Brief Announcement: Distributed Model Checking on Graphs of Bounded Treedepth

Fedor Fomin
fedor.fomin@uib.no
University of Bergen
Bergen, Norway

Pierre Fraigniaud
pierre.fraigniaud@irif.fr
Université Paris Cité and CNRS
Paris, France

Pedro Montealegre
p.montealegre@uai.cl
Universidad Adolfo Ibáñez
Santiago, Chile

Ivan Rapaport
rapaport@dim.uchile.cl
Universidad de Chile
Santiago, Chile

Ioan Todinca
ioan.todinca@univ-orleans.fr
Université d'Orléans and INSA
Centre-Val de Loire
Orléans, France

ABSTRACT

We establish that every monadic second-order logic (MSO) formula on graphs with bounded treedepth is decidable in a constant number of rounds within the CONGEST model. To our knowledge, this marks the first meta-theorem regarding distributed model-checking. Various optimization problems on graphs are expressible in MSO. Examples include determining whether a graph G has a clique of size k, whether it admits a coloring with k colors, whether it contains a graph H as a subgraph or minor, or whether terminal vertices in G could be connected via vertex-disjoint paths. Our meta-theorem significantly enhances the work of Bousquet et al. [PODC 2022], which was focused on *distributed certification* of MSO on graphs with bounded treedepth. Moreover, our results can be extended to solving optimization and counting problems expressible in MSO, in graphs of bounded treedepth.

CCS CONCEPTS

• **Theory of computation → Distributed algorithms**; **Verification by model checking**.

ACM Reference Format:

Fedor Fomin, Pierre Fraigniaud, Pedro Montealegre, Ivan Rapaport, and Ioan Todinca. 2024. Brief Announcement: Distributed Model Checking on Graphs of Bounded Treedepth. In *ACM Symposium on Principles of Distributed Computing (PODC '24), June 17–21, 2024, Nantes, France.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3662158.3662793

1 INTRODUCTION

Distributed *decision* and distributed *certification* are two complementary fields of distributed computing, closely associated with distributed fault-tolerant computing. Both fields are addressing the problem of checking whether a distributed system is in a legal state with respect to a given specification, or not. We examine this problem in the classical context of distributed computing in networks, under the standard CONGEST model. Recall that this model assumes networks modeled as simple connected n-node graphs, in which every node is provided with an identifier on $O(\log n)$ bits that is unique in the network. Computation proceeds synchronously as a sequence of *rounds*. At each round, every node sends a message to each of its neighbors in the graph, receives the messages sent by its neighbors, and performs some individual computation. A crucial point is that messages are restricted to be of size $O(\log n)$ bits.

Distributed Decision. Given a boolean predicate Π on graphs, e.g., whether the graph is H-free for some fixed graph H, a *decision* algorithm for Π takes as input a graph $G = (V, E)$, and outputs whether G satisfies Π or not. Specifically, every node v receives as input its identifier $\mathrm{id}(v)$, and, after a certain number of rounds of communication with its neighbors, it outputs *accept* or *reject*, under the constraint that G satisfies Π if and only if the output of each of the nodes $v \in V$ is accept. In other words,

$$G \models \Pi \iff \forall v \in V(G),\ \mathrm{out}(v) = \mathrm{accept}.$$

Some predicates are easy to decide *locally*, i.e., in a constant number of rounds. A canonical example is checking whether the (connected) graph G is regular, for which one round suffices. However, other predicates cannot be checked locally, with canonical example checking whether there is a unique node of degree 3 in the network. Indeed, checking this property requires $\Omega(n)$ rounds in networks of diameter $\Theta(n)$, as two nodes of degree 3 may be at arbitrarily large distances in the graph. Another example of a difficult problem is checking whether the graph is C_4-free, i.e., does not contain a 4-cycle as a subgraph, which requires $\tilde{\Omega}(\sqrt{n})$ rounds [4]. One way to circumvent the difficulty of local checkability, i.e., to address graph predicates requiring a large number of rounds for being decided, is to consider distributed *certification.*

Distributed Certification. A *certification scheme* for a boolean predicate Π is a pair *prover-verifier.* The prover is a centralized, computationally unbounded, non-trustable oracle. Given a graph $G = (V, E)$, the prover assigns a *certificate* $c(v) \in \{0, 1\}^\star$ to each node $v \in V$. These certificates are forged by the prover using the complete knowledge of the graph G. The verifier is a distributed 1-round algorithm. Each node v takes as sole input its identifier $\mathrm{id}(v)$ and its

