Structural Lemmas on Temporal Connectivity

(Revisiting gossip literature and beyond)

Arnaud Casteigts

University of Geneva

Dagstuhl seminar 26251

on Temporal Graphs

June 15, 2026

https://arxiv.org/abs/2606.15606

The world is dynamic...

In technologies

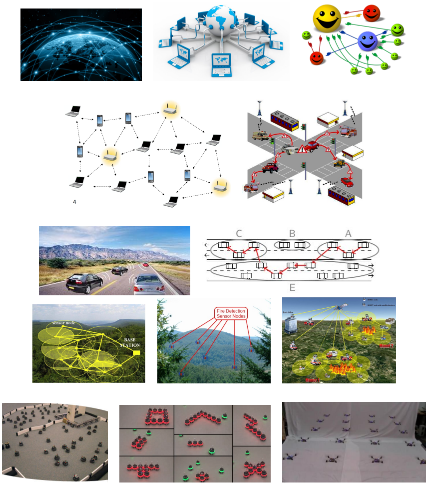

In nature

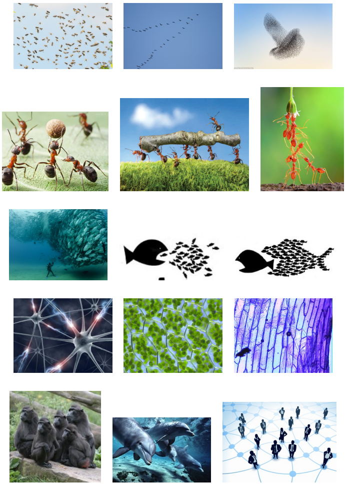

Temporal graphs

(*a.k.a.* time-varying, time-dependent, evolving, dynamic,...)

$\mathcal{G} = (\underline{V, E}, \lambda)$, where $\lambda : E \to 2^{\mathbb{N}}$ assigns *time labels* to edges. *Lifetime* L = largest label.

footprint

Example:

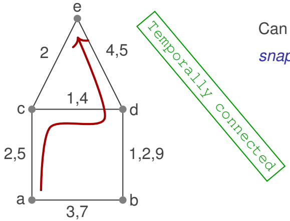

Can also be viewed as a sequence of *snapshots* $\{G_i = \{e \in E : i \in \lambda(e)\}\}$

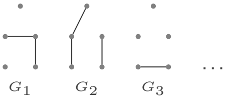

Temporal paths

- e.g. $\langle (a,c,2), (c,d,4), (d,e,4) \rangle$ (*non-strict*)
- e.g. $\langle (a,c,2), (c,d,4), (d,e,5) \rangle$ (*strict*)

Temporal connectivity: Temporal paths between all pairs (class TC).

$\to$ In general, reachability is non-symmetrical... and non-transitive!

Restrictions on the labeling

- Simple: one label per edge $(\lambda : E \to \mathbb{N})$
- Proper: no common label for adjacent edges (λ locally injective)

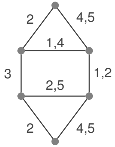

Non-proper & non-simple

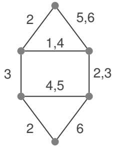

Proper & non-simple

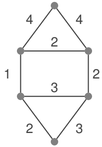

Simple & non-proper

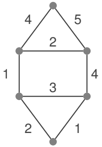

Simple & proper
(= Happy)

Expressiveness of the reachability relation (classes of possible digraphs $\mathcal{R}(\mathcal{G})$):

Reachability relation = digraph $\mathcal{R}(\mathcal{G})$ whose arcs represent the existence of temporal paths.

Classes of possible digraphs:

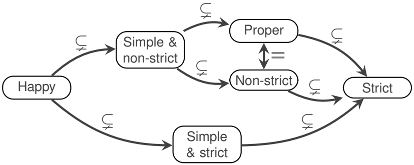

[C.,Corsini, Sarkar, TCS 2024]

How about directed temporal graphs?

