

A TIGHT META-THEOREM FOR LOCAL CERTIFICATION OF MSO₂ PROPERTIES WITHIN BOUNDED TREEWIDTH

Eunjung Kim

SoC KAIST and DIMAG IBS, Daejeon, South Korea & CNRS, France

Jointwork with Linda Cook and Tomáš Masařík

LOGALG 2025, 19-21 November 2025, Vienna, Austria

PROOF LABELING SCHEME (PLS)

A *certificate* on a graph G is a mapping φ_G from $V(G)$ to $\{0,1\}^*$. The size of a certificate on G is $\max_{v \in V(G)} |\varphi_G(v)|$.

PROOF LABELING SCHEME

A pair (P, V) is said to be a *proof labeling scheme* for a graph property \mathcal{P} if

- P is an algorithm which outputs a certificate φ_G for each graph G , and
- V is an algorithm whose outputs YES or NO on an input from $\{0,1\}^*$

for which the following holds:

- if $G \in \mathcal{P}$, then $V(\varphi_G(v), \bigcup_{w \in N_G(v)} \varphi_G(w)) = \text{YES}$ for every $v \in V(G)$, and
- if $G \notin \mathcal{P}$, then $V(\varphi_G(v), \bigcup_{w \in N_G(v)} \varphi_G(w)) = \text{NO}$ for some $v \in V(G)$.

PROOF LABELING SCHEME (PLS)

Intuitively, the prover P computes a certificate φ_G and assigns to each vertex $v \in V(G)$ the *local* certificate $\varphi_G(v)$.

After a one round of *synchronized* communication between neighboring vertices, the verifier V at v outputs YES or NO based on its computation on the instance consisting of $\varphi_G(v)$ and $(\varphi_G(w))_{w \in N_G(v)}$.

Each vertex is equipped with a unique *identifier* as a $\{0, 1\}$ -string of length $\text{poly}(n)$, where n is the number of vertices.

It is assumed that the underlying graph G is connected.

PLS for { all bipartite graphs }.

- P outputs a certificate $\varphi_G : V(G) \rightarrow \{\text{red}, \text{blue}\}$; φ_G is a proper 2-coloring of G if one exists, a random mapping otherwise.
- V at a vertex v outputs YES if $\varphi_G(w)$ differs from $\varphi_G(v)$ for every $w \in N_G(v)$; NO otherwise.

PLS for { all acyclic graphs }.

- If G is a tree, P chooses an arbitrary vertex r as the root and computes φ_G such that for each vertex v , $\varphi_G(v) = \text{dist}(v, r)$.
- V at v outputs YES if either $\text{dist}(v, r) = 0$ or the following holds:
 - 1 there is a unique neighbor w such that $\text{dist}(w, r) \leq \text{dist}(v, r)$ and it holds that $\text{dist}(w, r) = \text{dist}(v, r) + 1$, and
 - 2 for every other neighbor z , it holds that $\text{dist}(z, r) = \text{dist}(v, r) + 1$.

Property	upper bound	ref
H -minor-free for small H	$O(\log n)$	BFP'21
planarity	$O(\log n)$	FFMRRT'21
bounded genus	$O(\log n)$	EL'22
treedepth at most k	$O(\log n)$	FBP'22
treewidth at most k	$O(\log^2 n)$	FMRT'22
cographs	$O(\log n)$	FMMRT'23
cliquewidth at most k	$O(\log^2 n)$	FMMRT'23

OUR MAIN RESULT

THEOREM (COOK, K. MASAŘÍK 2025)

There is an approximate PLS for graphs of treewidth at most k of size $O(\log n)$; i.e. there exists a computable function f for which P computes φ_G of size $O(\log n)$ such that

- *if $\text{TW}(G) \leq k$, then each V outputs YES, and*
- *if $\text{TW}(G) > f(k)$, then some V outputs NO.*

THEOREM (FRAIGNIAUD, MONTEALEGRE, RAPAPORT, TODINCA 2022, COOK, K. MASAŘÍK 2025)

