraints are satisfied and the worst-case distance $\mathcal{D} \in \{\text{EA}, \text{LD}, \text{FT}, \text{ST}, \text{MH}, \text{MW}\}$ between a source and any other node is optimal. In other words, we assume a fixed infrastructure (a static graph with given traversal times) and an offline planner that assigns time labels to edges, subject to multiplicity constraints, to optimize the worst-case temporal distance from the sources.

The $\mathcal{D}$-Temporal Multi-Broadcast problem has several applications in logistics, multi-agent information spreading, and wireless networks. In a logistics setting, consider a set of different suppliers that distribute different kinds of goods to a set of distribution points or warehouses connected by a road network. This naturally leads to an offline scheduling problem where one schedules, over time, which road segments can be used for shipments, so that we have a guarantee on the time the last destination receives the last missing good under a chosen temporal distance measure. The time required to traverse a road may depend on the time at which we traverse it (captured by the traversal function), and each road segment may only support a bounded number of total shipments due to limited resources such as truck availability (captured by multiplicity). Depending on the chosen measure, we might want to minimize the time at which the last destination receives the last missing good (EA), or the maximum time a good spends in transit on roads (ST), or the total time (including waiting at a middle point) a good needs to go from its supplier to a distribution point (FT).

A similar problem arises in information spreading over a social network, where a set of agents, each holding different data, want to spread their data in such a way that all network users receive all the data as early as possible. Two users can exchange information through a *connection* between them (e.g., an online or in-person meeting) in which users share the data they received so far, see e.g. [6]. A connection can be scheduled at specific time slots, but a connection between the same two users cannot occur more than a fixed number of times. The aim is to schedule all connections in such a way that the time when the last user receives all the data is minimum. Similarly, in a wireless sensor network, the data gathered by a set of sensors must be broadcast to all the nodes as early as

*Work done while the author was a PhD student at the GSSI (L'Aquila, Italy).

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/ITHO9185

possible, but no more than a given number of connections can be established among nodes for energy-saving purposes.

The $\mathcal{D}$-Temporal Multi-Broadcast problem generalizes the ReachFast problem, introduced in [6], where we are given a temporal graph and are allowed to shift the temporal edges, i.e., to change the times at which edges appear, in such a way that the latest earliest arrival time from a source to a node is minimum. We prove that, when we consider only earliest arrival distances, problems ReachFast and EA-Temporal Multi-Broadcast are equivalent under value-preserving reductions. Deligkas et al. [6] proved that ReachFast is tractable for a single source, for two sources when the underlying graph consists of a number of parallel paths between them, and for multiple sources when the underlying graph is a tree. Moreover, they showed that it is NP-complete to decide whether ReachFast with two sources admits a solution of value 6, which implies that the problem is APX-hard.

1.1 Related Work

Delaying (i.e., postponement of edge availability) and shifting (i.e., postponement or advancement of edge availability) are common tools to influence the temporal reachability of a temporal graph. Both have been studied in several works as means to satisfy specific connectivity objectives. Motivated by application in epidemiology, Deligkas and Potapov [7] introduced the operation of delaying and studied the problem of minimizing reachability under a bounded number of delays. They proved that minimizing reachability, as well as minimizing the maximum or average reachability, is NP-hard when a given number of edges can each be delayed by at most a specified input parameter. The problems become tractable when an unbounded number of delays is allowed.

Kutner and Sommer [14] studied whether, given a temporal graph with passengers (i.e., vertices), each having a specified destination and deadline, it is possible to delay edges so that all passengers reach their targets within the required time. They proved that the problem is NP-complete both for graphs of lifetime 2 and for planar graphs with larger lifetime. On the positive side, the problem is tractable on trees and fixed-parameter tractable (FPT) when parametrized by the number of demands plus the size of the feedback edge set of the underlying graph. For a broader comparison with related delaying/shifting problems, we refer the interested reader to Table 1 in [14].