PODC '24, June 17–21, 2024, Nantes, France

ACM ISBN 979-8-4007-0668-4/24/06…$15.00
https://doi.org/10.1145/3662158.3662793

certificate $c(v)$. In particular, for distributed decisions, v is unaware of the graph G. Every node v just exchanges once its identifier and certificate with its neighbors, and then it must output *accept* or *reject.*

The certification scheme is correct if the following two conditions hold. The *completeness* condition states that if G satisfies Π, then the oracle can provide the nodes with certificates that they all accept. The *soundness* condition says that if G does not satisfy Π, then no matter the certificates assigned by the oracle to the nodes, at least one of them rejects. That is, the role of the verifier is to check that the collection of certificates assigned to the nodes by the prover is indeed a global proof that the graph satisfies the predicate. In other words,

$$G \models \Pi \iff \exists c : V(G) \to \{0,1\}^\star \ : \ \forall v \in V(G),\ \text{out}(v) = \text{accept}.$$

The main measure of complexity of a certification scheme is the maximum *size* of the certificates assigned by the prover to the nodes on legal instances, i.e., for graphs G satisfying the given predicate. Ideally, to be implemented in a single round under the CONGEST model, the certificates should be of size $O(\log n)$ bits. Interestingly, many graph properties can be certified with such short certificates, including acyclicity [15], planarity [8], bounded genus [6, 9], etc. On the other hand, basic graph properties require large certificates, including diameter 2 vs. 3 (requiring $\tilde{\Omega}(n)$-bit certificates [2]), non-3-colorability (requiring $\tilde{\Omega}(n^2)$-bit certificates [13]), C_4-freeness (requiring $\tilde{\Omega}(\sqrt{n})$-bit certificates [4]), etc. The following question was thus raised, under different formulations (see, e.g., [7]): *What are the graph properties that admit certification schemes with $O(\log n)$-bit certificates, or, to the least, certificates of polylogarithmic size?* Answering this question requires formalizing the notion of "graph predicate".

Monadic Second-Order Logic. Recall that, in the *first-order* logic (FO) of graphs, a graph property is expressed as a quantified logical sentence whose variables represent vertices, with predicates for equality (=) and adjacency (adj). An FO formula is therefore constructed according to the following set of rules, where x and y are vertices, and φ and ψ are FO formulas:

$$x = y \mid \mathtt{adj}(x,y) \mid \varphi \lor \psi \mid \varphi \land \psi \mid \neg\varphi \mid \exists x\varphi \mid \forall x\varphi.$$

The *monadic second-order* logic (MSO) extends FO by allowing quantification on *sets* of vertices and edges, with the incidence predicate $\mathtt{inc}(v,e)$ indicating whether vertex v is incident to edge e, and the membership ($\in$) predicate. Since FO can express properties such as C_4-freeness, which, as mentioned before, requires certificates on $\tilde{\Omega}(\sqrt{n})$ bits, there is no hope of establishing a *meta-theorem* about FO regarding compact certification in all graphs. Nevertheless, a breakthrough in the theory of distributed certification was recently obtained by Bousquet, Feuilloley, and Pierron [7], who showed that every MSO predicate admits a distributed certification scheme with $O(\log n)$-bit certificates in the family of graphs with bounded *treedepth.*

Algorithmic Meta-Theorems. A vibrant line of research in sequential computing is the development of algorithmic meta-theorems. According to Grohe and Kreutzer [14], algorithmic meta-theorems assert that certain families of algorithmic problems, typically defined by some logical and combinatorial conditions, can be solved efficiently under some suitable definition of this term. Such theorems play an essential role in the theory of algorithms as they reveal a profound interplay between algorithms, logic, and combinatorics. One of the most celebrated examples of a meta-theorem is Courcelle's theorem, which asserts that graph properties definable in MSO are decidable in linear time on graphs of bounded treewidth [3].