For each integer t and MSO_2 -sentence ϕ , there is a $O(\log n)$ -size PLS for an MSO_2 -definable property on graphs of bounded treewidth; with promise of $\text{TW}(G) \leq t$, the prover P computes φ_G such that

- *if $G \models \phi$, then each V outputs YES, and*
- *if $G \not\models \phi$, then some V outputs NO.*

ELIMINATION TREE AND WIDTH

A rooted tree F is an *elimination tree* of a (simple) graph G if $V(F) = V(G)$ and for every $uv \in E(G)$, u is a strict ancestor of v or vice versa.

The *width* of an elimination tree F of G is defined as

$$\text{WIDTH}(F) = \max_{v \in V(G)} \{w \in V(G) \mid w \text{ is an ancestor of } v \text{ and } w \text{ is adjacent with } F_v\} - 1$$


WITNESS FOR TREEWIDTH AT MOST t

ELIMINATION TREE AND WIDTH

A rooted tree F is an *elimination tree* of a (simple) graph G if $V(F) = V(G)$ and for every $uv \in E(G)$, u is a strict ancestor of v or vice versa.

The *width* of an elimination tree F of G is defined as

$$\text{WIDTH}(F) = \max_{v \in V(G)} \{w \in V(G) \mid w \text{ is an ancestor of } v \text{ and } w \text{ is adjacent with } F_v\} - 1$$



ELIMINATION TREE AND WIDTH

A rooted tree F is an *elimination tree* of a (simple) graph G if $V(F) = V(G)$ and for every $uv \in E(G)$, u is a strict ancestor of v or vice versa.

The *width* of an elimination tree F of G is defined as

$$\text{WIDTH}(F) = \max_{v \in V(G)} \{w \in V(G) \mid w \text{ is an ancestor of } v \text{ and } w \text{ is adjacent with } F_v\} - 1$$

ARNBORG 1985, BOADLAENDER ET AL. 1991, BOJAŃCZYK AND PILIPCZUK 2016

A graph G has treewidth at most t if and only if G admits an elimination tree of width at most t .

Toy case: suppose that G admits an elimination tree F of width at most t such that every edge of F is an (actual) edge of G .

CERTIFICATE φ_G ASSIGNED BY THE PROVER

If G admits an elimination tree F of width at most t , the local certificate $\varphi_G(v)$ carries the following information for each $v \in V(G)$.

- 1 the distance $\text{dist}_F(v, \text{root})$.
- 2 the list L_v of all strict ancestors of v with a neighbor in F_v , together with their distances to the root.

TOY CASE AND BOTTLENECK FOR TREEWIDTH AT MOST t

Toy case: suppose that G admits an elimination tree F of width at most t such that every edge of F is an (actual) edge of G .

VERIFIER'S ALGORITHM AT EACH $v \in V(G)$

Verifier V at each v checks the following.

- 1 Ensure that there is no neighbor u with $\text{dist}_F(u, \text{root}) = \text{dist}_F(v, \text{root})$.
- 2 Ensure that there is a unique neighbor w with $\text{dist}_F(w, \text{root}) = \text{dist}_F(v, \text{root}) + 1$ (or $\text{dist}_F(v, \text{root}) = 0$).
- 3 Whenever its neighbor z is an ancestor, ensure that z is the list L_v .
- 4 Ensure that $|L_v| \leq t$.
- 5 Ensure that \bigcup_w is a neighbor and strict descendant of v $L_w \subseteq L_v \cup \{v\}$.

TOY CASE AND BOTTLENECK FOR TREEWIDTH AT MOST t

Toy case: suppose that G admits an elimination tree F of width at most t such that every edge of F is an (actual) edge of G .

↪ We cannot such an elimination tree in general.

BLANCO, COOK, HATZEL, HILAIRE, ILLINGWORTH, MCCARTY 2024

For every function f , there exists a graph G which does not admit a tree-decomposition (T, β) of width $f(\text{TW}(G))$ such that T is a minor of G .