Temporal branching and temporal spanning subgraphs [1, 9, 10] are similar to our problem in the single-source case. A temporal branching is a directed out-tree rooted at a source that spans all vertices, whereas in a temporal spanning subgraph, the underlying graph is not required to be a tree but is still required to be of minimum size. Huang et al. [9] proved that, assuming the source can reach every other vertex, it is possible to compute both a temporal branching from the source and an EA-branching, which is a branching in which the source reaches every other vertex at the earliest possible arrival time. Bubboloni et al. [1] studied temporal branchings and temporal spanning subgraphs using the same six distance measures adopted in this work. For each distance $\mathcal{D}$, their objective is to find a temporal branching (or a spanning subgraph) rooted at a given vertex that reaches every other vertex via a temporal path whose value under distance $\mathcal{D}$ is minimized. They proved that finding $\mathcal{D}$-branchings and $\mathcal{D}$-spanning subgraphs is NP-complete for every temporal distance but EA. The main difference between their problem and our single-source case is that we only require the *maximum* distance between the source and any vertex to be minimized.

1.2 Our Results

We introduce the definitions of $\mathcal{D}$-ReachFast and $\mathcal{D}$-Temporal Multi-Broadcast, and show their equivalence for any $\mathcal{D}$, which allows us to focus exclusively on the latter. We then characterize the computational complexity and approximability of $\mathcal{D}$-Temporal Multi-Broadcast in terms of distance and number of sources.

In the single-source setting, the problem is tractable for EA (as established in [6]) and we show that it is also tractable for LD, while for the other temporal distances we prove that it is NP-complete to decide whether the problem admits a solution of a given value. This latter result implies that the problem cannot be approximated within a given factor that depends on the specific temporal distance measure, unless P = NP. In particular, for $\mathcal{D} \in \{\text{ST}, \text{MH}\}$, $\mathcal{D}$-Temporal Multi-Broadcast cannot be approximated within a factor better than 2, while, for $\mathcal{D} \in \{\text{FT}, \text{MW}\}$, it cannot be approximated within any factor smaller than an exponential in the number of nodes or a value depending on the input weights and the distance function. In the latter case, we further show that a simple approximation algorithm matches this lower bound.

For the case of multiple sources, we prove that even deciding whether a feasible solution exists is NP-complete, regardless of the adopted temporal distance function. Consequently, if feasibility is not assumed a priori, the problem cannot be approximated within any finite factor. Moreover, this result holds for any fixed number of sources greater than one. We complement these negative results by identifying structural conditions that guarantee tractability for EA and LD. We establish that for EA and LD, the problem admits a polynomial-time algorithm for any number of sources, when the multiplicity is at least equal to the number of sources or when the underlying static graph is a tree and the multiplicity of each edge is at least two.

A longer version is available on arXiv [2].

ACKNOWLEDGMENTS

This work was partially supported by: PNRR MUR project GAMING “Graph Algorithms and MinINg for Green agents” (PE0000013, CUP D13C24000430001); the European Union - NextGenerationEU under the Italian Ministry of University and Research National Innovation Ecosystem grant ECS00000041 - VITALITY - CUP: D13C21000430001; 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); The Italian Institute for Advanced Mathematics and the Italian National Group for Scientific Computation (GNCS-INdAM); MUR (Italy), Department of Excellence 2023–2027.

REFERENCES

[1] Daniela Bubboloni, Costanza Catalano, Andrea Marino, and Ana Silva. 2025. On computing optimal temporal branchings and spanning subgraphs. *J. Comput. System Sci.* 148 (2025), 103596. https://doi.org/10.1016/j.jcss.2024.103596

[2] Daniele Carnevale and Gianlorenzo D'Angelo. 2026. Optimizing Distances for Multi-Broadcast in Temporal Graphs. arXiv:2602.12126 [cs.DS] https://arxiv.org/abs/2602.12126

[3] Daniele Carnevale, Gianlorenzo D'Angelo, and Martin Olsen. 2025. Approximating Optimal Labelings for Temporal Connectivity. In *AAAI-25, Sponsored by the Association for the Advancement of Artificial Intelligence, February 25 - March 4, 2025, Philadelphia, PA, USA*, Toby Walsh, Julie Shah, and Zico Kolter (Eds.). AAAI Press, 26490–26497. https://doi.org/10.1609/AAAI.V39I25.34849

