



Span-core Decomposition for Temporal Networks: Algorithms and Applications

EDOARDO GALIMBERTI, ISI Foundation and University of Turin

MARTINO CIAPERONI, Aalto University

ALAIN BARRAT, Aix Marseille Univ, Université de Toulon, CNRS, CPT and ISI Foundation

FRANCESCO BONCHI, ISI Foundation and Eurecat

CIRO CATTUTO, University of Turin and ISI Foundation

FRANCESCO GULLO, UniCredit

2

When analyzing temporal networks, a fundamental task is the identification of dense structures (i.e., groups of vertices that exhibit a large number of links), together with their temporal span (i.e., the period of time for which the high density holds). In this article, we tackle this task by introducing a notion of temporal core decomposition where each core is associated with two quantities, its coreness, which quantifies how densely it is connected, and its span, which is a temporal interval: we call such cores *span-cores*.

For a temporal network defined on a discrete temporal domain T , the total number of time intervals included in T is quadratic in $|T|$, so that the total number of span-cores is potentially quadratic in $|T|$ as well. Our first main contribution is an algorithm that, by exploiting containment properties among span-cores, computes all the span-cores efficiently. Then, we focus on the problem of finding only the *maximal span-cores*, i.e., span-cores that are not dominated by any other span-core by both their coreness property and their span. We devise a very efficient algorithm that exploits theoretical findings on the maximality condition to directly extract the maximal ones without computing all span-cores.

Finally, as a third contribution, we introduce the problem of *temporal community search*, where a set of query vertices is given as input, and the goal is to find a set of densely-connected subgraphs containing the query vertices and covering the whole underlying temporal domain T . We derive a connection between this problem and the problem of finding (maximal) span-cores. Based on this connection, we show how temporal community search can be solved in polynomial-time via dynamic programming, and how the maximal span-cores can be profitably exploited to significantly speed-up the basic algorithm.

We provide an extensive experimentation on several real-world temporal networks of widely different origins and characteristics. Our results confirm the efficiency and scalability of the proposed methods. Moreover, we showcase the practical relevance of our techniques in a number of applications on temporal networks, describing face-to-face contacts between individuals in schools. Our experiments highlight the relevance of the notion of (maximal) span-core in analyzing social dynamics, detecting/correcting anomalies in the data, and graph-embedding-based network classification.

Authors' addresses: E. Galimberti, ISI Foundation, Via Chisola 5, 10126 Turin, Italy and University of Turin, Turin, Italy; email: edoardo.galimberti@isi.it; M. Ciaperoni, Aalto University Foundation sr, PO BOX 11000, FI-00076, Aalto, Finland; email: martino.ciaperoni@aalto.fi; A. Barrat, Centre de Physique Théorique, Campus de Luminy, Case 907, 163 Avenue de Luminy, 13288 Marseille Cedex 9, France and ISI Foundation, Turin, Italy; email: alain.barrat@cpt.univ-mrs.fr; F. Bonchi, ISI Foundation, Via Chisola 5, 10126 Turin, Italy and Eurecat, Barcelona, Spain; email: francesco.bonchi@isi.it; C. Cattuto, University of Turin, Corso Svizzera 185, 10149 Torino, Italy and ISI Foundation, Turin, Italy; email: ciro.cattuto@unito.it; F. Gullo, Unicredit, Via Molfetta, 101 - 00171 Rome, Italy; email: gullof@acm.org.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1556-4681/2020/12-ART2 \$15.00

<https://doi.org/10.1145/3418226>

CCS Concepts: • **Mathematics of computing** → **Graph algorithms**; *Graph theory*; • **Information systems** → **Temporal data**; Data mining; • **Theory of computation** → *Dynamic programming*;

Additional Key Words and Phrases: Temporal networks, core decomposition, maximal cores, community search, face-to-face interaction networks

ACM Reference format:

Edoardo Galimberti, Martino Ciaperoni, Alain Barrat, Francesco Bonchi, Ciro Cattuto, and Francesco Gullo. 2020. Span-core Decomposition for Temporal Networks: Algorithms and Applications. *ACM Trans. Knowl. Discov. Data* 15, 1, Article 2 (December 2020), 44 pages.
<https://doi.org/10.1145/3418226>

1 INTRODUCTION

A temporal network¹ is a representation of entities (vertices), their relations (links), and how these relations are established/broken over time. Notice that here we will consider discrete times, i.e., the temporal networks can be represented as a time-ordered series of snapshots (instantaneous graphs). Extracting dense structures (i.e., groups of vertices exhibiting a large number of links with each other), together with their temporal span (i.e., the period of time for which the high density is observed) is a key mining primitive to characterize such temporal networks and extract relevant structures. This type of pattern enables fine-grain analysis of the network dynamics and can be a building block towards more complex tasks and applications, such as finding temporally recurring subgraphs or anomalously dense ones. For instance, they can help in studying contact networks among individuals to quantify the transmission opportunities of respiratory infections in a population and uncover situations where the risk of transmission is higher, with the goal of designing mitigation strategies [43]. Anomalously dense temporal patterns among entities in a co-occurrence graph (e.g., extracted from the Twitter stream) have also been used to identify events and buzzing stories in real time [5, 15, 16]. Another example concerns scientific collaboration and citation networks, where these patterns can help understand the dynamics of collaboration in successful professional teams, study the evolution of scientific topics, and detect emerging technologies [32].

In this article, we adopt as a measure of density of a pattern, the *minimum degree* holding among the vertices in the subgraph during the pattern’s span. The problem of extracting *all* these patterns is tackled by introducing a notion of *temporal core decomposition* in which each core is associated with its *span*, i.e., an interval of *contiguous timestamps*, for which the coreness property holds. We term such a notion of temporal core *span-core*.

Moreover, in several application scenarios, it is typically required to identify only those dense patterns that contain a given set of query vertices. We therefore introduce the problem of *temporal community search*, whose goal is to find a set of cohesive temporal subgraphs containing the input query vertices and covering the whole temporal domain.

To the best of our knowledge, the problems of (efficient) *span-core* computation and temporal community search have never been studied so far.

1.1 Challenges and Contributions

As the number of possible time intervals is quadratic in the size of the input temporal domain T , the total number of span-cores is, in the worst case, quadratic in T too. The naive method to find all *span-cores*, which would be to operate a core decomposition for each of these time intervals,

¹We use “network” and “graph” interchangeably throughout the article.

would therefore be very time-consuming. This is a major challenge that we tackle by deriving containment properties between span-cores and by exploiting them to devise an algorithm for computing all the *span-cores* that is significantly more efficient than the naïve exhaustive method.

We then shift our attention to the problem of finding only the *maximal span-cores*, defined as the span-cores that are not dominated by any other span-core by both the coreness property and the span. A straightforward way of approaching this problem is to filter out non-maximal span-cores during the execution of an algorithm for computing the whole span-core decomposition. However, as the maximal ones are usually much less numerous than the overall span-cores, it would be desirable to have a method that effectively exploits the maximality property and extracts maximal span-cores directly, without computing the complete decomposition. The design of an algorithm of this kind is an interesting challenge, as it contrasts with the intrinsic conceptual properties of core decomposition, based on which a core of order k can be efficiently computed from the core of order $k-1$, of which it is a subset. For this reason, at first glance, the computation of the core of the highest order would seem as hard as computing the overall core decomposition. Instead, in this work, we derive a number of theoretical properties about the relationship among span-cores of different temporal intervals and, based on these findings, we show how such a challenging goal may be achieved.

Finally, we focus on the problem of *community search in temporal networks*. Community search has been extensively studied in static graphs. It requires to find a subgraph containing a given set of query vertices and maximizing a certain density measure [36, 51]. Here, we propose a formulation of the community-search problem in temporal networks as follows: given a set Q of query vertices, and a positive integer h , find a segmentation of the underlying temporal domain in h segments $\{\Delta_i\}_{i=1}^h$ and a subgraph S_i for every identified segment Δ_i such that each S_i contains the query vertices Q and the total density of the subgraphs is maximized. Following the bulk of the literature in community search on static networks, in our definition of temporal community search, we adopt the minimum degree as a density measure.

We show that, with some manipulations, temporal community search can be reformulated as an instance of the popular *sequence segmentation* problem, which asks for partitioning a sequence of numerical values into h segments so as to minimize the sum of the penalties (according to some penalty function) on the identified segments [10]. Therefore, the classical dynamic-programming (DP) algorithm for sequence segmentation by Bellman [10] can be easily adapted to solve temporal community search in polynomial time.

A criticality of this approach is that a naïve adaptation of the Bellman's algorithm takes quadratic time in the size of the input temporal domain T . As a major contribution in this regard, we prove that the set of maximal span-cores provide a sound and complete basis to still have an optimal solution to temporal community search, while at the same time leading to a significant speed-up with respect to the naïve method. In fact, let $T^* \subseteq T$ be the subset of timestamps that are covered by the span of at least one maximal span-core, together with the timestamps that immediately precede or succeed any of such spans. We show that considering T^* (instead of T) in the (adaptation of the) Bellman's algorithm is sufficient to optimally solve the underlying temporal-community-search problem instance. As, typically, $|T^*| \ll |T|$, this finding guarantees a considerable improvement in efficiency (as confirmed by our experiments).

A further challenge in our temporal-community-search problem is a typical one in community-search formulations based on minimum degree, namely, that the output subgraphs are typically large in size. We tackle this challenge by devising a method to reduce the size of the output subgraphs without affecting optimality. The proposed method is inspired by the one devised by Barbieri et al. [7] for the problem of minimum community search (in static graphs).

To summarize, the main contributions of this article are as follows:

- We introduce the notion of span-core decomposition and maximal span-core in temporal networks, characterizing structure and size of the search space and providing important containment properties (Section 3).
- We devise an algorithm for computing all span-cores that exploits the aforementioned containment properties and is orders of magnitude faster than a naïve method based on traditional core decomposition (Section 4).
- We study the problem of finding only the maximal span-cores. We derive several theoretical findings about the relationship between maximal span-cores and exploit these findings to devise an algorithm that is more efficient than computing all span-cores and discarding the non-maximal ones (Section 5).
- We introduce the problem of temporal community search and show how it can be solved in polynomial time via DP. We prove an important connection between temporal community search and maximal span-cores, which allows us to devise an algorithm that is considerably more efficient than the naïve DP one. We also propose a method to achieve the critical challenge of having too large communities as output (Section 6).
- We provide a comprehensive experimentation on several real-world temporal networks, with millions of vertices, tens of millions of edges, and hundreds of timestamps, which attests efficiency and scalability of our methods (Section 7).
- We present applications on face-to-face interaction networks that illustrate the relevance of the notions of (maximal) span-core and temporal community search in real-life analyses and applications (Section 8).

The next section provides an overview of the related literature, while Section 9 discusses future work and concludes the article.

An abridged version of this work, covering Sections 4 and 5, together with the corresponding experiments (i.e., parts of Sections 7 and 8), was presented in [37].

Reproducibility. For the sake of reproducibility, all our code and some of the datasets used in this article are available at github.com/egalimberti/span_cores.

2 BACKGROUND AND RELATED WORK

2.1 Core Decomposition

Given a simple (static) graph $G = (V, E)$, let $d(S, u)$ denote the degree of vertex $u \in V$ in the subgraph induced by vertex set $S \subseteq V$, i.e., $d(S, u) = |\{v \in S \mid (u, v) \in E\}|$. The notions of k -core and core decomposition are defined as follows:

Definition 1 (k -core and Core Decomposition [67]). The k -core (or core of order k) of G is a maximal set of vertices $C_k \subseteq V$ such that $\forall u \in C_k : d(C_k, u) \geq k$. The set of all k -cores $V = C_0 \supseteq C_1 \supseteq \dots \supseteq C_{k^*}$ ($k^* = \arg \max_k C_k \neq \emptyset$) is the core decomposition of G .

Core decomposition can be computed in linear time by iteratively removing the smallest-degree vertex and setting its core number as equal to its degree at the time of removal [9]. Among the many definitions of dense structures, core decomposition is particularly appealing as, among others, it is fast to compute, and can speed-up/approximate dense-subgraph extraction according to various other definitions. For instance, core decomposition allows for finding cliques more efficiently [31], as a k -clique is contained into a $(k-1)$ -core, which can be significantly smaller than the original graph. Moreover, core decomposition is at the basis of approximation algorithms for the densest-(at-least- k -)subgraph problem [3, 56], and betweenness centrality [47]. Core decomposition has also been recognized as an important tool to analyze and visualize complex networks [2, 8] in

several domains, e.g., bioinformatics [6, 87], software engineering [90], and social networks [41, 54]. It has been studied under various settings, such as distributed [69], streaming/maintenance [60, 77], and disk-based [23], and generalized to various types of static graphs, such as uncertain [19], directed [44], weighted [28, 40], bipartite graphs [61], or including attributes on the nodes [89]. For a comprehensive survey about theory, algorithms, and applications of core decomposition we refer to [18, 64].

Two types of extension of core decomposition bear some relation to our work. First, core decomposition in *multilayer networks*—i.e., networks that are composed of a superposition of networks—has been studied in [38, 39]. In the multilayer setting, a core is allowed to extend on any subset of layers, thus implying that the total number of multilayer cores is exponential in the number of layers. Although temporal networks can be viewed as a special case of multilayer networks (where each timestamp is interpreted as a layer), there is a fundamental difference: in a temporal network the “layers” are ordered, and the sequentiality of timestamps represents an important structural constraint. In other words, in the temporal setting we are interested in cores that span a temporal interval, and not simply any subset of (potentially non-contiguous) timestamps. This aspect has two critical consequences. First, the search space and the number of temporal cores are no longer exponential, unlike the multilayer case. Second, to guarantee an effective fulfilment of the constraint on temporal sequentiality, the requirements for the edges that contribute to the formation of a temporal core are stricter than the ones at the basis of the multilayer-core definition. A more detailed technical discussion of the relationship between multilayer core decomposition and the proposed temporal core decomposition is reported in Section 3.3.

The second extension of core decomposition that shares some relation to the one proposed in this work is due to Wu et al., who have proposed in [86] an alternative definition of temporal core decomposition. A major difference between Wu et al.’s definition and ours is that the former *does not take any kind of temporal constraint into account*. Indeed, Wu et al. define the (k, h) -core as the largest subgraph in which every vertex has at least k neighbors and there are at least h temporal edges between the vertex and its neighbors, without any restriction on when these h edges occur: the sequentiality of connections is not taken into account and non-contiguous timestamps can support the same core. In fact, the (k, h) -core decomposition can be seen as a kind of weighted static core decomposition on the weighted static network resulting from the aggregation of the temporal network. In contrast, our temporal cores have each a clear temporal collocation and continuous spans, so that our definition includes temporality in an explicit way and cannot be reduced to Wu et al.’s one. As we will see in Section 8, associating a temporal collocation to each core is important in applications.

2.2 Patterns in Temporal Networks

A number of works on extracting dense patterns from a temporal network focus on the well-established notion of *densest subgraph*, i.e., a subgraph maximizing the average-degree density. Jethava and Beerenwinkel [53] consider as input a set of graphs sharing the same vertex set, which can thus also be interpreted as a temporal network. On such an input, they study the *densest common subgraph* problem, i.e., the problem of finding a subgraph maximizing the minimum average degree over all graphs (timestamps), and devise a linear-programming formulation and a greedy heuristic algorithm for it. Further (mostly theoretical) advancements to the densest-common-subgraph problem have been provided by Reinthal et al. [71] and Charikar et al. [22]. Semertzidis et al. [78] instead introduce two more variants of the problem, where the goal is to maximize the average average degree and the minimum minimum degree, respectively. They show that the average–average variant easily reduces to the traditional densest-subgraph problem, and

that the minimum–minimum variant can be exactly solved by a simple adaptation of the classic algorithm for core decomposition.

Complementary works focus on variants of the densest-subgraph-discovery problem. Rozenstein et al. study the problem of discovering dense temporal subgraphs whose edges occur in short time intervals considering the exact timestamp of the occurrences [73], and the problem of partitioning the timeline of a temporal network into non-overlapping intervals, such that the intervals span subgraphs with maximum total density [72]. Epasto et al. [29] deal with the problem of maintaining the densest subgraph in a dynamic setting.

Attention in the literature has also been devoted to densities other than the average degree. The notion of Δ -clique, as a set of vertices in which each pair is in contact at least every Δ timestamps, has been proposed in [48, 85]. Bentert et al. [11] introduce the Δ - k -plex, a relaxation of Δ -clique in which each vertex has an edge to all but at most $k - 1$ vertices at least once every Δ consecutive timestamps. Li et al. [59] study the problem of finding the maximum (θ, Δ) -persistent k -core in a temporal network, i.e., the largest subgraph that is a connected k -core in all the subintervals of duration θ of a given temporal interval Δ .

A different, but still slightly related body of literature focuses on other definitions of temporal patterns, such as frequent evolution patterns in temporal attributed graphs [12, 26, 52], link-formation rules in temporal networks [20, 58], frequency-estimation algorithms for counting temporal motifs [57, 62], finding a small vertex set whose removal eliminates all temporal paths connecting two designated terminal vertices [91], finding a subgraph that maximizes the sum of edge weights in a network whose topology remains fixed but edge weights evolve over time [14, 63], and the discovery of dynamic relationships and events [25], or of correlated activity patterns [42].

This work differs from all the above ones as our notions of span-core and temporal core decomposition do not correspond (or are straightforwardly reducible) to any of those temporal patterns.

2.3 Community Search

Given a static graph and a set of query vertices, the *community search* problem aims at finding a cohesive subgraph *containing the query vertices*. Community search has attracted a great deal of attention in the last years [36, 51]. Sozio and Gionis [79] are the first to introduce this problem by employing the *minimum degree* as a cohesiveness measure. Their formulation can be solved by a simple (linear-time) greedy algorithm, which resembles the traditional two-approximation algorithm for densest subgraph proposed in [21]. More recently, Cui et al. [24] devise a local-search approach to improve the efficiency of the method defined in [79], but only for the special case of a single query vertex. The case of multiple query vertices has instead been addressed by Barbieri et al. [7], who exploit core decomposition as a preprocessing step to improve efficiency. They also tackle the problem of *minimum community search*, i.e., a variant of community search where the size of the output subgraph has to be minimized.

Community search has also been studied under different names and/or settings. Huang et al. [49] introduce a community-search model based on the k -truss notion. Andersen and Lang [4] and Kloumann and Kleinberg [55] study *seed set expansion* in social graphs, in order to find communities with small conductance or that are well-resemblant of the characteristics of the query vertices, respectively. Other works define connectivity subgraphs based on electricity analogues [33], random walks [84], the minimum-description-length principle [1], the Wiener index [74], and network efficiency [75]. Recent approaches also introduce the flexibility of having query vertices belonging to different communities [13, 88]. Finally, community search has been formalized for attributed graphs [34, 50] and spatial graphs [35] as well.

In this work, we study for the first time community search in temporal graphs. Specifically, we provide a novel definition of the problem by asking for a set of subgraphs containing the given

query vertices, along with their corresponding temporal intervals, such that the total minimum-degree density of the identified subgraphs is maximized and the union of the temporal intervals spanned by those subgraphs covers the whole underlying temporal domain. None of the above works deal with such a definition of temporal community search, not even the works by Rozenshtein et al. [72] and Li et al. [59] discussed in the previous subsection. In fact, although Rozenshtein et al. [72] and Li et al. [59] search for cohesive temporal subgraphs, they do not accept any query vertices in input. This is a fundamental feature, which makes those works actually solve a problem other than community search. Another key difference is that they focus on different notions of density.

3 TEMPORAL CORE DECOMPOSITION: PROBLEM STATEMENT

In this section, we provide preliminary definitions and the needed notations, introduce the problem of finding all span-cores and only the maximal ones, and prove containment properties among span-cores that are at the basis of our efficient algorithms.