Bousquet, Feuilloley, and Pierron in [7] introduced the exploration of algorithmic meta-theorems in distributed computing. Their primary result in this direction is that any MSO formula can be locally *certified* on graphs with bounded treedepth using a logarithmic number of bits per node, which represents the golden standard in certification. This theorem has numerous consequences for certification — for more details, we refer to [7]. Notably, the FO property C_4-freeness, and the MSO property non-3-colorability, which both necessitate certificates of polynomial size in general, can be certified with just $O(\log n)$-bit certificates in graphs of bounded treedepth. Bousquet et al.'s result has been extended to more comprehensive classes of graphs, including graphs excluding a small minor [1], as well as graphs of bounded *treewidth*, and graphs of bounded *cliquewidth.* However, the last two extensions come both at the cost of slightly larger certificates, of $O(\log^2 n)$ bits, as seen in [11] and [10], respectively.

With significant advances in developing meta-theorems for distributed *certification*, there's a notable absence of similar results for distributed *decision.* It prompts a natural question: could such results be obtained for the round-complexity of CONGEST? More concretely, the fundamental inquiry that remains unaddressed by Bousquet et al.'s paper, and by the subsequent works regarding distributed certification of MSO predicates is:

> **Question.** *What is the round-complexity in* CONGEST *of deciding MSO formulas in graphs of bounded treedepth?*

A first step in answering this question was proposed in [17] where it is stated that, in any graph class of treedepth at most d, for every fixed connected graph H, H-freeness can be decided in $O(1)$ rounds in CONGEST. In this paper, we offer a comprehensive answer to the question. To elucidate our results, we first need to define the treedepth of a graph.

Treedepth. For any non-negative integer d, a (connected) graph G has treedepth at most d if there exists a rooted tree T spanning the vertices of G, with depth at most d, such that, for every edge $\{u,v\}$ in G, u is an ancestor of v in T, or v is an ancestor of u in T. The treedepth of a graph G, denoted by $\mathtt{td}(G)$, is the smallest d for which such a tree exists.

The class of graphs with bounded treedepth, i.e., of treedepth d for some fixed $d \geq 0$, has strong connections with minor-closed families of graphs. Specifically, for any family $\mathcal{F}$ of graphs closed under taking graph minors, the graphs in $\mathcal{F}$ have bounded treedepth if and only if $\mathcal{F}$ does not include all the paths [16]. Similarly, the graphs with bounded treedepth have a finite set of forbidden induced subgraphs, and any property of graphs monotonic with respect to induced subgraphs can be tested in polynomial time on graphs of bounded treedepth [16]. Computing the treedepth of a graph is

NP-hard, but since treedepth is monotonic under graph minors, it is fixed-parameter tractable (FPT) [12]. Last but not least, MSO and FO have the same expressive power in graph classes of bounded treedepth [5].

2 OUR RESULTS

Distributed Model Checking. We prove that, for every MSO formula φ, there is an algorithm $\mathcal{A}$ that, for every n-node graph G, decides whether $G \models \varphi$ in $O(2^{2^{\mathrm{td}(G)}})$ rounds in the CONGEST model. That is, the round-complexity of $\mathcal{A}$ depends only on the treedepth of the input graph, and on the MSO formula, i.e., it does not depend on the *size* n of the graph. Thus $\mathcal{A}$ performs a constant number of rounds in any class of graphs with bounded treedepth. In particular, deciding non-3-colorability can be done in $O(1)$ rounds in graphs of bounded treedepth, in contrast to general graphs, for which deciding non-3-colorability requires a polynomial number of rounds by [13]. Our meta-theorem is essentially the best that one may hope regarding distributed model checking MSO formulas in a constant number of rounds in CONGEST. Indeed, the FO predicate "there is at least one vertex of degree > 2" requires $\Omega(n)$ rounds to be checked in this class. Hence our theorem cannot be extended to graphs of bounded treewidth or bounded cliquewith, actually not even to bounded pathwith, and not even to the class $\mathcal{P} \cup \mathcal{B}$ where $\mathcal{P}$ is the set of all paths, and $\mathcal{B}$ is the set of all graphs composed of a path to which is attached a claw at one of its endpoints.

Labeled Graphs. We also consider distributed model checking of *labeled* graphs. For instance, one can check whether a given set of vertices is a feedback vertex set, i.e., whether the graph obtained by removing this set of vertices is acyclic. For such a predicate, it is sufficient to add a unary predicate to the logical structure used to mark the nodes, say $\mathsf{mark}(x) = true$ means that vertex x is in the set. Using this unary predicate, φ can express the fact that there are no cycles in G passing only trough nodes x for which $\mathsf{mark}(x) = false$. Since we also deal with MSO, we can also label edges. For instance, one can check whether a given set of edges forms a spanning tree. Indeed, it is sufficient to introduce a unary predicate used to mark the edges: $\mathsf{mark}(e) = true$ means that edge e is in the set. As for feedback vertex set, using this unary predicate, φ can express the fact that the set of marked edges is a spanning tree (i.e., every node is incident to at least one marked edge, and any two vertices are connected by a path of marked edges). We show that deciding MSO formulas on *labeled* graphs of bounded treedepth can be done in $O(1)$ rounds in the CONGEST model.

