

RETHINKING THE EXPRESSIVENESS OF GNNs: A COMPUTATIONAL MODEL PERSPECTIVE

Anonymous authors

Paper under double-blind review

ABSTRACT

Graph Neural Networks (GNNs) are extensively employed in graph machine learning, with considerable research focusing on their expressiveness. Current studies often assess GNN expressiveness by comparing them to the Weisfeiler-Lehman (WL) tests or classical graph algorithms. However, we identify three key issues in existing analyses: (1) some studies use preprocessing to enhance expressiveness but overlook its computational costs; (2) some claim the limited power of the **identical-feature** WL test while enhancing expressiveness using **distinct** features, thus creating a mismatch; and (3) some characterize message-passing GNNs (MPGNNs) with the CONGEST model but make unrealistic assumptions about computational resources, allowing NP-Complete problems to be solved in $O(m)$ depth. We contend that a well-defined computational model is urgently needed to serve as the foundation for discussions on GNN expressiveness. To address these issues, we introduce the Resource-Limited CONGEST (RL-CONGEST) model, incorporating optional preprocessing and postprocessing to form a framework for analyzing GNN expressiveness **from an algorithmic alignment perspective**. Our framework sheds light on computational aspects, including the computational hardness of hash functions in the WL test and the role of virtual nodes in reducing network capacity. Additionally, we suggest that high-order GNNs correspond to first-order model-checking problems, offering new insights into their expressiveness.

1 INTRODUCTION

Graph Neural Networks (GNNs) have attracted widespread attention in the graph machine learning community due to their impressive performance in areas such as recommendation systems, drug discovery, and combinatorial optimization. One key area of research has focused on characterizing the expressive power of existing GNNs and developing new models with enhanced expressive power. Existing work in this area typically aligns GNNs with various algorithms. One line of research focuses on connecting GNNs to the Weisfeiler-Lehman (WL) graph isomorphism test and its variants. For instance, Xu et al. (2019) pioneered the exploration of the relationship between message-passing GNNs (MPGNNs) and the WL test. Several studies (Morris et al., 2019; Maron et al., 2019; Feng et al., 2023) have proposed high-order GNNs inspired by the k -WL test and the k -Folklore WL (FWL) test, showing that these models exhibit stronger power compared to standard MPGNNs. Additionally, works such as (Alsentzer et al., 2020; Cotta et al., 2021; Papp et al., 2021; Feng et al., 2022; Frasca et al., 2022; Zhou et al., 2023a) introduced subgraph GNNs, where subgraphs are obtained through sampling or partitioning, followed by message-passing on these subgraphs. Furthermore, Zhou et al. (2023b); Zhang et al. (2024) analyzed the counting capabilities of different GNN types. Other studies focus on aligning GNNs with traditional graph algorithms. For instance, Zhang et al. (2023) designed a GD-WL framework, which incorporates precomputed distance information as additional features in message-passing, enabling the detection of graph biconnectivity. Additionally, Loukas (2020) attempted to align MPGNNs with the CONGEST model in distributed computing. They used existing lower bounds on the communication complexity of graph algorithms in the CONGEST model to derive lower bounds on the width and depth of MPGNNs when simulating these algorithms.

We carefully revisit these works and identify **inconsistent or unreasonable results** among them:

- 054 • **Underestimated Preprocessing Time Complexity.** Some existing works employ preprocessing
055 techniques, such as substructure recognition or distance computation, to enhance their models
056 and show that the proposed models can perform algorithmic tasks beyond the capabilities of the
057 standard WL test. However, we observe that some of these works underestimate the computational
058 cost of preprocessing, sometimes resulting in a preprocessing time complexity that exceeds that of
059 the algorithmic task intended to show the model’s expressiveness, which [may lead to undesirable](#)
060 [results, such as allowing directly precomputing answers to the algorithmic task as features.](#)
- 061 • **Mismatch Between Identical-Feature WL Test and Distinct Features.** Existing works compar-
062 ing GNNs with WL tests typically derive negative results regarding the WL test, such as its
063 inability to distinguish certain toy example graphs. Some of these works then propose models
064 that incorporate additional features to enhance expressiveness. We identify that their analysis of
065 the WL test is based on the anonymous (**identical-feature**) setting, whereas the solution employs
066 **distinct** features. Therefore, directly comparing the **identical-feature** WL test with models that
067 violate this setting by incorporating **distinct** features creates a mismatch.
- 068 • **CONGEST Model Addresses Mismatch but Retains Unrealistic Assumptions.** A proposal
069 suggests resolving the mismatch between the **identical-feature** WL test and **distinct** features by
070 using the CONGEST model as a computational framework to characterize MPGNNs. However,
071 we argue that directly adopting the CONGEST model as a computational model for GNNs can
072 lead to unrealistic outcomes, such as enabling GNNs to solve many NP-Complete problems in
073 $O(m)$ rounds, due to its implicit assumption of unlimited computational resources.

075 We conclude that these inconsistent or unreasonable results stem from the ad-hoc settings in these
076 works, arising from the **lack of a well-defined computational model to characterize GNNs** and
077 analyze their expressive power. Motivated by this, we propose the **Resource-Limited CONGEST**
078 **(RL-CONGEST) model**, a simple and elegant computational framework for characterizing GNNs.
079 This model extends the standard CONGEST model by introducing constraints on computational
080 resources and incorporating optional preprocessing and postprocessing phases. It addresses key
081 issues by accounting for the complexity of preprocessing and postprocessing, explicitly allowing the
082 use of **distinct features as node IDs**, and imposing limitations on nodes’ computational resources.
083 Additionally, we present several **novel theoretical results using this model**:

- 084 • **WL Test Requires Large Networks to Compute.** We find that previous works have underesti-
085 mated the complexity of the HASH function in the WL problem. In the RL-CONGEST model,
086 if no preprocessing – such as graph modification or the use of additional features – is permitted,
087 we prove that the HASH function in the WL test typically requires the network capacity (depth
088 multiplied by width) to be linear in relation to the graph’s size for computation.
- 089 • **Virtual Nodes Reduce Network Size for WL Test.** We present evidence suggesting that virtual
090 nodes can enhance the performance of MPGNNs. Specifically, we prove that introducing a virtual
091 node can reduce the network capacity required to compute one iteration of the WL problem.
- 092 • **Aligning High-Order GNNs with Model Checking is Natural.** Additionally, we leverage ins-
093 ights from descriptive and fine-grained complexity theories to argue that aligning high-order
094 GNNs with the model checking problem is more natural.

095 The content of this paper is organized as follows: Section 2 introduces the notations used throughout
096 the paper and provides a brief overview of the relevant background knowledge. In Section 3, we
097 identify three key issues in existing works. In Section 4, we propose the RL-CONGEST model as
098 a canonical computational framework for analyzing GNNs and present our theoretical findings. In
099 Section 5, we present several open problems that may serve as directions for future work.

103 2 PRELIMINARIES

104 In this section, we first define the notations used throughout the paper. We then provide an overview
105 of the relevant background knowledge, including GNNs, various variants of the WL tests, distributed
106 computing models, and basic concepts in logic.

2.1 NOTATIONS

We use curly braces $\{\cdot\}$ to denote a set and double curly braces $\{\!\!\{\cdot\}\!\!\}$ for a multi-set where elements can appear multiple times. $[n]$ is shorthand for the set $\{0, 1, \dots, n-1\}$. Boldface lowercase letters, such as \mathbf{a} and \mathbf{b} , represent vectors, while boldface uppercase letters, such as \mathbf{A} and \mathbf{B} , represent matrices. For two vectors of the same length k , we define the Hamming distance between them, denoted as $d_H(\mathbf{x}, \mathbf{y})$, as the number of differing coordinates. We denote a graph by $G = (V, E)$, where V is the vertex set and $E \subseteq V \times V$ is the edge set. Unless otherwise specified, any graph G mentioned in this paper is undirected, meaning that for any two nodes $i, j \in V$, $(i, j) \in E$ if and only if $(j, i) \in E$. Given a node u in a graph G , the neighborhood of u , denoted by $N(u)$, is defined as $N(u) = \{v : (v, u) \in E\}$. We use $n := |V(G)|$ to denote the number of nodes and $m := |E(G)|$ to denote the number of edges when G can be inferred from context. We use D to denote the diameter of a graph, which is the length of the longest shortest path.