3.1 Span-cores

We are given a *temporal graph* $G = (V, T, \tau)$, where V is a set of vertices, $T = [0, 1, \dots, t_{max}] \subseteq \mathbb{N}$ is a discrete time domain, and $\tau : V \times V \times T \rightarrow \{0, 1\}$ is a function defining for each pair of vertices $u, v \in V$ and each timestamp $t \in T$ whether edge (u, v) exists in t . We denote $E = \{(u, v, t) \mid \tau(u, v, t) = 1\}$ the set of all temporal edges. Given a timestamp $t \in T$, $E_t = \{(u, v) \mid \tau(u, v, t) = 1\}$ is the set of edges existing at time t . A temporal interval $\Delta = [t_s, t_e]$ is contained into another temporal interval $\Delta' = [t'_s, t'_e]$, denoted $\Delta \sqsubseteq \Delta'$, if $t'_s \leq t_s$ and $t'_e \geq t_e$. Given an interval $\Delta \sqsubseteq T$, we denote $E_\Delta = \bigcap_{t \in \Delta} E_t$ the edges existing in *all timestamps*² of Δ . Given a subset $S \subseteq V$ of vertices, let $E_\Delta[S] = \{(u, v) \in E_\Delta \mid u \in S, v \in S\}$ and $G_\Delta[S] = (S, E_\Delta[S])$. Finally, the *temporal degree* of a vertex u within $G_\Delta[S]$ is denoted $d_\Delta(S, u) = |\{v \in S \mid (u, v) \in E_\Delta[S]\}|$.

Definition 2 ((k, Δ)-core). The (k, Δ) -core of a temporal graph $G = (V, T, \tau)$ is (when it exists) a maximal and non-empty set of vertices $\emptyset \neq C_{k, \Delta} \subseteq V$, such that $\forall u \in C_{k, \Delta} : d_\Delta(C_{k, \Delta}, u) \geq k$, where $\Delta \sqsubseteq T$ is a temporal interval and $k \in \mathbb{N}^+$.

A (k, Δ) -core is thus a set of vertices implicitly defining a cohesive subgraph (where k represents the cohesiveness constraint), together with its *temporal span*, i.e., the interval Δ for which the subgraph satisfies the cohesiveness constraint. In the remainder of the article, we refer to this type of temporal pattern as *span-core*.

The first problem we tackle in this work is to compute the *span-core decomposition* of a temporal graph G , i.e., all span-cores of G .

PROBLEM 1 (SPAN-CORE DECOMPOSITION). *Given a temporal graph G , find the set of all (k, Δ) -cores of G .*

²We remark that this is just one of the possible ways of defining the existence of an edge in a temporal domain. There are two basic semantics used in the literature: the “AND” semantics we employ here, where an edge is required to exist in all the timestamps of an interval, and an “OR” semantics, requiring that an edge appears in at least one of the timestamps. Although both semantics can be meaningful and there is no strong a-priori argument to prefer one over the other, the types of application and the desired semantics of the data analysis can dictate the choice. In this work, we are particularly interested in networks of social interactions (contacts, communications, etc.), and in exposing structures that are cohesive and stable, together with their duration. It seems then natural to consider the AND semantics, as an OR semantics would correspond to an aggregation on a temporal interval and would not constrain the simultaneity of interactions to define a structure. This simultaneity is crucial in applications such as the ones in Section 8, in which we show the relevance of our work in the analysis of contact networks among individuals recorded by an RFID-based proximity-sensing infrastructure.




Fig. 1. Search space: for a temporal span $\Delta = [t_s, t_e]$, the (k, Δ) -core is depicted as a node labeled “ $k, [t_s, t_e]$ ”. An arrow $C_1 \rightarrow C_2$ denotes $C_1 \supseteq C_2$ (the distinction between solid and dotted arrows is for visualization sake only).

Unlike standard cores of simple graphs, span-cores are not all nested into each other, due to their spans. However, they still exhibit containment properties. Indeed, it can be observed that a (k, Δ) -core is contained into any other (k', Δ') -core with less restrictive degree and span conditions, i.e., $k' \leq k$, and $\Delta' \sqsubseteq \Delta$. This property is depicted in Figure 1, and formally stated in the next proposition.

PROPOSITION 1 (SPAN-CORE CONTAINMENT). *For any two span-cores $C_{k,\Delta}$, $C_{k',\Delta'}$ of a temporal graph G it holds that*

$$k' \leq k \wedge \Delta' \sqsubseteq \Delta \Rightarrow C_{k,\Delta} \subseteq C_{k',\Delta'}$$

PROOF. The result can be proved by separating the two conditions in the hypothesis, i.e., by separately showing that (i) $k' \leq k \Rightarrow C_{k,\Delta} \subseteq C_{k',\Delta}$, and (ii) $\Delta' \sqsubseteq \Delta \Rightarrow C_{k,\Delta} \subseteq C_{k,\Delta'}$. The first point holds as, keeping the span Δ fixed, the maximal set of vertices C for which $d_\Delta(C, u) \geq k$ is clearly contained in the maximal set of vertices C' for which $d_\Delta(C', u) \geq k'$, if $k' \leq k$. To prove (ii), it can be noted that $\Delta' \sqsubseteq \Delta \Rightarrow E_\Delta \subseteq E_{\Delta'}$, which implies that $\forall u \in C_{k,\Delta} : d_\Delta(C_{k,\Delta}, u) \leq d_{\Delta'}(C_{k,\Delta}, u)$. Therefore, all vertices within $C_{k,\Delta}$ satisfy the condition to be part of $C_{k,\Delta'}$ too. \square

The following observation directly derives from Proposition 1 and states that finding all the span-cores having a fixed span Δ corresponds to computing the core decomposition of a simple graph.

OBSERVATION 1. *For a fixed temporal interval $\Delta \sqsubseteq T$, finding all span-cores that have Δ as their span is equivalent to computing the classic core decomposition [9] of the simple graph $G_\Delta = (V, E_\Delta)$.*

3.2 Maximal Span-cores

As the total number of temporal intervals that are contained into the whole time domain T is $|T|(|T|+1)/2$, the total number of span-cores is potentially $O(|T|^2 \times k_{max})$, where k_{max} is the largest value of k for which a (k, Δ) -core exists. It is thus quadratic in $|T|$, which may be too

large an output for human direct inspection. In this regard, it may be useful to focus only on the most relevant cores, i.e., the *maximal* ones, as defined next.

Definition 3 (Maximal Span-core). A span-core $C_{k,\Delta}$ of a temporal graph G is said *maximal* if there does not exist any other span-core $C_{k',\Delta'}$ of G such that $k \leq k'$ and $\Delta \subseteq \Delta'$.

Hence, a span-core is recognized as maximal if it is not dominated by another span-core both on the order k and the span Δ . Differently from the *innermost core* (i.e., the core of the highest order) in the classic core decomposition, which is unique, in our temporal setting, the number of maximal span-cores is $O(|T|^2)$, as, in the worst case, there may be one maximal span-core for every temporal interval. However, as observed in empirical temporal-network data, maximal span-cores are always much less than the overall span-cores: the difference is usually one order of magnitude or more. The second problem we tackle in this work is to compute the maximal span-cores of a temporal graph.

PROBLEM 2 (MAXIMAL SPAN-CORE MINING). *Given a temporal graph G , find the set of all maximal (k, Δ) -cores of G .*

Clearly, one could solve Problem 2 by solving Problem 1 and filtering out all the non-maximal span-cores. However, an interesting yet challenging question is whether one can exploit the maximality condition to develop faster algorithms that can directly extract the maximal ones, without computing all the span-cores. We provide a positive answer to this question in Section 5.

3.3 Relation to Multilayer Core Decomposition [38, 39]

Multilayer graphs are a representation paradigm of complex systems, where multiple relations of different types occur between the same pair of entities [17, 27, 82]. A multilayer graph is formally defined as a triple $G = (V, E, L)$, where V is a set of vertices, L is a set of layers, and $E \subseteq V \times V \times L$ is a set of edges. Given a multilayer graph $G = (V, E, L)$ and an $|L|$ -dimensional integer vector $\vec{k} = [k_\ell]_{\ell \in L}$, Galimberti et al. [38, 39] define the notion of *multilayer \vec{k} -core* of G as a *maximal* set $C \subseteq V$ of vertices such that, for all $\ell \in L$, the minimum degree of a vertex in C in layer ℓ is larger than or equal to k_ℓ . In other words, a \vec{k} -multilayer-core corresponds to a subgraph that satisfies the k_ℓ -core definition in layer ℓ , for all $\ell \in L$. For instance, for $|L| = 2$, a multilayer (k_1, k_2) -core is a subgraph that is simultaneously a k_1 -core in the first layer and a k_2 -core in the second layer.

Temporal graphs can be viewed as a special case of multilayer graphs where timestamps correspond to layers. Therefore, a natural question while introducing a notion of temporal core is how it relates to the definition of a core in the multilayer setting. A fundamental difference is that, unlike the multilayer context, in a temporal graph the “layers” are ordered, and the consecution of timestamps should be taken into account. As a result, the two definitions are not comparable and have different conceptual and computational properties in the general case. A major remark in this regard is that the multilayer cores, as defined in [38, 39], are exponential in the number $|L|$ of layers $|L|$, *thus the multilayer core decomposition takes (worst-case) exponential time, while temporal core decomposition is computable in polynomial time.*

Once having ascertained such a key difference, another meaningful investigation would be understanding whether the notion of multilayer core may still be exploited to define/compute span-cores, even if only in limited circumstances. In this regard, as formally shown in the next proposition and illustrated in the example in Figure 2, there exists a containment relationship between span-cores and multilayer cores.

PROPOSITION 2. *Let C be the (k, Δ) -span-core of a temporal graph $G = (V, T, \tau)$. Let also $\vec{k} = [k_t]_{t \in T}$ be a $|T|$ -dimensional integer vector such that $\forall t \in \Delta : k_t = k, \forall t \notin \Delta : k_t = 0$, and let C' be the*




Fig. 2. Relationship between multilayer cores defined in [38, 39] and span-cores introduced in this work (Proposition 2). The figure depicts a toy temporal graph G , with time domain $T = \{0, 1, 2\}$. Solid, dashed, and dotted edges refer to timestamp 0, 1, and 2, respectively. The $(2, [0, 1])$ -span-core of G corresponds to $C = \{A, D, E\}$. At the same time, $C' = \{A, B, D, E\}$ corresponds to the $(2, 2, 0)$ -multilayer-core of G , when G is interpreted as a (three-layer) multilayer graph with the first, second, and third layer corresponding to timestamps 0, 1, and 2, respectively. While there is no exact correspondence, it can be observed that span-core C is contained into multilayer core C' .

\vec{k} -multilayer-core extracted from G by interpreting it as a multilayer graph where layers correspond to timestamps. It holds that $C \subseteq C'$.

PROOF. According to Definition 2, every vertex $v \in C$, has a at least k neighbors within $C \setminus \{v\}$, for every timestamp $t \in \Delta$. This complies with the definition of \vec{k} -multilayer-core, $\vec{k} = [k_t]_{t \in T}, \forall t \in \Delta : k_t = k, \forall t \notin \Delta : k_t = 0$, meaning that all vertices in C are necessarily part of the \vec{k} -multilayer-core as well. \square

Proposition 2 suggests that, in principle, to compute span-cores, one may: (i) compute all multilayer cores, (ii) among all multilayer cores, retrieve the ones complying with Proposition 2, and (iii) post-process those multilayer cores in order to extract the actual span-cores. However, this strategy is not feasible, as, due to the aforementioned exponential-time computation, extracting multilayer cores from a temporal graph would be affordable only for very small values of $|T|$. In fact, Galimberti et al. [38, 39] show experiments on graphs with at most 10 layers, and in the 10-layer graphs computing the multilayer core decomposition takes more than 28 hours. In the temporal setting, we are interested in analyzing long-term, sometimes high-frequency, interactions, thus temporal graphs have typically much more than 10 timestamps. Indeed, all the datasets used in this article have many more timestamps, and the algorithms from [38, 39] cannot run in a reasonable amount of time. Moreover, even assuming to be able to compute multilayer cores on temporal networks, those cores have still to be filtered and post-processed, which makes this strategy meaningless with respect to methods that compute span-cores directly, as the ones introduced in the next section.

4 ALGORITHMS: COMPUTING ALL SPAN-CORES

In this section, we devise algorithms for computing a complete span-core decomposition of a temporal graph (Problem 1).

A naïve approach. As stated in Observation 1, for a fixed temporal interval $\Delta \subseteq T$, mining all span-cores $C_{k,\Delta}$ is equivalent to computing the classic core decomposition of the graph $G_\Delta = (V, E_\Delta)$. A naïve strategy is thus to run a core-decomposition subroutine [9] on graph G_Δ for each temporal interval $\Delta \subseteq T$. Such a method has time complexity $\mathcal{O}(\sum_{\Delta \subseteq T} (|\Delta| \times |E|))$, i.e., $\mathcal{O}(|T|^2 \times |E|)$.

A more efficient algorithm. Looking at Figure 1, one can observe that the naïve algorithm only exploits one dimension of the containment property: it starts from each point on the top

ALGORITHM 1: Span-cores

Input: A temporal graph $G = (V, T, \tau)$.
Output: The set C of all span-cores of G .

```

1  $C \leftarrow \emptyset$ ;  $Q \leftarrow \emptyset$ ;  $\mathcal{A} \leftarrow \emptyset$ 
2 forall  $t \in T$  do
3    $\lfloor$  enqueue  $[t, t]$  to  $Q$ ;  $\mathcal{A}[t, t] \leftarrow V$ 
4 while  $Q \neq \emptyset$  do
5   dequeue  $\Delta = [t_s, t_e]$  from  $Q$ 
6    $E_\Delta[\mathcal{A}[\Delta]] \leftarrow \{(u, v) \in E_\Delta \mid u \in \mathcal{A}[\Delta], v \in \mathcal{A}[\Delta]\}$ 
7   if  $|E_\Delta[\mathcal{A}[\Delta]]| > 0$  then
8      $C_\Delta \leftarrow \text{core-decomposition}(\mathcal{A}[\Delta], E_\Delta[\mathcal{A}[\Delta]])$ 
9      $C \leftarrow C \cup C_\Delta$ 
10     $\Delta_1 = [\max\{t_s - 1, 0\}, t_e]$ ;  $\Delta_2 = [t_s, \min\{t_e + 1, t_{max}\}]$ 
11    forall  $\Delta' \in \{\Delta_1, \Delta_2\} \mid \Delta' \neq \Delta$  do
12      if  $\mathcal{A}[\Delta'] \neq \text{NULL}$  then
13         $\mathcal{A}[\Delta'] \leftarrow \mathcal{A}[\Delta'] \cap C_{1,\Delta}$ 
14        enqueue  $\Delta'$  to  $Q$ 
15      else
16         $\mathcal{A}[\Delta'] \leftarrow C_{1,\Delta}$ 

```

level, i.e., from cores of order 1, and goes down vertically with the classic core decomposition. Based on Proposition 1, it is possible to design a more efficient algorithm that exploits also the “horizontal containment” relationships.

Example 1. Consider core $C_{1,[0,2]}$ in Figure 1: by Proposition 1 it holds that it is a subset of both $C_{1,[0,1]}$ and $C_{1,[1,2]}$. Therefore, to compute $C_{1,[0,2]}$, instead of starting from the whole V , one can start from $C_{1,[0,1]} \cap C_{1,[1,2]}$. Starting from a much smaller set of vertices can provide a substantial speed-up to the whole computation.

This observation, although simple, produces a speed-up of orders of magnitude as we will empirically show in Section 7. The next straightforward corollary of Proposition 1 states that, not only $C_{1,[0,2]} \subseteq C_{1,[0,1]} \cap C_{1,[1,2]}$, but this is the best one can get, meaning that intersecting these two span-cores is equivalent to intersecting all span-cores structurally containing $C_{1,[0,2]}$.

COROLLARY 1. *Given a temporal graph $G = (V, T, \tau)$, and a temporal interval $\Delta = [t_s, t_e] \sqsubseteq T$, let $\Delta_+ = [\min\{t_s + 1, t_e\}, t_e]$ and $\Delta_- = [t_s, \max\{t_e - 1, t_s\}]$. It holds that*

$$C_{1,\Delta} \subseteq (C_{1,\Delta_+} \cap C_{1,\Delta_-}) = \bigcap_{\Delta' \sqsubseteq \Delta} C_{1,\Delta'}$$

Example 2. Consider again $C_{1,[0,2]}$ in Figure 1: Proposition 1 states that it is a subset of $C_{1,[0,0]}$, $C_{1,[0,1]}$, $C_{1,[1,1]}$, $C_{1,[1,2]}$, $C_{1,[2,2]}$. Corollary 1 suggests that there is no need to intersect them all, but only $C_{1,[0,1]}$ and $C_{1,[1,2]}$: in fact, $C_{1,[0,1]} \subseteq C_{1,[0,0]} \cap C_{1,[1,1]}$ and $C_{1,[1,2]} \subseteq C_{1,[1,1]} \cap C_{1,[2,2]}$.

The main idea behind our efficient Span-cores algorithm (whose pseudocode is given as Algorithm 1) is to generate temporal intervals of increasing size (starting from size one) and, for each Δ of width larger than one, to initiate the core decomposition from $(C_{1,\Delta_+} \cap C_{1,\Delta_-})$, i.e., the smallest intersection of cores containing $C_{1,\Delta}$ (Corollary 1). The intervals to be processed are added to queue Q , which is initialized with the intervals of size one (Lines 2–3): these are the only intervals

for which no other interval can be used to reduce the set of vertices from which the core decomposition is started, thus they have to be initialized with the whole vertex set V . The algorithm utilizes a map \mathcal{A} that, given an interval Δ , returns the set of vertices to be used as a starting set of the core decomposition on Δ . The algorithm processes all intervals stored in Q , until Q has become empty (Lines 4–16). For every temporal interval Δ extracted from Q , the starting set of vertices is retrieved from $\mathcal{A}[\Delta]$ and the corresponding set of edges is identified (Line 6). Unless this is empty, the classic core-decomposition algorithm [9] is invoked over $(\mathcal{A}[\Delta], E_\Delta[\mathcal{A}[\Delta]])$ (Line 8) and its output (a set of span-cores of span Δ) is added to the ultimate output set C (Line 9).

Afterwards, the two intervals, denoted Δ_1 and Δ_2 , for which $C_{1,\Delta}$ can be used to obtain the smallest intersections of cores containing them (Corollary 1) are computed at Line 10. For Δ_1 (and analogously Δ_2), we check whether $\mathcal{A}[\Delta_1]$ has already been initialized (Line 12): this would mean that previously the other “father” (i.e., smallest containing core) of C_{1,Δ_1} has been computed, thus we can intersect $C_{1,\Delta}$ with $\mathcal{A}[\Delta_1]$ and enqueue Δ_1 to be processed (Lines 13–14). Instead, if $\mathcal{A}[\Delta_1]$ was not yet initialized, we initialize it with $C_{1,\Delta}$ (Line 16): in this case Δ_1 is not enqueued because it still lacks one father to be intersected before being ready for core decomposition. This procedural update of Q ensures that both fathers of every interval in Q exist and have been previously computed, thus no a-posteriori verification is needed.

Example 3. Consider again the search space in Figure 1. Algorithm 1 first processes the intervals $[0, 0]$, $[1, 1]$, $[2, 2]$, and $[3, 3]$. Then, it intersects $C_{1,[0,0]}$ and $C_{1,[1,1]}$ to initialize $C_{1,[0,1]}$, intersects $C_{1,[1,1]}$ and $C_{1,[2,2]}$ to initialize $C_{1,[1,2]}$, and intersects $C_{1,[2,2]}$ and $C_{1,[3,3]}$ to initialize $C_{1,[2,3]}$. Then, it continues with the intervals of size 3: it intersects $C_{1,[0,1]}$ and $C_{1,[1,2]}$ to initialize $C_{1,[0,2]}$ and so on.

The next theorem formally shows soundness and completeness of our Span-cores algorithm.

THEOREM 1. *Algorithm 1 is sound and complete for Problem 1.*

PROOF. The algorithm generates and processes a subset of temporal intervals $\mathcal{X} \subseteq \{\Delta \mid \Delta \sqsubseteq T\}$. For every interval $\Delta \subseteq \mathcal{X}$, it computes *all* span-cores $C_\Delta = \{C_{1,\Delta}, C_{2,\Delta}, \dots, C_{k,\Delta}\}$ defined on Δ by means of the core-decomposition subroutine on the graph $(\mathcal{A}[\Delta], E_\Delta[\mathcal{A}[\Delta]])$. The set of vertices $\mathcal{A}[\Delta]$ is equivalent to $(C_{1,\Delta_+} \cap C_{1,\Delta_-})$ because of Line 13 (Corollary 1) and the fact that Δ is enqueued (Line 14) only when both fathers have been processed and the intersection done. The correctness of doing the classic core decomposition is guaranteed by Observation 1.