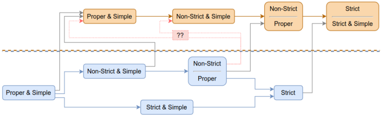

[Döring, ISAAC 2025]

Why caring?

Example: MINIMUM TEMPORAL SPANNER

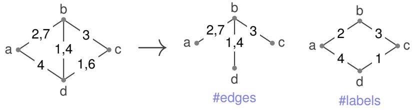

Input: $\mathcal{G} \in$ TC, $k \in \mathbb{N}$
Question: $\exists \mathcal{G}' \subseteq \mathcal{G}$, $|\mathcal{G}'| = k$ and $\mathcal{G}' \in$ TC?

At least 4 natural versions! {strict, non-strict} $\times$ {MIN-EDGE, MIN-LABEL}

Hardness of the problem:

version	because	article
strict $\times$ {MIN-LABEL}	–	[Akrida, Gasieniec, Mertzios, Spirakis, WAOA 2017]
non-strict $\times$ {MIN-EDGE, MIN-LABEL}	Simple	[Axiotis, Fotakis, ICALP 2016]
{strict, non-strict} $\times$ MIN-EDGE	Proper	[C., Corsini, SIROCCO 2024]
all versions at once	Happy	[C., Molter, Zehavi, arxiv 2026]

Key properties:

- Strict and non-strict coincide in proper temporal graphs
- MIN-EDGE and MIN-LABEL coincide in simple temporal graphs

Approved by Pharrell:

Good practice:

- Negative results: Try to show hardness/impossibility in happy graphs! Most general possible.
- Positive results: Start Solving the problem anywhere requires solving it in happy.

Structural Lemmas on Temporal Connectivity

https://arxiv.org/abs/2606.15606

With:

David Schindl

Daniele Carnevale

Main goals:

- Revisit gossip literature
- Extract relevant lemmas
- Make them accessible to our community
- Extend with new lemmas

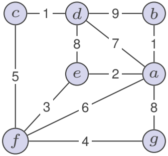

Featuring (among others):

- R Tijdeman - On a telephone problem (1971)
- B. Baker and R. Shostak - Gossips and Telephones (1972)
- D. West - A class of solutions to the gossip problem (1982)
- D. West - Gossiping without duplicate transmission (1982)
- A. Seress - Quick gossiping without duplicate transmissions (1986)
- **R. Labahn - Information flows on hypergraphs (1993)**
- D. Krumme - Reordered gossip schemes (1996)

From this point on, all graphs are **happy**.

(some lemmas hold more generally)

Preliminaries

Some notations:

$e^-(v)$: minimum edge of v. $\qquad E^- = \{e^-(v) : v \in V\}$

$e^+(v)$: maximum edge of v $\qquad E^+ = \{e^+(v) : v \in V\}$

$v \rightsquigarrow u$: $\exists$ temporal walk from v to u.

$v \rightsquigarrow e$: $\exists$ temporal walk starting at v that traverses e

$e \rightsquigarrow v$: $\exists$ temporal walk traversing e that ends at u

$e \rightsquigarrow e'$: $\exists$ temporal walk traversing e and e' in this order (not necessarily consecutively)

Extensions to sets, e.g. $X \rightsquigarrow Y$ means $x \rightsquigarrow y$ for all $x \in X$, $y \in Y$

Lemma 3.1: Basic features

(i) If $e_1 \rightsquigarrow e_2$ and $e_2 \rightsquigarrow e_3$, then $e_1 \rightsquigarrow e_3$. (transitivity works for edges!)

(ii) If $v \rightsquigarrow e_1$ and $e_1 \rightsquigarrow e_2$, then $v \rightsquigarrow e_2$.

(iii) If $v \rightsquigarrow e$ and $e \rightsquigarrow v'$, then $v \rightsquigarrow v'$.

