The Thirty-Ninth AAAI Conference on Artificial Intelligence (AAAI-25)

Temporal Triadic Closure: Finding Dense Substructures in Social Networks That Evolve Over Time

Tom Davot, Jessica Enright, Jayakrishnan Madathil, Kitty Meeks
University of Glasgow, Glasgow, United Kingdom
tom.davot@glasgow.ac.uk, jessica.enright@glasgow.ac.uk, jayakrishnan.madathil@glasgow.ac.uk, kitty.meeks@glasgow.ac.uk

Abstract

A graph G is c-closed if every two vertices with at least c common neighbors are adjacent to each other. This definition is an abstraction of the triadic closure property exhibited by many real-world social networks, namely, friends of friends tend to be friends themselves. Social networks, however, are often temporal rather than static—the connections change over a period of time. And hence temporal graphs, rather than static graphs, are often better suited to model social networks. Motivated by this, we introduce a definition of temporal c-closed graphs, in which if two vertices u and v have at least c common neighbors during a short interval of time, then u and v are adjacent to each other around that time. Our pilot experiments show that several real-world temporal networks are c-closed for rather small values of c. We also study the computational problems of enumerating maximal cliques and other dense subgraphs in temporal c-closed graphs. A clique in a temporal graph is a subgraph that lasts for a certain period of time, during which every possible edge in the subgraph becomes active often enough; other dense subgraphs are defined similarly. We bound the number of such maximal dense subgraphs in a temporal c-closed graph that evolves slowly, and thus show that the corresponding enumeration problems admit efficient algorithms; by slow evolution, we mean that between consecutive time-steps, the local change in adjacencies remains small. Our work also adds to a growing body of literature on defining suitable structural parameters for temporal graphs that can be leveraged to design efficient algorithms.

1 Introduction

Social networks evolve. Influencers gain and lose followers on social media; ants in a colony guide each other to food; scientists collaborate with their peers. These examples all involve networks in which connections are created and destroyed as time passes. Moreover, even when a relationship within a network is continuous, the interactions that provide evidence of that relationship—such as being in the same location as a friend, or exchanging an email—only happen at some points in time. All these considerations mean that temporal information must be taken into account to gain a full understanding of many social networks. Temporal graphs provide a useful formalism for modeling these evolving networks; in this work we adopt the widely-used model of Kempe, Kleinberg, and Kumar (2002), in which a temporal graph $\mathcal{G}$ consists of a static graph G called the *underlying graph* or the *footprint* of $\mathcal{G}$, together with a function (which we will denote by λ in this paper) assigning to each edge a (finite) subset of $\mathbb{N}$ representing the discrete timesteps at which the edge appears in the graph or is *active*. We can equivalently consider a temporal graph to be a sequence of static graphs or *snapshots*, where the t^{th} snapshot contains all vertices and only those edges that are active at time t.

Modeling social networks in this way, however, introduces a new level of algorithmic challenge. Even problems that are polynomial-time solvable on static graphs often become NP-hard when a temporal dimension is added; in some cases this holds even when very strong restrictions are placed on the footprint of the graph, such as requiring it to be a path or a star (Mertzios et al. 2023; Akrida et al. 2021). This has prompted significant recent interest in the design of temporal graph parameters that additionally take into account temporal information, with the hope of identifying situations in which natural problems admit efficient algorithms. None of these, however, seems particularly well suited for the design of efficient algorithms to solve problems on social networks. Some only offer a limited benefit over restricting the structure of the footprint: the *timed feedback vertex number* (Casteigts et al. 2021), *temporal feedback edge/connection number* (Haag et al. 2022), and the various temporal analogues of treewidth (Fluschnik et al. 2020) are guaranteed to take small values when the footprint is a tree, so they cannot hope to give tractability for those problems that remain NP-hard under this (or a stronger) restriction. Restricting others, such as the *vertex-interval-membership-width* (Bumpus and Meeks 2023) has proved more widely useful in making problems tractable, but this success comes at a price: the vertex-interval-membership-width of a temporal graph representing a social network will only be small if, at every timestep, most individuals have either (i) never yet formed any connections, or (ii) will never again form a connection, an assumption that seems unrealistic for the vast majority of social networks.

In general, most of the temporal graph parameters defined so far require that the graph be sparse (in the sense of having only a linear number of edges active) at every timestep. Very recent work (Enright et al. 2024) has tried to address this limitation by introducing temporal analogues of the parameter cliquewidth, modular width and neighborhood diversity, which can take small values on dense temporal graphs that are sufficiently highly structured, but these are not without their downsides. Computing the temporal cliquewidth of a temporal graph is an NP-hard problem in itself (Enright et al. 2024), making it hard to gain intuition about which real-world networks, if any, might have small values of this parameter. Meanwhile, for the temporal modular width or temporal neighborhood diversity of a social network to be small there must be large groups of vertices between which the interactions are uniform (i.e. every member of each group has identical interactions with each member of the other group at every time-step), which again does not seem a credible property for many social networks.

In this paper we introduce a new structural parameter for temporal graphs, namely the notion of temporal c-closed graphs. This is an adaptation of static c-closed graphs introduced by Fox et al. (2018; 2020); we discuss the usefulness of these graphs as a model for social networks, and present some preliminary results on real-world networks in in Section 2. We also show that in order to replicate many of the algorithmic results about static c-closed graphs we need to impose further restrictions, and in Section 3 we introduce an additional parameter that captures the extent to which the network changes locally between one timestep and the next. In Sections 4 and 5, we present our main theoretical results—bounds for the number of maximal cliques and similar dense subgraphs in a temporal c-closed graph that evolves sufficiently slowly, which imply the existence of efficient algorithms to find all such subgraphs. All omitted details can be found in the full version (Davot et al. 2024).

Terminology and notation. For a (static) graph G, we use $V(G)$ and $E(G)$ to denote the vertex set and edge set of G, respectively. A temporal graph $\mathcal{G}$ is a pair (G, λ), where G is a static graph and the function $\lambda : E(G) \to 2^{\mathbb{N}}$ specifies the discrete time-steps at which each edge e of G is active. We assume throughout that $\lambda(e)$ is finite. We also assume that the lifespan of a temporal graph $\mathcal{G}$ starts at time-step 1. We use $\Lambda_{\mathcal{G}}$ (or simply Λ when $\mathcal{G}$ is clear) to denote the maximum time-step at which any edge is active, and call $\Lambda_{\mathcal{G}}$ the lifetime of $\mathcal{G}$. We call G the footprint or the underlying graph of $\mathcal{G}$. For $\mathcal{G} = (G, \lambda)$ and a vertex $v \in V(G)$, we use $\mathcal{G} - v$ to denote the temporal graph obtained from $\mathcal{G}$ by deleting v, i.e., $\mathcal{G} - v = (G - v, \lambda')$, where $G - v$ is the subgraph of G induced by $V(G) \setminus \{v\}$ and λ' is the restriction of λ to $E(G - v)$. By a time-interval $I \subseteq \mathbb{N}$ we mean a set of consecutive time-steps, i.e., $I = [a, b] = \{a, a+1, a+2, \dots, b\}$ for some $a, b \in \mathbb{N}$, where $a \leq b$. The length of the time-interval $[a, b]$ is $b-a+1$, i.e., the number of time-steps in $[a, b]$. For a time-interval I and a temporal graph $\mathcal{G} = (G, \lambda)$, we use G_I to denote the subgraph of G that consists of all the edges of G that are active at some time-step in I, i.e., $V(G_I) = V(G)$ and $E(G_I) = \{e \in E(G) \mid \lambda(e) \cap I \neq \emptyset\}$. For $u, v \in V(G)$, we say that u and v are adjacent to each other *during* I if $uv \in E(G_I)$; to emphasize, u and v are adjacent during I if u and v are adjacent to each other at *some* time-step in I, and not necessarily at *every* time-step in I. For $v \in V(G)$, we use $N_I(v)$ to denote the set of neighbors of v in the graph G_I; when $I = \{t\}$, we omit the braces and simply write $N_t(v)$. For $u, v \in V(G)$, we use $CN_I(u, v)$ to denote the common neighborhood of u and v in the graph G_I, i.e., $CN_I(u, v) = N_I(u) \cap N_I(v)$. Notice that this definition does not require that a vertex $w \in CN_I(u, v)$ be adjacent to both u and v at the same time-step; we have $w \in CN_I(u, v)$ if the edge uw is active at time-step t and the edge vw is active at time-step t' for possibly distinct $t, t' \in I$.

