

10. Basics Concepts of Temporal Graphs

Teacher: Arnaud Casteigts

Assistant: Matteo De Francesco

A temporal graph is a graph whose set of vertices and edges may vary over time. These graphs model various phenomena in the real world, such as social interaction, communication in mobile networks (robots, drones, etc.), scheduling problems, and evolving networks.

In this course, we will focus on a basic model, where only the edges vary. Furthermore, the edges are undirected, and time is discrete. This setting is sufficient for discussing most aspects of temporal graphs. An example of temporal graph is given in fig. 1.




Figure 1: A temporal graph.

The edge labels indicate when the edges are available. For example, the edge ac is present only at time 3, the edge ab is present at times 2 and 5, and so on.

10.1 Definitions

Formally, a temporal graph is defined as an **edge-labeled graph**: $\mathcal{G} = (V, E, \lambda)$, where V is a set of vertices, E is a set of edges ($E \subseteq V \times V$) and $\lambda : E \rightarrow 2^{\mathbb{N}}$ is a function that assigns a nonempty set of time labels to every edge of E . The graph $G = (V, E)$ is called the **footprint** of \mathcal{G} . Equivalently, \mathcal{G} can also be seen as a **sequence of snapshots** $\mathcal{G} = \{G_1, G_2, \dots, G_k\}$, where each snapshot G_i is the graph (V, E_i) with $E_i = \{e \in E \mid i \in \lambda(e)\}$. In other words, G_i contains the edges available at time i . This point of view is illustrated in fig. 2.

Finally, one may also consider \mathcal{G} as a time-ordered **sequence of contacts**, where each contact is a pair (uv, t) such that $t \in \lambda(uv)$, i.e. the edge uv is present at time t . The same graph again corresponds to the following sequence:

$$\mathcal{G} = \{(cd, 1), (bd, 1), (ab, 2), (bd, 2), (cd, 2), (ac, 3), (cd, 4), (de, 4), (ab, 5), (de, 5)\}$$




Figure 2: A temporal graph, seen as a sequence of snapshots.

For most uses, all these representations are equivalent, and we will switch among them depending on the context. In some cases, choosing one or the other may impact the complexity of a problem, but let's ignore this for now.

10.2 Temporal paths and reachability

In a temporal graph, reachability takes place over time. It is possible that two nodes are never connected in a snapshot (e.g. none of the snapshots have a path between a and e), and yet, that they can reach each other using **temporal paths**. Formally, a temporal path is a sequence of contacts $\langle (e_i, t_i) \rangle$ such that $\langle e_i \rangle$ is a path in the footprint of \mathcal{G} , $\langle t_i \rangle$ is non-decreasing¹, and for all i , $t_i \in \lambda(e_i)$. If $\langle t_i \rangle$ is increasing, then the temporal path is **strict**. Here are a few examples:

- $\langle (ec, 2), (ca, 3), (ab, 5) \rangle$ is a strict temporal path from e to b
- $\langle (ac, 3), (cd, 4), (de, 4) \rangle$ is a (non-strict) temporal path from a to e
- $\langle (ac, 3), (ce, 2) \rangle$ is not a temporal path

If a temporal path exists from a node u to a node v , we can write $u \rightsquigarrow v$ to indicate that u can reach v .

Given a temporal graph \mathcal{G} , one can construct a directed graph R such that an arc exists from u to v in R if and only if $u \rightsquigarrow v$. In other words, the graph R encodes the reachability relation among vertices of \mathcal{G} . One may consider two versions of this graph, depending on whether the temporal paths must be strict or not, as illustrated on fig. 3.

Connectivity. If all the nodes can reach each other, i.e. the reachability graph R is a complete directed graph, then \mathcal{G} is **temporally connected**. The class of temporal graphs that are temporally connected is called **TC**. The version of this class for strict temporal connectivity can be noted **TC[>]**. In the above example, $\mathcal{G} \notin \text{TC}$ and $\mathcal{G} \notin \text{TC}^>$. However, the graph in fig. 1 is both in **TC** and in **TC[>]**.

¹In French, non-decreasing means “croissant”, and increasing means “strictement croissant”.




Figure 3: A temporal graph and its reachability.