2.2 GRAPH NEURAL NETWORKS AND WEISFEILER-LEHMAN TESTS

Graph Neural Networks (GNNs) are neural network models defined on graphs. The most prominent and widely used framework for implementing GNNs, as found in libraries like PyTorch-Geometric (Fey & Lenssen, 2019) and DGL (Wang et al., 2019), is the message-passing GNN (MPGNN) framework proposed by Gilmer et al. (2017). The MPGNNs can be formulated as:

$$\mathbf{h}_u^{(\ell+1)} = \text{UPD}^{(\ell)} \left(\mathbf{h}_u^{(\ell)}, \left\{ \left\{ \text{MSG}^{(\ell)} \left(\mathbf{h}_u^{(\ell)}, \mathbf{h}_v^{(\ell)}, e_{(v,u)} \right) : v \in N(u) \right\} \right\} \right), \forall u \in V, \quad (1)$$

where $\mathbf{h}_u^{(\ell)}$ is the feature of node u in the ℓ -th layer, $e_{(v,u)}$ is the edge feature on (v, u) , $\text{UPD}^{(\ell)}$ is the updating function in the ℓ -th layer, and $\text{MSG}^{(\ell)}$ is the message function in the ℓ -th layer, which maps the features of a pair of adjacent nodes and the edge feature to another vector called a message.

Xu et al. (2019) claim that the expressive power of MPGNNs is bounded by the Weisfeiler-Lehman (WL) test, which was proposed by Weisfeiler and Lehman in (Weisfeiler & Lemman, 1968) as a graph isomorphism test. Initially, each node is assigned a natural number, called a color, from $[n]$ (usually, all nodes are assigned 0). The iteration formula of the WL test is as follows:

$$C^{(\ell+1)}(u) = \text{HASH}^{(\ell)} \left(C^{(\ell)}(u), \left\{ \left\{ C^{(\ell)}(v) : v \in N(u) \right\} \right\} \right), \forall u \in V, \quad (2)$$

where $C^{(\ell)}(u)$ is the color of node u in the ℓ -th iteration, and $\text{HASH}^{(\ell)}$ is a perfect hashing function mapping a multi-set of colors to a new color. It can be observed that the iteration formula of the WL test can be regarded as a special case of MPGNNs, where the message function outputs only the features of the neighboring nodes, and the updating function is a hashing function.

There are several variants of the standard WL test, and we will introduce some of them that will appear in our discussions later. A generalization is the higher-order WL tests, such as k -WL or k -FWL, which are defined on k -tuples of nodes in G . The updating formula for k -WL is described in (Huang & Villar, 2021) as:

$$C^{(\ell+1)}(\mathbf{u}) = \text{HASH}^{(\ell)} \left(C^{(\ell)}(\mathbf{u}), \left\{ \left\{ C^{(\ell)}(\mathbf{v}) : \mathbf{v} \in N_1(\mathbf{u}) \right\}, \dots, \left\{ C^{(\ell)}(\mathbf{v}) : \mathbf{v} \in N_k(\mathbf{u}) \right\} \right\} \right), \forall \mathbf{u} \in V^k, \quad (3)$$

where $\mathbf{u} = (\mathbf{u}_1, \dots, \mathbf{u}_k) \in V^k$ is a k -tuple of nodes, and the i -th neighborhood of \mathbf{u} is defined as $N_i(\mathbf{u}) = \{(\mathbf{u}_1, \dots, \mathbf{u}_{i-1}, v, \mathbf{u}_{i+1}, \dots, \mathbf{u}_k) : v \in V\}$, consisting of all k -tuples in which the i -th coordinate is substituted with each node v . Meanwhile, the updating formula for k -FWL is described in (Huang & Villar, 2021) as:

$$C^{(\ell+1)}(\mathbf{u}) = \text{HASH}^{(\ell)} \left(C^{(\ell)}(\mathbf{u}), \left\{ \left\{ C^{(\ell)}(\mathbf{u}_{[1] \leftarrow w}), \dots, C^{(\ell)}(\mathbf{u}_{[k] \leftarrow w}) \right\} : w \in V \right\} \right), \forall \mathbf{u} \in V^k, \quad (4)$$

where $\mathbf{u}_{[i] \leftarrow w} = (\mathbf{u}_1, \dots, \mathbf{u}_{i-1}, w, \mathbf{u}_{i+1}, \dots, \mathbf{u}_k)$ is the k -tuple of nodes where the i -th coordinate in \mathbf{u} is substituted with node w .

Another variant is the GD-WL framework proposed by Zhang et al. (2023), which is defined as:

$$C^{(\ell+1)}(u) = \text{HASH}^{(\ell)} \left(\left\{ \left\{ \left(d_G(u, v), C^{(\ell)}(v) \right) : v \in V \right\} \right\} \right), \forall u \in V, \quad (5)$$

where $d_G(u, v)$ is a distance, such as Shortest Path Distance (SPD) or Resistance Distance (RD).

High-order GNNs relate to high-order WL tests in the same way that MPGNNs relate to the standard WL test. In other words, if we replace the HASH function in the updating formula of a variant of the WL test with another updating function UPD, we obtain a corresponding GNN model. Therefore, we sometimes use the terms WL tests and their corresponding GNN models interchangeably.

2.3 DISTRIBUTED COMPUTING MODELS

Distributed computing involves multiple processors collaborating to compute a common result. A distributed computing model is an abstract framework used to characterize this process. LOCAL and CONGEST proposed by (Linial, 1987; 1992; Peleg, 2000) are two classic distributed computing models based on synchronous message-passing between processors. In this paper, we follow the model definitions from (Ghaffari, 2022). These models are based on an n -node graph $G = (V = [n], E)$, where G is assumed to be simple and undirected unless stated otherwise. Each node in the network hosts a processor. Initially, each processor knows the total number of nodes n , its unique identifier in $[n]$, and its initial features. In each round, a node computes based on its knowledge and sends messages to its neighbors, which may differ for each. By the end of the round, it receives all messages from its neighbors. In this model, each node must determine its own portion of the output. This process is described in (Loukas, 2020) as:

$$s_u^{(\ell+1)} = \text{UPD}_u^{(\ell)} \left(s_u^{(\ell)}, \left\{ \left\{ \text{MSG}_{v \rightarrow u}^{(\ell)} \left(s_v^{(\ell)}, v, u \right) : v \in N(u) \right\} \right\} \right), \forall u \in V, \quad (6)$$

where $s_u^{(\ell)}$ is the internal state (which may not be a vector) of the processor at node u . The primary difference between the LOCAL and CONGEST models is that, in each communication round, the LOCAL model permits nodes to exchange messages of unbounded length, while the CONGEST model restricts messages to a bounded length, typically $O(\log n)$.

2.4 BASIC CONCEPTS IN FIRST-ORDER LOGIC

First-Order Logic (FOL) is a formal system widely used in mathematics and various fields of computer science. An formula in FOL is composed of variable symbols such as x, y, z , and so on; punctuation symbols like parentheses and commas; relation symbols or predicates such as P, Q, R , and so forth; logical connectives including $\vee, \wedge, \neg, \rightarrow$, and \leftrightarrow ; and logical quantifiers, specifically the universal quantifier \forall and the existential quantifier \exists . A sentence is a special case of a formula where all variables are quantified; in other words, there are no free variables. We also introduce an extension to standard FOL called First-Order Logic with Counting (FOLC), which incorporates additional counting quantifiers. Specifically, for any natural number $i \in \mathbb{N}$, we define the counting quantifiers $\exists^{\geq i}$, $\exists^{\leq i}$, and $\exists^{=i}$. The expression $\exists^{\geq i} x \varphi(x)$ ($\exists^{\leq i} x \varphi(x)$, $\exists^{=i} x \varphi(x)$) means that there exist at least (or at most, exactly, respectively) i elements that satisfy the property φ . We use \mathcal{L}^k and \mathcal{C}^k to denote the sets of FOL and FOLC sentences, respectively, that use no more than k variables.

3 ISSUES DUE TO ABSENCE OF WELL-DEFINED COMPUTATIONAL MODEL

Many studies have analyzed the expressiveness of GNNs, but they lack a well-defined computational model as a foundation, often relying on ad-hoc methods that lead to unreasonable outcomes.