2 Temporal c-Closed Graphs

What structural properties could we reasonably expect to see in temporal graphs derived from social networks? Ideally we want to identify deterministic properties of such networks, rather than rely on any of the random models of temporal social networks in the literature (Chakrabarti and Faloutsos 2006; Takaguchi 2015), as deterministic properties are needed to design algorithms with theoretical guarantees on the worst-case running time. Moreover, as noted by Koana, Komusiewicz, and Sommer (2022), an ideal structural parameter should be computable in polynomial-time and also easy to understand. With these considerations in mind, an excellent candidate for such a property is that of triadic closure: This formalizes the idea that people with many friends in common are likely to be friends themselves.

This notion has been leveraged very effectively in the static setting via the notion of *c-closed graphs*. Introduced by Fox et al. (2018; 2020) as a deterministic model for social networks, c-closed graphs are those with the property that every two vertices that have at least c common neighbors are adjacent to each other. The *closure number* of a graph is the least integer c for which it is c-closed. In a series of papers, Koana, Komusiewicz, and Sommer (2022; 2023a; 2023b) demonstrated that the closure number and a related parameter called the weak closure number, which we will discuss shortly, are extremely useful graph parameters that can be exploited to design fixed-parameter tractable (FPT) algorithms for several problems, including classic problems such as INDEPENDENT SET and DOMINATING SET. These parameters have since then received considerable attention from the (parameterized) algorithms community (Behera et al. 2022; Koana et al. 2022; Koana and Nichterlein 2021; Lokshtanov and Surianarayanan 2021; Kanesh et al. 2023).

A temporal analogue of the closure number should take the timing of interactions into account; informally, we might expect that two individuals that both interact with many of the same individuals during some short time-interval are likely to interact with each other either during this same interval or shortly before or afterwards (see Figure 1 for empirical evidence). This leads us to the following definition (in which the notions of "short time-interval", "shortly before" and "shortly afterwards" can be tuned by setting appropriate values of Δ_0, Δ_1 and Δ_2).

Definition 2.1 ($\Delta_0, \Delta_1, \Delta_2, c$)**-closed graphs.** For integers $\Delta_0, \Delta_1, \Delta_2 \geq 0$ and $c \geq 1$, a temporal graph $\mathcal{G} = (G, \lambda)$ is $(\Delta_0, \Delta_1, \Delta_2, c)$-closed if the following condition holds: for every two distinct vertices $u, v \in V(G)$ and any time-interval $[a, b]$ of length at most $\Delta_1 + 1$ (i.e., $b - a \leq \Delta_1$)

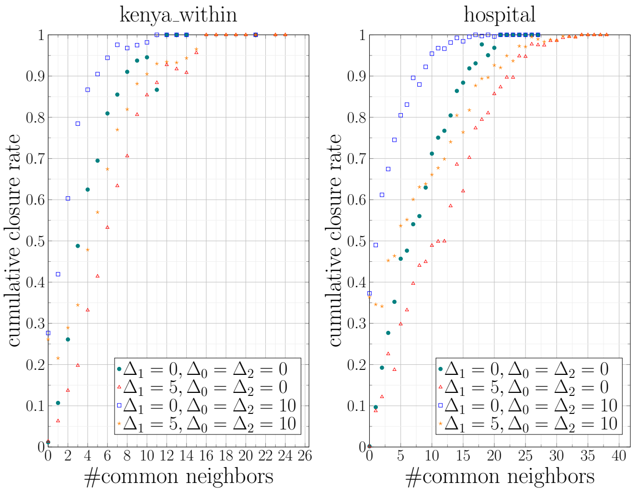

Figure 1: Cumulative closure rate of two real-world temporal networks. Each color corresponds to one choice of $(\Delta_0, \Delta_1, \Delta_2)$. For each x-value, the corresponding y-value is the cumulative closure rate, i.e., the fraction of tuples $([a, a+\Delta_1], u, v)$ such that $[a, a+\Delta_1] \subseteq [1+\Delta_0, \Lambda - \Delta_2]$ and u and v are distinct vertices that have x common neighbors during $[a, a+\Delta_1]$ and are adjacent to each other during $[a - \Delta_0, a + \Delta_1 + \Delta_2]$.

with $a \geq 1 + \Delta_0$ and $b \leq \Lambda_{\mathcal{G}} - \Delta_2$, if u and v have at least c common neighbors during $[a, b]$, then u and v are adjacent to each other during $[a - \Delta_0, b + \Delta_2]$.

We must note that Definition 2.1 is one of the several plausible adaptations of static c-closed graphs to the temporal setting. But it is more general than some of the more simplistic adaptations, say requirements such as (a) every snapshot be c-closed, or (b) the footprint be c-closed, or (c) the static graph during any interval of length Δ_1+1 be c-closed. By appropriately choosing Δ_0, Δ_1 and Δ_2, we can derive each of these three requirements as a special case of Definition 2.1. For example, fixing $\Delta_0 = \Delta_1 = \Delta_2 = 0$ leads to the first requirement, i.e., every snapshot be c-closed.

Along with c-closed graphs, Fox et al. (2020) had also defined a more general class of graphs called weakly γ-closed graphs;[1] we extend this too to the temporal setting. For $\gamma \geq 1$, a (static) graph G is weakly γ-closed if every induced subgraph H of G contains a vertex v such that the number of common neighbors v has with any non-neighbor (in H) is at most $\gamma - 1$. Motivated by weakly γ-closed graphs, we define weakly $(\Delta_0, \Delta_1, \Delta_2, \gamma)$-closed temporal graphs, and for that, we first define the $(\Delta_0, \Delta_1, \Delta_2)$-closure number of a vertex as follows. Recall that $CN_{[a,b]}(u, v)$ denotes the common neighborhood of u and v during the time-interval $[a, b]$.

Definition 2.2 closure number of a vertex. Consider integers $\Delta_0, \Delta_1, \Delta_2 \geq 0$ and a temporal graph $\mathcal{G} = (G, \lambda)$. For a vertex $v \in V(G)$, the $(\Delta_0, \Delta_1, \Delta_2)$-closure of v in $\mathcal{G}$, denoted by $\mathrm{cl}_{\mathcal{G}}(v, (\Delta_0, \Delta_1, \Delta_2))$, is defined as follows:

$$\mathrm{cl}_{\mathcal{G}}(v, (\Delta_0, \Delta_1, \Delta_2)) = \max_{(u,[a,b])} \left\{0, |CN_{[a,b]}(u, v)|\right\},$$

where the maximum is over all vertices u and time-intervals $[a, b]$ such that $u \in V(G) \setminus \{v\}$, $[a, b] \subseteq [1+\Delta_0, \Lambda_{\mathcal{G}} - \Delta_2]$, $b - a \leq \Delta_1$ and $uv \notin E(G_{[a-\Delta_0, b+\Delta_2]})$.