If a node can reach all the others, then this node is called a **temporal source** (e.g. node a). If a node can be reached by all others, then it is a **temporal sink** (e.g. node e). Clearly, a graph is temporally connected if and only if all its vertices are temporal sources (or equivalently, all of them are temporal sinks). These notions can of course be adapted for strict temporal paths as well.

10.3 Restricted families of labelings

The labeling function λ can be restricted in various ways. The following two restrictions are commonly used.

- A temporal graph $\mathcal{G} = (V, E, \lambda)$ is **proper** if λ is locally injective, which means that two edges incident to a same node cannot have a common label.
- A temporal graph $\mathcal{G} = (V, E, \lambda)$ is **simple** if λ is single-valued, i.e. an edge can have only one label.

Proper temporal graphs are realistic in many contexts. For example, when the graph represents pairwise interactions (e.g. phone calls). It also has the great advantage that the distinction between strict and non-strict temporal paths vanishes (all temporal paths are strict). Simple temporal graphs are harder to motivate practically, but they are convenient for studying the basic features of temporal graphs. Indeed, temporal graphs are typically hard to analyse mathematically, and restricting the labeling helps develop a good understanding of the simpler cases. An example of each kind is shown in fig. 4.




Figure 4: Properness and simpleness.

10.4 Some facts

10.4.1 Time matters more than structure

Consider the temporal graph given in fig. 5 and suppose that its schedule repeats periodically (say, every 5 time units). Looking at the structure of the graph, i.e. its footprint only, node c is clearly the most central, and nodes a and e appear to be equally excentered. However, the role of these nodes changes dramatically when time is added: node a can reach every other node within one full period, whereas node e may take up to 4 periods to reach node a , and node c has an intermediate score at this game.




Figure 5: A temporal graph that repeats every 5 time units.

The conclusion of this simple observation is that classical centrality measures based on the structure of the graph only have very limited relevance. What matters is time!

10.4.2 Reachability is non-symmetric and non-transitive

In contrast to static graphs, the reachability relation in a temporal graph is *non-symmetric*: $u \rightsquigarrow v$ does not imply $v \rightsquigarrow u$. In this respect, temporal graphs are closer in spirit to directed graphs. However, unlike directed graphs, the reachability relation is also *non-transitive*: having $u \rightsquigarrow v$ and $v \rightsquigarrow w$ does not imply $u \rightsquigarrow w$, as illustrated in fig. 6.




Figure 6: Reachability is non-transitive (e.g. $a \rightsquigarrow c$ and $c \rightsquigarrow d$, but $a \not\rightsquigarrow d$)

10.5 Temporal versus structural diameter

The **diameter** of a static graph G is the maximum distance between any pair of vertices. In temporal graphs, this notion can be extended in various ways. Two naive options are the diameter of the *footprint* and the diameter of the *snapshots*. A third one is to define a notion of **temporal diameter** as the maximum time needed for any vertex to reach all the others (e.g., starting at time 0). Clearly, the temporal diameter may differ significantly from the diameter of the footprint. Interestingly, it can also differ significantly from the diameter of the snapshots. Consider a temporal graph $\mathcal{G} = \langle G_1, G_2, G_3, \dots \rangle$ whose snapshots are all connected (class **AC**, for “always connected” temporal graphs). For example:




Figure 7: A temporal graph $\mathcal{G} \in \text{AC}$.

What is the temporal diameter if non-strict temporal paths are allowed? Right. So let's consider strict temporal paths instead. We will show that one can build a graph that prevents at least one node u from reaching all the others until $n - 1$ time steps, and so, even if all the snapshots have diameter 3. The graph is built as follows. At each time step, V is split into two parts: one part contains all the vertices that have already been reached by u (initially, u alone), and the other part contains the others. Each part is a clique, and a single edge is added between the two parts (see fig. 8). Then, the newly informed vertex




Figure 8: Small diameter vs. large temporal diameter.

after G_1 joins the informed part in G_2 , and so on. More generally, each G_i contains a clique of i informed vertices, connected by a single edge to a vertex in a clique of size $n - i$. Clearly, such a process will inform exactly one new vertex in each time step, and terminate only after $n - 1$ time steps.