(iv) If $v \rightsquigarrow u$, then $e^-(v) \rightsquigarrow u$; $v \rightsquigarrow e^+(u)$; and $e^-(v) \rightsquigarrow e^+(u)$.

(v) If $\mathcal{G} \in$ TC, then for all $u, v \in V$, $e^-(u) \rightsquigarrow v$; $u \rightsquigarrow e^+(v)$; and $e^-(u) \rightsquigarrow e^+(v)$.

(vi) If e_1 and e_2 are distinct edges, then either $e_1 \not\rightsquigarrow e_2$ or $e_2 \not\rightsquigarrow e_1$.

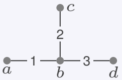

We will often drop the "temporal" adjective.

Time compression

Some graphs are equivalent in a strong sense (temporal walks in bijection).

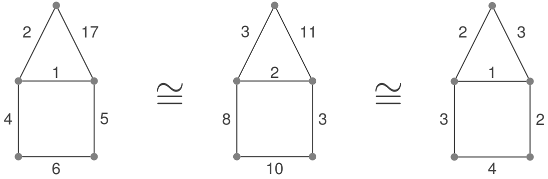

The graph on the right satisfies a contiguity property: If an edge is labeled $t > 1$, then at least one adjacent edge is labeled $t - 1$. We call such graphs time-compressed (or just compressed).

We will consider **compressed** graphs only (without loss of generality)

Some definitions in compressed graphs:

- A label $\lambda(e)$ is unique if for all other edges e', $\lambda(e) \neq \lambda(e')$.

- A graph is full-range if all the labels are unique ($L = |E|$).

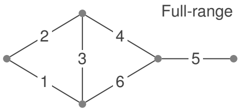

Contiguous caterpillars

Unfold the contiguity property from a given edge (choosing arbitrarily when two choices exist).

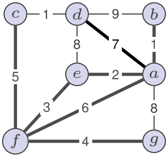

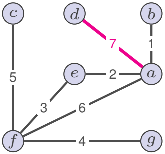

Caterpillar from $(ad, 7)$:

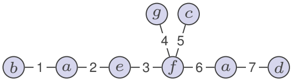

In general, the caterpillar of an edge is not unique (here it was). Vertices may be duplicated (e.g. vertex a).

Lemma 2.4: Basic properties of caterpillars

(i) A caterpillar from an edge e contains exactly one edge for each label in $[1, \lambda(e)]$.

(ii) For all edges e_1, e_2 in a caterpillar, with $\lambda(e_1) < \lambda(e_2)$, we have $e_1 \rightsquigarrow e_2$.

(iii) The lifetime L is the number of edges in the caterpillar of the largest edge $e = \arg\max(\lambda)$.

Example of use:

Lemma 3.2: If $\lambda(e)$ is unique in $\mathcal{G}$ and $\lambda(e) < \lambda(e')$ for some e', then $e \rightsquigarrow e'$

The caterpillar of e' must have an edge with label $\lambda(e)$. As this label is unique, the edge is e, thus $e \rightsquigarrow e'$.

Three interesting properties

Properties:

- Acyclic: no temporal cycle (nontrivial temporal walk from a node to itself).
- Minimal: no edge can be removed without impacting $\mathcal{R}(\mathcal{G})$.
- Two-path-free: For all u, v, at most one temporal walk from u to v.

Each played an important role in gossip theory.

A lot of interesting structure for each $+$ some inclusions:

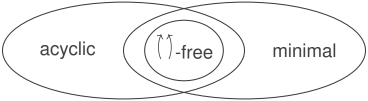

Remaining plan (mostly)

Acyclic | Acyclic TC | Two-path-free | Two-path-free TC | TC | Minimal TC

Remember that all graphs are **happy** (with loss o.g.) and **compressed** (without loss o.g.).