As for completeness, it suffices to show that the intervals $\Delta \notin \mathcal{X}$ that have not been processed by the algorithm do not yield any span-core. The algorithm generates all temporal intervals size by size, starting from those of size one and then going to larger sizes. This is done by maintaining the queue Q . As said above, an interval Δ is enqueued as soon as both C_{1,Δ_+} and C_{1,Δ_-} have been processed. Thus, an interval Δ is not in \mathcal{X} only if either C_{1,Δ_+} or C_{1,Δ_-} does not exist. In this case $C_{1,\Delta}$ and all other $C_{k,\Delta}$ do not exist as well. \square

Discussion. Algorithm 1 exploits the “horizontal containment” relationships only at the first level of the search space. For a given Δ , once the restricted starting set of vertices has been defined for $k = 1$, the traditional core decomposition is started to produce all the span-cores of span Δ . In other words, for $k > 1$ only the “vertical containment” is exploited. Consider the span-core $C_{3,[1,2]}$ in Figure 1: we know that it is a subset of $C_{2,[1,2]}$ (“vertical”) and of $C_{3,[1,1]}$ and $C_{3,[2,2]}$ (“horizontal”). One could consider intersecting all these three span-cores before computing $C_{3,[1,2]}$. We tested this alternative approach, but concluded that the overhead of computing intersections and data-structure maintenance was outweighing the benefit of starting from a smaller vertex set.

The worst-case time complexity of Algorithm 1 is equal to the naïve approach, however, in practice, it is orders of magnitude faster, as shown in Section 7.

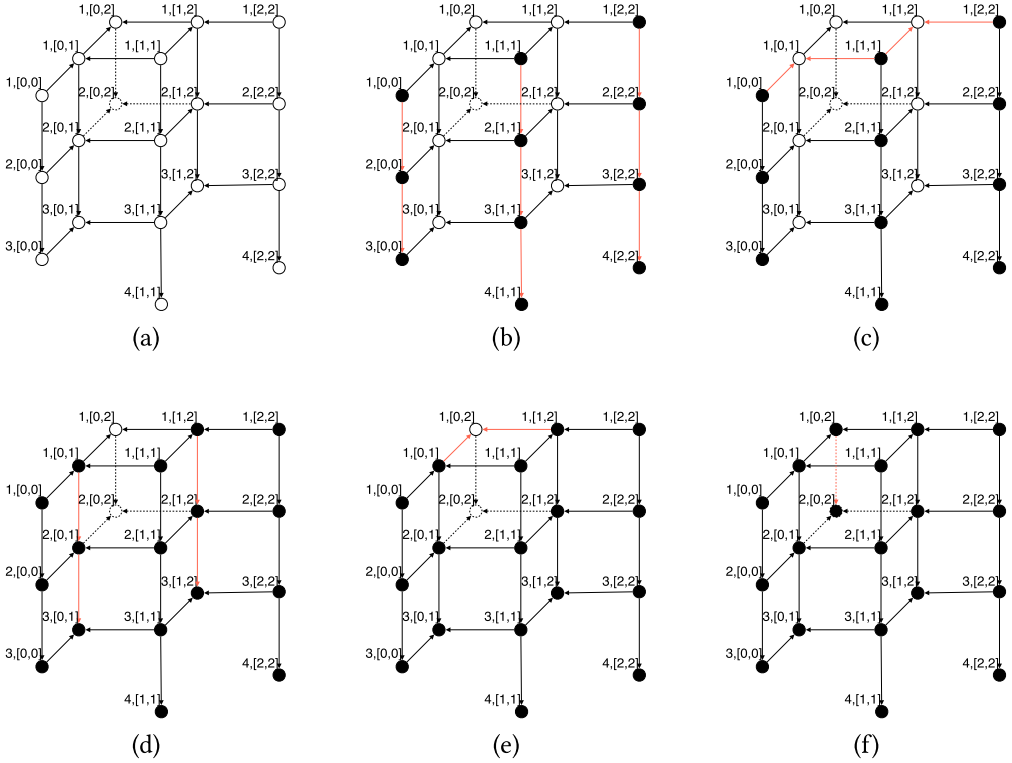


Fig. 3. Run-through example of the execution of Span-cores (Algorithm 1) over the search space of a temporal graph having $T = [0, 2]$. Full nodes represent computed span-cores, while empty nodes are span-cores that will be visited in the next steps of the algorithm. Red arrows highlight the containment relationships exploited during the current step.

Example 4. Figure 3 reports a run-through example, illustrating the execution of Span-cores (Algorithm 1) over the search space of a toy temporal graph having $T = [0, 2]$ (shown in Figure 3(a)). The algorithm starts by computing all the span-cores having span of size 1 (Figure 3(b)); in this case, only the “vertical containment” is exploited by the core-decomposition subroutine. Then, Span-cores proceeds with the computation of the span-cores having span of size 2. At first, the algorithm exploits the “horizontal containment” relationships at the first level of the search space to restrict the starting set of vertices for computing the span-cores of $k = 1$ (Figure 3(c)). Afterwards, the core-decomposition subroutine computes all the span-cores with span of size 2, by following the “vertical containment” (Figure 3(d)). Finally, the same method is applied for visiting the span-cores with span of size 3 (Figure 3(e)-(f)).

5 ALGORITHMS: COMPUTING MAXIMAL SPAN-CORES

In this section, we focus on Problem 2: computing the *maximal* span-cores of a temporal graph.

A filtering approach. As anticipated above, a straightforward way of solving this problem consists in filtering the span-cores computed during the execution of Algorithm 1, so as to ultimately output only the maximal ones. This can easily be accomplished by equipping Algorithm 1 with a data structure \mathcal{M} that stores the span-core of the highest order for every temporal interval $\Delta \subseteq T$ that has been processed by the algorithm. Moreover, at the storage of a span-core $C_{k,\Delta}$ in

\mathcal{M} , the span-cores previously stored in \mathcal{M} for subintervals of the temporal interval Δ and with the same order k are removed from \mathcal{M} . This removal operation, together with the order in which span-cores are processed, ensures that \mathcal{M} eventually contains only the maximal span-cores.

Efficient maximal-span-core finding. Our next goal is to design a more efficient algorithm that extracts maximal span-cores directly, without computing complete core decompositions, passing over more peripheral ones, and without generating all temporal cores. This is a quite challenging design principle, as it contrasts the intrinsic structural properties of core decomposition, based on which a core of order k is usually computed from the core of order $k-1$, thus making the computation of the core of the highest order as hard as computing the overall decomposition. Nevertheless, thanks to theoretical properties that relate the maximal span-cores to each other, in the temporal context such a challenge can be achieved. In the following we discuss such properties in detail, by starting from a result that has already been discussed above, but only informally.

Consider the classic core decomposition in a standard (non-temporal) graph G (Definition 1) and let $C_{k^*}[G]$ denote the *innermost* core of G , i.e., the non-empty k -core of G with the largest k .

LEMMA 1. *Given a temporal graph $G = (V, T, \tau)$, let C_M be the set of all maximal span-cores of G , and $C_{\text{inner}} = \{C_{k^*}[G_\Delta] \mid \Delta \sqsubseteq T\}$ be the set of innermost cores of all graphs G_Δ . It holds that $C_M \subseteq C_{\text{inner}}$.*

PROOF. Every $C_{k,\Delta} \in C_M$ is the innermost core of the non-temporal graph G_Δ : else, there would exist another core $C_{k',\Delta} \neq \emptyset$ with $k' > k$, implying that $C_{k,\Delta} \notin C_M$. \square

Lemma 1 states that each maximal span-core is an innermost core of a G_Δ , for some temporal interval $\Delta \sqsubseteq T$. Hence, there can exist at most one maximal span-core for every $\Delta \sqsubseteq T$ (while an interval Δ may not yield any maximal span-core). The key question to design an efficient maximal-span-core-mining algorithm thus becomes how to extract innermost cores of the graphs G_Δ more efficiently than by computing the full core decompositions of all G_Δ . The answer to this question comes from the result stated in the next two lemmas (with Lemma 2 being auxiliary to Lemma 3).

LEMMA 2. *Given a temporal graph $G = (V, T, \tau)$, and three temporal intervals $\Delta = [t_s, t_e] \sqsubseteq T$, $\Delta' = [t_s-1, t_e] \sqsubseteq T$, and $\Delta'' = [t_s, t_e+1] \sqsubseteq T$. The innermost core $C_{k^*}[G_\Delta]$ is a maximal span-core of G if and only if $k^* > \max\{k', k''\}$ where k' and k'' are the orders of the innermost cores of $G_{\Delta'}$ and $G_{\Delta''}$, respectively.*

PROOF. The “ \Rightarrow ” part comes directly from the definition of maximal span-core (Definition 3): if k^* were not larger than $\max\{k', k''\}$, then $C_{k^*}[G_\Delta]$ would be dominated by another span-core both on the order and on the span (as both Δ' and Δ'' are superintervals of Δ). For the “ \Leftarrow ” part, from Lemma 1 and Proposition 1 it follows that $\max\{k', k''\}$ is an upper bound on the maximum order of a span-core of a superinterval of Δ . Therefore, $k^* > \max\{k', k''\}$ implies that there cannot exist any other span-core that dominates $C_{k^*}[G_\Delta]$ both on the order and on the span. \square

LEMMA 3. *Given G , Δ , Δ' , Δ'' , k' , and k'' defined as in Lemma 2, let $\tilde{V} = \{u \in V \mid d_\Delta(V, u) > \max\{k', k''\}\}$, and let $C_{k^*}[G_\Delta[\tilde{V}]]$ be the innermost core of $G_\Delta[\tilde{V}]$. If $k^* > \max\{k', k''\}$, then $C_{k^*}[G_\Delta[\tilde{V}]]$ is a maximal span-core; otherwise, no maximal span-core exists for Δ .*

PROOF. Lemma 2 states that, to be recognized as a maximal span-core, the innermost core of G_Δ should have order larger than $\max\{k', k''\}$. This means that, if the innermost core of G_Δ is a maximal span-core, all vertices $u \notin \tilde{V}$ cannot be part of it. Therefore, G_Δ yields a maximal span-core only if the innermost core of subgraph $G_\Delta[\tilde{V}]$ has order $k^* > \max\{k', k''\}$. \square

Lemma 3 provides the basis of our efficient method for extracting maximal span-cores. Basically, it states that, to verify whether a certain temporal interval $\Delta = [t_s, t_e]$ yields a maximal span-core (and, if so, compute it), there is no need to consider the whole graph G_Δ , rather it suffices

ALGORITHM 2: Maximal-span-cores

Input: A temporal graph $G = (V, T, \tau)$.
Output: The set C_M of all maximal span-cores of G .

```

1  $C_M \leftarrow \emptyset$ 
2  $\mathcal{K}'[t] \leftarrow 0, \forall t \in T$ 
3 forall  $t_s \in [0, 1, \dots, t_{max}]$  do
4    $t^* \leftarrow \max\{t_e \in [t_s, t_{max}] \mid E_{[t_s, t_e]} \neq \emptyset\}$ 
5    $k'' \leftarrow 0$ 
6   forall  $t_e \in [t^*, t^*-1, \dots, t_s]$  do
7      $\Delta \leftarrow [t_s, t_e]$ 
8      $lb \leftarrow \max\{\mathcal{K}'[t_e], k''\}$ 
9      $V_{lb} \leftarrow \{u \in V \mid d_\Delta(V, u) > lb\}$ 
10     $E_\Delta[V_{lb}] \leftarrow \{(u, v) \in E_\Delta \mid u \in V_{lb}, v \in V_{lb}\}$ 
11     $C \leftarrow \text{innermost-core}(V_{lb}, E_\Delta[V_{lb}])$ 
12     $k^* \leftarrow \text{order of } C$ 
13    if  $k^* > lb$  then
14       $C_M \leftarrow C_M \cup \{C\}$ 
15     $k'' \leftarrow \max\{k'', k^*\}; \mathcal{K}'[t_e] \leftarrow \max\{\mathcal{K}'[t_e], k''\}$ 

```

to start from a smaller subgraph, which is given by all vertices whose temporal degree is larger than the maximum between the orders of the innermost cores of intervals $\Delta' = [t_s - 1, t_e]$ and $\Delta'' = [t_s, t_e + 1]$. This finding suggests a strategy that is opposite to the one used for computing the overall span-core decomposition: a *top-down* strategy that processes temporal intervals starting from the larger ones. Indeed, in addition to exploiting the result in Lemma 3, this way of exploring the temporal-interval space allows us to skip the computation of complete core decompositions of the whole “singleton-interval” graphs $\{G_{[t, t']}\}_{t \in T}$, which may easily become a critical bottleneck, as they are the largest ones among the graphs induced by temporal intervals.

The Maximal-span-cores algorithm. Algorithm 2 iterates over all timestamps $t_s \in T$ in *increasing order* (Line 3), and for each t_s it first finds all the maximal span-cores that have span starting in t_s . This way of proceeding *ensures that a span-core that is recognized as maximal will not be later dominated by another span-core*. Indeed, an interval $[t_s, t_e]$ can never be contained in another interval $[t'_s, t'_e]$ with $t_s < t'_s$. For a given t_s , all maximal span-cores are computed as follows. First, the maximum timestamp $\geq t_s$ such that the corresponding edge set $E_{[t_s, t_e]}$ is not empty is identified as t^* (Line 7). Then, all intervals $\Delta = [t_s, t_e]$ are considered one by one in *decreasing order* of t_e (Lines 6–7): this again *guarantees that a span-core that is recognized as maximal will not be later dominated by another span-core, as the intervals are processed from the largest to the smallest*. At each iteration of the internal cycle, the algorithm resorts to Lemma 3 and computes the lower bound lb on the order of the innermost core of G_Δ to be recognized as maximal, by taking the maximum between $\mathcal{K}'[t_e]$ and k'' (Line 8). \mathcal{K}' is a map that maintains, for every timestamp $t \in [t_s, t^*]$, the order of the innermost core of graph $G_{\Delta'}$, where $\Delta' = [t_s - 1, t]$ (i.e., $\mathcal{K}'[t]$ stores what in Lemmas 2–3 is denoted as k'). Whereas k'' stores the order of the innermost core of $G_{\Delta''}$, where $\Delta'' = [t_s, t_e + 1]$. Afterwards, the sets of vertices V_{lb} and of edges $E_\Delta[V_{lb}]$ that comply with this lower-bound constraint are built (Lines 9–10), and the innermost core of the subgraph $(V_{lb}, E_\Delta[V_{lb}])$ is extracted (Lines 11–12). Ultimately, based again on Lemma 3, such a core is added to the output set of maximal span-cores only if its order is actually larger than lb (Lines 13–14), and the values of k'' and $\mathcal{K}'[t_e]$ are updated (Line 15). Specifically, note that the order k^* of core C

may in principle be less than k'' , as C is extracted from a subgraph of G_Δ . If this happens, it means that the actual order of the innermost core of G_Δ is equal to k'' . This motivates the update rules (and their order) reported in Line 15.

THEOREM 2. *Algorithm 2 is sound and complete for Problem 2.*

PROOF. The algorithm processes all temporal intervals $\Delta \sqsubseteq T$ yielding a non-empty edge set E_Δ , in an order such that no interval is processed before one of its superintervals: this guarantees that a span-core recognized as maximal will not be dominated by another span-core found later on. For every Δ it extracts a core C that is used as a proxy of the innermost core of graph G_Δ . C is added to the output set C_M only if Lemma 3 recognizes it as a maximal span-core, otherwise it is discarded. This proves the soundness of the algorithm. Completeness follows from Lemma 1, which states that to extract all maximal span-cores it suffices to focus on the innermost cores of graphs $\{G_\Delta \mid \Delta \sqsubseteq T\}$, and Lemma 3 again, which states the condition for a proxy core C to be safely discarded because it is a non-maximal span-core. \square

Discussion. The worst-case time complexity of Algorithm 2 is the same as the algorithm for computing the overall span-core decomposition, i.e., $O(|T|^2 \times |E|)$. It is worth mentioning that it is not possible to do better than this, as the output itself is potentially quadratic in $|T|$. However, as we will show in Section 7, the proposed algorithm is in practice much more efficient than computing the overall span-core decomposition and filtering out the non-maximal span-cores as, in this case, we avoid the visit of portions of the span-coresearch space and the computations are run over subgraphs of reduced dimensions.

To conclude, we discuss how the crucial operation of building the subgraph $(V_{lb}, E_\Delta[V_{lb}])$ may be carried out efficiently in terms of both time and space. Consider a fixed timestamp $t_s \in [0, \dots, t_{max}]$. The following reasoning holds for every t_s . Let $E^-(t_e) = E_{[t_s, t_e]} \setminus E_{[t_s, t_{e+1}]}$ be the set of edges that are in $E_{[t_s, t_e]}$ but not in $E_{[t_s, t_{e+1}]}$, for $t_e \in [t_s, \dots, t^* - 1]$. As a first general step, for each t_s , we compute and store *all* edge sets $\{E^-(t_e)\}_{t_e \in [t_s, t^* - 1]}$. These operations can be accomplished in $O(|T| \times |E|)$ overall time, because every $E^-(t_e)$ can be computed incrementally from $E_{[t_s, t_e]}$ as $E^-(t_e) = \{(u, v) \in E_{[t_s, t_e]} \mid \tau(u, v, t_e + 1) = 0\}$. Moreover, for any timestamp t_e , we keep a map \mathcal{D} storing all vertices of $G_{[t_s, t_e]}$ organized by degree. Specifically, the set $\mathcal{D}[k]$ contains all vertices having degree $> k$ in $G_{[t_s, t_e]}$. Every vertex in \mathcal{D} is thus replicated a number of times equal to its degree. This way, the overall space taken by \mathcal{D} is $O(|E|)$, i.e., as much space as G . \mathcal{D} is initialized as empty (when $t_e = t^*$) and repeatedly augmented as t_e decreases, by a linear scan of the various $E^-(t_e)$. The overall filling of \mathcal{D} (for all t_e) therefore takes $O(|T| \times |E|)$ time. Then, the desired V_{lb} can be computed in constant time simply as $V_{lb} = \mathcal{D}[lb]$.

As for $E_\Delta[V_{lb}]$, for any t_e , we first reconstruct $E_{[t_s, t_e]}$ as $E_{[t_s, t_{e+1}]} \cup E^-(t_e)$, having previously computed $E_{[t_s, t_{e+1}]}$. Note that storing all $E^-(t_e)$ takes $O(|E|)$ space. That is why we store all $E^-(t_e)$ and reconstruct $E_{[t_s, t_e]}$ afterward (instead of storing the latter, which would take $O(|T| \times |E|)$ space). $E_\Delta[V_{lb}]$ is ultimately derived by a linear scan of $E_{[t_s, t_e]}$, taking all edges in $E_{[t_s, t_e]}$ having both endpoints in V_{lb} . This way, the step of building $E_\Delta[V_{lb}]$ for all t_e takes again $O(|T| \times |E|)$ overall time.