Definition 2.3 weakly $(\Delta_0, \Delta_1, \Delta_2, \gamma)$-closed graphs. For integers $\Delta_0, \Delta_1, \Delta_2 \geq 0$ and $\gamma \geq 1$, a temporal graph $\mathcal{G} = (G, \lambda)$ is weakly $(\Delta_0, \Delta_1, \Delta_2, \gamma)$-closed if there exists an ordering $v_1, v_2, \ldots, v_n$ of the vertices of $\mathcal{G}$ such that $\mathrm{cl}_{\mathcal{G}_i}(v_i) \leq \gamma - 1$ for every $i \in [n]$, where $\mathcal{G}_i$ is the subgraph of $\mathcal{G}$ induced by $\{v_i, v_{i+1}, \ldots, v_n\}$.

We define the $(\Delta_0, \Delta_1, \Delta_2)$-closure number of a temporal graph $\mathcal{G}$ to be the least c for which $\mathcal{G}$ is $(\Delta_0, \Delta_1, \Delta_2, c)$-closed; we define the weak $(\Delta_0, \Delta_1, \Delta_2)$-closure number analogously.

Empirical Results. We now present a dataset of real-world temporal graphs that we have used to observe how well these parameters worked in practice. We have taken the data available on Sociopatterns[2] which brings together a large number of real-world proximity networks. As we wanted to be able to calculate the values of the parameters on a personal computer, we selected only the modest-sized instances (*i.e.* those with at most 250 vertices). The time-steps were also selected to reduce the computation time. The detailed list of the selected instances is the following:

- *baboons*: Interactions between guinea baboons living in an enclosure of a primate center in France (Gelardi et al. 2020). As time step, we chose one day.
- *hospital*: Co-presence in a workplace in a hospital (Génois and Barrat 2018). As time step, we chose one hour.
- *kenya_across* and *kenya_within*: Contact networks between members of households of rural Kenya (Kiti et al. 2016). As time step, we chose one hour.
- *malawi*: Contact patterns in a village in rural Malawi (Ozella et al. 2021). As time step, we chose one day.
- *workplace13* and *workplace15*: Co-presence in a workplace in two different years (Génois and Barrat 2018). As time step, we chose one hour.

Some general statistics are depicted in Table 1. As it was not possible to test every possible values of Δ_0, Δ_1 and Δ_2, we chose to test four configurations for which $\Delta_0 + \Delta_1 + \Delta_2$ does not exceed 10% of the lifetime of the longest instance (*i.e. workspace15*): (a) $\Delta_1 = 0$ and $\Delta_0 = \Delta_2 = 0$, (b) $\Delta_1 = 0$ and $\Delta_0 = \Delta_2 = 10$, (c) $\Delta_1 = 5$ and $\Delta_0 = \Delta_2 = 0$, (d) $\Delta_1 = 5$ and $\Delta_0 = \Delta_2 = 10$.

Our empirical results indicate that for various of choices of $(\Delta_0, \Delta_1, \Delta_2)$, the $(\Delta_0, \Delta_1, \Delta_2)$-closure number and particularly the weak $(\Delta_0, \Delta_1, \Delta_2)$-closure number are often small compared to the number of vertices and edges in the network. More telling is the fact that these values are considerably smaller than their static counterparts on the network's footprint (in which we ignore the time-steps and treat the network as static). See Table 1 for an overview.

[1] Fox et al. (2020) use the letter c rather than γ for weakly closed graphs as well. But subsequent literature on the topic (Koana et al. 2022; Koana, Komusiewicz, and Sommer 2023b) use γ for the weak version, and we defer to this trend.

[2] Available at http://www.sociopatterns.org/datasets/

	(A) General statistics					**(B)** Temporal c and γ				**(C)** Instabilities		
Instance	$\lvert V\rvert$	$\lvert E\rvert$	Λ	Degree	c, γ	$\Delta_1/\Delta_0(\Delta_2 = \Delta_0)$				loc-η	pair-η	
Δ_i				max, min		$0/0$	$0/10$	$5/0$	$5/10$		0	5
baboons	21	162	27	19, 1	15, 4	8, 5	5, 3	12, 11	12, 9	11	8	7
hospital	73	1381	71	61, 2	45, 25	20, 15	20, 9	33, 21	33, 18	26	24	36
kenya_across	21	54	45	14, 1	10, 4	8, 3	8, 3	8, 3	8, 3	13	6	8
kenya_within	47	479	61	39, 6	27, 11	11, 7	10, 6	15, 10	15, 10	19	14	14
malawi	86	347	30	31, 1	10, 5	5, 4	4, 2	6, 5	6, 4	15	5	7
workplace13	95	3915	275	93, 17	70, 45	36, 15	14, 7	41, 20	24, 14	78	77	77
workplace15	217	4274	275	84, 1	41, 20	19, 12	19, 11	30, 16	30, 16	33	23	23

Table 1: Statistics for the instances used for numerical experiments. The Δ_i line contains the values of Δ_0, Δ_1 and Δ_2, when they are relevant. **Part (A):** The columns $|V|$, $|E|$, c, γ and Deg max, min contain contain the number of vertices, the number of edges, the (static) closure and weak closure numbers, and the maximum and minimum degrees of the footprint, respectively, and Λ is the lifetime of the temporal network. **Part (B):** Contains the closure and the weak closure numbers for different values of Δ_0, Δ_1 and Δ_2; for an entry a, b in a column, a is the closure number (i.e., c) and b is the weak closure number (i.e., γ). **Part (C):** The column "loc-η" contains the minimum value η for which the graph is locally η-unstable. The columns "pair-η" contain the minimum value η for which the graph is pairwise η-unstable with the following restriction; to make the computation feasible, instead of considering all intervals $[a, b]$ as required by Definition 3.4, we only considered intervals $[a, a + \Delta_1]$ for $\Delta_1 = 0$ and $\Delta_1 = 5$.

3 The Need for Neighborhood Stability

In the previous section we have introduced our notion of triadic closure for temporal graphs, and have provided some evidence that this definition is useful in describing structural properties of some social networks. Our primary goal in defining this new parameter, however, is to capture realistic structural properties that can be leveraged to design efficient algorithms.

The most obvious computational problems to tackle with our new parameter are the temporal analogues of those known to admit efficient algorithms on static c-closed graphs. The canonical such problem is clique enumeration (in which the goal is to output all maximal subsets of vertices that induce complete subgraphs). Fox et al. (2018; 2020) showed that the number of maximal cliques in an n-vertex c-closed graphs is at most $3^{(c-1)/3} \cdot n^2$; this bound leads to an algorithm for enumerating all maximal cliques in time $3^{(c-1)/3} \cdot \text{poly}(n)$. We must note here that the number of maximal cliques in an arbitrary n-vertex graph could be as large as $3^{n/3}$ (Moon and Moser 1965) (which necessarily implies that any algorithm enumerating all maximal cliques requires time $\Omega(3^{n/3})$ in the worst case). Enumeration of cliques and similar dense subgraphs is also a natural problem to consider in the context of social networks, as such structures correspond to very highly connected communities within the network (a clique represents a community in which every two members interact directly).