Acyclic graphs

Some examples:

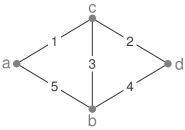

With cycles

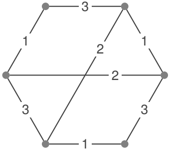

Acyclic

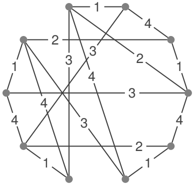

Acyclic

Lemma 3.3: Acyclic $\implies$ footprint is triangle free

Trivial. Any triangle induces a temporal cycle.

Lemma 3.4: Acyclic $\implies$ lifetime $L < n$

Recall that a caterpillar from $e = \arg\max(\lambda)$ has L edges.

If $L \geq n$, then this implies $\geq n + 1$ vertices.

One vertex must repeat.

This vertex has a temporal cycle.

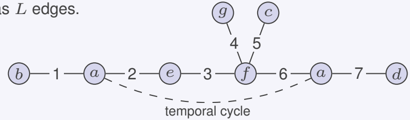

Acyclic TC graphs

Some examples:

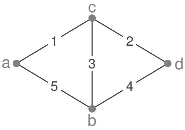

TC

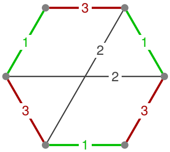

Acyclic TC

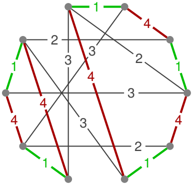

Acyclic TC

Lemma 3.5: Acyclic TC $\implies$ E^- and E^+ are perfect matchings (disjoint if $n > 2$)

All vertices have a minimum edge, so if E^- is a matching, it must be perfect.

Now, if E^- is not a matching, then two min-edges exist as in this picture $\rightarrow$

As the graph is TC, $c \rightsquigarrow a$, either through b or not through b.

In both cases, some composition implies a temporal cycle for a.

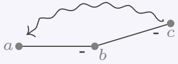

(symmetric arguments for E^+)

$+$ Lemma 3.6 (corollary): Acyclic TC $\implies$ n is even.

Lemma 3.7: Acyclic TC $\implies$ footprint is biconnected (no articulation point)

Let u be an articulation point in the footprint. Let v be the max neighbor of u.

Node v can reach the other side (because TC),

so it can reach u earlier than some left edge of u (themselves $< \lambda(uv)$).

But then it can also return to itself from u using uv.

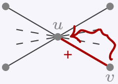

Acyclic TC graphs (2)

We already know that acyclic implies $L < n$. With TC, we get something stronger:

Lemma 3.8: Acyclic TC $\implies$ lifetime $L \leq n/2$

If not, a caterpillar $\mathcal{C}$ exists with at least $n/2 + 1$ edges and $n/2 + 2$ vertices.

As $\mathcal{G}$ is acyclic, the vertices of $\mathcal{C}$ are also distinct in $\mathcal{G}$.

Now augment the caterpillar with the edges from E^- incident to at least one vertex of $\mathcal{C}$ (picture).

At most $n/2$ edges from E^- (because matching) are attached to at least $n/2 + 2$ vertices of $\mathcal{C}$.

At least one edge $e \in E^-$ appears twice in the augmented caterpillar.

The white endpoint of the “first” occurrence of e the blue vertex of the “second” occurrence.

This vertex can reach itself through the caterpillar.

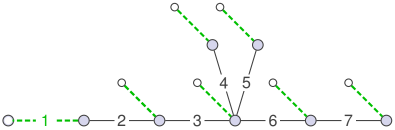

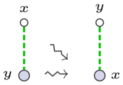

Two-path-free graphs

A graph is two-path-free if for all u, v, at most one temporal walk exists from u to v.

Lemma 3.9: Two-path-free graphs are acyclic

If not acyclic, there is a temporal cycle $v_0, v_1, \ldots, v_0$.

Then v_1 can reach v_0 using either their common edge or the rest of the cycle (contradiction).

Lemma 3.10: Two-path-free graphs are minimal

If not minimal, there is an edge $e = uv$ such that $\mathcal{R}(\mathcal{G}) = \mathcal{R}(\mathcal{G} \setminus e)$.