Example 5. We report here a run-through example of the execution of Maximal-span-cores (Algorithm 4) over the search space of a temporal graph having $T = [0, 2]$ (the same shown in Figure 3(a)). Maximal-span-cores starts by identifying the span-core of highest order in the largest possible temporal interval Δ , i.e., $\Delta = T$ (Figure 4 (a)). Such a span-core is guaranteed to be maximal, since the span-core of highest order with span T cannot be dominated in terms of span by any other span-core. The algorithm then processes interval $[0, 1]$ (Figure 4(b)): here $lb = 2$, since the only constraint derives from the identification of core $C_{2, [0, 2]}$ as maximal, therefore core $C_{3, [0, 1]}$

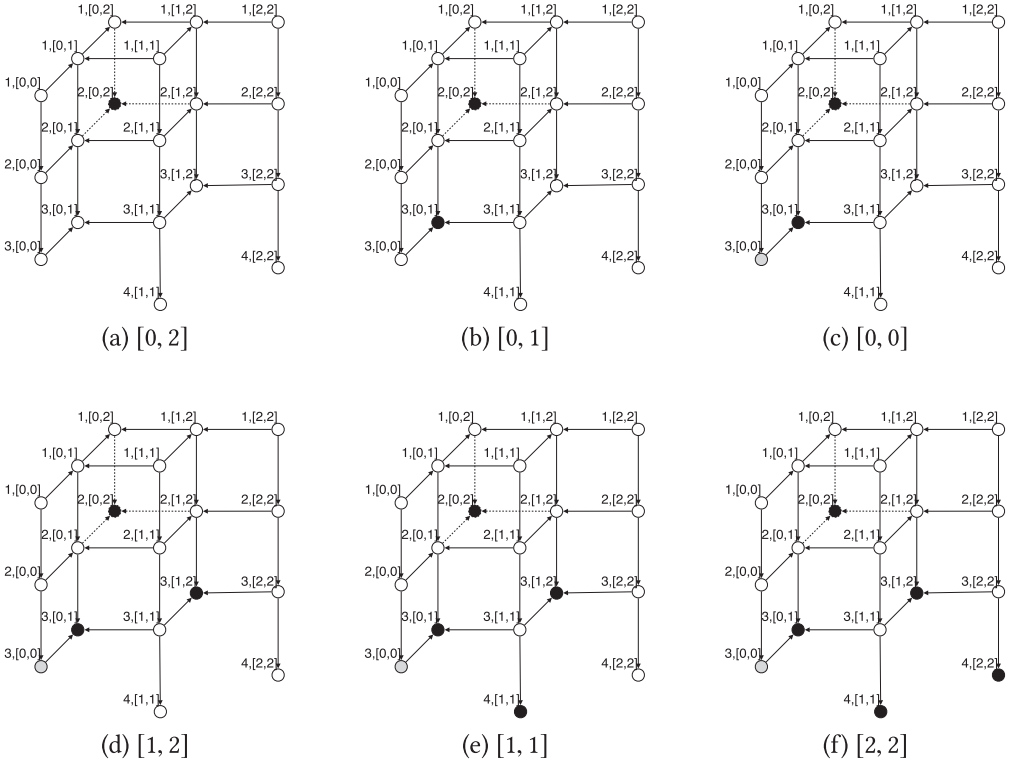


Fig. 4. Run-through example of the execution of Maximal-span-cores (Algorithm 4) over the search space of a temporal graph having $T = [0, 2]$ (same as Figure 3(a)). Full nodes represent computed span-cores: span-cores in black are recognized as maximal, while those in gray are discarded from the set of the maximal span-cores.

is recognized as maximal. Next, the algorithm searches for the last possible maximal span-core having span $\Delta = [t_s, t_e]$ such that $t_s = 0$, i.e., $\Delta = [0, 0]$ (Figure 4(c)). Core $C_{3,[0,0]}$ is computed, but discarded from the solution maximal, because it has order equal to the lower bound $lb = 3$, derived from core $C_{3,[0,1]}$. The algorithm proceeds in a similar way, by finding a maximal span-core in all the remaining intervals, i.e., $[1, 2]$, $[1, 1]$, and $[2, 2]$ (Figures 4(d)–(f)). It is important to note that in all such cases, the lower bound lb for the existence of a maximal span-core in a given temporal interval accounts for two factors. For example, consider interval $[1, 1]$. The innermost core in $G_{[1,1]}$ is a maximal span-core if it has order greater than both cores $C_{3,[0,1]}$ and $C_{3,[1,2]}$, that is $lb = 3$.

6 TEMPORAL COMMUNITY SEARCH

Community search in static graphs aims at finding a dense subgraph (community) containing a set of input query vertices [36, 51]. In the temporal setting, it is very likely that the communities spanning the query vertices change over time. To be more precise, it may happen that a certain subgraph S is a well-representative community for the given query vertices Q , but only for a certain time interval Δ . Instead, for another time interval Δ' , a relevant community for Q might correspond to a completely different subgraph S' . For this reason, we formulate community search on temporal networks as the problem of finding h subgraphs (with $h > 0$ being an input parameter) containing the query vertices, together with their temporal span, such that the sum of the density

of those subgraphs is maximized and the union of their temporal spans corresponds to the whole input temporal domain. Among the many densities proposed in the literature, here we follow the bulk of the literature on community search, and adopt the minimum-degree density [36, 51]. In fact, as well-discussed, among others, by Sozio and Gionis [79] in their seminal work, unlike other density notions, including the popular average degree, the minimum-degree density has the capability of mitigating the so-called “free-rider” effect, i.e., the fact that (large) subgraphs may be arbitrarily added to a community-search solution to artificially increase the objective-function value, and thus lead to unintuitive yet unnecessarily large output solutions. Formally, the problem we study in this work is:

PROBLEM 3 (TEMPORAL COMMUNITY SEARCH). *Given a temporal graph $G = (V, T, \tau)$, a set $Q \subseteq V$ of query vertices, and a positive integer $h \in \mathbb{N}^+$, find a set $\{(S_i, \Delta_i)\}_{i=1}^h$ of h pairs such that (i) $\forall 1 \leq i \leq h : Q \subseteq S_i \subseteq V$, (ii) $\bigcup_{1 \leq i \leq h} \Delta_i = T$, and (iii) the following is maximized:*

$$\sum_{i=1}^h \min_{u \in S_i} d_{\Delta_i}(S_i, u). \quad (1)$$

The input integer h is a user-defined parameter that gives the analyst the flexibility of requiring a specific number of output temporal communities, which might vary from application to application.

6.1 Connection with SEQUENCE SEGMENTATION

Here we provide some theoretical insights into the TEMPORAL COMMUNITY SEARCH problem. The main result we provide at the end of this subsection is an interesting connection with the well-established SEQUENCE SEGMENTATION problem [10]. As shown in the next subsections, such a result forms the basis for algorithmic design.

Let us first consider a single-interval variant of Problem 3: for a fixed temporal interval Δ , find a subgraph containing the input set Q of query vertices that maximizes the minimum temporal degree within Δ . Formally:

PROBLEM 4 (SINGLE TEMPORAL COMMUNITY SEARCH). *Given a temporal graph $G = (V, T, \tau)$, a set $Q \subseteq V$ of query vertices, and an interval $\Delta \sqsubseteq T$, find*

$$S^* = \operatorname{argmax}_{Q \subseteq S \subseteq V} \min_{u \in S} d_{\Delta}(S, u).$$

It is easy to see that solving Problem 4 corresponds to solving minimum-degree-based community search on graph G_{Δ} . Therefore, a solution to Problem 4 can straightforwardly be computed by applying a standard result on minimum-degree-based community search, which states that the highest-order core containing all query vertices is a solution to that problem [7]. This finding is formalized next.

Definition 4 ((Q, Δ)-highest-order-span-core). Given a temporal graph $G = (V, T, \tau)$, a set $Q \subseteq V$ of query vertices, and an interval $\Delta \sqsubseteq T$, the (Q, Δ)-highest-order-span-core of G , denoted $C_{Q, \Delta}^*$, is defined as the highest-order span-core among all span-cores of G with temporal span Δ and containing all query vertices in Q . Let also $v_{Q, \Delta}^*$ denote the order of $C_{Q, \Delta}^*$.

FACT 1. *Given a temporal graph $G = (V, T, \tau)$, a set $Q \subseteq V$ of query vertices, and an interval $\Delta \sqsubseteq T$, the (Q, Δ)-highest-order-span-core of G is a solution to Problem 4 on input $\langle G, Q, \Delta \rangle$.*

Note that Problem 4 may have multiple solutions: $C_{Q, \Delta}^*$ is only one of those possibly many ones. $C_{Q, \Delta}^*$ can be computed by running a core decomposition on (static) graph G_{Δ} , and stopping it

when the first core that does not contain all query vertices in Q has been encountered. Therefore, Problem 4 can be solved in $O(|\Delta| \times |E|)$ time.

In light of the above findings, an alternative yet equivalent way of formulating our TEMPORAL COMMUNITY SEARCH problem is to ask for a *segmentation* (i.e., a partition) of the time domain T into a set $\{\Delta_i\}_{i=1}^h$ of h intervals so as to maximize the sum $\sum_{i=1}^h v_{Q, \Delta_i}^*$ of the orders of the (Q, Δ) -highest-order-span-cores of those identified intervals. Once such an optimal segmentation of T has been computed, the ultimate $\{(S_i, \Delta_i)\}_{i=1}^h$ pairs are derived by simply setting $S_i = C_{Q, \Delta_i}^*, \forall 1 \leq i \leq h$. Formally:

PROBLEM 5 (ALTERNATIVE FORMULATION OF PROBLEM 3). *Given a temporal graph $G = (V, T, \tau)$, a set $Q \subseteq V$ of query vertices, and a positive integer $h \in \mathbb{N}^+$, find a set $\{(S_i, \Delta_i)\}_{i=1}^h$ of h pairs such that (i) $\forall 1 \leq i \leq h : S_i = C_{Q, \Delta_i}^*$, (ii) $\{\Delta_i\}_{i=1}^h$ is a partition of T , and (iii) the following is maximized:*

$$\sum_{i=1}^h v_{Q, \Delta_i}^*. \quad (2)$$

Correspondence between Problems 3 and 5 easily follows from Fact 1 and from the observation that for any feasible solution $\{(S_i, \Delta_i)\}_{i=1}^h$ to Problem 3 with overlapping intervals, there exists an overlapping-interval-free feasible solution with not smaller objective-function value. To see the latter, for any two overlapping intervals Δ_i and Δ_j , simply replace one of the two intervals, say Δ_i , with $\Delta'_i = \Delta_i \setminus (\Delta_i \cap \Delta_j)$. As $\Delta'_i \sqsubseteq \Delta_i$, it holds that $v_{Q, \Delta'_i}^* \geq v_{Q, \Delta_i}^*$, therefore the resulting overlapping-interval-free solution will have objective-function value greater than or equal to the objective-function value of the starting solution with overlapping intervals.

Thanks to the reformulation in Problem 5, it is immediate to observe that our TEMPORAL COMMUNITY SEARCH problem is an instance of the well-established SEQUENCE SEGMENTATION problem, which asks for partitioning a sequence of numerical values into b segments so as to minimize the sum of the penalties (according to some penalty function) on each identified segment [10]:

PROBLEM 6 (SEQUENCE SEGMENTATION [10]). *Given a sequence $X = (x_0, x_1, \dots, x_{max})$ of numerical values, and a function $p : \{Y\}_{Y \subseteq X} \rightarrow \mathbb{R}$ that assigns a penalty score to every subsequence Y of X , partition X into a set $\{X_i\}_{i=1}^b$ of b subsequences such that $\sum_{i=1}^b p(X_i)$ is minimized.*

FACT 2. *TEMPORAL COMMUNITY SEARCH (Problem 3) on input $\langle G = (V, T, \tau), Q, h \rangle$ is an instance of SEQUENCE SEGMENTATION (Problem 6) with $X = T$, $b = h$, and $\forall \Delta \sqsubseteq T : p(\Delta) = -v_{Q, \Delta}^*$.*

In the following two subsections we show how to exploit the result in Fact 2 (and a further important finding about maximal span-cores) to design efficient algorithms for our TEMPORAL COMMUNITY SEARCH problem.

6.2 A Basic Algorithm (Based on All Span-cores)

SEQUENCE SEGMENTATION can be solved in $O(|X|^2 \times h + \tau_p)$ time via DP [10], where τ_p is the overall time spent for computing the penalty score of all subsequences of the input sequence X (according to the given penalty function p). Thanks to the connection shown in Fact 2, the DP algorithm for SEQUENCE SEGMENTATION can be easily adapted to solve TEMPORAL COMMUNITY SEARCH as well. The pseudocode of this algorithm – termed Temporal-community-search – is reported as Algorithm 3, and described next.

The Temporal-community-search algorithm makes use of two $(|T| \times h)$ -dimensional matrices, i.e., \mathbf{P} and \mathbf{R} . Matrix \mathbf{P} represents the *penalty matrix*. It contains, $\forall t \in T, \forall i \in [0, h)$, the minimum cost of segmenting the sequence corresponding to the first t timestamps of T into $i + 1$ segments. As a result, $\mathbf{P}[t_{max}, h - 1]$ contains the objective-function value of the ultimate optimal solution

ALGORITHM 3: Temporal-community-search

Input: A temporal graph $G = (V, E, T)$, a set $Q \subseteq V$ of query vertices, an integer $h \in \mathbb{N}^+$.
Output: A set $\{\{S_i, \Delta_i\}_{i=1}^h\}$, where $Q \subseteq S_i \subseteq V, \forall 1 \leq i \leq h$, and $\{\Delta_i\}_{i=1}^h$ is a partition of T .

```

/* Initialization */
1 Compute  $v_{Q,\Delta}^*$  and  $C_{Q,\Delta}^*, \forall \Delta \sqsubseteq T$ , via  $Q$ -constrained span-core decomposition
2  $\mathbf{P} \leftarrow$  an empty  $(|T| \times h)$ -dimensional matrix // Penalty matrix
3  $\mathbf{R} \leftarrow$  an empty  $(|T| \times h)$ -dimensional matrix // Reconstruction matrix
4 forall  $t \in T$  do
5    $\mathbf{P}[t, 0] \leftarrow -v_{Q,[0,t]}^*$ 
6    $\mathbf{R}[t, 0] \leftarrow 0$ 
/* Dynamic-programming step */
7 forall  $t \in T$  do
8   forall  $i \in [1, h]$  do
9      $\mathbf{P}[t, i] \leftarrow \min_{\ell \in [0, t]} \mathbf{P}[\ell, i-1] - v_{Q, [\ell+1, t]}^*$ 
10     $\mathbf{R}[t, i] \leftarrow \operatorname{argmin}_{\ell \in [0, t]} \mathbf{P}[\ell, i-1] - v_{Q, [\ell+1, t]}^*$ 
/* Reconstruction of the solution */
11  $ub \leftarrow t_{max}$ 
12 forall  $i \in (h, 0]$  do
13    $lb \leftarrow \mathbf{R}[ub, i]$ 
14    $\Delta_i \leftarrow [lb, ub]$ 
15    $ub \leftarrow lb - 1$ 
16 forall  $i \in (h, 0]$  do
17    $S_i \leftarrow C_{Q, \Delta_i}^*$ 

```

to Problem 5. Matrix \mathbf{R} is the *reconstruction matrix*. It provides information about the optimal segmentation, and is used at the end of the algorithm to reconstruct the output $\{\Delta_i\}_{i=1}^h$. Note that the algorithm does not explicitly compute the S_i subgraphs corresponding to the optimal Δ_i intervals. In fact, as discussed above, each S_i can be easily retrieved at the end of the algorithm, by simply setting it equal to the corresponding (Q, Δ_i) -highest-order-span-core C_{Q, Δ_i}^* . According to Fact 2, the penalty score of an interval $\Delta \sqsubseteq T$ corresponds to $-v_{Q, \Delta}^*$, i.e., the negative of the order of the (Q, Δ) -highest-order-span-core $C_{Q, \Delta}^*$. All individual $v_{Q, \Delta}^*$ values, for all $\Delta \sqsubseteq T$, are efficiently computed altogether, at the beginning of the algorithm, via a “ Q -constrained” variant of span-core decomposition (an alternative, but much less efficient strategy consists in computing every single $v_{Q, \Delta}^*$ from scratch, on the fly). Specifically, a simple (yet more efficient) variant of the span-core decomposition algorithm (Algorithm 1) is employed for this purpose, which outputs only those span-cores containing all the vertices in Q . This is easily achievable by stopping the core-decomposition subroutine, for every interval $\Delta \sqsubseteq T$, as soon as a core not containing all query vertices in Q has been encountered.

The time complexity of Algorithm 3 is $O(|T|^2 \times h + \tau_{sc})$, where τ_{sc} is the time spent for computing the Q -constrained span-core decomposition of the input graph G .

6.3 A More Efficient Algorithm (Based on Maximal Span-cores)

A more efficient algorithm can be designed by noticing that, actually, one does not need to consider all timestamps in T in the DP step. Rather, focusing on a subset $T^* \subseteq T$ – which is properly defined

ALGORITHM 4: Efficient-temporal-community-search

Input: A temporal graph $G = (V, E, T)$, a set $Q \subseteq V$ of query vertices, an integer $h \in \mathbb{N}^+$.
Output: A set $\{\{S_i, \Delta_i\}\}_{i=1}^h$, where $Q \subseteq S_i \subseteq V, \forall 1 \leq i \leq h$, and $\{\Delta_i\}_{i=1}^h$ is a partition of T .

/* Identification of T^* */

- 1 Compute the set $C_M(Q)$ of Q -constrained maximal span-cores of G
- 2 $\mathbf{D} \leftarrow \{\Delta \sqsubseteq T \mid C_{k,\Delta} \in C_M(Q)\}$
- 3 $T_{\mathbf{D}} \leftarrow \bigcup_{\Delta \in \mathbf{D}} \Delta$; $T_{\mathbf{D}}^+ \leftarrow \{\min\{t_e+1, t_{max}\} \mid [t_s, t_e] \in \mathbf{D}\}$; $T_{\mathbf{D}}^- \leftarrow \{\max\{t_s-1, 0\} \mid [t_s, t_e] \in \mathbf{D}\}$
- 4 $T_{sup} \leftarrow \{t_i \in T \setminus (T_{\mathbf{D}} \cup T_{\mathbf{D}}^- \cup T_{\mathbf{D}}^+ \cup \{t_{max}\}) \mid i \in [1, h+1 - |T_{\mathbf{D}} \cup T_{\mathbf{D}}^- \cup T_{\mathbf{D}}^+ \cup \{t_{max}\}|]\}$
- 5 $T^* \leftarrow T_{\mathbf{D}} \cup T_{\mathbf{D}}^+ \cup T_{\mathbf{D}}^- \cup \{t_{max}\} \cup T_{sup}$

/* Initialization */

- 6 Compute $v_{Q,\Delta}^*, \forall \Delta \sqsubseteq T$
- 7 $\mathbf{M} \leftarrow$ mapping function $[0, |T^*|) \rightarrow T^*$
- 8 $\mathbf{P} \leftarrow$ an empty $(|T^*| \times h)$ -dimensional matrix // Penalty matrix
- 9 $\mathbf{R} \leftarrow$ an empty $(|T^*| \times h)$ -dimensional matrix // Reconstruction matrix
- 10 **forall** $r \in [0, |T^*|)$ **do**
- 11 $\mathbf{P}[r, 0] \leftarrow -v_{Q, [0, \mathbf{M}[r]]}^*$
- 12 $\mathbf{R}[r, 0] \leftarrow 0$

/* Dynamic-programming step */

- 13 **forall** $r \in [0, |T^*|)$ **do**
- 14 **forall** $i \in [1, h)$ **do**
- 15 $\mathbf{P}[r, i] \leftarrow \min_{\ell \in [0, r]} \mathbf{P}[\ell, i-1] - v_{Q, [\mathbf{M}[\ell+1], \mathbf{M}[r]]}^*$
- 16 $\mathbf{R}[r, i] \leftarrow \operatorname{argmin}_{\ell \in [0, r]} \mathbf{P}[\ell, i-1] - v_{Q, [\mathbf{M}[\ell+1], \mathbf{M}[r]]}^*$

/* Reconstruction of the solution */

- 17 $ub \leftarrow |T^*| - 1$
- 18 **forall** $i \in (h, 0)$ **do**
- 19 $lb \leftarrow \mathbf{R}[ub, i]$
- 20 $\Delta_i \leftarrow [\mathbf{M}[lb], \mathbf{M}[ub]]$
- 21 $ub \leftarrow lb - 1$
- 22 **forall** $i \in (h, 0)$ **do**
- 23 $S_i \leftarrow C_{Q, \Delta_i}^*$

based on the maximal span-cores of the input graph, see next – allows for significantly reducing the dimensionality of the penalty matrix P and the reconstruction matrix R , hence the overall time complexity of the algorithm, without affecting optimality of the output solution. The following fact provides the theoretical basis for defining such a reduced temporal domain T^* .