To address these algorithmic problems on $(\Delta_0, \Delta_1, \Delta_2, c)$-closed temporal graphs, we first need a notion of temporal cliques. We use a notion of temporal clique that has been studied in the literature, particularly in the context of algorithms for enumerating maximal cliques (Viard, Latapy, and Magnien 2016; Himmel et al. 2017). According to this notion, a clique in a temporal graph is a subgraph in which every possible edge appears every so often. The formal definition is as follows.

Definition 3.1 Δ-clique (Viard, Latapy, and Magnien 2016). Consider a temporal graph $\mathcal{G} = (G, \lambda)$. For a nonnegative integer Δ, a Δ-clique in $\mathcal{G}$ is a pair $(X, [a, b])$, where $X \subseteq V(G)$ and $[a, b]$ is a time-interval, with the following properties: $b - a \geq \Delta$, and for any two distinct vertices $u, v \in X$ and for every $\tau \in [a, b-\Delta]$, the edge uv is active during $[\tau, \tau+\Delta]$, i.e., there exists $t \in \lambda(uv) \cap [\tau, \tau+\Delta]$.

Informally, a maximal Δ-clique is one which is not contained in any other Δ-clique. As a Δ-clique has two constituent parts—a set of vertices and a time interval—the "not contained in any other" needs to be defined with respect to both. We define this formally now. Consider a Δ-clique $(X, [a, b])$. We say that $(X, [a, b])$ is *vertex-maximal* if there does not exist a vertex subset $X' \subseteq V(G)$ such that $X \subsetneq X'$ and $(X', [a, b])$ is a Δ-clique. Similarly, we say that $(X, [a, b])$ is *time-maximal* if there does not exist a time-interval $[a', b']$ such that $[a, b] \subsetneq [a', b']$ and $(X, [a', b'])$ is a Δ-clique. And we say that $(X, [a, b])$ is a *maximal Δ-clique* if it is both vertex-maximal and time-maximal. Notice that if a Δ-clique $(X, [a, b])$ is not vertex-maximal, then $(X \cup \{u\}, [a, b])$ is a Δ-clique for some vertex $u \notin X$. And if $(X, [a, b])$ is not time-maximal, then either $(X, [a-1, b])$ is a Δ-clique or $(X, [a, b+1])$ is a Δ-clique.

The number of maximal Δ-cliques in an n-vertex temporal graph $\mathcal{G}$ could be as large as $2^n \cdot \Lambda_{\mathcal{G}}$ (see Example 3.2 below). Himmel et al. (2017) introduced a parameter called Δ-slice degeneracy, denoted by d, which is the maximum degeneracy of the underlying graph during any time-interval of length Δ, and showed that the number of maximal Δ-cliques is at most $3^{d/3} \cdot 2^{d+1} \cdot n \cdot \Lambda$. A bound for maximal cliques w.r.t. a parameter called vertex deletion to order preservation is implicit in the work of Hermelin et al. (2023).

Ideally, we would like to extend the results of Fox et al. (2020) to the temporal setting; that is, bound the number of maximal Δ-cliques in a $(\Delta_0, \Delta_1, \Delta_2, c)$-closed temporal graph by $f(c, \Delta_0, \Delta_1, \Delta_2, \Delta) \cdot \text{poly}(n, \Lambda)$ for some function f, where n and Λ respectively are the number of vertices and the lifetime of the temporal graph. But this is not possible: It is not difficult to construct pathological examples of temporal graphs that are $(\Delta_0, \Delta_1, \Delta_2, c)$-closed but have $\Omega(2^n)$ maximal Δ-cliques; see Example 3.2.

Example 3.2. Consider the temporal graph $\mathcal{G}' = (G', \lambda')$ defined as follows. The footprint G' is a clique on n vertices, and we will assign time-steps to the edges in such a way that for each non-empty $X \subseteq V(G')$, we will have a maximal Δ-clique $(X, [a, b])$ for an appropriate time-interval $[a, b]$. Let $X_1, X_2, \ldots, X_{2^n-1}$ be the non-empty subsets of $V(G')$ (ordered arbitrarily). For $i \in [2^n-1]$, let $t_i = (\Delta+2)i-1$. Now, for each $uv \in E(G')$, we define $\lambda'(uv) = \{t_i \mid u, v \in X_i\}$ so that each X_i induces a clique at time-step t_i. Notice also that $t_{i+1}-t_i = [(\Delta+2)(i+1)-1]-[(\Delta+2)i-1] = \Delta+2$, and thus there is a gap of $\Delta + 2$ time-steps between the cliques induced by X_i and X_{i+1}. Hence, for each X_i with $|X_i| \geq 2$, $(X_i, [t_i - \Delta, t_i + \Delta])$ is a maximal Δ-clique. As for X_is with $|X_i| = 1$, notice that $(X_i, [1, \Lambda])$ is trivially a maximal Δ-clique. It is straightforward to verify that $\mathcal{G}'$ is $(\Delta_0, \Delta_1, \Delta_2, 1)$-closed for any $\Delta_0, \Delta_1, \Delta_2 \geq 0$ such that $\Delta_1 \leq \Delta$; for any interval $[a, b]$ with $b-a \leq \Delta_1 \leq \Delta$, if two vertices u and v have at least one common neighbor during $[a, b]$, then $t_i \in [a, b]$ and $u, v \in X_i$ for some $i \in [2^n - 1]$, and hence u and v are adjacent to each other during $[a, b]$.

Example 3.2 works because of all the cliques that appear suddenly only to disappear in an instant. Or at a local level, neighborhoods of vertices change too drastically and too suddenly. We refer to such changes as local instability, and try to limit them, which leads us to the following definition.

Definition 3.3 locally η-unstable graphs. For a non-negative integer η, a temporal graph $\mathcal{G} = (G, \lambda)$ is locally η-unstable if $\max\{|N_t(v) \setminus N_{t+1}(v)|, |N_{t+1}(v) \setminus N_t(v)|\} \leq \eta$ for every vertex $v \in V(G)$ and every time-step $t \leq \Lambda - 1$.

We may think of the measure of η as capturing the rate of local change in adjacencies. And when η is small, the graph evolves slowly over time. In particular, when $\eta = 0$, the graph does not evolve at all; so a locally 0-unstable temporal graph is essentially a static graph. Unfortunately, initial investigations on real datasets indicate that requiring this parameter to be small may often be too restrictive; see Table 1. This leads us to the following weaker (and less intuitive) version of instability, in which common neighborhoods of pairs of vertices change slowly over time, which is nevertheless sufficient for our purposes. We define this below and observe from Table 1 that in some cases it takes much smaller values than our first version.

Definition 3.4 pairwise η-unstable graphs. For a non-negative integer η, a temporal graph $\mathcal{G} = (G, \lambda)$ is pairwise η-unstable if for any two distinct vertices u and v and any time-interval $[a, b]$ with $b - a = \Delta_1$ and non-negative integers ℓ, ℓ' such that $[a - \ell, b + \ell'] \subseteq [1, \Lambda_{\mathcal{G}}]$, we have $|CN_{[a-\ell,b+\ell']}(u, v) \setminus CN_{[a,b]}(u, v)| \leq \eta(\ell + \ell')$.

It is not difficult to prove that if a temporal graph $\mathcal{G}$ is locally η-unstable, then it is also pairwise-2η-unstable; if $\mathcal{G}$ is locally η-unstable, then for any time-interval $[a, b]$ and ℓ, ℓ', each of the $\ell+\ell'$ time-steps in $[a-\ell', b+\ell] \setminus [a, b]$ contributes at most 2η vertices to $CN_{[a-\ell',b+\ell]}(u, v) \setminus CN_{[a,b]}(u, v)$. The same argument, when applied to a $(\Delta_0, \Delta_1, \Delta_2, c)$-closed graph $\mathcal{G}$ and vertices u and v that are not adjacent during an interval of length at least $\Delta_0 + \Delta_1 + \Delta_2 + 1$, implies the following lemma, which we will use in Section 4 to bound the number of maximal Δ-cliques.