3.1 UNDERESTIMATED PREPROCESSING TIME COMPLEXITY

Many GNNs fit into a “preprocessing-then-message-passing” framework. Given an input graph G with features \mathbf{X} , they first perform preprocessing to build a new graph G' with updated features \mathbf{X}' , followed by message passing on G' . For example, high-order GNNs based on k -WL and k -FWL tests construct graphs $G' = (V^k, E')$ on k -tuples of nodes, where $E' = \{(\mathbf{u}, \mathbf{v}) \in (V^k)^2 : d_H(\mathbf{u}, \mathbf{v}) = 1\}$. Subgraph GNNs and GNNs with additional precomputed features naturally fit this framework. These models typically target algorithmic tasks beyond the capabilities of the standard WL test; however, in some cases, the time complexity of the preprocessing phase exceeds that of the algorithmic task used to show the model’s superior expressiveness, which is unreasonable from a complexity alignment perspective.

Finding Pattern Graphs is Computationally Expensive. One example comes from subgraph GNNs, which identify pattern subgraphs H in the input graph G . We note that without restrictions on H , it implies overly powerful preprocessing capabilities. For instance, Thiede et al. (2021); Bouritsas et al. (2023); Wollschläger et al. (2024) proposed variants of GNNs and WL tests that utilize hand-crafted features by recognizing subgraphs. They argued that certain subgraph schemes could make models more powerful than k -WL for any k . However, these models achieve full expressiveness only when no constraints are imposed on H , implicitly assuming the existence of oracles

capable of counting isomorphic subgraphs. Such preprocessing requirements are overly strong from a theoretical perspective. The counting version of subgraph isomorphism is a #P-Complete problem, as its decision version is NP-Complete (Karp, 1972) and by the definition of #P-completeness. By Toda’s Theorem (Toda, 1991) $\text{PH} \subseteq \text{P}^{\#\text{P}}$, we conclude that a polynomial number of queries to a subgraph isomorphism counting oracle can solve any problem in the polynomial hierarchy PH, rendering such requirements implausible.

Computing Distances Exceeds Biconnectivity Detection Time. Another example can be found in distance-based GNNs. Zhang et al. (2023) introduced the GD-WL framework and proved that the WL test cannot recognize the biconnectivity of a graph, while the proposed framework can. However, this framework encounters the issue where the computation time for all-pair distances exceeds that of biconnectivity detection. Specifically, the biconnectivity of a graph can be determined in linear time, $O(m)$, using Tarjan’s algorithm (Tarjan, 1972); whereas the worst-case time complexities for exactly computing all-pairs shortest path distances (SPDs) and resistance distances (RDs) are both $\tilde{O}(\min\{nm, n^\omega\})^1$. Additionally, we point out through the following theorem that the “existence of more efficient approximate algorithms” claim remains debatable due to the lack of sensitivity analysis on the expressiveness with respect to error in the precomputed RDs.

Theorem 1. *For any integer $n \geq 3$, there exist two graphs with n nodes, P_n and C_n , such that for any two adjacent nodes in P_n , the resistance distance is 1, while in C_n , it is $1 - 1/n$. Furthermore, P_n is neither vertex- nor edge-biconnected, whereas C_n is both vertex- and edge-biconnected.*

We defer the proofs of all theorems to the appendices to save space. Therefore, if the RD-WL framework requires the error to be less than $\Theta(1/n)$ to fully show the expressiveness, to the best of our knowledge, no existing approximation algorithms can compute all-pairs RDs within this error threshold in $o(\min\{nm, n^\omega\})$ time. Allowing higher preprocessing costs to solve a lower-complexity problem can lead to undesirable outcomes, such as using direct answers as features. The following theorem show this by demonstrating that precomputed RDs can be used to directly identify whether an edge is a cut edge, rendering subsequent message-passing phase unnecessary.

Theorem 2. *Given any edge $(u, v) \in E$ in an undirected, unweighted graph G , (u, v) is a cut edge if and only if the resistance distance $R(u, v) = 1$.*

3.2 MISMATCH BETWEEN IDENTICAL-FEATURE WL TEST AND DISTINCT FEATURES

We observe that many existing works (Bouritsas et al., 2023; Zhang et al., 2023; Wollschläger et al., 2024) which compare GNNs’ expressiveness to the WL test handling the anonymous (identical-feature) setting either in an ad-hoc manner. These works equate the expressiveness of MPGNNs with the WL test, which fails to distinguish toy examples. To address this, they propose adding features to enhance expressiveness. However, this approach may conflict with the identical-feature setting, and their models’ expressiveness are not compared to distinct-feature MPGNNs (equivalent to CONGEST, as discussed in Section 3.3). Moreover, equating MPGNNs with the WL test overlooks the fact that real-world graphs often have features, making the claim that MPGNNs are weak based solely on the WL test questionable.

WL Test is Weak in the Identical-Feature Setting. As described by Huang & Villar (2021), the WL test operates under an identical-feature setting, where all nodes are initially assigned the same color and the process relies solely on node colors, making it unable to distinguish between nodes with identical colors during its iterations. This limited expressive power is often exemplified by its inability to distinguish certain graph pairs, such as two disjoint triangle graphs ($C_3 \cup C_3$) and a six-node cycle (C_6). As shown in Figure 1, the “type” of each node, which consists of its color and the multiset of its neighbors’ colors, is the same ($\circ, \{\{\circ, \circ\}\}$) for all nodes in both graphs. After applying the HASH function, all node colors remain identical, so the WL test cannot distinguish between the graphs due to the identical multisets.

¹The notation $\tilde{O}(\cdot)$ hides polylogarithmic factors. We note that authors were not aware of more efficient algorithms; however, this does not affect our conclusion. For details on the complexity of computing all-pairs distances, please refer to Appendix D and E.




Figure 1: A running example of the anonymous WL test on $C_3 \cup C_3$ and C_6 .

Additional Features Empower Models by Breaking Anonymity? Motivated by this limitation, many works such as (Zhang et al., 2023; Wollschläger et al., 2024) have proposed **introducing features** to enhance the expressiveness of GNNs. We speculate that these features may enhance the expressiveness by breaking the **identical-feature** setting. For example, in the C_6 graph in Figure 1, if nodes are assigned the same initial color, nodes B and C are indistinguishable from A 's perspective. However, if distances are included as features, B and C will not be equivalent. **The distinct-feature setting is commonly used in existing models.** For example, GCN (Kipf & Welling, 2017) and other models applied to real-world datasets use node features that are distinct with high probability. The LINKX model proposed by Lim et al. (2021) employs $MLP(A)$, which can be reformulated as $MLP'(\sigma(AIW))$ and uses the identity matrix as features. In the GD-WL framework, the distance matrix serves as distinct features, as only the diagonal elements are zero and each row is unique. Additionally, works on GNN expressiveness (Loukas, 2020; Sato et al., 2021) adopt random features, which are also unique w.h.p.. As a result, previous works that equate GNNs with identical-feature WL tests to claim GNNs' weak expressiveness, while claiming to enhance expressiveness by distinct features, actually create a mismatch. We also include a discussion on the relationship between distinct features and permutation invariance or equivariance in Appendix A.

3.3 CONGEST ADDRESSES MISMATCH BUT RETAINS UNREALISTIC ASSUMPTIONS

We observe that one existing work (Loukas, 2020) attempts to align MPGNNs with the CONGEST model, which unintentionally resolves the inconsistency in the **identical-feature** setting. However, directly using the CONGEST model as a computational model introduces a problem: nodes are assumed to have unlimited computational resources, which leads to impractical outcomes.

Breaking Anonymity Empowers Models! In the previous section, we noted that many variants enhance the expressive power by adding precomputed features, which may implicitly rely on node IDs to distinguish nodes. This raises a natural question: can we directly improve the expressive power by explicitly incorporating node IDs into the framework of the WL test?

When comparing the equation of MPGNNs with the WL test, it becomes clear that MPGNNs can be viewed as the WL test without the constraints of **identical node features** and the HASH function. Meanwhile, MPGNNs are similar to the CONGEST model if we compare Equation 1 and 6. Thus, MPGNNs and CONGEST models are expected to have stronger expressive power than the WL tests once these limitations are removed. Loukas (2020) provides evidence for this through the following theorem by aligning MPGNNs with LOCAL and CONGEST models:

Theorem 3 ((Loukas, 2020)). *MPGNN can compute any computable function over connected graphs if the conditions are jointly met: (1) each node is uniquely identified; (2) the message and update functions are Turing-complete for every ℓ ; and (3) the depth and width are sufficiently large.*

Moreover, Pritchard & Thurimella (2011) proposed a CONGEST algorithm that solves the edge-biconnectivity problem in $O(D)$ rounds, challenging the claim that MPGNNs are weak.