The edge e is a temporal walk from u to v in $\mathcal{G}$.

The fact that $\mathcal{R}(\mathcal{G}) = \mathcal{R}(\mathcal{G} \setminus e)$ implies $u \rightsquigarrow v$ in $\mathcal{G} \setminus e$ (different walk).

Lemma 3.11: Two-path free have at most $O(n \log n)$ edges (tight on hypercubes)

At any point, define the *predecessors* of a node u as the nodes that already reached u (initially u itself).

Whenever an edge uv occurs, the predecessors or u and v just before must be disjoint. (why?)

Thus every edge doubles the number of predecessors of at least one node.

If the graph has $\omega(n \log n)$ edges, at least one set of predecessors will double $\omega(\log n)$ times (impossible).

Two-path-free TC graphs

Inherits all properties of acyclic TC graphs, plus of course the ones of two-path-free graphs.

In addition, there is a nice *lower* bound on the number of edges if TC by Akos Seress (the bound is tight in the sense that graphs of that size exist).

Lemma 3.12 (Seress'89): Two-path-free TC graphs have at least $2.25n - 6$ edges

The proof is the entire paper, not presented here.

Regarding lifetime, Douglas West (1982) construct an infinite family of two-path-free graphs which (if examined as a compressed happy graph) has lifetime $L = n/4 + 1$.

We conjecture that this is maximum:

Conjecture 3.13: The lifetime of a two-path-free TC graph is at most $L \leq n/4 + 1$

Exhaustively verified up to $n = 12$ (included).

We know that two-path-free graphs are both minimal and acyclic, but the converse does not hold. Still, we conjecture:

Conjecture 3.26: Graphs that are both TC minimal and acyclic have $O(n \log n)$ edges

(Ask me why at the end.)

TC graphs

Two basic facts before we start:

Lemma 3.14 (folklore in gossip): $L \geq \lceil \log_2 \rceil n$.

Seeing $\mathcal{G}$ as a sequence of snapshots $G_1, G_2, \ldots, G_L$. The edges of each G_i are a matching (proper).
The set of nodes reached by any fixed node at most doubles in each snapshot.

Lemma x.y (observation): all nodes are simultaneously alive at some point. (VIM width $= n$.)

TC implies $\max\{\lambda(e) : e \in E^-\} \leq \min\{\lambda(e) : e \in E^+\}$.

In the rest of this section, we dissect TC graphs based mostly on:

- Information flows on hypergraphs, Labahn (1993) (setting $k = 2$)

Reachability among edges (as a poset)

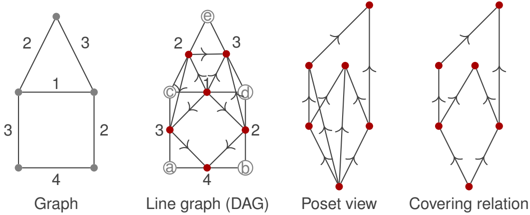

Example:

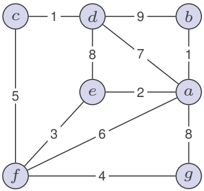

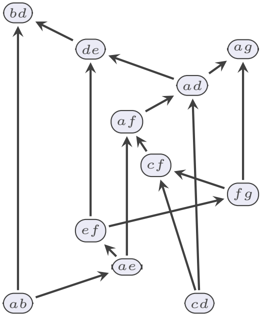

Important types of edges

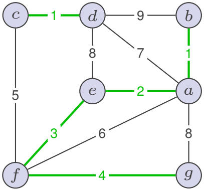

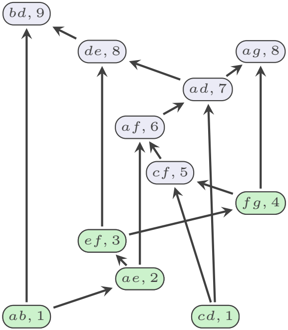

- Minimum edges E^-

Important types of edges

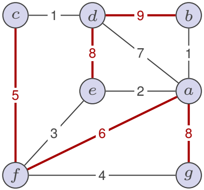

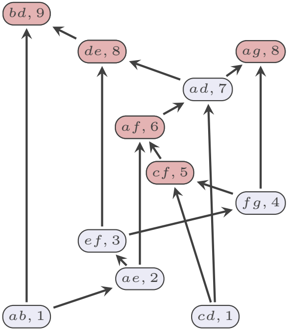

- Minimum edges E^-
- Maximum edges E^+

Important types of edges

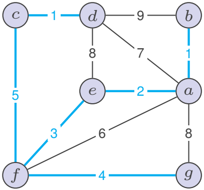

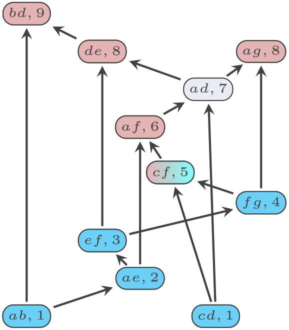

- Minimum edges E^-
- Maximum edges E^+
- Source edges: $\{e \in E \mid e \rightsquigarrow E^+\} \supseteq E^-$

Important types of edges

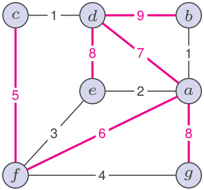

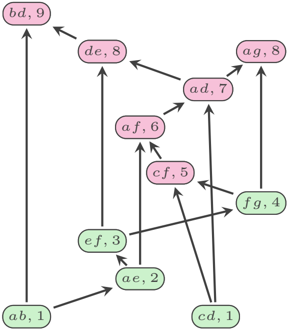

- Minimum edges E^-
- Maximum edges E^+
- Source edges: $\{e \in E \mid e \rightsquigarrow E^+\} \supseteq E^-$
- Sink edges: $\{e \in E \mid E^- \rightsquigarrow e\} \supseteq E^+$

Important types of edges

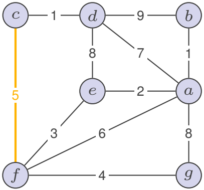

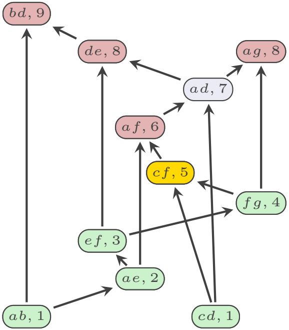

- Minimum edges E^-
- Maximum edges E^+
- Source edges: $\{e \in E \mid e \rightsquigarrow E^+\} \supseteq E^-$
- Sink edges: $\{e \in E \mid E^- \rightsquigarrow e\} \supseteq E^+$
- Pivot edges: $\{e \in E \mid E^- \rightsquigarrow e \rightsquigarrow E^+\}$ (also $= Source \cap Sink$)

Temporal connectivity

(Labahn's decomposition)

- Divergent edges (●)
  (maximal source edges)
- Div-sink edges
  (sink edges that *all divergent* edges can reach)
- Convergent edges (●):
  (minimal div-sink edges)
- Kernel:
  - edges between ● and ● (both included)
- Inner kernel:
  - same (excluding ● and ●)

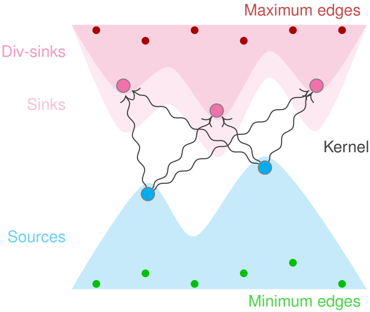

Lemma 3.19 (characterization of TC): A graph is TC if and only if:

- Each vertex (in fact, each source edge, including E^-) can reach at least one ●
- All ● can reach all ●
- Each vertex (in fact, each div-sink edge, including E^+) can be reached by at least one ●

Lemma 3.22 (decomposition): The edges of a TC graph can be partitioned into

- Lower part: edges that can reach at least one ● (them included)
- Inner kernel: edges between a ● and a ●
- Upper part: edges that can be reached by at least one ● (them included)

Minimal TC graphs

Lemma 3.23: Every node reaches exactly one divergent (+ symmetric fact for convergent).