Lemma 3.5. *Let $\mathcal{G}$ be a locally η-unstable, $(\Delta_0, \Delta_1, \Delta_2, c)$-closed temporal graph. For any two distinct vertices $u, v \in V(G)$ such that $uv \notin E(G_{[a,b]})$ for some time-interval $[a, b]$ where $b-a \geq \Delta_0+\Delta_1+\Delta_2$, it holds that $|CN_{[a,b]}(u, v)| \leq c - 1 + 2\eta(b - a - \Delta_1)$.*

4 Bound for Maximal Cliques

We now bound the number of maximal Δ-cliques in a $(\Delta_0, \Delta_1, \Delta_2, c)$ closed, locally η-unstable temporal graph.

Theorem 4.1. *For every $c \geq 1$ and $\eta, \Delta_0, \Delta_1, \Delta_2 \geq 0$, where $\Delta \geq \Delta_0 + \Delta_1 + \Delta_2$, every locally η-unstable, $(\Delta_0, \Delta_1, \Delta_2, c)$-closed temporal graph with n vertices and lifetime Λ has at most $3 \cdot 2^{c-1+2\eta(\Delta+1-\Delta_1)} \cdot n^2 \cdot \Lambda$ maximal Δ-cliques.*

We need the following observation to prove Theorem 4.1.

Observation 4.1. *Consider a temporal graph $\mathcal{G} = (G, \lambda)$ and a non-negative integer Δ. For each $X \subseteq V(G)$ and each positive integer $\tau \leq \Lambda_{\mathcal{G}}$, there exists at most one time-interval $[a, b]$ such that $\tau \in [a, b]$ and $(X, [a, b])$ is a maximal Δ-clique. Notice that this holds for* all *temporal graphs, and not just $(\Delta_0, \Delta_1, \Delta_2, c)$-closed or locally (or pairwise) η-unstable temporal graphs.*

Proof of Theorem 4.1. Let $F(\eta, \Delta_0, \Delta_1, \Delta_2, c, \Delta, n, \Lambda)$ be the maximum number of maximal Δ-cliques in an η-unstable $(\Delta_0, \Delta_1, \Delta_2, c)$-closed temporal graph with n vertices and lifetime Λ; for convenience, we use ρ as a shorthand for $(\eta, \Delta_0, \Delta_1, \Delta_2, c)$, and write $F(\rho, \Delta, n, \Lambda)$. Let $\mathcal{G} = (G, \lambda)$ be an η-unstable, $(\Delta_0, \Delta_1, \Delta_2, c)$-closed temporal graph with n vertices and lifetime Λ, and let v be an arbitrary vertex of $\mathcal{G}$. Then every maximal Δ-clique $(X, [a, b])$ in $\mathcal{G}$ is of one (or more) of the following five types.

- Type 1: X does not contain v and hence $(X, [a, b])$ is a maximal Δ-clique in $\mathcal{G} - v$.
- Type 2: X contains v and $(X \setminus \{v\}, [a, b])$ is a maximal Δ-clique in $\mathcal{G} - v$.
- Type 3: X contains v and $(X \setminus \{v\}, [a, b])$ is not a vertex-maximal Δ-clique in $\mathcal{G} - v$.
- Type 4a: X contains v, $(X \setminus \{v\}, [a, b])$ is not a time-maximal Δ-clique in $\mathcal{G} - v$, but $(X \setminus \{v\}, [a - 1, b])$ is a Δ-clique in $\mathcal{G} - v$.
- Type 4b: X contains v, $(X \setminus \{v\}, [a, b])$ is not a time-maximal Δ-clique in $\mathcal{G} - v$, but $(X \setminus \{v\}, [a, b + 1])$ is a Δ-clique in $\mathcal{G} - v$.

We will bound the number of maximal Δ-cliques of each type. First, types 1 and 2. As they are maximal Δ-cliques in $\mathcal{G} - v$, and since $\mathcal{G} - v$, being an induced subgraph of $\mathcal{G}$, is an η-unstable and $(\Delta_0, \Delta_1, \Delta_2, c)$-closed temporal graph with $n - 1$ vertices and lifetime at most Λ, the number of maximal Δ-cliques in $\mathcal{G} - v$ is at most $F(\rho, \Delta, n - 1, \Lambda)$.

Therefore the number of maximal Δ-cliques of types 1 and 2 in $\mathcal{G}$ is at most $F(\rho, \Delta, n-1, \Lambda)$.

Let us now bound the number of Δ-cliques of type 3. Consider such a Δ-clique $(X, [a, b])$. Then $(X \setminus \{v\}, [a, b])$ is not a vertex-maximal Δ-clique in $\mathcal{G} - v$, which implies that there exists a vertex $u \in V(G) \setminus X$ such that $((X \setminus \{v\}) \cup \{u\}, [a, b])$ is a Δ-clique. As $(X, [a, b])$ is a maximal Δ-clique and $u \notin X$, we can conclude that there exists a time-interval $[\tau, \tau + \Delta] \subseteq [a, b]$ such that u and v are not adjacent to each other during $[\tau, \tau + \Delta]$, i.e., $uv \notin E(G_{[\tau,\tau+\Delta]})$. But notice that every vertex in $X \setminus \{v\}$ is adjacent to v at some time-step in $[\tau, \tau + \Delta]$ (as $(X, [a, b])$ is a Δ-clique); that is, $X \setminus \{v\} \subseteq N_{[\tau,\tau+\Delta]}(v)$. Similarly, every vertex in $X \setminus \{v\}$ is also adjacent to u at some time step in $[\tau, \tau + \Delta]$ (as $((X \setminus \{v\}) \cup \{u\}, [a, b])$ is a Δ-clique); that is, $X \setminus \{v\} \subseteq N_{[\tau,\tau+\Delta]}(u)$. We thus have $X \setminus \{v\} \subseteq N_{[\tau,\tau+\Delta]}(v) \cap N_{[\tau,\tau+\Delta]}(u) = CN_{[\tau,\tau+\Delta]}(u, v)$. Therefore, the number of choices for $X \setminus \{v\}$ is at most $2^{|CN_{\tau,\tau+\Delta}(u,v)|}$. Also, by Observation 4.1, for each subset Y of $CN_{[\tau,\tau+\Delta]}(u, v)$ there exists at most one time-interval $[a', b']$ such that $[\tau, \tau+\Delta] \subseteq [a', b']$ and $(Y \cup \{v\}, [a', b'])$ is a maximal Δ-clique in $\mathcal{G}$. Thus, by summing over all choices for u and τ, we get that the number of maximal Δ-cliques of type 3 is at most $\sum_{(u,\tau)} 2^{|CN_{[\tau,\tau+\Delta]}(u,v)|}$, where the summation is over all pairs (u, τ) such that $u \in V(G) \setminus \{v\}$, $1 \le \tau \le \Lambda - \Delta$ and $uv \notin E(G_{[\tau,\tau+\Delta]})$. Now, as $\mathcal{G}$ is locally η-unstable and $(\Delta_0, \Delta_1, \Delta_2, c)$-closed, and as $uv \notin E(G_{[\tau,\tau+\Delta]})$ with $\Delta \ge \Delta_0 + \Delta_1 + \Delta_2$, we can apply Lemma 3.5, by which we have $|CN_{[\tau,\tau+\Delta]}(u, v)| \le c - 1 + 2\eta(\Delta - \Delta_1)$. Then, as the pair (u, τ) has at most $n \cdot \Lambda$ choices, we get that the number of maximal Δ-cliques of type 3 is at most $2^{c-1+2\eta(\Delta-\Delta_1)} \cdot n \cdot \Lambda$.