[4] Argyrios Deligkas, Michelle Döring, Eduard Eiben, Tiger-Lily Goldsmith, George Skretas, and Georg Tennigkeit. 2025. How Many Lines to Paint the City: Exact Edge-Cover in Temporal Graphs. In *AAAI-25, Sponsored by the Association for the Advancement of Artificial Intelligence, February 25 - March 4, 2025, Philadelphia, PA, USA*, Toby Walsh, Julie Shah, and Zico Kolter (Eds.). AAAI Press, 26498–26506. https://doi.org/10.1609/AAAI.V39I25.34850

[5] Argyrios Deligkas, Eduard Eiben, Tiger-Lily Goldsmith, and George Skretas. 2023. Being an Influencer is Hard: The Complexity of Influence Maximization in Temporal Graphs with a Fixed Source. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2023, London, United Kingdom, 29 May 2023 - 2 June 2023*, Noa Agmon, Bo An, Alessandro Ricci, and William Yeoh (Eds.). ACM, 2222–2230. https://doi.org/10.5555/3545946.3598898

[6] Argyrios Deligkas, Eduard Eiben, and George Skretas. 2023. Minimizing Reachability Times on Temporal Graphs via Shifting Labels. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*. International Joint Conferences on Artificial Intelligence Organization, 5333–5340. https://doi.org/10.24963/ijcai.2023/592 Main Track.

[7] Argyrios Deligkas and Igor Potapov. 2020. Optimizing Reachability Sets in Temporal Graphs by Delaying. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. AAAI Press, 9810–9817. https://doi.org/10.1609/AAAI.V34I06.6533

[8] Petter Holme and Jari Saramäki. 2012. Temporal networks. *Physics Reports* 519, 3 (2012), 97–125. https://doi.org/10.1016/j.physrep.2012.03.001

[9] Silu Huang, Ada Wai-Chee Fu, and Ruifeng Liu. 2015. Minimum Spanning Trees in Temporal Graphs. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data* (Melbourne, Victoria, Australia) *(SIGMOD '15)*. Association for Computing Machinery, New York, NY, USA, 419–430. https://doi.org/10.1145/2723372.2723717

[10] Naoyuki Kamiyama and Yasushi Kawase. 2015. On packing arborescences in temporal networks. *Inf. Process. Lett.* 115, 2 (2015), 321–325. https://doi.org/10.1016/J.IPL.2014.10.005

[11] David Kempe, Jon Kleinberg, and Amit Kumar. 2002. Connectivity and Inference Problems for Temporal Networks. *J. Comput. System Sci.* 64, 4 (2002), 820–842. https://www.sciencedirect.com/science/article/pii/S0022000002918295

[12] Nina Klobas, George B. Mertzios, Hendrik Molter, Rolf Niedermeier, and Philipp Zschoche. 2022. Interference-free walks in time: temporally disjoint paths. *Autonomous Agents and Multi-Agent Systems* 37, 1 (2022), 1. https://doi.org/10.1007/S10458-022-09583-5

[13] Pascal Kunz, Hendrik Molter, and Meirav Zehavi. 2023. In Which Graph Structures Can We Efficiently Find Temporally Disjoint Paths and Walks?. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023, 19th-25th August 2023, Macao, SAR, China*. ijcai.org, 180–188. https://doi.org/10.24963/IJCAI.2023/21

[14] David C. Kutner and Anouk Sommer. 2025. Better Late, Then? The Hardness of Choosing Delays to Meet Passenger Demands in Temporal Graphs. In *4th Symposium on Algorithmic Foundations of Dynamic Networks, SAND 2025, June 9-11, 2025, Liverpool, UK (LIPIcs, Vol. 330)*, Kitty Meeks and Christian Scheideler (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 7:1–7:18. https://doi.org/10.4230/LIPICS.SAND.2025.7