FACT 3. *Given a temporal graph $G = (V, T, \tau)$ and a set $Q \subseteq V$ of query vertices, let $C_M(Q)$ be the set of all Q -constrained maximal span-cores of G . For a temporal interval $\Delta \sqsubseteq T$, it holds that $v_{Q,\Delta}^* = \max\{0, \max\{k \mid C_{k,\Delta'} \in C_M(Q), \Delta \sqsubseteq \Delta'\}\}$.*

Fact 3 states that the penalty score $v_{Q,\Delta}^*$ of an interval Δ corresponds to the maximum among the orders of the Q -constrained maximal span-cores whose span includes Δ , if some exist. If an interval Δ is not a subset of any span of a Q -constrained maximal span-core, then $v_{Q,\Delta}^* = 0$. In that case, therefore, Δ can be safely discarded, as it cannot be part of the optimal solution of the given TEMPORAL COMMUNITY SEARCH problem instance (unless it is needed to fill possible “holes”, see below). The ultimate consequence of this finding is that the aforementioned reduced

temporal domain T^* is identified by the timestamps covered by the spans of the maximal span-cores, along with auxiliary timestamps, which are needed to ensure a smooth execution of the DP step, as well as a correct handling of some extreme cases. Specifically, let $\mathbf{D} = \{\Delta \sqsubseteq T \mid C_{k,\Delta} \in C_M(Q)\}$ be the set of the spans of the Q -constrained maximal span-cores of the input graph, and $T_{\mathbf{D}} = \bigcup_{\Delta \in \mathbf{D}} \Delta$ be the set of timestamps that are part of a span of a Q -constrained maximal span-core. The first two sets of auxiliary timestamps correspond to the timestamps that immediately precede and succeed the intervals in \mathbf{D} , i.e., the sets $T_{\mathbf{D}}^+ = \{\min\{t_e + 1, t_{max}\} \mid [t_s, t_e] \in \mathbf{D}\}$ and $T_{\mathbf{D}}^- = \{\max\{t_s - 1, 0\} \mid [t_s, t_e] \in \mathbf{D}\}$, respectively. The timestamps in $T_{\mathbf{D}}^+$ and $T_{\mathbf{D}}^-$ (along with the last timestamp t_{max} of the input temporal domain T) are needed to allow the DP step to identify a solution that actually covers the whole temporal domain T (as per Condition (ii) of Problem 3). In particular, such timestamps may be interpreted as a trick to give the DP step the flexibility to select “holes” (i.e., time intervals in-between two consecutive but not necessarily contiguous timestamps in $T_{\mathbf{D}}$). Moreover, we define T_{sup} as the set of the first $h + 1 - |T_{\mathbf{D}} \cup T_{\mathbf{D}}^- \cup T_{\mathbf{D}}^+ \cup \{t_{max}\}|$ timestamps of T not contained in $T_{\mathbf{D}} \cup T_{\mathbf{D}}^- \cup T_{\mathbf{D}}^+ \cup \{t_{max}\}$, i.e., $T_{sup} = \{t_i \in T \setminus (T_{\mathbf{D}} \cup T_{\mathbf{D}}^- \cup T_{\mathbf{D}}^+ \cup \{t_{max}\}) \mid i \in [1, h + 1 - |T_{\mathbf{D}} \cup T_{\mathbf{D}}^- \cup T_{\mathbf{D}}^+ \cup \{t_{max}\}|\}$. The timestamps in T_{sup} are further auxiliary timestamps that are needed to return a correct h -sized solution when the timestamps in $T_{\mathbf{D}} \cup T_{\mathbf{D}}^- \cup T_{\mathbf{D}}^+ \cup \{t_{max}\}$ are less than $h + 1$ (the minimum number of timestamps required in T^* to have a solution of size h). Note that T_{sup} is nonempty only if $|T_{\mathbf{D}} \cup T_{\mathbf{D}}^- \cup T_{\mathbf{D}}^+ \cup \{t_{max}\}| < h + 1$. Ultimately, T^* is defined as

$$T^* = T_{\mathbf{D}} \cup T_{\mathbf{D}}^+ \cup T_{\mathbf{D}}^- \cup \{t_{max}\} \cup T_{sup}.$$

The proposed more efficient method for TEMPORAL COMMUNITY SEARCH, termed Efficient-temporal-community-search, is summarized in Algorithm 4 and described next. The first five lines of the algorithm are devoted to the identification of T^* . As said above, matrices \mathbf{P} and \mathbf{R} have here reduced dimensionality with respect to Algorithm 3: they are $(|T^*| \times h)$ -dimensional matrices, where $|T^*| \leq |T|$. A mapping function \mathbf{M} is used to assign an index within $[0, |T^*|)$ to every timestamp in $|T^*|$ (Line 6). Such a mapping is needed to have every timestamp in $|T^*|$ logically assigned to a row of matrices \mathbf{P} and \mathbf{R} . The rest of the algorithm resembles Algorithm 3, except for the fact that \mathbf{M} is used every time that a row index has to be mapped to its corresponding timestamp (e.g., during the reconstruction of the solution).

An important point to clarify is that, during the execution of the Efficient-temporal-community-search algorithm, we might need the penalty score $v_{Q,\Delta}^*$ of intervals $\Delta \sqsubseteq T$ corresponding to *non-maximal* (Q -constrained) span-cores. Therefore, the algorithm needs the $v_{Q,\Delta}^*$ score of *all* intervals $\Delta \sqsubseteq T$. To compute these $v_{Q,\Delta}^*$ scores (and, related to this, the set $C_M(Q)$ of Q -constrained maximal span-cores, at Line 1), there are two main options. The first one consists in computing the whole Q -constrained span-core decomposition (as done in Algorithm 3), keep the $v_{Q,\Delta}^*$ scores of all such cores, and eventually compute $C_M(Q)$ by simply filtering out non-maximal span-cores. The second option corresponds instead to compute $C_M(Q)$ directly, without passing through the whole Q -constrained span-core decomposition. This may be carried out by running a simple variant of the algorithm for computing maximal span-cores (Algorithm 2), where containment of query vertices is added as a further constraint. The computation of all the $v_{Q,\Delta}^*$ scores comes for free during the execution of this algorithm for Q -constrained maximal span-cores: these scores can therefore be retained by adding a few straightforward (constant-time) instructions to that algorithm. In our implementation we stick to the latter, as the Maximal-span-cores algorithm has been experimentally recognized as faster than the naïve filtering approach in all tested datasets.

The time complexity of the proposed Efficient-temporal-community-search algorithm is $O(|T^*|^2 \times h + \tau_{msc})$, with τ_{msc} being the time spent in computing the Q -constrained maximal span-cores and the penalty scores $v_{Q,\Delta}^*$. As in practice (attested by our experiments) $|T^*| \ll |T|$, the

proposed Efficient-temporal-community-search algorithm is expected to be much more efficient than its naïve counterpart, i.e., Algorithm 3.

6.4 Minimum Community Search

An instance of TEMPORAL COMMUNITY SEARCH may admit several optimal solutions which might differ either in terms of output intervals $\{\Delta_i\}_{i=1}^h$, or in terms of subgraphs assigned to the various identified intervals. More precisely, the latter refers to the fact that two optimal solutions might find the same segmentation $\{\Delta_i\}_{i=1}^h$ of the input temporal domain, but select different subgraphs S_i for any interval Δ_i . Therefore, if the communities S_i are not chosen carefully, they may result to be excessively large, not really cohesive, and containing redundant/outlying vertices. This is a well-recognized issue of minimum-degree-based community search [79]. At the same time, large communities might include more cohesive and denser subgraphs that still exhibit optimality. Motivated by this, in this subsection we devise a method to refine the communities originally found by our algorithms for TEMPORAL COMMUNITY SEARCH, specifically attempting to minimize their size while preserving optimality. The main idea behind our refinement method is based on the following result:

PROPOSITION 3 (COMMUNITY CONTAINMENT). *Given a temporal graph $G = (V, T, \tau)$, a set $Q \subseteq V$ of query vertices, and a positive integer $h \in \mathbb{N}^+$, let $\{\langle S_i, \Delta_i \rangle\}_{i=1}^h$ be a solution to Problem 3 on input $\langle G, Q, h \rangle$ with S_i corresponding to the (Q, Δ_i) -highest-order-span-core of G , $\forall i \in [1, h]$. For every other solution $\{\langle S'_i, \Delta_i \rangle\}_{i=1}^h$ (referring to the same segmentation $\{\Delta_i\}_{i=1}^h$) to Problem 3 on input $\langle G, Q, h \rangle$ it holds that $S'_i \subseteq S_i, \forall i \in [1, h]$.*

PROOF. Let k_i be the minimum degree of S_i , i.e., $k_i = v_{Q, \Delta_i}^*$ is the order of the (Q, Δ_i) -highest-order-span-core. Assume that there exists a solution S'_i to Problem 4 that is not contained in S_i . This implies that (i) the minimum degree of a vertex of S'_i in Δ_i is k_i , and (ii) the minimum degree of a vertex of $S_i \cup S'_i$ in Δ_i is k_i as well. This violates the maximality condition of the definition of span-core, since, by hypothesis, S_i corresponds to the (Q, Δ_i) -highest-order-span-core of G . \square

The above proposition states that, given a solution $\{\langle S_i, \Delta_i \rangle\}_{i=1}^h$ to the TEMPORAL COMMUNITY SEARCH problem where every S_i corresponds to the (Q, Δ_i) -highest-order-span-core of the input graph, one can focus on the various S_i solely to refine the output communities, as such S_i are guaranteed to contain *all* optimal solutions of the underlying problem instance (while keeping the segmentation $\{\Delta_i\}_{i=1}^h$ fixed). Within this view, we formulate the following optimization problem (which is a variant of Problem 4, with the additional constraint of requiring a smallest-sized solution):

PROBLEM 7. *Given a temporal graph $G = (V, T, \tau)$, a set $Q \subseteq V$ of query vertices, and an interval $\Delta \subseteq T$, let $S^* \subseteq V$ be the subset of vertices containing all the solutions to Problem 4 on input $\langle G, Q, \Delta \rangle$ (according to what stated in Proposition 3). Find*

$$S_{min}^* = \operatorname{argmin}_{\{S \mid Q \subseteq S \subseteq S^*, \min_{u \in S} d_{\Delta}(S, u) \geq \min_{u \in S^*} d_{\Delta}(S^*, u)\}} |S|.$$

THEOREM 3. *Problem 7 in NP-hard.*

PROOF. Consider (the optimization version of) the NP-hard mCST problem introduced by Cui et al. [24]: given a graph $H = (V_H, E_H)$ and a query vertex $q \in V_H$, find a minimum-sized subgraph that contains q , is connected, and maximizes the minimum degree. Given an instance $\langle H, q \rangle$ of the mCST problem, construct an instance $\langle G, Q, \Delta \rangle$ of Problem 7 by defining G as composed by a single temporal snapshot corresponding to graph H , Δ as a singleton interval composed of the single timestamp of G , and setting $Q = \{q\}$. It is straightforward to see that solving Problem 7 on input

ALGORITHM 5: Greedy-minimum-community-search

Input: A temporal graph $G = (V, E, T)$, a set $Q \subseteq V$ of query vertices, an interval $\Delta \sqsubseteq T$, a subset of vertices $S^* \subseteq V$ containing all the solutions to Problem 4 on input $\langle G, Q, \Delta \rangle$.

Output: A subset S_{min}^* of vertices such that $Q \subseteq S_{min}^* \subseteq S^*$ and $\min_{u \in S_{min}^*} d_{\Delta}(S_{min}^*, u) \geq \min_{u \in S^*} d_{\Delta}(S^*, u)$.

```

1  $S_{min}^* \leftarrow \emptyset$ ;  $P \leftarrow \emptyset$ ;  $\mathcal{A} \leftarrow \emptyset$ 
2 add every  $q \in Q$  to  $P$  with priority  $+\infty$ 
3  $k^* \leftarrow \min_{u \in S^*} d_{\Delta}(S^*, u)$ ;  $k_{min}^* \leftarrow 0$ 
4 while  $k_{min}^* < k^*$  or  $Q \not\subseteq S_{min}^*$  do
5   dequeue  $u$  from  $P$ 
6    $S_{min}^* \leftarrow S_{min}^* \cup \{u\}$ 
7   forall  $v \in \text{neigh}_{\Delta}(S^*, u) \setminus S_{min}^* \setminus P$  do
8      $\mathcal{A}[v] \leftarrow \text{score}(v)$ 
9     add  $v$  to  $P$  with priority  $\mathcal{A}[v]$ 
10  forall  $v \in \text{neigh}_{\Delta}(S_{min}^*, u)$  do
11    if  $d_{\Delta}(S_{min}^*, v) = k^*$  then
12      forall  $w \in \text{neigh}_{\Delta}(P, v)$  do
13         $\mathcal{A}[w] \leftarrow \mathcal{A}[w] - 1$ 
14   $k_{min}^* \leftarrow \min_{v \in S_{min}^*} d_{\Delta}(S_{min}^*, v)$ 

```

$\langle G, Q, \Delta \rangle$ is equivalent to solving mCST on input $\langle H, q \rangle$, as the constraint about connectedness is automatically satisfied in Problem 7 for the special case of a single query vertex. \square

As Problem 7 is NP-hard, we devise a heuristic that is inspired to the greedy one proposed for the MINIMUM COMMUNITY SEARCH problem in [7]. The proposed heuristic is outlined in Algorithm 5 and described next. In the pseudocode and in the following we denote as k^* and k_{min}^* the minimum degree of S^* and S_{min}^* , respectively, and as $\text{neigh}_{\Delta}(S, u)$ the neighbors of a vertex $u \in V$ in the subgraph induced by $S \subseteq V$ and $\Delta \sqsubseteq T$. Algorithm 5 iteratively adds vertices to the solution S_{min}^* according to a priority queue P . Priorities of vertices in P are defined based on a score that measures how promising a vertex is for making the current solution S_{min}^* reach the optimal minimum degree. Specifically, the score of a vertex $u \in S^*$ is defined as:

$$\text{score}(u) = \text{score}^+(u) - \text{score}^-(u),$$

where

$$\text{score}^+(u) = |\{v \in \text{neigh}_{\Delta}(S_{min}^*, u) \mid d_{\Delta}(S_{min}^*, v) < k^*\}|;$$

$$\text{score}^-(u) = \max\{0, k^* - d_{\Delta}(S_{min}^*, u)\}.$$

$\text{score}^+(u)$ is the gain effect of adding u to S_{min}^* , while $\text{score}^-(u)$ is the penalty effect. In particular, $\text{score}^+(u)$ counts the number of neighbors of u in S_{min}^* that would benefit from the inclusion of u to S_{min}^* , i.e., that have degree less than k^* . On the other hand, $\text{score}^-(u)$ represents the number of neighbors of u still required in S_{min}^* so that u has degree at least k^* . The algorithm starts by adding the query vertices to the queue P with priority $+\infty$, in order to ensure that they will be selected at the very beginning. At each iteration of the main cycle of the algorithm (starting at Line 4), the vertex u exhibiting the highest priority is dequeued from P and is added to the solution S_{min}^* . As a consequence, a couple of updates are performed. First, u 's neighbors not in the priority queue P are added to it (Lines 8–9). Note that this is the only step of the algorithm where the score of a vertex

Table 1. Temporal Graphs Used in the Experiments

dataset	$ V $	$ E $	$ T $	window size	domain
HighSchool	327	47k	1, 212	5 mins	face-to-face
PrimarySchool	242	55k	390	5 mins	face-to-face
HongKong	806	2M	2, 976	5 mins	face-to-face
ProsperLoans	89k	3M	307	7 days	economic
Last.fm	992	4M	77	21 days	co-listening
WikiTalk	2M	10M	192	28 days	communication
DBLP	1M	11M	80	366 days	co-authorship
StackOverflow	2M	16M	51	56 days	question answering
Wikipedia	343k	18M	101	56 days	co-editing
Amazon	2M	22M	115	28 days	co-rating
Epinions	120k	33M	25	21 days	co-rating

is computed from scratch and stored in \mathcal{A} , a map that keeps the scores of all vertices in P up-to-date during the whole execution of the algorithm. The second update consists in recomputing the score of every v 's neighbor w in the queue, if a vertex $v \in S_{min}^*$ has reached the desired minimum degree k^* after the addition of u .

7 EXPERIMENTS

In this section, we present an experimental evaluation to empirically assess the performance of all the proposed methods. Specifically, we focus on whole span-core decomposition (Section 7.1), maximal span-cores (Section 7.2), characterization of the extracted span-cores (Section 7.3), and temporal community search (Section 7.5).

Datasets. We use 11 real-world datasets recording timestamped interactions between entities. For each dataset, we select a window size to define a discrete time domain, composed of contiguous timestamps of the same duration, and build the corresponding temporal graph. If multiple interactions occur between two entities during the same discrete timestamp, they are counted as one. The characteristics of the resulting temporal graphs, along with the selected window sizes, are reported in Table 1.

The three smallest datasets were gathered by using wearable proximity sensors in schools, with a temporal resolution of 20 seconds. PrimarySchool³ contains the contact events between 242 volunteers (232 children and 10 teachers) in a primary school in Lyon, France, during 2 days [81]. HighSchool³ describes the close-range proximity interactions between students and teachers (327 individuals overall) of nine classes during 5 days in a high school in Marseilles, France [66]. HongKong reports the same kind of interactions for a primary school in Hong Kong, whose population consists of 709 children and 65 teachers divided into 30 classes, for 11 consecutive days [76].

ProsperLoans⁴ represents the network of loans between the users of Prosper, a marketplace of loans between privates. Last.fm⁴ records the co-listening activity of the Last.fm streaming platform: an edge exists between two users if they listened to songs of the same band within the same discrete timestamp. WikiTalk⁴ is the communication network of the English Wikipedia. DBLP⁴ is the co-authorship network of the authors of scientific papers from the DBLP computer

³sociopatterns.org.

⁴konect.cc.

science bibliography. StackOverflow⁵ includes the answer-to-question interactions on the stack exchange of the stackoverflow.com website. Wikipedia⁴ connects users of the Italian Wikipedia that co-edited a page during the same discrete timestamp. Finally, for both Amazon⁴ and Epinions⁴, vertices are users and edges represent the rating of at least one common item within the same discrete timestamp.

Implementation. All methods are implemented in Python (v. 2.7.16) and compiled by Cython. All the experiments were run on a machine equipped with Intel Xeon CPU at 2.1 GHz. The experiments reported in Sections 7.1 and 7.2 used 64 GB RAM, while the ones in Section 7.5 used 32 GB RAM.

7.1 Span-core Decomposition

We compare the two methods to compute a complete decomposition described in Section 4, i.e., the baseline Naïve-span-cores and the proposed Span-cores, in terms of execution time, memory, and total number of vertices input to the core-decomposition subroutine. We report these measures, together with the number of span-cores and maximal span-cores of each dataset, in Table 2.

In terms of execution time, Span-cores considerably outperforms Naïve-span-cores in all datasets, achieving a speed-up from 2.1 up to two orders of magnitude. The speed-up is explained by the number of vertices processed by the core-decomposition subroutine, which is the most time-consuming step of the algorithms albeit linear in the size of the input subgraph. The difference of this quantity between Span-cores and Naïve-span-cores reaches over an order of magnitude in the WikiTalk, Wikipedia, and Epinions dataset, confirming the effectiveness of the “horizontal containment” relationships. The memory required by the two procedures is comparable in all cases since the largest structures needed in memory are the temporal graph itself and the set C of all span-cores.

7.2 Maximal Span-cores

We compare our Maximal-span-cores algorithm to the naïve approach, described at the beginning of Section 5, based on running the Span-cores algorithm and filtering out the non-maximal span-cores, which we refer to as Naïve-maximal-span-cores. The results are again reported in Table 2.

Naïve-maximal-span-cores behaves very similarly to Span-cores: they only differ for the filtering mechanism which requires a few additional seconds in most cases. Maximal-span-cores is much faster than Naïve-maximal-span-cores for all datasets, with a speed-up from 1.3 for the Epinions dataset to one order of magnitude for the HongKong dataset. Except for the school datasets and Last.fm, the difference in terms of number of processed vertices is between one and three orders of magnitude, attesting the advantages of the top-down strategy of Maximal-span-cores, which avoids the visit of portions of the span-coresearch space and handles the overhead of reconstructing graphs, i.e., $(V_{lb}, E_{\Delta}[V_{lb}])$, efficiently. Finally, the memory requirements of the two methods are comparable for all datasets.

7.3 Span-cores Characterization

We compare and characterize all span-cores against maximal span-cores. At first, Table 2 shows that span-cores are at least one order of magnitude more numerous than maximal span-cores for all datasets, with the maximum difference of three orders of magnitude for the HongKong dataset.

In Figure 5, we show the number (top) and the average size (bottom) of span-cores and maximal span-cores as a function of the order k for the DBLP and Epinions datasets. For both datasets, the number of maximal span-cores is at least one order of magnitude lower than the total number of

⁵snap.stanford.edu.