Optimization. More generally, we also consider the *optimization* variants of decision problems expressible in MSO on graphs of bounded treedepth. For instance, an independent set can be expressed as an MSO formula with a free variable S, such as $\varphi(S) = \forall x \in S\ \forall y \in S\ \neg\mathsf{adj}(x, y)$. Then, $max\varphi$, i.e., maximum independent set, consists in, given any graph $G = (V, E)$, finding the largest set $S \subseteq V$ such that $G \models \varphi(S)$. We show that, for every MSO formula $\varphi(S)$ with free variable $S \subseteq V$ or $S \subseteq E$, there is an algorithm for graphs of bounded treedepth solving $max\varphi$ (and $min\varphi$) in a constant number of rounds in the CONGEST model. This constant is of the form $O(g(\mathrm{td}(G), \varphi))$ for some function g. Due to the expressive power of MSO, our results yield algorithms with a constant number of rounds in the CONGEST model on graphs of bounded treedepth for numerous popular optimization problems including minimum vertex cover, minimum feedback vertex set, minimum dominating set, maximum independent set, maximum induced forest, maximum clique, maximum matching, minimum spanning tree, Hamiltonian cycle, cubic subgraph, planar subgraph, Eulerian subgraph, Steiner tree, disjoint paths, min-cut, minor and topological minor containment, rural postman, k-colorability, edge k-colorability, partition into k cliques, and covering by k cliques. Importantly, we also extend our results to *counting* problems, such as counting triangles or perfect matchings.

Graphs with Bounded Expansion. Finally, we obtained some applications of our results to much larger classes of graphs, namely graphs of bounded *expansion* (see [16] for an extended introduction). Graphs of bounded expansion include planar graphs, and more generally, all classes of graphs defined from excluding minor. It was shown [17] that, for every class $\mathcal{G}$ of graphs with bounded expansion, and every positive integer p, there is an algorithms performing in $O(\log n)$ rounds under the CONGEST model that partitions the vertex set V of any graph $G = (V, E) \in \mathcal{G}$ into $f(p)$ parts $V_1, \dots, V_{f(p)}$ such that every collection $V_{i_1}, \dots, V_{i_q}$ of at most p parts, $1 \le q \le p$, $\{i_1, \dots, i_q\} \subseteq \{1, \dots, f(p)\}$, induces a (not necessarily connected) subgraph of G with treedepth at most p. The function f solely depends on the considered class $\mathcal{G}$ of bounded expansion. The vertex partitioning $V_1, \dots, V_{f(p)}$ is called a low treedepth decomposition with parameter p. Plugging in our techniques into this framework, we show that, for every connected graph H, H-freeness can be decided in $O(\log n)$ rounds under the CONGEST model in any class of graphs with bounded expansion. This result was claimed in [17] with no proofs. We provide that claim with a complete formal proof.

3 OPEN PROBLEM

There might exist some fragments of FO that could be tractable on graphs of bounded expansion in the distributed setting. It would be interesting to identify the exact boundaries of intractability in this context, regarding both distributed decision, and distributed certification. An initial step in this direction was taken in [17], resulting in a distributed algorithm for computing a low treedepth decomposition of graphs of bounded expansion, running in $O(\log n)$ rounds under CONGEST. As we pointed out, this result enables to efficiently decide FO-expressible decision problems (such as H-freeness, for H connected) in classes of graphs with bounded expansion, in $O(\log n)$ rounds. We restate a question stated in [17]: Given a *local* FO formula $\varphi(x)$, i.e., a formula where $\varphi(x)$ depends on a fixed-radius neighborhood of vertex x only, can we mark all vertices satisfying φ in $O(\log n)$ rounds?

ACKNOWLEDGMENTS

Fedor Fomin received funding from the Research Council of Norway via the project BWCA (grant no. 314528); Pierre Fraigniaud received additional supports from ANR projects DUCAT (ANR-20-CE48-0006) and QuDATA (ANR-18-CE47-0010); Pedro Montealegre received funding from FONDECYT 1230599, and Ivan Rapaport from FB210005, BASAL funds for centers of excellence from ANID-Chile, and FONDECYT 1220142.