We now bound the number of maximal Δ-cliques of type 4a. Consider a Δ-clique $(X, [a, b])$ of type 4a. Then $(X \setminus \{v\}, [a, b])$ is not time-maximal, and $(X \setminus \{v\}, [a-1, b])$ is a Δ-clique. As $(X, [a, b])$ is a maximal Δ-clique, and in particular a time-maximal Δ-clique, $(X, [a-1, b])$ is not a Δ-clique, which, along with the fact that $(X \setminus \{v\}, [a-1, b])$ is a Δ-clique, implies that there exists a vertex $u \in X \setminus \{v\}$ such that u and v are not adjacent to each other during the interval $[a-1, a-1+\Delta]$. That is, $uv \notin E(G_{[a-1,a-1+\Delta]})$. Then, as before, Lemma 3.5 applies, and we thus have $|CN_{[a-1,a+\Delta]}(u, v)| \le c - 1 + 2\eta(\Delta + 1 - \Delta_1)$. Now, as $(X, [a, b])$ is a Δ-clique, and $u, v \in X$, every vertex $w \in X \setminus \{u, v\}$ is adjacent to both u and v during the interval $[a, a + \Delta]$. That is, $uw, vw \in E(G_{[a,a+\Delta]})$ for every $w \in X \setminus \{u, v\}$, which implies that $X \setminus \{u, v\} \subseteq CN_{[a,a+\Delta]}(u, v) \subseteq CN_{[a-1,a+\Delta]}(u, v)$. Hence, $|X \setminus \{u, v\}| \le |CN_{a-1,a-1+\Delta}(u, v)| \le c-1+2\eta(\Delta+1-\Delta_1)$, and thus the number of choices for $X \setminus \{u, v\}$ is at most $2^{c-1+2\eta(\Delta+1-\Delta_1)}$. By summing over all choices of u and a, we get that the number of maximal Δ-cliques of type 4a is at most $\sum_{(u,a)} 2^{c-1+2\eta(\Delta+1-\Delta_1)}$, where the summation is over all pairs (u, a) such that $u \in V(G) \setminus \{v\}$, and $2 \le a \le \Lambda - \Delta$, and $uv \notin E(G_{[a-1,a-1+\Delta]})$. As (u, a) has at most $n \cdot \Lambda$ choices, we get that the number of maximal Δ-cliques of type 4a is at most $2^{c-1+2\eta(\Delta+1-\Delta_1)} \cdot n \cdot \Lambda$. By symmetric arguments, we can also conclude that the number of maximal Δ-cliques of type 4b is at most $2^{c-1+2\eta(\Delta+1-\Delta_1)} \cdot n \cdot \Lambda$.

We have thus shown that the number of maximal Δ-cliques of types 3, 4a and 4b together is at most $2^{c-1+2\eta(\Delta-\Delta_1)} \cdot n \cdot \Lambda + 2 \cdot 2^{c-1+2\eta(\Delta+1-\Delta_1)} \cdot n \cdot \Lambda$, which is at most $3 \cdot 2^{c-1+2\eta(\Delta+1-\Delta_1)} \cdot n \cdot \Lambda$. Recall that the number of maximal Δ-cliques of types 1 and 2 together is at most $F(\rho, \Delta, n-1, \Lambda)$. Thus, taking into account all five types, the number of maximal Δ-cliques in $\mathcal{G}$ is governed by the recursive inequality $F(\rho, \Delta, n, \Lambda) \le F(\rho, \Delta, n-1, \Lambda) + 3 \cdot 2^{c-1+2\eta(\Delta+1-\Delta_1)} \cdot n \cdot \Lambda$. By induction on n with base case $F(\rho, \Delta, 1, \Lambda) = 1$, we get that $F(\rho, \Delta, n, \Lambda)$ is at most $3 \cdot 2^{c-1+2\eta(\Delta+1-\Delta_1)} \cdot n^2 \cdot \Lambda$. This completes the proof. ☐

Extensions and Implications of Theorem 4.1. We now note down a number of results that can be derived by using the same arguments as in the proof of Theorem 4.1. Notice first that in the proof of Theorem 4.1, the vertex v was chosen arbitrarily, and the only property of v that we used was $\mathrm{cl}_{\mathcal{G}}(v, (\Delta_0, \Delta_1, \Delta_2)) \le c - 1$. Hence the same proof will still go through so long as there exists a vertex v with this property at each recursive level; that is, there exists an ordering $v_1, v_2, \ldots, v_n$ of the vertices of $\mathcal{G}$ such that $\mathrm{cl}_{\mathcal{G}_i}(v_i, (\Delta_0, \Delta_1, \Delta_2)) \le c - 1$, where $\mathcal{G}_i$ is the subgraph of $\mathcal{G}$ induced by $v_i, v_{i+1}, \ldots, v_n$. Recall that this is precisely the definition of weakly $(\Delta_0, \Delta_1, \Delta_2, c)$-closed graphs (Definition 2.3). We thus have the following result.

Theorem 4.2. *For every $\gamma \ge 1$ and $\eta, \Delta_0, \Delta_1, \Delta_2 \ge 0$, where $\Delta \ge \Delta_0 + \Delta_1 + \Delta_2$, every locally η-unstable, weakly $(\Delta_0, \Delta_1, \Delta_2, \gamma)$-closed temporal graph with n vertices and lifetime Λ has at most $3 \cdot 2^{\gamma-1+2\eta(\Delta+1-\Delta_1)} \cdot n^2 \cdot \Lambda$ maximal Δ-cliques.*

Observe also that the proof of Theorem 4.1 can naturally be turned into an algorithm that enumerates all maximal Δ-cliques. But we may instead use any known algorithm for enumerating maximal Δ-cliques. And by using an algorithm due to Himmel et al. (2017),[3] we get the following result.

Theorem 4.3. *There is an algorithm that given a locally η-unstable, weakly $(\Delta_0, \Delta_1, \Delta_2, \gamma)$-closed temporal graph $\mathcal{G} = (G, \lambda)$ and $\Delta \ge \Delta_0 + \Delta_1 + \Delta_2$, runs in time $\mathcal{O}(2^{\gamma-1+2\eta(\Delta+1-\Delta_1)} \cdot n^2 \cdot m \cdot \Lambda)$, and enumerates all maximal Δ-cliques in $\mathcal{G}$, where n and m respectively are the number of vertices and edges of the footprint G.*

Theorem 4.3 says that the maximal clique enumeration problem and consequently the decision problem of checking if a temporal graph contains a clique of a given size are fixed-parameter tractable, when parameterized by $\gamma + \eta + \Delta$.

Remark 4.2 Pairwise η-instability is sufficient. In Theorem 4.1, we used the local η-instability of $\mathcal{G}$ only when we invoked Lemma 3.5; for example, in the type 3 case,

[3]We use Theorem 2 in (Himmel et al. 2017), which says that all maximal Δ-cliques can be enumerated in time $\mathcal{O}(x \cdot m + m \cdot \Lambda)$, where x is the number of *time-maximal* Δ-cliques, and m is the number of edges in the footprint. To use this result, we therefore need to bound the number of time-maximal Δ-cliques. But it is easy to adapt our proof of Theorem 4.1 to bound the number of time-maximal Δ-cliques; we simply omit Type 3 cliques.

to bound $|CN_{[\tau,\tau+\Delta]}(u,v)|$. For this, pairwise η-instability is sufficient. Thus we may replace local η-instability with pairwise 2η-instability in Theorems 4.1, 4.2 and 4.3, and the same bounds will still hold; the same goes for Theorem 5.1.