For an edge uv in the elimination treewidth (v being the parent of u), we want to have a (u, v) -path in G so as to use it for “channeling” the oriented edge from u to v .

TOY CASE AND BOTTLENECK FOR TREEWIDTH AT MOST t

Toy case: suppose that G admits an elimination tree F of width at most t such that every edge of F is an (actual) edge of G .

\leadsto We cannot such an elimination tree in general.

BLANCO, COOK, HATZEL, HILAIRE, ILLINGWORTH, MCCARTY 2024

For every function f , there exists a graph G which does not admit a tree-decomposition (T, β) of width $f(\text{TW}(G))$ such that T is a minor of G .

For an edge uv in the elimination treewith (v being the parent of u), we want to have a (u, v) -path in G so as to use it for “channeling” the oriented edge from u to v .

ORIENTED PATH SYSTEM WITNESSING AN ELIMINATION TREE

Let F be an elimination tree of G . A collection \mathcal{P} of directed path system is said to witness F if, for every vertex u and its parent v in F , there is some (u, v) -path of G oriented from u toward v in \mathcal{P}

The *congestion* of \mathcal{P} is

$$\max_{v \in V(G)} \# \text{ of oriented paths in } \mathcal{P} \text{ using } v \text{ as a start or internal vertex.}$$

KEY TECHNICAL LEMMA [COOK, K. MASAŘÍK 2024 + BOJAŃCZYK, PILIPCZUK 2015]

For any graph G of treewidth at most t , there exists

- an elimination tree F of G of width at most $f(t)$, and
- an oriented path system \mathcal{P} of congestion $f(t)$ witnessing F .

Suppose that G admits an elimination tree F of width $f(t)$ and an oriented path system \mathcal{P} of congestion at most $f(t)$ witnessing F .

CERTIFICATE φ_G ASSIGNED BY THE PROVER

$\varphi_G(v)$ carries the following information for each $v \in V(G)$.

- ① the distance $\text{dist}_F(v, \text{root})$.
- ② the parent u of v .
- ③ the list L_v of all strict ancestors of v with a neighbor in F_v , together with their distances to the root.
- ④ **channel for (s, t) -path in \mathcal{P} :** when v is on an oriented path $P \in \mathcal{P}$ as a start or internal vertex: the start s and final vertex t of P (together with their distances to the root), the predecessor and successor of v , the list L_s of the start vertex.

VERIFIER'S ALGORITHM AT EACH $v \in V(G)$: CHANNEL WORKS PROPERLY

Ensure the **channel for (s, t) -path P in \mathcal{P} is working properly**

- 1 Ensure that the (s, t) -channel is proper: $\text{dist}_F(s, \text{root}) = \text{dist}_F(t, \text{root}) - 1$.
- 2 Ensure that the predecessor / successor of v are neighbors of v and their channel information is consistent with what v knows.
- 3 If $v = s$ for some (s, t) -channel, ensure there is only one such channel.
- 4 If v has s as a predecessor, ensure $\text{dist}_F(s, \text{root})$ matches what v knows.
- 5 If v has t as a successor, ensure $\text{dist}_F(t, \text{root})$ matches what v knows.
- 6 If v has a neighbor which carries an (s, t) -channel with $v = t$, ensure that $\text{dist}_F(s, \text{root}) = \text{dist}_F(v) - 1$.

VERIFIER'S ALGORITHM AT EACH $v \in V(G)$: ELIMINATION TREE

Assuming that the **channel for (s, t) -path P in \mathcal{P} is working properly**, the presumed elimination tree is verified.

- 1 Ensure that there is no neighbor u with $\text{dist}_F(u, \text{root}) = \text{dist}_F(v, \text{root})$.
- 2 Ensure that there is a unique **target of a channel starting with v** (or $\text{dist}_F(v, \text{root}) = 0$).
- 3 Whenever its neighbor (including its "parent" as a target of a channel starting with v) z is an ancestor, ensure that z is the list L_v
- 4 Ensure that $|L_v| \leq t$.
- 5 Ensure that \bigcup_w is a neighbor and strict descendant of v $L_w \subseteq L_v \cup \{v\}$.