Table 2. Evaluation of the Proposed Algorithms: Number of Output Span-cores, Running Time, Memory, and Number of Processed Vertices

dataset	method	# output span-cores	running time (s)	memory (GB)	# processed vertices
HighSchool	Naïve-span-cores	12,320	18	0.1	3M
	Span-cores		1	0.1	581k
	Naïve-maximal-span-cores	450	1	0.1	581k
	Maximal-span-cores		0.3	0.1	181k
PrimarySchool	Naïve-span-cores	4, 703	4	0.1	818k
	Span-cores		0.6	0.1	174k
	Naïve-maximal-span-cores	409	0.6	0.1	174k
	Maximal-span-cores		0.1	0.1	63k
HongKong	Naïve-span-cores	2, 367, 743	85, 180	1	819M
	Span-cores		18, 389	0.8	216M
	Naïve-maximal-span-cores	1, 807	18, 641	0.8	216M
	Maximal-span-cores		339	0.5	212M
ProsperLoans	Naïve-span-cores	4, 273	101	2	55M
	Span-cores		46	2	27M
	Naïve-maximal-span-cores	293	48	2	27M
	Maximal-span-cores		8	2	980k
Last.fm	Naïve-span-cores	126, 819	707	0.5	2M
	Span-cores		199	0.5	531k
	Naïve-maximal-span-cores	1, 670	202	0.5	531k
	Maximal-span-cores		57	0.5	271k
WikiTalk	Naïve-span-cores	19, 693	322, 302	36	25B
	Span-cores		1, 084	36	555M
	Naïve-maximal-span-cores	632	1, 194	36	555M
	Maximal-span-cores		126	35	2M
DBLP	Naïve-span-cores	6, 135	10, 506	11	1B
	Span-cores		278	11	150M
	Naïve-maximal-span-cores	268	292	11	150M
	Maximal-span-cores		116	11	620k
StackOverflow	Naïve-span-cores	1, 238	5, 360	10	1B
	Span-cores		245	10	127M
	Naïve-maximal-span-cores	129	245	10	127M
	Maximal-span-cores		128	10	3M
Wikipedia	Naïve-span-cores	125 191	17, 155	4	1B
	Span-cores		522	4	35M
	Naïve-maximal-span-cores	2, 147	537	4	35M
	Maximal-span-cores		201	4	320k
Amazon	Naïve-span-cores	29, 318	10, 415	18	2B
	Span-cores		409	18	247M
	Naïve-maximal-span-cores	303	580	18	247M
	Maximal-span-cores		123	18	688k
Epinions	Naïve-span-cores	63, 111	699	4	39M
	Span-cores		186	4	3M
	Naïve-maximal-span-cores	320	201	4	3M
	Maximal-span-cores		154	5	129k

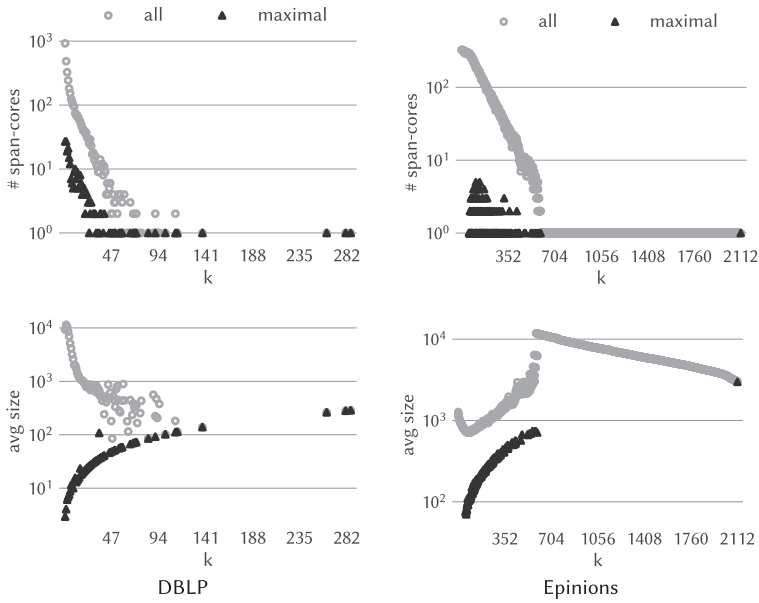


Fig. 5. Top plots: number of span-cores and maximal span-cores as a function of the order k . Bottom plots: average size of all span-cores and maximal span-cores as a function of the order k .

span-cores up to a quarter of the k domain, where the span-cores are more numerous. Instead, in the rest of the domain, they mostly coincide due to the maximality condition over $|\Delta|$. The average size is also smaller for maximal span-cores, difference that wears thin when the gap between the numbers of span-cores and maximal span-cores starts decreasing since, for high values of k , most (or all) span-cores are maximal.

Figure 6 shows a different picture when numbers and average sizes of span-cores are shown as a function of the size of the span $|\Delta|$. For both datasets, the number of span-cores and maximal span-cores is decreasing—which is expected since the number of intervals decreases when $|\Delta|$ increases—with a constant gap close to one and two orders of magnitude, respectively. On the other hand, the behavior of the average size is quite different between the two datasets. For low values of $|\Delta|$, the average size of span-cores of the DBLP dataset is much higher than the average size of maximal span-cores, then the difference decreases and vanishes at the end of domain where a maximal span-core of $|\Delta| = 37$ dominates all other span-cores with $|\Delta| \geq 20$. Instead, for the Epinions dataset, the average size of all span-cores and of maximal span-cores follow the same behavior, with a difference of less than an order of magnitude, because the maximality condition over k excludes the largest span-cores from the set of maximal span-cores.

Figure 7 yields some additional insights by showing the whole distribution of sizes of the span-cores: these distributions are very skewed and span several orders of magnitude. The figure also shows the size distributions of span-cores with a given order or duration: all are broad, becoming narrower as the order or duration increase. We have also considered randomized versions of the datasets, in which edges are reshuffled at random at each timestamp. In this case as well, the distributions of the sizes of the span-cores are found to be broad, showing that the heterogeneity in span-coresizes can also be obtained in largely random data. However, the cohesive temporal structures are destroyed in the reshuffled data (see also Section 8.1): all the span-cores have then very small order and span.

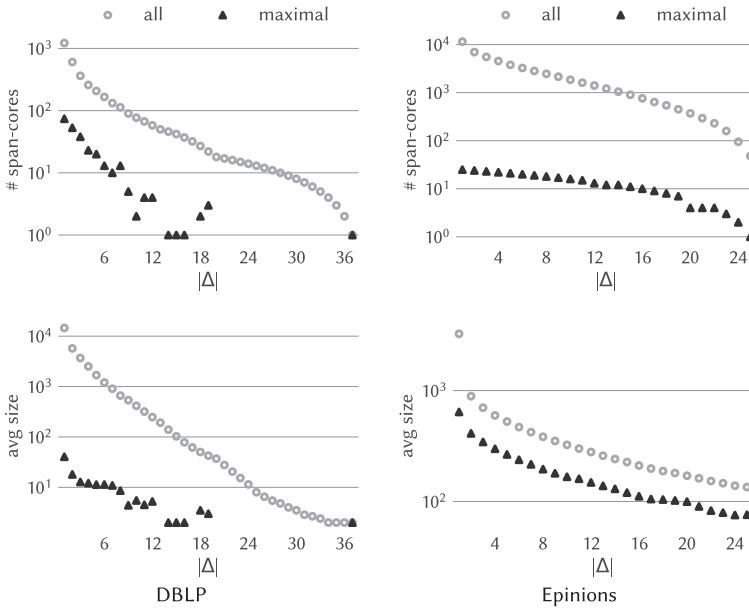


Fig. 6. Top plots: number of span-cores and maximal span-cores as a function of the size of the temporal span $|\Delta|$. Bottom plots: average size of all span-cores and maximal span-cores as a function of the size of the temporal span $|\Delta|$.

7.4 Scalability Analysis

Here we evaluate the scalability of Algorithms 1 and 2. To this aim, we consider temporal subgraphs derived from the Amazon dataset with increasing number of vertices and timestamps. To obtain temporal subgraphs with varying number of timestamps, we simply consider the temporal graphs associated with the first 20, 40, 60, 80, and 100 timestamps, while considering all the vertices and edges existing in these time-frames. For what concerns the temporal subgraphs with controlled number of vertices, we consider the entire temporal domain and sample sets of vertices of size $0.2M$, $0.4M$, $0.6M$, $0.8M$, $1M$, $1.2M$, $1.4M$, $1.6M$, and $1.8M$. Vertices are sampled according to the following simple procedure:

- Select a vertex uniformly at random from the whole V and add such a vertex to the set of sampled nodes S
- Starting from the first timestamp of the temporal domain T , iteratively
 - For each vertex v in S , select a neighbor of v uniformly at random and add it to S
 - Move to the next timestamp
- If the last timestamp is reached, restart from the first one

The results of this scalability experiment are reported in Figure 8. It can be observed that the trends shown in the figure comply with the time complexities discussed in Sections 4 and -5: running times are quadratic in the number of timestamps and linear in the number of vertices/edges. Another relevant consideration is that the difference between the algorithm to compute maximal span-cores and the one for computing all the span-cores gets larger as the number of timestamps or vertices increase. This further attests the usefulness of introducing an algorithm that is specifically devoted to maximal-span-core computation.

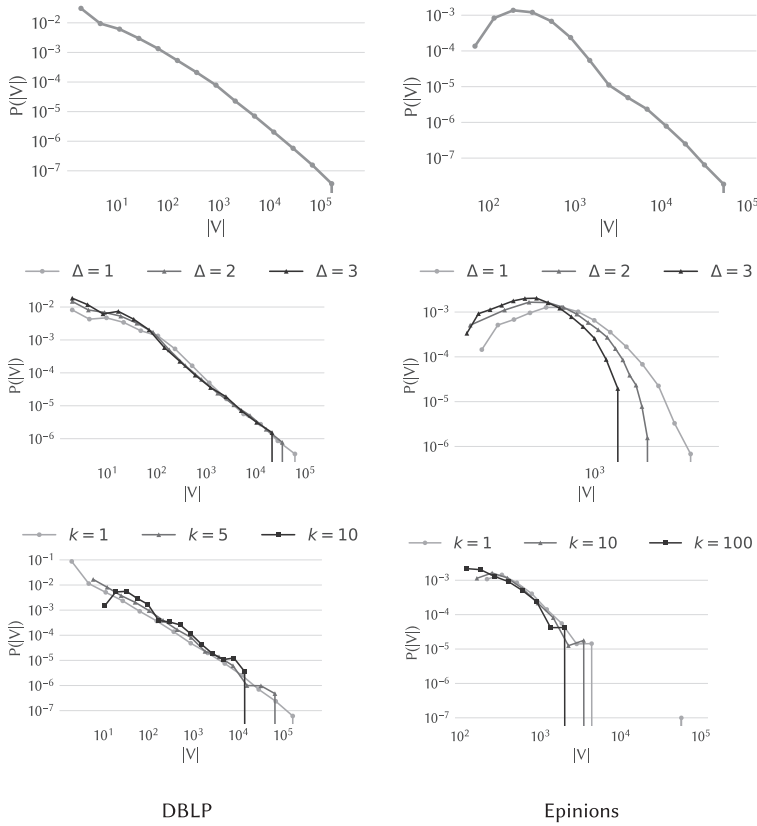


Fig. 7. Distribution of sizes of the span-cores. Top plots: overall distributions. Middle plots: distributions of the sizes of span-cores with fixed temporal span. Bottom plots: distributions of the sizes of span-cores with fixed order.

7.5 Temporal Community Search

In this subsection, we assess the performance of the proposed algorithms for temporal community search (presented in Sections 6.2 and 6.3), as well as the greedy procedure for reducing the size of the output communities (presented in Section 6.4). In the remainder of this subsection, we refer to our basic algorithm (i.e., Algorithm 3, which precomputes the penalty scores via span-core decomposition) as SC-TCS, and to our more efficient algorithm (i.e., Algorithm 4, which exploits maximal span-cores to reduce the number of timestamps to be considered) as MSC-TCS. We also involve in the comparison a naïve version of Algorithm 3, where the penalty scores of the various intervals are computed from scratch during the execution of the algorithm, instead of precomputing them all via span-core decomposition. We refer to such a naïve method as Naïve-TCS.

The experimental setting we consider here is as follows. We vary the number $|Q|$ of query vertices from 1 to 3. In particular, when $|Q| = 1$, we sample the single query vertex uniformly at random from the whole vertex set V . Instead, for $|Q| > 1$, we employ a more sophisticated sampling strategy that aims at finding meaningful query-vertex sets, i.e., vertices interacting with each other during the temporal observations, and, at the same time, independent from the specific form of the resulting span-core decomposition. Specifically, the sampling strategy we use is based on an adaptation of random walk to the temporal settings:




Fig. 8. Scalability analysis: running time of the algorithms for extracting all the span-cores (SC, Algorithm 1) and only the maximal span-cores (MSC, Algorithm 2) as a function of the number timestamps and vertices, on the Amazon dataset.

- Select a vertex uniformly at random from the whole V and add such a vertex to the set Q_{visited} of visited vertices
- Starting from the first timestamp of the temporal domain T , iteratively:
 - With probability p , move the random walker to a neighbor of the current vertex and add the neighbor to Q_{visited} . If the current vertex has no neighbors in a given timestamp, the random walker jumps to the first next timestamp in which that vertex has at least one neighbor
 - With probability $1 - p$, keep the random walker at the current vertex, but go to the next timestamp
- Restart if the last timestamp of T is reached
- Stop when $|Q_{\text{visited}}|$ reaches a proper (user-defined) size ν
- Sample $|Q|$ query vertices from Q_{visited} with probability proportional to the frequency of the visits during the random walk

In our experiments, we set $p = 0.8$ and $\nu = 3|Q|$. As far as the number h of output communities, we consider the range $h \in [10, 20, 30, 40, 50, 60]$ on all datasets, with the exception of StackOverflow, for which we discard $h = 60$, and Epinions, for which we consider $h \in [4, 8, 12, 16, 20, 24]$. For every parameter configuration, we perform five runs of every algorithm (in every run we sample a different query-vertex set). Note that we were not able to run the algorithms for temporal community search on the WikiTalk dataset due to memory constraints.

Running time. In Figure 9, we show the running time of the proposed algorithms as a function of the number h of output communities, for the HighSchool, DBLP, Wikipedia, and Amazon datasets. The first general observation we make is that the running time of all algorithms increases as h gets higher. This is in accordance with the time-complexity analysis reported in Section 6. Also, running times are independent of the selected query-vertex set Q . Looking at the individual performance, we notice that, as expected, the Naïve-TCS method has severe limitations in terms of efficiency: it takes hours to run on the HighSchool and Wikipedia datasets, while it is not able to terminate in less than 10 days on the remaining datasets. SC-TCS and MSC-TCS are much faster than Naïve-TCS, achieving a speedup of up to more than four orders of magnitude. MSC-TCS is in most cases faster than SC-TCS, with speedup up to one order of magnitude (on HighSchool, for $h = 60$). This confirms that the exploitation of the maximal span-cores is effective in both shortening the precomputation time and reducing the temporal domain considered in the DP step. The only exception is the Wikipedia dataset. To dive deeper into the motivations of this exception, we report in Figure 10 the split of the average running time of SC-TCS and MSC-TCS into the

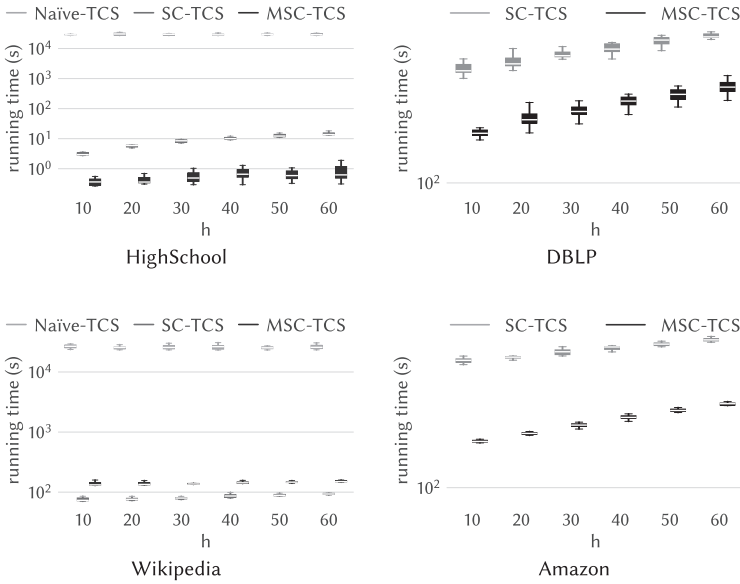


Fig. 9. Running time of the algorithms for TEMPORAL COMMUNITY SEARCH, as a function of the number h of output communities. Each boxplot corresponds to 15 data points.




Fig. 10. Split of the average running time of the SC-TCS and MSC-TCS algorithms into DP and precomputation, for the Wikipedia and Last.fm datasets.

time spent in the DP step (which also includes the identification of the reduced temporal domain T^* for MSC-TCS), and the precomputation time (i.e., the time required for computing all penalty scores via span-core decomposition or maximal span-cores). Interestingly, what affects the most the running time is the precomputation of the scores. Apparently, the Q -constrained version of Span-cores is more efficient than Maximal-span-cores in some datasets, which we believe is due to the structure of the search space. On the other hand, these results confirm that the reduction of the temporal domain considered by the DP step is actually effective since the DP running time of MSC-TCS is always less than (or equal to) the DP running time of SC-TCS.

Greedy-minimum-community-search. Here we evaluate the performance of the proposed Greedy-minimum-community-search algorithm (Algorithm 5) for reducing the size of the output communities. We recall that the proposed algorithms for TEMPORAL COMMUNITY SEARCH (evaluated above) output communities corresponding to the (Q, Δ_i) -highest-order-span-cores for all $\{\Delta_i\}_{i=1}^h$ temporal intervals identified. The Greedy-minimum-community-search algorithm takes every (Q, Δ_i) -highest-order-span-core and attempts to reduce its size, while preserving optimality. Thus, the ultimate goal of the evaluation presented next is to show how well

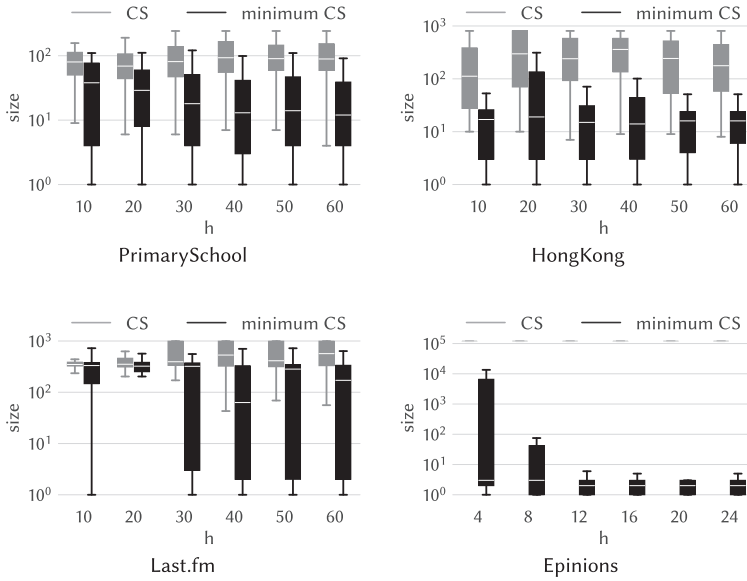


Fig. 11. Comparison of the size of the communities in the solutions to TEMPORAL COMMUNITY SEARCH: original output of the algorithms for TEMPORAL COMMUNITY SEARCH (CS) and after running the Greedy-minimum-community-search algorithm on top of them (minimum CS). Each boxplot corresponds to 15 data points.

Greedy-minimum-community-search is able to reduce the size of the original span-cores, and what is its overhead in terms of running time.

Figure 11 compares the size of the starting (Q, Δ_i) -highest-order-span-cores and the size of the corresponding reduced community yielded by the Greedy-minimum-community-search algorithm, for the PrimarySchool, HongKong, Last.fm, and Epinions datasets. It can be easily observed that, as a general trend, the reduced communities are much smaller than the original ones, in all datasets, up to four orders of magnitude. The results on the Epinions dataset are a bit different than the other three datasets. In fact, on that dataset, the original communities (CS) always include the whole 120k vertices of the graph, while the communities found by Greedy-minimum-community-search (minimum CS) have median size smaller than 10, and, in many cases, they correspond to communities composed of the query vertices only. This means that, on the Epinions dataset, for our tested queries, the algorithms for TEMPORAL COMMUNITY SEARCH do not extract communities that are really cohesive around the query vertices. This way, the benefits of exploiting an a-posteriori community-size-reduction step are less evident. Also, we do not notice any evident pattern as a function of h , for any dataset.