Direct Use of CONGEST Is Inappropriate. In the aforementioned paper (Loukas, 2020), the author proposes using the CONGEST model from distributed computing to characterize MPGNNs, as it permits **distinct features** and supports more complex update functions compared to hash functions, making it a closer representation of the real-world implementation of MPGNNs than the WL test. However, we argue that directly using the CONGEST model as a computational model for MPGNN is not entirely appropriate, as unlimited computational resources assumption can lead to unrealistic and surprising results, as stated in the following theorem:

Theorem 4. *If we allow a single node to have unbounded computational power to solve any computable problem, then every NP-Complete decision problems on undirected unweighted **connected** graphs can be solved by the CONGEST model in $O(m)$ rounds.*

This unreasonable outcome shows that directly using the CONGEST model as the computational model for MPGNNs is inappropriate due to the unlimited computational resources for nodes.

4 PROPOSED COMPUTATIONAL MODEL AND OUR RESULTS

In the previous section, we discussed the inconsistent or unreasonable results in existing studies on the expressiveness of GNNs. We argue that these problems primarily arise from the lack of a well-defined computational model for GNNs, leading researchers to propose various ad-hoc solutions, some of which are inconsistent or unreasonable. In this section, we introduce a computational model to characterize GNNs, introduce how it addresses the aforementioned issues, and prove several interesting conclusions using this model.

We propose the Resource-Limited CONGEST (RL-CONGEST) model, an extension of the CONGEST model with constrained computational resources at each node, to serve as a computational framework for characterizing GNNs.

Definition 1 (RL-CONGEST Model and Computation Process). *Given a model width $w \in \mathbb{N}$ and a complexity class \mathcal{C} (e.g., TC^{0^2} , P), the RL-CONGEST model with width w and computational resource \mathcal{C} is defined as a CONGEST model where message sizes are limited to $w \lceil \log |V(G)| \rceil$ bits, and nodes can solve any problems in \mathcal{C} .*

For an attributed graph $G = (V, E, \mathbf{X}, \mathbf{E})$, where \mathbf{X} and \mathbf{E} represent node and edge features, the computation of a GNN using the RL-CONGEST model involves the following phases:

1. **(Optional) Preprocessing:** Operations such as building a hypergraph (for higher-order GNNs), extracting subgraphs (for subgraph GNNs), or computing additional features (for distance-based GNNs) occur in this phase, resulting in a new attributed graph $G' = (V', E', \mathbf{X}', \mathbf{E}')$. The time complexity of this step must be explicitly provided.
2. **Message-Passing with Limited Computational Resources:** Each node $u \in V'$ starts with its node features \mathbf{X}'_u and the edge features of its incident edges $\{(v, \mathbf{E}'_{(u,v)}) : v \in N_{G'}(u)\}$. The message-passing proceeds as in the standard CONGEST model, but with each node allowed to update its internal state using computations in \mathcal{C} . The total number of communication rounds corresponds to the GNN model’s depth d .
3. **(Optional) Postprocessing:** Additional computations, such as a READOUT operation, can be performed after message-passing. The time complexity must be explicitly stated.

In Figure 2, we present a diagram illustrating the three phases of the computation process for a GNN using the RL-CONGEST model.

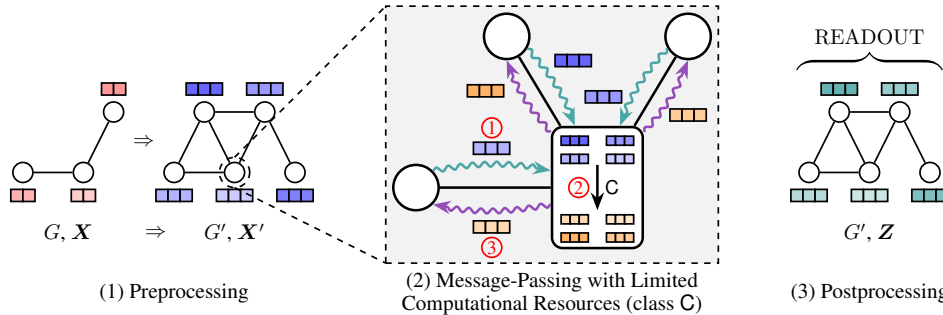


Figure 2: The three phases of computation in GNNs using the RL-CONGEST model.

We have essentially proposed a framework based on the RL-CONGEST model to characterize GNNs, which effectively addresses the three key issues identified earlier. Specifically:

- The RL-CONGEST model does not directly circumvent the “Underestimated Preprocessing Time Complexity” issue. **Our framework requires that future works report the preprocessing complex-**

²This is a complexity class of circuits. For further details on circuits, please refer to Appendix C.

ity. If it exceeds that of the algorithmic task, researchers should reassess whether preprocessing implicitly solves the task, as such results would not reflect the true expressiveness of the model.

- We address the “Mismatch Between WL Test and Features” issue by permitting nodes to know their own unique IDs. It is recommended that future works analyze the tasks GNNs can solve under this setting, rather than confining themselves to alignment with WL tests.
- We mitigate the “CONGEST Retains Unrealistic Assumptions” issue by allowing flexible configurations for the computational resources available to nodes. For instance, assuming each node operates as a Turing machine yields the standard CONGEST model. Alternatively, each node can be restricted to solving problems within P , the same class of problems that LLM-based agents enhanced by Chain-of-Thought (CoT) reasoning can address Li et al. (2024); Merrill & Sabharwal (2024); or be modeled as a TC^0 circuit, similar to MLPs (Shawe-Taylor et al., 1992; Beiu & Taylor, 1996), as observed in real-world GNN models. These different settings may lead to various interesting and independent results, and could even be extended to graph agents.

By using the RL-CONGEST model with preprocessing and postprocessing as our analysis framework, we can establish several interesting results and offer guidelines for future exploration of the expressiveness of high-order GNNs. These will be elaborated in the following subsections.

4.1 WL TEST REQUIRES LARGE NETWORKS TO COMPUTE

In previous work aligning MPGNNs with the WL test to study their expressive power, researchers aligned the update function of MPGNNs directly with the HASH function in the WL test. Aamand et al. (2022) noted the challenges in constructing the HASH function for the WL test but did not establish a lower bound on the trade-off between network depth and width. Within the RL-CONGEST framework, we rigorously prove the relationship between the depth and width required for an MPGNN to simulate one iteration of color-refinement in the WL test. This enables us to prove that the HASH function in the WL test is computationally hard, as shown by the following theorem.

Theorem 5. *If an MPGNN can simulate **one** iteration of the WL test without preprocessing, either deterministically or randomly with zero error, regardless of the computational power available to each node, the model’s width w and depth d must satisfy $d = \Omega\left(D + \frac{m}{w \log n}\right)$, given that $w = o\left(\frac{n}{\log n}\right)$.*

We defer the formal definition of the problem concerning one iteration of the WL test, along with the proof of the above theorem, to Appendix I. Notably, in our proof, we employed techniques from communication complexity without making any assumptions about the complexity class required for computational resources in the RL-CONGEST model. Therefore, the result also holds for the general CONGEST model, indicating that our findings – showing that WL-like HASH functions are hard to compute – are of independent interest to the field of distributed computing.

Furthermore, we design a deterministic RL-CONGEST algorithm with a round complexity that nearly matches the lower bound, indicating that the algorithm is near-optimal.

Theorem 6. *There exists a deterministic RL-CONGEST algorithm that can simulate **one** iteration of the WL test without preprocessing, with width w and depth d satisfying $d = O\left(D + \frac{m}{w}\right)$. Additionally, it is sufficient to set the nodes’ computational resource class to $\mathsf{C} = \text{DTIME}(n^2 \log n)$.*

4.2 VIRTUAL NODES REDUCE NETWORK SIZE FOR WL TEST

Several works have attempted to enhance the performance of GNNs by introducing a virtual node that connects to all or some nodes in the original graph (Gilmer et al., 2017; Hwang et al., 2022). Subsequent studies have analyzed the impact of this node. For instance, Barceló et al. (2020) show that virtual nodes can bring GNNs closer to a \mathcal{C}^2 classifier, while Rosenbluth et al. (2024) compare MPGNNs with virtual nodes and graph transformers. To the best of our knowledge, no prior work

has explored how a virtual node helps reduce the network’s capacity when simulating one iteration of the WL test. We address this in the following theorem:

Theorem 7. *There exists a deterministic RL-CONGEST algorithm that can simulate **one** iteration of the WL test by adding a virtual node, which connects to other nodes, as preprocessing. The algorithm operates with width w and depth d satisfying $dw = O(\Delta)$, where Δ is the maximum degree of the graph before the addition of the virtual node. Additionally, it is sufficient to set the nodes’ computational resource class to $\mathcal{C} = \text{DTIME}(n^2 \log n)$.*

Some studies suggest that virtual nodes do not enhance expressive power (Zhu et al., 2023), which contrasts with empirical evidence showing improvements in model performance. We find this is because they often equate “expressive power” with the ability to compute specific functions, akin to computability. Through our analysis, we introduce a computational model that provides a more refined view of expressive power by examining problem complexity and focusing on resource usage. This approach shows that virtual nodes can reduce the network size required to simulate the WL test, deepening our understanding of their impact.