Remark 4.3 Necessity of the assumptions in Theorem 4.1. Theorem 4.1 crucially relies on three requirements: η-instability, $(\Delta_0, \Delta_1, \Delta_2, c)$-closure and $\Delta \geq \Delta_0+\Delta_1+\Delta_2$. Each of these requirements is necessary to yield our bound for the number of maximal Δ-cliques. Without any one of them, the number of maximal Δ-cliques may blow up to $2^{\Omega(n)}$. We already saw the necessity of η-instability in Example 3.2. The necessity of $(\Delta_0, \Delta_1, \Delta_2, c)$ closure is even more straightforward, because for every $n \geq 1$, there exists a static graph with $3n$ vertices and 3^n maximal cliques; this is the classic Moon-Moser graph, the complete n-partite graph with exactly 3 vertices in each part (Moon and Moser 1965). To adapt it to the temporal setting, we only need to assign all time-steps from 1 to $\Delta + 1$ to every edge; the resulting temporal graph is 0-unstable, has lifetime $\Delta + 1$ and contains 3^n maximal Δ-cliques, but not $(\Delta_0, \Delta_1, \Delta_2, c)$-closed any $\Delta_0, \Delta_1, \Delta_2 \geq 0$ with $\Delta_0 + \Delta_1 + \Delta_2 \leq \Delta$ and $c \leq 3n - 3$. Finally, to see that $\Delta \geq \Delta_0 + \Delta_1 + \Delta_2$ is also necessary, notice that for any $c \geq 1$ and any temporal graph $\mathcal{G}$ (including the temporal adaptation of the Moon-Moser graph that we just saw), we can always choose $(\Delta_0, \Delta_1, \Delta_2)$ in such a way that $\mathcal{G}$ is $(\Delta_0, \Delta_1, \Delta_2, c)$-closed; for example, if we choose $\Delta_0 = \Delta_2 = \Lambda_{\mathcal{G}}$, then $\mathcal{G}$ would be vacuously $(\Delta_0, \Delta_1, \Delta_2, c)$ closed for any c.

Remark 4.4 Weak versions of pairwise-instability. Just like the weak γ-closure, it is possible to define a weak version of the pairwise instability. A temporal graph is weakly pairwise η-unstable if it is possible to find an ordering $v_1, \ldots, v_n$ of the vertices such that for pair v_i and v_j with $i < j$, any time interval $[a, b]$ with $b - a = \Delta_1$ and any two non-negative integers ℓ, ℓ' such that $[a - \ell, b + \ell'] \subseteq [1, \Lambda]$, we have $|CN_{[a-\ell,b+\ell']}(u, v) \setminus CN_{[a,b]}(u, v)| \leq \eta(\ell + \ell')$ in the subgraph induced by $\{v_i, \ldots, v_n\}$. Another particularly interesting "weak version" consists in finding an order of the vertices that minimizes both γ and η. That is, an ordering $v_1, \ldots, v_n$ that minimizes the maximum between the temporal closure of v_i and the pairwise instability of v_i in the graph induced by $\{v_i, \ldots, v_n\}$.

5 Bounds for Other Dense Subgraphs

In Section 4, we bounded the number of maximal Δ-cliques. Using the same ideas, we can prove similar bounds for other dense subgraphs such as the temporal adaptations of $(k+1)$-plexes and k-defective cliques. For $k \geq 0$, a $(k + 1)$-plex is a static graph in which every vertex has at most k non-neighbors, and a k-defective clique is one which has at least $\binom{n}{2} - k$ edges (where n is the number of vertices). Thus a static clique is a 1-plex and a 0-defective clique.

Behera et al. (2022) showed that various dense subgraphs also admit bounds of the form $f(c) \cdot \text{poly}(n)$ in a static c-closed graph. In particular, they showed that the number of maximal $(k + 1)$-plexes in an n-vertex c-closed graph is at most $n^{2k} r_k^c c^{\mathcal{O}(1)}$, where $r_k < 2$ is a constant. The other dense subgraphs they considered were complements of forests, bounded treewidth graphs and bounded degeneracy graphs. Koana, Komusiewicz, and Sommer (2023a) proved similar bounds on weakly γ-closed graphs; they showed that in an n-vertex weakly γ-closed graph, the number of maximal $(k + 1)$-plexes is at most $2^\gamma n^{2k-1}$ and the number of maximal k-defective cliques is at most $2^\gamma n^{k+1}$.

Bentert et al. (2019) adapted $(k+1)$-plexes to the temporal setting as follows. For non-negative integers Δ and k, a $(\Delta, k + 1)$-plex in a temporal graph $\mathcal{G}$ is a pair $(X, [a, b])$, where $X \subseteq V(G)$ and $[a, b]$ is a time-interval such that $b - a \geq \Delta$, and for every vertex $v \in X$ and for every $\tau \in [a, b - \Delta]$, there exist at most k vertices $u \in X \setminus \{v\}$ such that $uv \notin E(G_{[\tau,\tau+\Delta]})$. We similarly adapt k-defective cliques as follows. A (Δ, k)-defective clique in $\mathcal{G}$ is a pair $(X, [a, b])$, where $X \subseteq V(G)$ and $[a, b]$ is a time-interval such that $b - a \geq \Delta$, and for every $\tau \in [a, b - \Delta]$, the static graph $G[X]$ has at least $\binom{|X|}{2} - k$ edges that are active during the interval $[\tau, \tau + \Delta]$. The definitions of a maximal $(\Delta, k + 1)$-plex and a maximal (Δ, k)-defective cliques are analogous to the definition of a maximal Δ-clique. And we prove the following result.

Theorem 5.1. *For $\gamma \geq 1$, $\eta, \Delta_0, \Delta_1, \Delta_2, k \geq 0$, where $\Delta \geq \Delta_0 + \Delta_1 + \Delta_2$, every η-unstable, weakly $(\Delta_0, \Delta_1, \Delta_2, \gamma)$-closed temporal graph with n vertices and lifetime Λ has at most $4 \cdot 2^{\gamma-1+2\eta(\Delta+1-\Delta_1)} \cdot n^{\max\{2k,k+2\}} \cdot \Lambda$ maximal $(\Delta, k+1)$-plexes, and at most $4 \cdot 2^{\gamma-1+2\eta(\Delta+1-\Delta_1)} \cdot n^{k+2} \cdot \Lambda$ maximal (Δ, k)-defective cliques.*

6 Concluding Remarks

We introduced $(\Delta_0, \Delta_1, \Delta_2, c)$-closed temporal graphs, which formalizes the triadic closure property of social networks. Our empirical results suggest that temporal versions of c and γ could be meaningful parameters in the study of real-world networks. But the networks we looked at are all relatively small contact networks, and further evaluation of various types of large real-world networks is necessary to confirm this. Our theoretical results demonstrate the usefulness of these parameters in designing algorithms. We hope that these parameters could be further exploited, and that temporal c and γ will prove to be just as useful as their static counterparts in designing algorithms for a variety of problems. While (local or pairwise) η-instability is one sufficient condition that, when coupled with $(\Delta_0, \Delta_1, \Delta_2, c)$-closure, can yield a non-trivial bound for the number of maximal Δ-cliques, it will be worth investigating whether there are more reasonable conditions that can yield a similar bound. We must also add a note of caution here. These parameters are all rather crude abstractions of properties exhibited by real-world networks; they are based on how adjacencies behave and evolve over time in an idealized temporal social network. They can nonetheless illuminate the structure of real-world networks. Also, less-than-realistic abstraction is often the price we must pay for algorithms with provable worst-case guarantees. It is not our case that these parameters will be sufficiently small for practical purposes for all real-world temporal networks. Practical applications will require further refinement of these parameters, and we hope that our work will trigger such inquiries.