In Table 3, we report the average running time of an execution of Greedy-minimum-community-search, for all datasets. Note that this is the average time required to process one of the h communities in a solution to TEMPORAL COMMUNITY SEARCH. Greedy-minimum-community-search runs in 8 seconds or less in all tested datasets. Therefore, the additional running time required by the algorithm is rather negligible.

To summarize, Greedy-minimum-community-search is empirically recognized as a powerful post-processing method for improving the quality of the solutions to TEMPORAL COMMUNITY SEARCH: it finds much smaller communities at a very small additional computational cost.

Table 3. Average Running Time of an Execution of the Greedy-minimum-community-search Algorithm

	HighSchool	PrimarySchool	HongKong	ProsperLoans	Last.fm
running time (s)	0.003	0.001	0.02	0.3	0.06
	DBLP	StackOverflow	Wikipedia	Amazon	Epinions
running time (s)	7	8	1	7	6

8 APPLICATIONS

In this section, we illustrate applications of (maximal) span-cores in the analysis of face-to-face interaction networks, and how the methods for TEMPORAL COMMUNITY SEARCH can be profitably exploited in a task of graph classification. For these applications we use the three networks gathered in schools, i.e., PrimarySchool, HighSchool, and HongKong, which are described above, at the beginning of Section 7. We use a window size of 5 minutes and, in the analysis, we discard span-cores of $|\Delta| = 1$, i.e., having span of 5 minutes, since they represent short interactions, not significant for our purposes. In the following we show (i) three types of interesting temporal patterns (Section 8.1), i.e., social activities of groups of students within a school day, mixing of gender and class, and length of social interactions in groups; (ii) a procedure to detect anomalous contacts and intervals that exploits maximal span-cores (Section 8.2); and, (iii) an approach to graph classification based on temporal community search (Section 8.3).

8.1 Temporal Patterns

Temporal activity. We first show how span-cores afford a simple temporal analysis of social activities of groups of people within a school day. The left side of Figure 12 reports colormaps of the order k of the span-cores as a function of their starting time t_s (x -axis) and of the size of their temporal span $|\Delta|$ (y -axis), for a school day of the PrimarySchool and HighSchool datasets. Darker gray indicates span-cores of high order and slots located in the upper part of the plots refer to span-cores of long span. It is important to notice that the linear decay in span duration is naturally due to the definition of span-core and to the shifting of the starting time t_s ; therefore, it is not a distinguishing feature of the activity patterns found in the analyzed data. In both datasets, fluctuations of k and $|\Delta|$ are observed along the day, which can be related to school events. Around 10 a.m., the size of the span $|\Delta|$ reaches a local maximum in correspondence to the morning break, which means that students establish long-lasting interactions that hold beyond the break itself. Moreover, when classes gather for the lunch break, the order k reaches its maximum value since students tend to form larger and more cohesive groups.

In order to verify that these results are not trivially derived from the general temporal activity, as simply given by the number of interactions in each timestamp, we compare our findings to a null model. At each timestamp of the temporal graphs, we reshuffle the edges by the Maslov–Sneppen algorithm [65] which consists in repeating the following operations up to when all edges have been processed: select at random two edges with no common vertices, e.g., (u, v) and (w, z) , and transform them into (u, z) and (w, v) , if neither (u, z) and (w, v) existed in the original timestamp. This reshuffling preserves the degree of each vertex in each timestamp and the global activity (i.e., the number of contacts per timestamp), but destroys correlations between edges of successive timestamps. In the right side of Figure 12, we show the results of the temporal analysis described above for the reshuffled datasets. In both, the values of $|\Delta|$ and k reached are much smaller than in the original datasets. The size of the span $|\Delta|$ is always shorter than 20 minutes, while in the

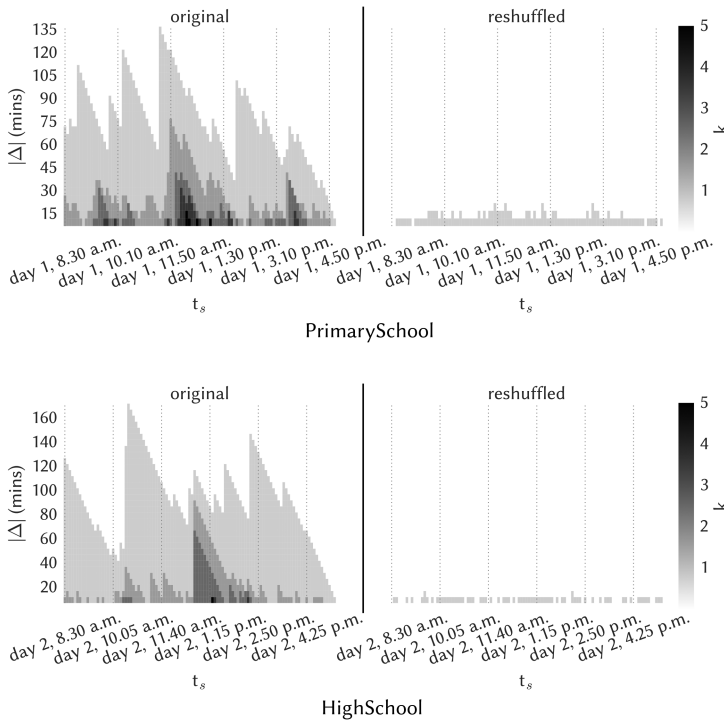


Fig. 12. Temporal activity of a school day of the PrimarySchool and HighSchool datasets: the x -axis reports the hour of the day at which the span of a span-core starts, the y -axis specifies the size of the span (in minutes), and the color scale shows the order k . At a glance, it can be observed that the temporal structure of the span-core decomposition detects time-evolving cohesive structures in the original datasets (left plots) that completely disappear in the reshuffled datasets (right plots).

original datasets it is much longer, up to 170 minutes, and the order k is always equal to 1, compared to the original maximum of 5. The time-evolving cohesive structures detected by the temporal core decomposition in the original datasets are completely lost on reshuffling, since only span-cores of short span and low coreness are observed in the latter case. This shows that the temporal structure exposed by the span-core decomposition is not simply a consequence of temporal patterns of global activity but that span-cores represent a concrete method to detect complex cohesive structures and their temporal evolution.

Mixing patterns. We now show an analysis of mixing patterns of students with respect to gender and class. Such vertex attributes are indeed available for the individuals of the PrimarySchool dataset. We define as *gender purity* of a span-core the fraction of individuals of the most represented gender within the span-core. *Class purity* is analogously defined. The left plot of Figure 13 reports the temporal evolution of the average gender and class purity of the maximal span-cores spanning each timestamp, during the first school day of the PrimarySchool dataset. During lessons, when students are in their own classes, class purity has naturally very high values, very close to 1. Gender purity is instead rather low. On the other hand, when students are gathered together, during the morning break at 10 a.m. and the lunch break between 12 a.m. and 2 p.m., the situation is overturned: gender purity reaches large values while class purity drastically decreases. This shows that primary school students group with individuals of the same class, disregarding the gender, only when they are forced by the schedule of the lessons, but prefer on average to form cohesive




Fig. 13. Temporal evolution (time on the x -axis) of average gender purity and average class purity (y -axis) of the maximal span-cores of the PrimarySchool dataset. Original data on the left, reshuffled data on the right.




Fig. 14. Distribution of the size of the span $|\Delta|$ of the maximal span-cores. The x -axis reports the size of the span (in minutes), while the y -axis the percentage of maximal span-cores having a given size of the span.

groups with students of the same gender during breaks. This is in agreement and complements a previous study of the same dataset focusing on single interactions in the static aggregated network [80].

The right plot of Figure 13 shows the temporal evolution of the average gender and class purity for a null model in which gender and class are randomly reshuffled among individuals. The two curves are more flat and the anti-correlation between them completely vanishes. This testifies that the results on the original dataset are not simply due to the relative abundance of individuals of each type interacting at each time, but reflect genuine mixing patterns and their temporal evolution.

Interaction length. Finally, we analyze the duration of interactions of social groups in schools by studying the distribution of the size of the span of the maximal span-cores of the three datasets (Figure 14). All distributions are extremely skewed with broad tails: most maximal span-cores have duration less than 1 hour, but durations much larger than the average can also be observed. Interestingly, the three datasets at hand all exhibit the same functional shape, confirming a robust statistical behavior. We also note that similar robust broad distributions have been observed for simpler characteristics of human interactions such as the statistics of contact durations [66, 81]. Outliers appear also at very large durations, especially for the HongKong dataset that has maximal span-cores lasting up to 83 hours. Group interactions of such long span are clearly abnormal and represent outliers in the distributions. We will show, in the following of this section, how to exploit such outliers to detect both irregular interactions and anomalous temporal intervals.

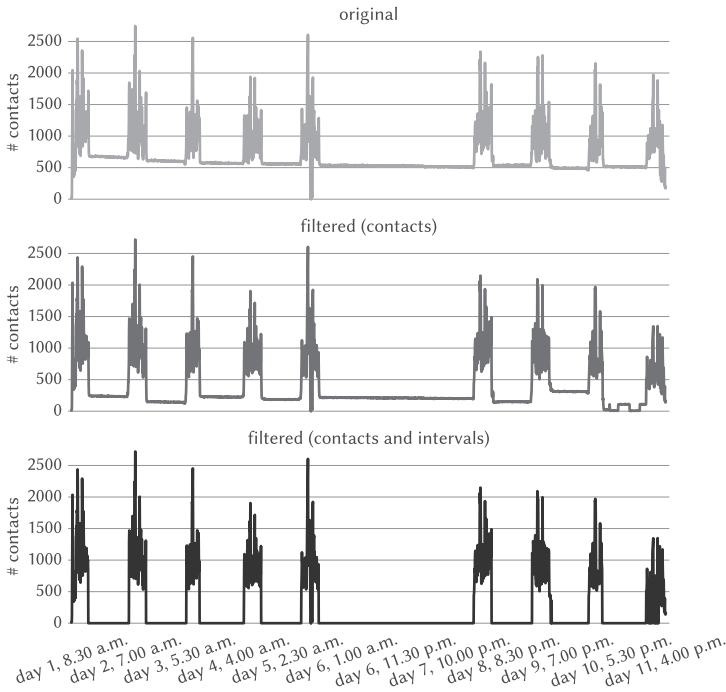


Fig. 15. HongKong dataset: number of edges per timestamp in the original data (top), after filtering anomalous edges (middle), and after filtering anomalous edges and intervals (bottom). Days 6 and 7 are weekend.

8.2 Anomaly Detection

The identification of anomalous behaviors in temporal networks has been the focus of several studies in the last few years [68, 76]. Based on the above findings, we devise a simple procedure to detect anomalous edges and intervals of the HongKong dataset that exploits maximal span-cores. The topmost plot of Figure 15 reports the number of edges for each timestamp of the original HongKong dataset. It is easy to notice that there is a lot of constant anomalous activity between school days and during the weekend, i.e., days six and seven: unexpectedly, the number of interactions per timestamp does not drop to zero. This happened in fact because proximity sensors were left in each class and close to each other, at the end of the lessons. In order to automatically detect these steady activity patterns that do not correspond to any genuine social dynamics, we apply the following procedure: (i) find a set of anomalously long temporal intervals supporting maximal span-cores, (ii) identify anomalous vertices, and (iii) filter out anomalous edges.

The first step of this procedure requires to find the set of temporal intervals $\mathcal{I} = \{\Delta \subseteq T \mid C_{k,\Delta} \in C_M \wedge |\Delta| > tr\}$ that are the span of a maximal span-core $C_{k,\Delta}$ with size longer than a certain threshold tr . Then, for each timestamp $t \in T$, select as anomalous all those vertices that appear in the span-cores $\{C_{1,\Delta} \mid \Delta \in \mathcal{I} \wedge t \in \Delta\}$, i.e., the span-cores of $k = 1$ whose span is in \mathcal{I} and contains t . Finally, at each timestamp $t \in T$, remove edges that are incident to at least a vertex that has been marked as anomalous at time t . Consistently with the distribution of the span durations of the maximal span-cores, we select the threshold $tr = 22$ (110 minutes). The results of this filtering procedure are shown in the middle plot of Figure 15. The number of edges during school days remains approximately unchanged, while the activity noticeably decreases in-between. Identifying as positives the spurious interactions occurring when the school is closed and as negatives the

genuine interactions observed when the school is open, this approach achieves a precision of 0.91 and a recall of 0.64.

We can refine this anomaly detection process by identifying, in addition to anomalous edges, also anomalous temporal intervals. We define a timestamp $t \in T$ as anomalous if the ratio between the number of original edges (top plot of Figure 15) and the number of filtered edges (middle plot of Figure 15) exceeds a given threshold. We apply this further filtering to the HongKong dataset with a threshold of 1.5 and report the results in the bottommost plot of Figure 15. The number of edges when the school is closed drops to zero, while the activity during school days is not modified, except for the last one, which is affected by the proximity to the end of the time domain. The overall procedure yields a slightly higher value of precision, 0.93, and substantially improves the recall to 0.99.

8.3 Graph Embedding and Vertex Classification

In this subsection, we show how TEMPORAL COMMUNITY SEARCH can be profitably exploited for classifying the vertices of a temporal graph. Specifically, the classification framework we set up is based on the paradigm of *graph embedding*, which has attracted a great deal of attention in the last few years, and whose goal is to assign to every vertex of a graph a numerical vector (i.e., an *embedding*) such that structurally similar vertices are represented by similar vectors, and vice versa [30, 45, 46]. Here, our framework simply consists in learning suitable embeddings for the vertices of the input graph, and then give them as input to some (well-established) classifier to ultimately accomplish the desired classification task. Thus, the main goal is to learn embeddings that are well-representative of the relationships among vertices, so as to help the classifier perform accurately. As our main result here, we show how an embedding strategy based on a simple exploitation of the output of TEMPORAL COMMUNITY SEARCH achieves results comparable to well-established vertex-embedding methods such as DeepWalk [70], LINE [83], and node2vec [46].

Method. For every vertex of the input temporal graph, we build an embedding as an h -dimensional vector conveying the information provided by a solution to the TEMPORAL COMMUNITY SEARCH problem on the same graph. Specifically, consider a vertex $u \in V$ and a solution $\{\langle S_i, \Delta_i \rangle\}_{i=1}^h$ to TEMPORAL COMMUNITY SEARCH on query-vertex set $Q = \{u\}$. We define u 's embedding as

$$\mathbf{X}_u = [v_{Q, \Delta_1}^*, v_{Q, \Delta_2}^*, \dots, v_{Q, \Delta_h}^*],$$

which corresponds to the temporally-ordered sequence of minimum degrees of the h communities identified by the temporal-community-search solution. Below we show that this simple approach is sufficient to achieve interesting experimental results. Clearly, more sophisticated methods are possible, e.g., by simultaneously exploiting information from the S_i communities. However, our main goal here is to give an idea of how the TEMPORAL COMMUNITY SEARCH problem can be successfully leveraged in a relevant application scenario, rather than devise the best temporal-community-search-based graph-embedding method.

Evaluation. We assess the performance of our method on the PrimarySchool and HighSchool datasets. In these datasets vertices correspond to students, and vertex labels (to be predicted) are the classes that every student belongs to. We involve in the comparison the following state-of-the-art vertex-embedding methods:

- DeepWalk [70], a method that preserves the proximity between vertices by running a set of random walks and maximizing the sum of the log-likelihood of a set of vertices for each walk.




Fig. 16. Vertex classification: Macro F_1 -score of the proposed temporal-community-search-based graph-embedding method TCS and the competing methods, with varying the dimensionality h of the output embeddings, on the PrimarySchool and HighSchool datasets.

- LINE [83], which optimizes a suitable objective function preserving both first-order (one-hop) and second-order (two-hop) proximities. Neighborhoods are not explored via random walk, but in a breadth-first fashion.
- node2vec [46], which is based on the same idea underlying DeepWalk, but allowing more flexibility on how random walks explore and leave the neighborhood of the current vertex.

These three methods consider non-temporal graphs. Therefore, we feed them with aggregated graphs in which every edge exists if it exists in at least one timestamp.

Exhaustive grid search is carried out to tune parameters of node2vec [46] and DeepWalk [70]. In particular, for both methods, we tune the neighborhood parameters of a vertex, i.e., number of walks r , and walk length l , while the neighborhood size k is set to 10. Furthermore, for node2vec we tune the return and in-out parameters p and q . For each parameter, we use the same grid of values as the one considered in the parameter sensitivity analysis reported in the original node2vec paper [46]. Specifically, we consider the following parameter space:

- Number of walks, $r = \{6, 8, 10, 12, 14, 16, 18, 20\}$;
- Walk length, $l = \{30, 40, 50, 60, 70, 80, 90, 100, 110\}$;
- Return parameter, $p = \{0.25, 0.5, 1, 2, 4\}$;
- In-out parameter, $q = \{0.25, 0.5, 1, 2, 4\}$.

We select the combination of parameters maximizing the Macro F_1 -score averaged over a range of numbers of latent dimensions $d = \{16, 32, 64, 128, 256\}$.

After filtering out those vertices representing the teachers, we partition the remaining vertices (i.e., the students) into training and test sets with an 80–20 split. A standard scaler is applied to the features extracted by each embedding method and, then, a penalized logistic-regression classifier is trained.

In Figure 16, we report classification results in terms of Macro F_1 -score, with varying the dimensionality h of the embeddings. On the PrimarySchool dataset, for $h \geq 200$, our TCS has performance close to 1 in terms Macro F_1 -score, similarly to the three baselines. It can be observed that the TCS results are better as h gets higher; in particular, TCS is even better than node2vec for $h = |T|$. This is expected and is motivated as, for higher h , TCS is allowed to rely on more temporal information about the vertices. On the HighSchool dataset, TCS is outperformed by all methods for smaller h . However, again, the performance of TCS becomes competitive for larger h , up to achieving comparable results to the best method(s) for $h = |T|$.

9 CONCLUSIONS

Temporal networks are a powerful representation of how relations are established and interrupted along time among a given population of entities. An interesting primitive for analyzing this type of networks is the extraction of relevant patterns, such as dense subgraphs, together with their time interval of existence (or span). Following this idea, we introduced in this article, a notion of temporal core decomposition where each core is associated with its span. Exploiting containment properties among cores, we developed efficient algorithms for computing all the span-cores, and also only the maximal ones. We then introduced the problem of temporal community search and showed how it can be solved in polynomial time via DP. We also proved an interesting connection between temporal community search and maximal span-cores, which made it possible to devise a considerably more efficient algorithm than the naïve DP one. Finally, we presented applications on empirical networks of human close-range proximity, that illustrate the relevance of the notions of (maximal) span-core and temporal community search in a variety analyses and applications.

In future work, we will study the role of maximal span-cores with large coreness and/or $|\Delta|$ in spreading processes on temporal networks. Furthermore, span-cores represent features that can be used for network fingerprinting and classification, as well as for model validation, and that could provide support for new ways of visualizing large-scale time-varying graphs. We also plan to investigate different semantics of temporal patterns and the corresponding notions of core decompositions extracted from temporal networks, such as, for instance, the ones that might arise by considering a temporal edge existing in a given interval if it appears in at least one of the timestamps of the interval, or in a fraction of timestamps larger than a given threshold. Finally, investigation in temporal community search has just started: we plan to study different notions of community search and their corresponding extraction problems. As an example, a variant of the notion adopted in this article, which is worth to be investigated further, would correspond to relaxing the requirement of covering the whole temporal domain T , and instead looking for a set of communities giving a good enough temporal coverage.

REFERENCES

- [1] Leman Akoglu, Duen Horng Chau, Christos Faloutsos, Nikolaj Tatti, Hanghang Tong, and Jilles Vreeken. 2013. Mining connection pathways for marked nodes in large graphs. In *Proceedings of the 2013 SIAM International Conference on Data Mining (SDM'13)*. 37–45.
- [2] J. Ignacio Alvarez-Hamelin, Luca Dall'Asta, Alain Barrat, and Alessandro Vespignani. 2006. Large scale networks fingerprinting and visualization using the k-core decomposition. In *Proceedings of the 18th International Conference on Advances in Neural Information Processing Systems (NeurIPS'06)*. 41–50.
- [3] Reid Andersen and Kumar Chellapilla. 2009. Finding dense subgraphs with size bounds. In *Proceedings of the 6th International Workshop on Algorithms and Models for the Web-Graph (WAW'09)*. 25–37.