4.3 ALIGNING HIGH-ORDER GNNs WITH MODEL CHECKING IS NATURAL

In this section, we show from both fine-grained and descriptive complexity perspectives that it is more natural to align higher-order GNNs with the \mathcal{C}^k model checking problem. We will begin by introducing the model checking problem and the related model equivalence problem.

The Model Checking (MC) problem asks whether, given a model A and a logic sentence φ , the sentence φ holds in A (i.e., $A \models \varphi$). In this paper, we focus on cases where the model is a graph G , and φ uses only the edge predicate $E(x, y)$ and the equality predicate $=(x, y)$ ³. The \mathcal{L}^k MC problem is highly expressive, capturing many key problems. For instance, deciding whether $G \models \varphi_\Delta$, where $\varphi_\Delta := \exists x \exists y \exists z (E(x, y) \wedge E(y, z) \wedge E(z, x))$, determines whether G contains a triangle subgraph. Due to its expressiveness, the \mathcal{L}^k MC problem has been widely studied. It applies to database queries like SQL (Gao et al., 2017), formal verification (Godefroid, 1997), and is central to fine-grained complexity in P (Puatracscu & Williams, 2010; Williams, 2014; Gao et al., 2017). In Appendix L, we provide evidence from theoretical computer science to support the classification of the PNF \mathcal{L}^k model checking problem⁴ in the $\tilde{\Theta}(\min\{n^k, m^{k-1}\})$ complexity class.

Another related problem is the Model Equivalence (ME) problem. Given two models A and B , and a class of logic sentences, the task is to determine whether for any sentence φ in that class, $A \models \varphi$ if and only if $B \models \varphi$. In other words, the goal is to check whether that logic cannot distinguish between the two models. Two important results that connect descriptive complexity and WL tests were proven by Cai et al. (1989) and Grohe (2017), showing that, for any $k \geq 3$, the expressiveness of $(k-1)$ -FWL and k -WL is equivalent to \mathcal{C}^k ME problem. Another result by Grohe (1998) shows that the expressiveness of both the standard WL test and the 2-WL test is equivalent to \mathcal{C}^2 ME problem. This means that the output colors from WL tests provide only a “type” of the graph, and we cannot directly interpret it for specific tasks such as determining whether a graph contains a triangle or is biconnected. The most we can infer is that if two graphs produce the same color multisets, then either both contain a triangle (or are biconnected, respectively), or neither does.

Therefore, we argue that from a computational model perspective, it is more meaningful to discuss the expressiveness of GNNs in terms of solving problems, such as model checking, rather than limiting the discussion to model equivalence, which only determines whether graph pairs are indistinguishable. A natural approach is to align higher-order GNNs, inspired by k -WL or $(k-1)$ -WL tests, with the \mathcal{C}^k MC problem. We support this claim by proving the following weaker theorem:

Theorem 8 (Informal). *If constructing the k -WL graph and additional features as preprocessing is allowed, the RL-CONGEST model can solve the PNF \mathcal{C}^k model checking problem in $O(k^2)$ rounds. Additionally, the computational resources required by each node are $\mathcal{C} = \text{DTIME}(k^2 n)$.*

³We use $x = y$ and $x \neq y$ as abbreviations for $=(x, y)$ and $\neg=(x, y)$.

⁴A PNF sentence is of the form $(Q_1 x_1)(Q_2 x_2) \cdots (Q_k x_k) \phi(x_1, \dots, x_k)$, where Q_i are quantifiers and $\phi(x_1, \dots, x_k)$ is a quantifier-free formula. Since variables can be reused, as in programming languages, a non-PNF sentence in \mathcal{L}^k may not always be convertible to an equivalent PNF sentence that remains in \mathcal{L}^k .

Therefore, using the RL-CONGEST model as the computational framework, GNNs can go beyond WL tests, which only yield hard-to-interpret graph classification types rather than addressing specific problems like model checking.

As a supplement, we list several variants of the WL test, their corresponding ME problems, and the relationships between their expressive power on the left side of Figure 3. On the right side, we present the corresponding \mathcal{C}^k MC problems, PNF \mathcal{C}^k MC problems, and PNF \mathcal{L}^k MC problems. On the far-right, the time complexity of the PNF \mathcal{L}^k MC problems is displayed. It can be observed that the expressiveness and time complexity of the WL test variants, ME problems, and MC problems form a hierarchical structure. Our Theorem 8 positions k -WL graph-based higher-order GNNs within the category of PNF \mathcal{C}^k MC problems. This is a weak result because the time complexity of constructing a k -WL graph, $O(kn^{k+1})^5$, already exceeds the time complexity of solving PNF \mathcal{C}^k MC problems using non-distributed algorithms. We conjecture that higher-order GNNs, with k -WL graph construction as preprocessing, may have the potential to solve PNF \mathcal{C}^{k+1} or general \mathcal{C}^k MC problems. We present this as an open problem in Section 5.

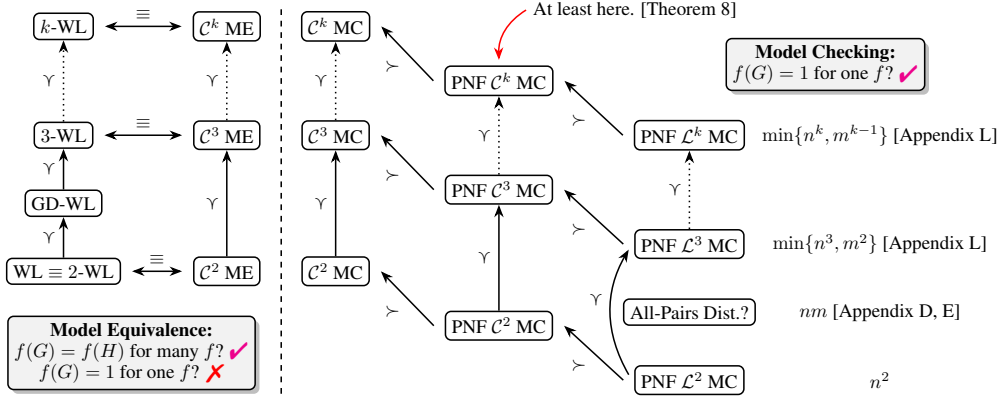


Figure 3: WL tests, model equivalence problems, model checking problems, their relationships, and the time complexity of PNF \mathcal{L}^k model checking problems. The notation $A \prec B$ means A is less powerful than B , while $A \succ B$ indicates that A is more powerful than B . $A \equiv B$ signifies that A and B have the same expressive power.

5 SOME OPEN PROBLEMS

Although we present some interesting results with our RL-CONGEST model and analysis framework, many open problems still remain, which are valuable for further research. We outline a few of them: (1) Can we establish a non-trivial trade-off between computational resources and round complexity in the RL-CONGEST model? (2) Is there a corresponding model equivalence problem, or other logic-related problems, for the GD-WL framework? (3) Do higher-order GNNs have the capability to solve the PNF \mathcal{C}^{k+1} model checking problem or general \mathcal{C}^k model checking problems?

6 CONCLUSIONS

In this paper, we identify three common issues in existing analyses of GNNs' expressive power, stemming from the absence of a well-defined computational model. To address this, we introduce the RL-CONGEST model, which includes optional preprocessing and postprocessing phases, as a standard framework for analyzing GNNs. Our framework addresses these issues and produces several noteworthy results, including the hardness of the WL problem, which may be of independent interest to the field of distributed computing. Additionally, we outline some open problems for potential future research.