Acknowledgments

Tom Davot and Jessica Enright are supported by EPSRC grant EP/T004878/1. Jayakrishnan Madathil and Kitty Meeks are supported by EPSRC grants EP/T004878/1 and EP/V032305/1.

References

Akrida, E. C.; Mertzios, G. B.; Spirakis, P. G.; and Raptopoulos, C. L. 2021. The temporal explorer who returns to the base. *J. Comput. Syst. Sci.*, 120: 179–193.

Behera, B.; Husic, E.; Jain, S.; Roughgarden, T.; and Seshadhri, C. 2022. FPT Algorithms for Finding Near-Cliques in c-Closed Graphs. In *ITCS*, volume 215 of *LIPIcs*, 17:1–17:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Bentert, M.; Himmel, A.; Molter, H.; Morik, M.; Niedermeier, R.; and Saitenmacher, R. 2019. Listing All Maximal k-Plexes in Temporal Graphs. *ACM J. Exp. Algorithmics*, 24(1): 1.13:1–1.13:27.

Bumpus, B. M.; and Meeks, K. 2023. Edge Exploration of Temporal Graphs. *Algorithmica*, 85(3): 688–716.

Casteigts, A.; Himmel, A.; Molter, H.; and Zschoche, P. 2021. Finding Temporal Paths Under Waiting Time Constraints. *Algorithmica*, 83(9): 2754–2802.

Chakrabarti, D.; and Faloutsos, C. 2006. Graph mining: Laws, generators, and algorithms. *ACM Comput. Surv.*, 38(1): 2.

Davot, T.; Enright, J. A.; Madathil, J.; and Meeks, K. 2024. Temporal Triadic Closure: Finding Dense Structures in Social Networks That Evolve. *CoRR*, abs/2412.09567.

Enright, J. A.; Hand, S. D.; Larios-Jones, L.; and Meeks, K. 2024. Structural Parameters for Dense Temporal Graphs. In *MFCS*, volume 306 of *LIPIcs*, 52:1–52:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Fluschnik, T.; Molter, H.; Niedermeier, R.; Renken, M.; and Zschoche, P. 2020. As Time Goes By: Reflections on Treewidth for Temporal Graphs. In *Treewidth, Kernels, and Algorithms*, volume 12160 of *Lecture Notes in Computer Science*, 49–77. Springer.

Fox, J.; Roughgarden, T.; Seshadhri, C.; Wei, F.; and Wein, N. 2018. Finding Cliques in Social Networks: A New Distribution-Free Model. In *ICALP*, volume 107 of *LIPIcs*, 55:1–55:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Fox, J.; Roughgarden, T.; Seshadhri, C.; Wei, F.; and Wein, N. 2020. Finding Cliques in Social Networks: A New Distribution-Free Model. *SIAM J. Comput.*, 49(2): 448–464.

Gelardi, V.; Godard, J.; Paleressompoulle, D.; Claidiere, N.; and Barrat, A. 2020. Measuring social networks in primates: wearable sensors versus direct observations. *Proceedings of the Royal Society A*, 476(2236): 20190737.

Génois, M.; and Barrat, A. 2018. Can co-location be used as a proxy for face-to-face contacts? *EPJ Data Science*, 7(1): 11.

Haag, R.; Molter, H.; Niedermeier, R.; and Renken, M. 2022. Feedback edge sets in temporal graphs. *Discret. Appl. Math.*, 307: 65–78.

Hermelin, D.; Itzhaki, Y.; Molter, H.; and Niedermeier, R. 2023. Temporal interval cliques and independent sets. *Theor. Comput. Sci.*, 961: 113885.

Himmel, A.; Molter, H.; Niedermeier, R.; and Sorge, M. 2017. Adapting the Bron-Kerbosch algorithm for enumerating maximal cliques in temporal graphs. *Soc. Netw. Anal. Min.*, 7(1): 35:1–35:16.

Kanesh, L.; Madathil, J.; Roy, S.; Sahu, A.; and Saurabh, S. 2023. Further Exploiting c-Closure for FPT Algorithms and Kernels for Domination Problems. *SIAM J. Discret. Math.*, 37(4): 2626–2669.

Kempe, D.; Kleinberg, J. M.; and Kumar, A. 2002. Connectivity and Inference Problems for Temporal Networks. *J. Comput. Syst. Sci.*, 64(4): 820–842.

Kiti, M. C.; Tizzoni, M.; Kinyanjui, T. M.; Koech, D. C.; Munywoki, P. K.; Meriac, M.; Cappa, L.; Panisson, A.; Barrat, A.; Cattuto, C.; et al. 2016. Quantifying social contacts in a household setting of rural Kenya using wearable proximity sensors. *EPJ data science*, 5: 1–21.

Koana, T.; Komusiewicz, C.; Nichterlein, A.; and Sommer, F. 2022. Covering Many (Or Few) Edges with k Vertices in Sparse Graphs. In *STACS*, volume 219 of *LIPIcs*, 42:1–42:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Koana, T.; Komusiewicz, C.; and Sommer, F. 2022. Exploiting c-Closure in Kernelization Algorithms for Graph Problems. *SIAM J. Discret. Math.*, 36(4): 2798–2821.

Koana, T.; Komusiewicz, C.; and Sommer, F. 2023a. Computing Dense and Sparse Subgraphs of Weakly Closed Graphs. *Algorithmica*, 85(7): 2156–2187.

Koana, T.; Komusiewicz, C.; and Sommer, F. 2023b. Essentially Tight Kernels for (Weakly) Closed Graphs. *Algorithmica*, 85(6): 1706–1735.

Koana, T.; and Nichterlein, A. 2021. Detecting and enumerating small induced subgraphs in c-closed graphs. *Discret. Appl. Math.*, 302: 198–207.

Lokshtanov, D.; and Surianarayanan, V. 2021. Dominating Set in Weakly Closed Graphs is Fixed Parameter Tractable. In *FSTTCS*, volume 213 of *LIPIcs*, 29:1–29:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Mertzios, G. B.; Molter, H.; Niedermeier, R.; Zamaraev, V.; and Zschoche, P. 2023. Computing maximum matchings in temporal graphs. *J. Comput. Syst. Sci.*, 137: 1–19.

Moon, J. W.; and Moser, L. 1965. On cliques in graphs. *Israel journal of Mathematics*, 3(1): 23–28.

Ozella, L.; Paolotti, D.; Lichand, G.; Rodríguez, J. P.; Haenni, S.; Phuka, J.; Leal-Neto, O. B.; and Cattuto, C. 2021. Using wearable proximity sensors to characterize social contact patterns in a village of rural Malawi. *EPJ Data Science*, 10(1): 46.

Takaguchi, T. 2015. Analyzing dynamical social interactions as temporal networks. *IFAC-PapersOnLine*, 48(18): 169–174.

Viard, T.; Latapy, M.; and Magnien, C. 2016. Computing maximal cliques in link streams. *Theor. Comput. Sci.*, 609: 245–252.