HOW TO GET A LOW-CONGESTION PATH SYSTEM

KEY TECHNICAL LEMMA

For any graph G of treewidth at most t , there exists an elimination tree F of G of width at most $f(t)$, and an oriented path system \mathcal{P} of congestion $f(t)$ witnessing F .

The proof relies on some key technical results from Bojańczyk and Pilipczuk (2016): sane tree-decomposition and some consequence of Simon's factorization forest theorem.

SANE TREE-DECOMPOSITION

If G has treewidth at most t , then it admits a **sane** tree-decomposition (T, β) of width at most t . That is, for every node t with parent t' ,

- $\beta(t) \setminus \beta(t') \neq \emptyset$,
- $Y_t := \bigcup_{b \in T_t} \beta(b) \setminus \beta(t')$ is connected, and
- every vertex in $\beta(t) \cap \beta(t')$ is adjacent with some vertex of Y_t .


From a sane tree-decomposition of width at most t , we can construct an elimination tree F of width at most t together with an oriented path system witnessing F .

HOW TO GET A LOW-CONGESTION PATH SYSTEM

FROM SANE TREE-DECOMPOSITION TO ELIMINATION TREE

Observe that for each node t of sane (T, β)

- the graph $G[\beta(t) - \text{adh}(t)]$ obtained by “torsofying with lower bags” makes each $\text{adh}(t')$ a clique for each child t' of t .
- Build an elimination tree F by combining an elimination tree F_t as DFS tree for each the above “marginal graph” for each t .
- Choose the root of F_t as a neighbor of the “largest vertex in $\text{adh}(t)$ ”.




HOW TO GET A LOW-CONGESTION PATH SYSTEM

FROM SANE TREE-DECOMPOSITION TO ELIMINATION TREE

Observe that for each node t of sane (T, β)

- the graph $G[\beta(t) - \text{adh}(t)]$ obtained by “torsofying with lower bags” makes each $\text{adh}(t')$ a clique for each child t' of t .
- Build an elimination tree F by combining an elimination tree F_t as DFS tree for each the above “marginal graph” for each t .
- Choose the root of F_t as a neighbor of the “largest vertex in $\text{adh}(t)$ ”.




HOW TO GET A LOW-CONGESTION PATH SYSTEM

FROM SANE TREE-DECOMPOSITION TO ELIMINATION TREE

Observe that for each node t of sane (T, β)

- the graph $G[\beta(t) - \text{adh}(t)]$ obtained by “torsofying with lower bags” makes each $\text{adh}(t')$ a clique for each child t' of t .
- Build an elimination tree F by combining an elimination tree F_t as DFS tree for each the above “marginal graph” for each t .
- Choose the root of F_t as a neighbor of the “largest vertex in $\text{adh}(t)$ ”.



HOW TO GET A LOW-CONGESTION PATH SYSTEM

From a sane tree-decomposition of width at most t , we can construct an elimination tree F of width at most t together with an oriented path system witnessing F .

Reducing the congestion needs much more work: core technical work done in Bojańczyk and Pilipczuk (2016); the main result was to MSO-transduce tree-decompositions of bounded width from graphs of bounded treewidth.

It was done with a nice combinatorial analysis of sane tree-decomposition, and **Simon's factorization theorem** applied to graphs of bounded pathwidth.

Our work for reducing the congestion builds on this.

Proof labeling scheme for the following properties are open, among others.

- $O(\log n)$ -sized PLS for cliquewidth / rankwidth at most t ? Open for linear cliquewidth at most t as well.
- $O(\log n)$ -sized PLS for H -minor-freeness, for any fixed H ?
Known: H planar, of size at most K , H -minor-free being bounded genus, etc.