REFERENCES

- [1] Nicolas Bousquet, Laurent Feuilloley, and Théo Pierron. 2021. Local Certification of Graph Decompositions and Applications to Minor-Free Classes. In *25th International Conference on Principles of Distributed Systems (OPODIS) (LIPIcs, Vol. 217)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 22:1–22:17.
- [2] Keren Censor-Hillel, Ami Paz, and Mor Perry. 2020. Approximate proof-labeling schemes. *Theor. Comput. Sci.* 811 (2020), 112–124.
- [3] Bruno Courcelle. 1990. The Monadic Second-Order Logic of Graphs. I. Recognizable Sets of Finite Graphs. *Inf. Comput.* 85, 1 (1990), 12–75. https://doi.org/10.1016/0890-5401(90)90043-H
- [4] Andrew Drucker, Fabian Kuhn, and Rotem Oshman. 2014. On the power of the congested clique model. In *33rd ACM Symposium on Principles of Distributed Computing (PODC)*. 367–376.
- [5] Michael Elberfeld, Martin Grohe, and Till Tantau. 2016. Where First-Order and Monadic Second-Order Logic Coincide. *ACM Trans. Comput. Log.* 17, 4 (2016), 25.
- [6] Louis Esperet and Benjamin Lévêque. 2022. Local certification of graphs on surfaces. *Theor. Comput. Sci.* 909 (2022), 68–75.
- [7] Laurent Feuilloley, Nicolas Bousquet, and Théo Pierron. 2022. What Can Be Certified Compactly? Compact local certification of MSO properties in tree-like graphs. In *41st ACM Symposium on Principles of Distributed Computing (PODC)*. 131–140.
- [8] Laurent Feuilloley, Pierre Fraigniaud, Pedro Montealegre, Ivan Rapaport, Éric Rémila, and Ioan Todinca. 2021. Compact Distributed Certification of Planar Graphs. *Algorithmica* 83, 7 (2021), 2215–2244.
- [9] Laurent Feuilloley, Pierre Fraigniaud, Pedro Montealegre, Ivan Rapaport, Éric Rémila, and Ioan Todinca. 2023. Local certification of graphs with bounded genus. *Discret. Appl. Math.* 325 (2023), 9–36.
- [10] Pierre Fraigniaud, Frédéric Mazoit, Pedro Montealegre, Ivan Rapaport, and Ioan Todinca. 2023. Distributed Certification for Classes of Dense Graphs. In *37th International Symposium on Distributed Computing (DISC) (LIPIcs, Vol. 281)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 20:1–20:17.
- [11] Pierre Fraigniaud, Pedro Montealegre, Ivan Rapaport, and Ioan Todinca. 2024. A Meta-Theorem for Distributed Certification. *Algorithmica* 86, 2 (2024), 585–612. https://doi.org/10.1007/S00453-023-01185-1
- [12] Jakub Gajarský and Petr Hlinený. 2015. Kernelizing MSO Properties of Trees of Fixed Height, and Some Consequences. *Log. Methods Comput. Sci.* 11, 1 (2015), 1–26.
- [13] Mika Göös and Jukka Suomela. 2016. Locally Checkable Proofs in Distributed Computing. *Theory Comput.* 12, 1 (2016), 1–33.
- [14] Martin Grohe and Stephan Kreutzer. 2009. Methods for Algorithmic Meta Theorems. In *Model Theoretic Methods in Finite Combinatorics - AMS-ASL Joint Special Session*, Vol. 558. AMS, 181–206. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.395.8282&rep=rep1&type=pdf
- [15] Amos Korman, Shay Kutten, and David Peleg. 2010. Proof labeling schemes. *Distributed Comput.* 22, 4 (2010), 215–233.
- [16] Jaroslav Nešetřil and Patrice Ossona de Mendez. 2012. *Sparsity - Graphs, Structures, and Algorithms*. Algorithms and combinatorics, Vol. 28. Springer. https://doi.org/10.1007/978-3-642-27875-4
- [17] Jaroslav Nešetřil and Patrice Ossona de Mendez. 2016. A distributed low tree-depth decomposition algorithm for bounded expansion classes. *Distributed Comput.* 29, 1 (2016), 39–49.