Let's prove something stronger: each node has a single walk reaching a divergent edge.

If not, some u has two choices of local edges to continue (picture).

Let uv be the lowest of the two edges.

Both u and v can reach a divergent without uv, starting *after* $\lambda(uv)$.

Other nodes that wanted to use uv can use either walk instead $\Longrightarrow$ not minimal.

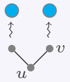

Lemma 3.24: If k divergent, lower part is a forest of $n - k$ edges (same for conv & upper part)

By Lemma 3.23, divergent edges partition V into $V_1, V_2, \ldots, V_k$

It partitions source edges into $E_1, E_2, \ldots, E_k$ (otherwise endpoints of share edge reaches several div).

If (V_i, E_i) has a cycle C, the smallest edge of C can be removed without consequence ($\Rightarrow$ not minimal).

Thus each E_i is a disjoint tree, $n - k$ edges in total.

In fact, LP and UP are forests both in the graph and in the cov. rel.

Lemma 3.25: These edges are needed only in one direction

Towards the divergent edge (resp. from the convergent edge).

Kernel under study... to be continued.

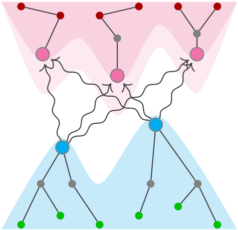

Additional lemmas

Lemma 3.27: Acyclic TC $\Rightarrow$ the divergents are E^-, the convergents are E^+ (the entire graph is the kernel!)

Lemma 3.28: Acyclic TC graphs cannot have a pivot edge (for $n > 2$)

Lemma 3.29: Full-range graphs must have a pivot edge

Lemma 3.30: Minimal TC graphs have at most one pivot edge

Lemma 3.29: In a TC graph, if $e = \min(E^+)$ is unique, then e is a pivot edge

Lemma 3.30: If a TC graph has only one divergent edge or one convergent edge, then it is a pivot edge

+ a few minor ones.

Summary of the main properties

Acyclic:

- Footprint is triangle-free
- Lifetime $L < n$

Acyclic TC:

- n is even
- Footprint is biconnected
- Lifetime $L \leq n/2$
- E^- and E^+ are perfect matchings
- Full kernel (Div = E^-, Conv = E^+)
- (more to come...)

Two-path free:

- Minimal
- Acyclic
- $O(n \log n)$ edges

Two-path free TC:

- $L \leq n/4 + 1$?

TC:

- $L \geq \log_2 n$
- Decomposition Lower part + Inner kernel + Upper part

Minimal TC:

- Lower and upper parts are forests

Minimal TC + acyclic:

- $O(n \log n)$ edges?

TC + acyclic + $2n - 4$: (West'82)

- Footprint is bipartite
- ... regular
- ... hamiltonian

TC + acyclic + minimal:

- Footprint is ~~bipartite~~
- ... ~~regular~~
- ... hamiltonian?

Open questions

Previous conjectures +

- Complexity of deciding acyclic spanner?
- Complexity of deciding two-path free spanner?
- Do cliques admit acyclic or two-path free spanner? (If not, complexity of deciding?)
- Worst size of TC acyclic minimal ? If $\Theta(n^2)$, admit spanners of $o(n^2)$ size?
- Minimum spanner in acyclic TC graphs?
- Minimum spanner in two-path free TC graphs?
- Related parameters ? (some of you already started...)
- Fast way to decide minimal? (deciding acyclic or two-path-free is linear)

One of these questions is pointless, which one? (Exercise :-))

Danke!