⁵The time complexity of treating k -tuples as new nodes is $O(n^k)$, with each node connected to $k(n-1)$ neighbors, resulting in a total time complexity of $O(kn^{k+1})$.

REFERENCES

- Anders Aamand, Justin Y. Chen, Piotr Indyk, Shyam Narayanan, Ronitt Rubinfeld, Nicholas Schiefer, Sandeep Silwal, and Tal Wagner. Exponentially improving the complexity of simulating the weisfeiler-lehman test with graph neural networks. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/af0ad514b9cda46bd49e14ee11e2672f-Abstract-Conference.html.
- Amir Abboud, Arturs Backurs, Karl Bringmann, and Marvin Künnemann. Fine-grained complexity of analyzing compressed data: Quantifying improvements over decompress-and-solve. In Chris Umans (ed.), *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pp. 192–203. IEEE Computer Society, 2017. doi: 10.1109/FOCS.2017.26. URL <https://doi.org/10.1109/FOCS.2017.26>.
- Amir Abboud, Nick Fischer, and Yarín Shechter. Faster combinatorial k-clique algorithms. In José A. Soto and Andreas Wiese (eds.), *LATIN 2024: Theoretical Informatics - 16th Latin American Symposium, Puerto Varas, Chile, March 18-22, 2024, Proceedings, Part I*, volume 14578 of *Lecture Notes in Computer Science*, pp. 193–206. Springer, 2024. doi: 10.1007/978-3-031-55598-5_13. URL https://doi.org/10.1007/978-3-031-55598-5_13.
- Emily Alsentzer, Samuel G. Finlayson, Michelle M. Li, and Marinka Zitnik. Subgraph neural networks. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/5bca8566db79f3788be9efd96c9ed70d-Abstract.html>.
- Pablo Barceló, Egor V. Kostylev, Mikhaël Monet, Jorge Pérez, Juan L. Reutter, and Juan Pablo Silva. The logical expressiveness of graph neural networks. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL <https://openreview.net/forum?id=r11z7AEKvB>.
- Valeriu Beiu and John G. Taylor. On the circuit complexity of sigmoid feedforward neural networks. *Neural Networks*, 9(7):1155–1171, 1996. doi: 10.1016/0893-6080(96)00130-X. URL [https://doi.org/10.1016/0893-6080\(96\)00130-X](https://doi.org/10.1016/0893-6080(96)00130-X).
- Giorgos Bouritsas, Fabrizio Frasca, Stefanos Zafeiriou, and Michael M. Bronstein. Improving graph neural network expressivity via subgraph isomorphism counting. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(1):657–668, 2023. doi: 10.1109/TPAMI.2022.3154319. URL <https://doi.org/10.1109/TPAMI.2022.3154319>.
- Karl Bringmann, Allan Grønlund, and Kasper Green Larsen. A dichotomy for regular expression membership testing. In Chris Umans (ed.), *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pp. 307–318. IEEE Computer Society, 2017. doi: 10.1109/FOCS.2017.36. URL <https://doi.org/10.1109/FOCS.2017.36>.
- Jin-yi Cai, Martin Fürer, and Neil Immerman. An optimal lower bound on the number of variables for graph identification. In *30th Annual Symposium on Foundations of Computer Science, Research Triangle Park, North Carolina, USA, 30 October - 1 November 1989*, pp. 612–617. IEEE Computer Society, 1989. doi: 10.1109/SFCS.1989.63543. URL <https://doi.org/10.1109/SFCS.1989.63543>.
- Michael B. Cohen, Rasmus Kyng, Gary L. Miller, Jakub W. Pachocki, Richard Peng, Anup B. Rao, and Shen Chen Xu. Solving SDD linear systems in nearly $m \log^{1/2} n$ time. In David B. Shmoys (ed.), *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pp. 343–352. ACM, 2014. doi: 10.1145/2591796.2591833. URL <https://doi.org/10.1145/2591796.2591833>.

- 594 Leonardo Cotta, Christopher Morris, and Bruno Ribeiro. Reconstruction for powerful graph rep-
595 resentations. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang,
596 and Jennifer Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems 34:
597 Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December
598 6-14, 2021, virtual*, pp. 1713–1726, 2021. URL [https://proceedings.neurips.cc/
599 paper/2021/hash/0d8080853a54f8985276b0130266a657-Abstract.html](https://proceedings.neurips.cc/paper/2021/hash/0d8080853a54f8985276b0130266a657-Abstract.html).
- 600 Rajat Vadiraj Dwaraknath, Ishani Karmarkar, and Aaron Sidford. Towards optimal
601 effective resistance estimation. In Alice Oh, Tristan Naumann, Amir Globerson,
602 Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), *Advances in Neural Infor-
603 mation Processing Systems 36: Annual Conference on Neural Information Process-
604 ing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16,
605 2023, 2023*. URL [http://papers.nips.cc/paper_files/paper/2023/hash/
606 b8e2046160a568145af6d42eeef199f4-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2023/hash/b8e2046160a568145af6d42eeef199f4-Abstract-Conference.html).
- 607 Friedrich Eisenbrand and Fabrizio Grandoni. On the complexity of fixed parameter clique and
608 dominating set. *Theor. Comput. Sci.*, 326(1-3):57–67, 2004. doi: 10.1016/J.TCS.2004.05.009.
609 URL <https://doi.org/10.1016/j.tcs.2004.05.009>.
- 610 Jiarui Feng, Yixin Chen, Fuhai Li, Anindya Sarkar, and Muhan Zhang. How power-
611 ful are k-hop message passing graph neural networks. In Sanmi Koyejo, S. Mo-
612 hamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural
613 Information Processing Systems 35: Annual Conference on Neural Information Process-
614 ing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9,
615 2022, 2022*. URL [http://papers.nips.cc/paper_files/paper/2022/hash/
616 1ece70d2259b8e9510e2d4ca8754cecf-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/1ece70d2259b8e9510e2d4ca8754cecf-Abstract-Conference.html).
- 617 Jiarui Feng, Lecheng Kong, Hao Liu, Dacheng Tao, Fuhai Li, Muhan Zhang, and Yixin Chen.
618 Extending the design space of graph neural networks by rethinking folklore weisfeiler-lehman.
619 In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine
620 (eds.), *Advances in Neural Information Processing Systems 36: Annual Conference on Neural
621 Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 -
622 16, 2023, 2023*. URL [http://papers.nips.cc/paper_files/paper/2023/hash/
623 1cac8326ce3fbe79171db9754211530c-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2023/hash/1cac8326ce3fbe79171db9754211530c-Abstract-Conference.html).
- 624 Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric.
625 *CoRR*, abs/1903.02428, 2019. URL <http://arxiv.org/abs/1903.02428>.
- 626 Fabrizio Frasca, Beatrice Bevilacqua, Michael M. Bronstein, and Haggai Maron. Under-
627 standing and extending subgraph gnns by rethinking their symmetries. In Sanmi Koyejo,
628 S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neu-
629 ral Information Processing Systems 35: Annual Conference on Neural Information Process-
630 ing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9,
631 2022, 2022*. URL [http://papers.nips.cc/paper_files/paper/2022/hash/
632 cb2a4cc70db72ea779abd01107782c7b-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/cb2a4cc70db72ea779abd01107782c7b-Abstract-Conference.html).
- 633 Jiawei Gao, Russell Impagliazzo, Antonina Kolokolova, and R. Ryan Williams. Completeness for
634 first-order properties on sparse structures with algorithmic applications. In Philip N. Klein (ed.),
635 *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA
636 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pp. 2162–2181. SIAM, 2017. doi: 10.
637 1137/1.9781611974782.141. URL [https://doi.org/10.1137/1.9781611974782.
638 141](https://doi.org/10.1137/1.9781611974782.141).
- 639 Mohsen Ghaffari. Distributed graph algorithms. Course Notes for Distributed Algorithms, 2022.
640 URL <https://people.csail.mit.edu/ghaffari/DA22/Notes/DGA.pdf>. Ac-
641 cessed: September, 2024.
- 642 Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural
643 message passing for quantum chemistry. In Doina Precup and Yee Whye Teh (eds.), *Proceedings
644 of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia,
645 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1263–1272.
646 PMLR, 2017. URL <http://proceedings.mlr.press/v70/gilmer17a.html>.