- [4] Reid Andersen and Kevin J. Lang. 2006. Communities from seed sets. In *Proceedings of the 15th International Conference on World Wide Web (WWW'06)*. 223–232.
- [5] Albert Angel, Nikos Sarkas, Nick Koudas, and Divesh Srivastava. 2012. Dense subgraph maintenance under streaming edge weight updates for real-time story identification. *Proceedings of the VLDB Endowment* 5, 6 (2012), 574–585.
- [6] Gary D. Bader and Christopher W. V. Hogue. 2003. An automated method for finding molecular complexes in large protein interaction networks. *BMC Bioinformatics* 4, 2 (2003).
- [7] Nicola Barbieri, Francesco Bonchi, Edoardo Galimberti, and Francesco Gullo. 2015. Efficient and effective community search. *Data Mining and Knowledge Discovery* 29, 5 (2015), 1406–1433.
- [8] Vladimir Batagelj, Andrej Mrvar, and Matjaz Zaversnik. 1999. Partitioning approach to visualization of large graphs. In *Proceedings of the International Symposium on Graph Drawing (GD'99)*. 90–97.
- [9] Vladimir Batagelj and Matjaz Zaversnik. 2011. Fast algorithms for determining (generalized) core groups in social networks. *Advances in Data Analysis and Classification* 5, 2 (2011), 129–145.
- [10] R. Bellman. 1961. On the approximation of curves by line segments using dynamic programming. *Communications of the ACM* 4, 6 (1961), 284.
- [11] Matthias Bentert, Anne-Sophie Himmel, Hendrik Molter, Marco Marik, Rolf Niedermeier, and René Saitenmacher. 2018. Listing all maximal k -plexes in temporal graphs. In *Proceedings of the 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM'18)*. 41–46.
- [12] Michele Berlingerio, Francesco Bonchi, Björn Bringmann, and Aristides Gionis. 2009. Mining graph evolution rules. In *Proceedings of the Joint European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD'09)*. 115–130.
- [13] Yuchen Bian, Yaowei Yan, Wei Cheng, Wei Wang, Dongsheng Luo, and Xiang Zhang. 2018. On multi-query local community detection. In *Proceedings of the 2018 IEEE International Conference on Data Mining (ICDM'18)*. 9–18.
- [14] Petko Bogdanov, Misael Mongiovi, and Ambuj K Singh. 2011. Mining heavy subgraphs in time-evolving networks. In *Proceedings of the IEEE International Conference on Data Mining (ICDM'11)*. 81–90.
- [15] Francesco Bonchi, Ilaria Bordino, Francesco Gullo, and Giovanni Stilo. 2016. Identifying buzzing stories via anomalous temporal subgraph discovery. In *Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence (WI'16)*. 161–168.
- [16] Francesco Bonchi, Ilaria Bordino, Francesco Gullo, and Giovanni Stilo. 2019. The importance of unexpectedness: Discovering buzzing stories in anomalous temporal graphs. *Web Intelligence* 17, 3 (2019), 177–198.
- [17] Francesco Bonchi, Aristides Gionis, Francesco Gullo, Charalampos E. Tsourakakis, and Antti Ukkonen. 2015. Chromatic correlation clustering. *ACM Transactions on Knowledge Discovery from Data* 9, 4 (2015), 34:1–34:24.
- [18] Francesco Bonchi, Francesco Gullo, and Andreas Kaltenbrunner. 2018. Core decomposition of massive, information-rich graphs. In *Encyclopedia of Social Network Analysis and Mining, 2nd Edition*, Reda Alhajj and Jon G. Rokne (Eds.). Springer.
- [19] Francesco Bonchi, Francesco Gullo, Andreas Kaltenbrunner, and Yana Volkovich. 2014. Core decomposition of uncertain graphs. In *Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (KDD'14)*. 1316–1325.
- [20] Björn Bringmann, Michele Berlingerio, Francesco Bonchi, and Aristides Gionis. 2010. Learning and predicting the evolution of social networks. *IEEE Intelligent Systems* 25, 4 (2010), 26–35.
- [21] Moses Charikar. 2000. Greedy approximation algorithms for finding dense components in a graph. In *Proceedings of the 3rd International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX'00)*. 84–95.
- [22] Moses Charikar, Yonatan Naamad, and Jimmy Wu. 2018. On finding dense common subgraphs. *CoRR* abs/1802.06361 (2018).
- [23] James Cheng, Yiping Ke, Shumo Chu, and M. Tamer Özsu. 2011. Efficient core decomposition in massive networks. In *Proceedings of the IEEE International Conference on Data Engineering (ICDE'11)*. 51–62.
- [24] Wanyun Cui, Yanghua Xiao, Haixun Wang, and Wei Wang. 2014. Local search of communities in large graphs. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data (SIGMOD'14)*. 991–1002.
- [25] Anish Das Sarma, Alpa Jain, and Cong Yu. 2011. Dynamic relationship and event discovery. In *Proceedings of the International Conference on Web Search and Data Mining (WSDM'11)*. 207–216.
- [26] Elise Desmier, Marc Plantevit, Céline Robardet, and Jean-François Boulicaut. 2012. Cohesive co-evolution patterns in dynamic attributed graphs. In *Proceedings of the International Conference on Discovery Science (DS'12)*. 110–124.
- [27] Mark E. Dickison, Matteo Magnani, and Luca Rossi. 2016. *Multilayer Social Networks*. Cambridge University Press.
- [28] Marius Eidsaa and Eivind Almaas. 2013. s -core network decomposition: A generalization of k -core analysis to weighted networks. *Physical Review E* 88, 6 (2013), 062819.
- [29] Alessandro Epasto, Silvio Lattanzi, and Mauro Sozio. 2015. Efficient densest subgraph computation in evolving graphs. In *Proceedings of the World Wide Web Conference (WWW'15)*. 300–310.

- [30] Alessandro Epasto and Bryan Perozzi. 2019. Is a single embedding enough? Learning node representations that capture multiple social contexts. In *Proceedings of the World Wide Web Conference (WWW'19)*. 394–404.
- [31] David Eppstein, Maarten Löffler, and Darren Strash. 2010. Listing all maximal cliques in sparse graphs in near-optimal time. In *Proceedings of the International Symposium on Algorithms and Computation (ISAAC'10)*. 403–414.
- [32] Péter Erdi, Kinga Makovi, Zoltán Somogyvári, Katherine Strandburg, Jan Tobochnik, Péter Volf, and László Zalányi. 2013. Prediction of emerging technologies based on analysis of the US patent citation network. *Scientometrics* 95, 1 (2013), 225–242.
- [33] Christos Faloutsos, Kevin S. McCurley, and Andrew Tomkins. 2004. Fast discovery of connection subgraphs. In *Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (KDD'04)*. 118–127.
- [34] Yixiang Fang, Reynold Cheng, Yankai Chen, Siqiang Luo, and Jiafeng Hu. 2017. Effective and efficient attributed community search. *The VLDB Journal* 26, 6 (2017), 803–828.
- [35] Yixiang Fang, Reynold Cheng, Xiaodong Li, Siqiang Luo, and Jiafeng Hu. 2017. Effective community search over large spatial graphs. *Proceedings of the VLDB Endowment* 10, 6 (2017), 709–720.
- [36] Yixiang Fang, Xin Huang, Lu Qin, Ying Zhang, Wenjie Zhang, Reynold Cheng, and Xuemin Lin. 2020. A survey of community search over big graphs. *The VLDB Journal* 29, 1 (2020), 353–392.
- [37] Edoardo Galimberti, Alain Barrat, Francesco Bonchi, Ciro Cattuto, and Francesco Gullo. 2018. Mining (maximal) span-cores from temporal networks. In *Proceedings of the International Conference on Information and Knowledge Management (CIKM'18)*. 107–116.
- [38] Edoardo Galimberti, Francesco Bonchi, and Francesco Gullo. 2017. Core decomposition and densest subgraph in multilayer networks. In *Proceedings of the International Conference on Information and Knowledge Management (CIKM'17)*. 1807–1816.
- [39] Edoardo Galimberti, Francesco Bonchi, Francesco Gullo, and Tommaso Lanciano. 2020. Core decomposition in multilayer networks: Theory, algorithms, and applications. *ACM Transactions on Knowledge Discovery from Data* 14, 1 (2020), 11:1–11:40.
- [40] Antonios Garas, Frank Schweitzer, and Shlomo Havlin. 2012. A k-shell decomposition method for weighted networks. *New Journal of Physics* 14, 8 (2012), 083030.
- [41] David García, Pavlin Mavrodiev, and Frank Schweitzer. 2013. Social resilience in online communities: The autopsy of Friendster. In *Proceedings of the ACM International Conference on Online Social Networks (COSN'13)*. 39–50.
- [42] Laetitia Gauvin, André Panisson, and Ciro Cattuto. 2014. Detecting the community structure and activity patterns of temporal networks: A non-negative tensor factorization approach. *PLOS ONE* 9, 1 (2014), e86028.
- [43] Valerio Gemmetto, Alain Barrat, and Ciro Cattuto. 2014. Mitigation of infectious disease at school: Targeted class closure vs school closure. *BMC Infectious Diseases* 14, 1 (2014), 695.
- [44] Christos Giatsidis, Dimitrios M. Thilikos, and Michalis Vazirgiannis. 2013. D-cores: Measuring collaboration of directed graphs based on degeneracy. *Knowledge and Information Systems* 35, 2 (2013), 311–343.
- [45] Palash Goyal and Emilio Ferrara. 2018. Graph embedding techniques, applications, and performance: A survey. *Knowledge-Based Systems* 151 (2018), 78–94.
- [46] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (KDD'16)*. 855–864.
- [47] John Healy, Jeannette Janssen, Evangelos E. Milios, and William Aiello. 2006. Characterization of graphs using degree cores. In *Proceedings of the 4th International Workshop on Algorithms and Models for the Web-Graph (WAW'06)*. 25–37.
- [48] Anne-Sophie Himmel, Hendrik Molter, Rolf Niedermeier, and Manuel Sorge. 2016. Enumerating maximal cliques in temporal graphs. In *Proceedings of the IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM'16)*. 337–344.
- [49] Xin Huang, Hong Cheng, Lu Qin, Wentao Tian, and Jeffrey Xu Yu. 2014. Querying k-truss community in large and dynamic graphs. In *Proceedings of the ACM International Conference on Management of Data (SIGMOD'14)*. 1311–1322.
- [50] Xin Huang and Laks VS Lakshmanan. 2017. Attribute-driven community search. *Proceedings of the VLDB Endowment* 10, 9 (2017), 949–960.
- [51] Xin Huang, Laks VS Lakshmanan, and Jianliang Xu. 2017. Community search over big graphs: Models, algorithms, and opportunities. In *Proceedings of the IEEE International Conference on Data Engineering (ICDE'17)*. 1451–1454.
- [52] Akihiro Inokuchi and Takashi Washio. 2010. Mining frequent graph sequence patterns induced by vertices. In *Proceedings of the SIAM International Conference on Data Mining (SDM'10)*. 466–477.
- [53] Vinay Jethava and Niko Beerenwinkel. 2015. Finding dense subgraphs in relational graphs. In *Proceedings of the Joint European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD'15)*. 641–654.
- [54] Maksim Kitsak, Lazaros K. Gallos, Shlomo Havlin, Fredrik Liljeros, Lev Muchnik, H. Eugene Stanley, and Hernán A. Makse. 2010. Identifying influential spreaders in complex networks. *Nature Physics* 6, 11 (2010), 888–893.

- [55] Isabel M. Kloumann and Jon M. Kleinberg. 2014. Community membership identification from small seed sets. In *Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (KDD'14)*. 1366–1375.
- [56] Guy Kortsarz and David Peleg. 1994. Generating sparse 2-spanners. *Journal of Algorithms* 17, 2 (1994), 222–236.
- [57] Lauri Kovanan, Márton Karsai, Kimmo Kaski, János Kertész, and Jari Saramäki. 2011. Temporal motifs in time-dependent networks. *Journal of Statistical Mechanics* 2011, 11 (2011), P11005.
- [58] Cane Wing-ki Leung, Ee-Peng Lim, David Lo, and Jianshu Weng. 2010. Mining interesting link formation rules in social networks. In *Proceedings of the International Conference on Information and Knowledge Management (CIKM'10)*. 209–218.
- [59] Rong-Hua Li, Jiao Su, Lu Qin, Jeffrey Xu Yu, and Qiangqiang Dai. 2018. Persistent community search in temporal networks. In *Proceedings of the IEEE International Conference on Data Engineering (ICDE'18)*. 797–808.
- [60] Rong-Hua Li, Jeffrey Xu Yu, and Rui Mao. 2014. Efficient core maintenance in large dynamic graphs. *IEEE Transactions on Knowledge and Data Engineering* 26, 10 (2014), 2453–2465.
- [61] Boge Liu, Long Yuan, Xuemin Lin, Lu Qin, Wenjie Zhang, and Jingren Zhou. 2019. Efficient (α, β) -core computation: An index-based approach. In *Proceedings of the World Wide Web Conference (WWW'19)*. 1130–1141.
- [62] Paul Liu, Austin R. Benson, and Moses Charikar. 2019. Sampling Methods for Counting Temporal Motifs. In *Proceedings of the International Conference on Web Search and Data Mining (WSDM'19)*. 294–302.
- [63] Shuai Ma, Renjun Hu, Luoshu Wang, Xuelian Lin, and Jimpeng Huai. 2017. Fast computation of dense temporal subgraphs. In *Proceedings of the IEEE International Conference on Data Engineering (ICDE'17)*. 361–372.
- [64] Fragkiskos D. Malliaros, Christos Giatsidis, Apostolos N. Papadopoulos, and Michalis Vazirgiannis. 2020. The core decomposition of networks: Theory, algorithms and applications. *The VLDB Journal* 29, 1 (2020), 61–92.
- [65] Sergei Maslov and Kim Sneppen. 2002. Specificity and stability in topology of protein networks. *Science* 296, 5569 (2002), 910–913.
- [66] Rossana Mastrandrea, Julie Fournet, and Alain Barrat. 2015. Contact patterns in a high school: A comparison between data collected using wearable sensors, contact diaries and friendship surveys. *PLOS ONE* 10, 9 (09 2015), 1–26.
- [67] David W. Matula and Leland L. Beck. 1983. Smallest-last ordering and clustering and graph coloring algorithms. *Journal of the ACM* 30, 3 (1983), 417–427.
- [68] Misael Mongiovi, Petko Bogdanov, Razvan Ranca, Evangelos E. Papalexakis, Christos Faloutsos, and Ambuj K. Singh. 2013. Netspot: Spotting significant anomalous regions on dynamic networks. In *Proceedings of the SIAM International Conference on Data Mining (SDM'13)*. 28–36.
- [69] Alberto Montesor, Francesco De Pellegrini, and Daniele Miorandi. 2013. Distributed k-core decomposition. *IEEE Transactions on Parallel and Distributed Systems* 24, 2 (2013), 288–300.
- [70] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In *Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (KDD'14)*. 701–710.
- [71] Alexander Reinthal, Arvid Andersson, Erik Norlander, Philip Stålhammar, and Sebastian Norlin. 2016. Finding the densest common subgraph with linear programming. Technical Report. Department of Computer Science and Engineering, Chalmers University of Technology.
- [72] Polina Rozenshtein, Francesco Bonchi, Aristides Gionis, Mauro Sozio, and Nikolaj Tatti. 2018. Finding events in temporal networks: Segmentation meets densest-subgraph discovery. In *Proceedings of the IEEE International Conference on Data Mining (ICDM'18)*. 397–406.
- [73] Polina Rozenshtein, Nikolaj Tatti, and Aristides Gionis. 2017. Finding dynamic dense subgraphs. *ACM Transactions on Knowledge Discovery from Data* 11, 3 (2017), 27.
- [74] Natali Ruchansky, Francesco Bonchi, David Garcia-Soriano, Francesco Gullo, and Nicolas Kourtellis. 2015. The minimum wiener connector problem. In *Proceedings of the ACM International Conference on Management of Data (SIGMOD'15)*. 1587–1602.
- [75] Natali Ruchansky, Francesco Bonchi, David Garcia-Soriano, Francesco Gullo, and Nicolas Kourtellis. 2017. To be connected, or not to be connected: That is the minimum inefficiency subgraph problem. In *Proceedings of the International Conference on Information and Knowledge Management (CIKM'14)*. 879–888.
- [76] Anna Sapienza, André Panisson, Joseph Wu, Laetitia Gauvin, and Ciro Cattuto. 2015. Detecting anomalies in time-varying networks using tensor decomposition. In *Proceedings of the IEEE International Conference on Data Mining - Workshops (ICDMW'15)*. 516–523.
- [77] Ahmet Erdem Sariyüce, Bugra Gedik, Gabriela Jacques-Silva, Kun-Lung Wu, and Ümit V. Çatalyürek. 2013. Streaming algorithms for k-core decomposition. *Proceedings of the VLDB Endowment* 6, 6 (2013), 433–444.
- [78] Konstantinos Semertzidis, Evaggelia Pitoura, Evimaria Terzi, and Panayiotis Tsaparas. 2018. Finding lasting dense subgraphs. *Data Mining and Knowledge Discovery* 33, 5 (2018), 1417–1445.
- [79] Mauro Sozio and Aristides Gionis. 2010. The community-search problem and how to plan a successful cocktail party. In *Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (KDD'10)*. 939–948.

- [80] Juliette Stehlé, François Charbonnier, Tristan Picard, Ciro Cattuto, and Alain Barrat. 2013. Gender homophily from spatial behavior in a primary school: A sociometric study. *Social Networks* 35, 4 (2013), 604–613.
- [81] Juliette Stehlé, Nicolas Voirin, Alain Barrat, Ciro Cattuto, Lorenzo Isella, Jean-François Pinton, Marco Quaggiotto, Wouter Van den Broeck, Corinne Régis, Bruno Lina, and Philippe Vanhems. 2011. High-resolution measurements of face-to-face contact patterns in a primary school. *PLOS ONE* 6, 8 (2011), e23176.
- [82] Andrea Tagarelli, Alessia Amelio, and Francesco Gullo. 2017. Ensemble-based community detection in multilayer networks. *Data Mining and Knowledge Discovery* 31, 5 (2017), 1506–1543.
- [83] Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. 2015. Line: Large-scale information network embedding. In *Proceedings of the World Wide Web Conference (WWW'15)*. 1067–1077.
- [84] Hanghang Tong and Christos Faloutsos. 2006. Center-piece subgraphs: Problem definition and fast solutions. In *Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (KDD'06)*. 404–413.
- [85] Tiphaine Viard, Matthieu Latapy, and Clémence Magnien. 2016. Computing maximal cliques in link streams. *Theoretical Computer Science* 609 (2016), 245–252.
- [86] Huanhuan Wu, James Cheng, Yi Lu, Yiping Ke, Yuzhen Huang, Da Yan, and Hejun Wu. 2015. Core decomposition in large temporal graphs. In *Proceedings of the IEEE International Conference on Big Data*. 649–658.
- [87] Stefan Wuchty and Eivind Almaas. 2005. Peeling the yeast protein network. *Proteomics* 5, 2 (2005), 444–449.
- [88] Yaowei Yan, Yuchen Bian, Dongsheng Luo, Dongwon Lee, and Xiang Zhang. 2019. Constrained local graph clustering by colored random walk. In *Proceedings of the World Wide Web Conference (WWW'19)*. 2137–2146.
- [89] Fan Zhang, Ying Zhang, Lu Qin, Wenjie Zhang, and Xuemin Lin. 2017. When engagement meets similarity: Efficient (k, r)-core computation on social networks. *Proceedings of the VLDB Endowment* 10, 10 (2017), 998–1009.
- [90] Haohua Zhang, Hai Zhao, Wei Cai, Jie Liu, and Wanlei Zhou. 2010. Using the k-core decomposition to analyze the static structure of large-scale software systems. *The Journal of Supercomputing* 53, 2 (2010), 352–369.
- [91] Philipp Zschoche, Till Fluschnik, Hendrik Molter, and Rolf Niedermeier. 2020. The complexity of finding small separators in temporal graphs. *Journal of Computer and System Sciences* 107 (2020), 72–92.

Received September 2019; revised May 2020; accepted July 2020