- 648 Patrice Godefroid. Model checking for programming languages using verisoft. In Peter Lee,
649 Fritz Henglein, and Neil D. Jones (eds.), *Conference Record of POPL'97: The 24th ACM*
650 *SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Papers Presented at*
651 *the Symposium, Paris, France, 15-17 January 1997*, pp. 174–186. ACM Press, 1997. doi:
652 10.1145/263699.263717. URL <https://doi.org/10.1145/263699.263717>.
- 653 Martin Grohe. Finite variable logics in descriptive complexity theory. *Bull. Symb. Log.*, 4(4):345–
654 398, 1998. doi: 10.2307/420954. URL <https://doi.org/10.2307/420954>.
- 655 Martin Grohe. *Descriptive Complexity, Canonisation, and Definable Graph Structure Theory*, vol-
656 ume 47 of *Lecture Notes in Logic*. Cambridge University Press, 2017. ISBN 9781139028868.
657 doi: 10.1017/9781139028868. URL <https://doi.org/10.1017/9781139028868>.
- 658 Takatori Hayashi, Takuya Akiba, and Yuichi Yoshida. Efficient algorithms for spanning tree central-
659 ity. In Subbarao Kambhampati (ed.), *Proceedings of the Twenty-Fifth International Joint Confer-*
660 *ence on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pp. 3733–3739.
661 IJCAI/AAAI Press, 2016. URL <http://www.ijcai.org/Abstract/16/525>.
- 662 Ningyuan Teresa Huang and Soledad Villar. A short tutorial on the weisfeiler-lehman test and
663 its variants. In *IEEE International Conference on Acoustics, Speech and Signal Processing,*
664 *ICASSP 2021, Toronto, ON, Canada, June 6-11, 2021*, pp. 8533–8537. IEEE, 2021. doi: 10.
665 1109/ICASSP39728.2021.9413523. URL [https://doi.org/10.1109/ICASSP39728.](https://doi.org/10.1109/ICASSP39728.2021.9413523)
666 2021.9413523.
- 667 EunJeong Hwang, Veronika Thost, Shib Sankar Dasgupta, and Tengfei Ma. An analysis of virtual
668 nodes in graph neural networks for link prediction (extended abstract). In *The First Learning on*
669 *Graphs Conference, 2022*. URL <https://openreview.net/forum?id=dI6KBKNRp7>.
- 670 Arun Jambulapati and Aaron Sidford. Ultrasparse ultrasparsifiers and faster laplacian system solvers.
671 In Dániel Marx (ed.), *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms,*
672 *SODA 2021, Virtual Conference, January 10 - 13, 2021*, pp. 540–559. SIAM, 2021. doi: 10.1137/
673 1.9781611976465.33. URL <https://doi.org/10.1137/1.9781611976465.33>.
- 674 Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller and
675 James W. Thatcher (eds.), *Proceedings of a symposium on the Complexity of Computer Com-*
676 *putations, held March 20-22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown*
677 *Heights, New York, USA*, The IBM Research Symposia Series, pp. 85–103. Plenum Press, New
678 York, 1972. doi: 10.1007/978-1-4684-2001-2_9. URL [https://doi.org/10.1007/](https://doi.org/10.1007/978-1-4684-2001-2_9)
679 978-1-4684-2001-2_9.
- 680 Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional net-
681 works. In *5th International Conference on Learning Representations, ICLR 2017, Toulon,*
682 *France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=SJU4ayYgl>.
- 683 D. J. Klein and M. Randić. Resistance distance. *Journal of Mathematical Chemistry*, 12(1):81–
684 95, Dec 1993. ISSN 1572-8897. doi: 10.1007/BF01164627. URL [https://doi.org/10.](https://doi.org/10.1007/BF01164627)
685 1007/BF01164627.
- 686 Ioannis Koutis, Gary L. Miller, and Richard Peng. Approaching optimality for solving SDD linear
687 systems. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010,*
688 *October 23-26, 2010, Las Vegas, Nevada, USA*, pp. 235–244. IEEE Computer Society, 2010. doi:
689 10.1109/FOCS.2010.29. URL <https://doi.org/10.1109/FOCS.2010.29>.
- 690 Ioannis Koutis, Gary L. Miller, and Richard Peng. A nearly-m log n time solver for SDD linear
691 systems. In Rafail Ostrovsky (ed.), *IEEE 52nd Annual Symposium on Foundations of Computer*
692 *Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pp. 590–598. IEEE Com-
693 puter Society, 2011. doi: 10.1109/FOCS.2011.85. URL [https://doi.org/10.1109/](https://doi.org/10.1109/FOCS.2011.85)
694 FOCS.2011.85.
- 695 Eyal Kushilevitz and Noam Nisan. *Communication complexity*. Cambridge University Press, 1997.
696 ISBN 978-0-521-56067-2. URL <https://doi.org/10.1017/CBO9780511574948>.

- 702 Zhiyuan Li, Hong Liu, Denny Zhou, and Tengyu Ma. Chain of thought empowers transform-
703 ers to solve inherently serial problems. In *The Twelfth International Conference on Learning*
704 *Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL
705 <https://openreview.net/forum?id=3EWTEy9MTM>.
- 706 Derek Lim, Felix Hohne, Xiuyu Li, Sijia Linda Huang, Vaishnavi Gupta, Omkar Bhalerao, and
707 Ser-Nam Lim. Large scale learning on non-homophilous graphs: New benchmarks and strong
708 simple methods. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang,
709 and Jennifer Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems 34:*
710 *Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December*
711 *6-14, 2021, virtual*, pp. 20887–20902, 2021. URL [https://proceedings.neurips.cc/
712 paper/2021/hash/ae816a80e4c1c56caa2eb4e1819cbb2f-Abstract.html](https://proceedings.neurips.cc/paper/2021/hash/ae816a80e4c1c56caa2eb4e1819cbb2f-Abstract.html).
- 713 Nathan Linial. Distributive graph algorithms-global solutions from local data. In *28th Annual*
714 *Symposium on Foundations of Computer Science, Los Angeles, California, USA, 27-29 October*
715 *1987*, pp. 331–335. IEEE Computer Society, 1987. doi: 10.1109/SFCS.1987.20. URL [https:
716 //doi.org/10.1109/SFCS.1987.20](https://doi.org/10.1109/SFCS.1987.20).
- 717 Nathan Linial. Locality in distributed graph algorithms. *SIAM J. Comput.*, 21(1):193–201, 1992.
718 doi: 10.1137/0221015. URL <https://doi.org/10.1137/0221015>.
- 719 Andreas Loukas. What graph neural networks cannot learn: depth vs width. In *8th International*
720 *Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.
721 OpenReview.net, 2020. URL <https://openreview.net/forum?id=B1l2bp4YwS>.
- 722 László Lovász. Random walks on graphs. *Combinatorics, Paul erdos is eighty*, 2(1-46):4, 1993.
723
- 724 Haggai Maron, Heli Ben-Hamu, Hadar Serviansky, and Yaron Lipman. Provably power-
725 ful graph networks. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Flo-
726 rence d’Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural In-*
727 *formation Processing Systems 32: Annual Conference on Neural Information Process-*
728 *ing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp.
729 2153–2164, 2019. URL [https://proceedings.neurips.cc/paper/2019/hash/
730 bb04af0f7eacee4aae62035497da1387-Abstract.html](https://proceedings.neurips.cc/paper/2019/hash/bb04af0f7eacee4aae62035497da1387-Abstract.html).
- 731 William Merrill and Ashish Sabharwal. The expressive power of transformers with chain of thought.
732 In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Aus-*
733 *tria, May 7-11, 2024*. OpenReview.net, 2024. URL [https://openreview.net/forum?
734 id=NjNGlPh8Wh](https://openreview.net/forum?id=NjNGlPh8Wh).
- 735 Christopher Morris, Martin Ritzert, Matthias Fey, William L. Hamilton, Jan Eric Lenssen, Gaurav
736 Rattan, and Martin Grohe. Weisfeiler and leman go neural: Higher-order graph neural networks.
737 In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Inno-*
738 *vative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium*
739 *on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, Jan-*
740 *uary 27 - February 1, 2019*, pp. 4602–4609. AAAI Press, 2019. doi: 10.1609/AAAI.V33I01.
741 33014602. URL <https://doi.org/10.1609/aaai.v33i01.33014602>.
- 742 Pál András Papp, Karolis Martinkus, Lukas Faber, and Roger Wattenhofer. Dropgnn: Ran-
743 dom dropouts increase the expressiveness of graph neural networks. In Marc’Aurelio Ran-
744 zato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan
745 (eds.), *Advances in Neural Information Processing Systems 34: Annual Conference on Neu-*
746 *ral Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*,
747 pp. 21997–22009, 2021. URL [https://proceedings.neurips.cc/paper/2021/
748 hash/b8b2926bd27d4307569ad119b6025f94-Abstract.html](https://proceedings.neurips.cc/paper/2021/hash/b8b2926bd27d4307569ad119b6025f94-Abstract.html).
- 749 David Peleg. *Distributed computing: a locality-sensitive approach*. Society for Industrial and
750 Applied Mathematics, USA, 2000. ISBN 978-0-89871-464-7. URL [https://doi.org/10.
751 1137/1.9780898719772](https://doi.org/10.1137/1.9780898719772).
- 752 David Pritchard and Ramakrishna Thurimella. Fast computation of small cuts via cycle space sam-
753 pling. *ACM Trans. Algorithms*, 7(4):46:1–46:30, 2011. doi: 10.1145/2000807.2000814. URL
754 <https://doi.org/10.1145/2000807.2000814>.
- 755

- 756 Mihai Pătrăscu and Ryan Williams. On the possibility of faster SAT algorithms. In Moses
757 Charikar (ed.), *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algo-*
758 *rithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, pp. 1065–1075. SIAM, 2010. doi:
759 10.1137/1.9781611973075.86. URL [https://doi.org/10.1137/1.9781611973075.](https://doi.org/10.1137/1.9781611973075.86)
760 86.
- 761 Eran Rosenbluth, Jan Tönshoff, Martin Ritzert, Berke Kisin, and Martin Grohe. Distinguished in
762 uniform: Self-attention vs. virtual nodes. In *The Twelfth International Conference on Learning*
763 *Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL
764 <https://openreview.net/forum?id=AcSChDWL6V>.
- 765 Ryoma Sato, Makoto Yamada, and Hisashi Kashima. Random features strengthen graph neural
766 networks. In Carlotta Demeniconi and Ian Davidson (eds.), *Proceedings of the 2021 SIAM In-*
767 *ternational Conference on Data Mining, SDM 2021, Virtual Event, April 29 - May 1, 2021*, pp.
768 333–341. SIAM, 2021. doi: 10.1137/1.9781611976700.38. URL [https://doi.org/10.](https://doi.org/10.1137/1.9781611976700.38)
769 [1137/1.9781611976700.38](https://doi.org/10.1137/1.9781611976700.38).
- 770 Raimund Seidel. On the all-pairs-shortest-path problem in unweighted undirected graphs. *J. Com-*
771 *put. Syst. Sci.*, 51(3):400–403, 1995. doi: 10.1006/JCSS.1995.1078. URL [https://doi.](https://doi.org/10.1006/jcss.1995.1078)
772 [org/10.1006/jcss.1995.1078](https://doi.org/10.1006/jcss.1995.1078).
- 773 John S. Shawe-Taylor, Martin H.G. Anthony, and Walter Kern. Classes of feedforward neural net-
774 works and their circuit complexity. *Neural Networks*, 5(6):971–977, 1992. ISSN 0893-6080. doi:
775 [https://doi.org/10.1016/S0893-6080\(05\)80093-0](https://doi.org/10.1016/S0893-6080(05)80093-0). URL [https://www.sciencedirect.](https://www.sciencedirect.com/science/article/pii/S0893608005800930)
776 [com/science/article/pii/S0893608005800930](https://www.sciencedirect.com/science/article/pii/S0893608005800930).
- 777 Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning,
778 graph sparsification, and solving linear systems. In László Babai (ed.), *Proceedings of the 36th*
779 *Annual ACM Symposium on Theory of Computing, Chicago, IL, USA, June 13-16, 2004*, pp.
780 81–90. ACM, 2004. doi: 10.1145/1007352.1007372. URL [https://doi.org/10.1145/](https://doi.org/10.1145/1007352.1007372)
781 [1007352.1007372](https://doi.org/10.1145/1007352.1007372).
- 782 Robert Endre Tarjan. Depth-first search and linear graph algorithms. *SIAM J. Comput.*, 1(2):146–
783 160, 1972. doi: 10.1137/0201010. URL <https://doi.org/10.1137/0201010>.
- 784 Erik H. Thiede, Wenda Zhou, and Risi Kondor. Autobahn: Automorphism-based graph neural nets.
785 In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wort-
786 man Vaughan (eds.), *Advances in Neural Information Processing Systems 34: Annual Confer-*
787 *ence on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021,*
788 *virtual*, pp. 29922–29934, 2021. URL [https://proceedings.neurips.cc/paper/](https://proceedings.neurips.cc/paper/2021/hash/faf02b2358de8933f480a146f4d2d98e-Abstract.html)
789 [2021/hash/faf02b2358de8933f480a146f4d2d98e-Abstract.html](https://proceedings.neurips.cc/paper/2021/hash/faf02b2358de8933f480a146f4d2d98e-Abstract.html).
- 790 Seinosuke Toda. PP is as hard as the polynomial-time hierarchy. *SIAM J. Comput.*, 20(5):865–877,
791 1991. doi: 10.1137/0220053. URL <https://doi.org/10.1137/0220053>.
- 792 Minjie Wang, Lingfan Yu, Da Zheng, Quan Gan, Yu Gai, Zihao Ye, Mufei Li, Jinjing Zhou,
793 Qi Huang, Chao Ma, Ziyue Huang, Qipeng Guo, Hao Zhang, Haibin Lin, Junbo Zhao, Jinyang Li,
794 Alexander J. Smola, and Zheng Zhang. Deep graph library: Towards efficient and scalable deep
795 learning on graphs. *CoRR*, abs/1909.01315, 2019. URL [http://arxiv.org/abs/1909.](http://arxiv.org/abs/1909.01315)
796 [01315](http://arxiv.org/abs/1909.01315).
- 797 Boris Weisfeiler and Andrei Leman. The reduction of a graph to canonical form and the algebra
798 which appears therein. *nti, Series*, 2(9):12–16, 1968.
- 799 Ryan Williams. Faster decision of first-order graph properties. In Thomas A. Henzinger and Dale
800 Miller (eds.), *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science*
801 *Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science*
802 *(LICS), CSL-LICS ’14, Vienna, Austria, July 14 - 18, 2014*, pp. 80:1–80:6. ACM, 2014. doi:
803 10.1145/2603088.2603121. URL <https://doi.org/10.1145/2603088.2603121>.
- 804 Tom Wollschläger, Niklas Kemper, Leon Hetzel, Johanna Sommer, and Stephan Günnemann. Ex-
805 pressivity and generalization: Fragment-biases for molecular gnns. In *Forty-first International*
806 *Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenRe-
807 view.net, 2024. URL <https://openreview.net/forum?id=rPm5cKb1VB>.

- 810 Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural
811 networks? In *7th International Conference on Learning Representations, ICLR 2019, New Or-*
812 *leans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL [https://openreview.net/](https://openreview.net/forum?id=ryGs6iA5Km)
813 [forum?id=ryGs6iA5Km](https://openreview.net/forum?id=ryGs6iA5Km).
- 814 Andrew Chi-Chih Yao. Some complexity questions related to distributive computing(preliminary
815 report). In *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing,*
816 *STOC '79*, pp. 209–213, New York, NY, USA, 1979. Association for Computing Machinery.
817 ISBN 9781450374385. doi: 10.1145/800135.804414. URL [https://doi.org/10.1145/](https://doi.org/10.1145/800135.804414)
818 [800135.804414](https://doi.org/10.1145/800135.804414).
- 819 Bohang Zhang, Shengjie Luo, Liwei Wang, and Di He. Rethinking the expressive power of gnns
820 via graph biconnectivity. In *The Eleventh International Conference on Learning Representa-*
821 *tions, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL [https://](https://openreview.net/forum?id=r9hNv76KoT3)
822 openreview.net/forum?id=r9hNv76KoT3.
- 823 Bohang Zhang, Jingchu Gai, Yiheng Du, Qiwei Ye, Di He, and Liwei Wang. Beyond weisfeiler-
824 lehman: A quantitative framework for GNN expressiveness. In *The Twelfth International Con-*
825 *ference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenRe-
826 view.net, 2024. URL <https://openreview.net/forum?id=HSKaGOi7Ar>.
- 827 Cai Zhou, Xiyuan Wang, and Muhan Zhang. From relational pooling to subgraph gnns: A univer-
828 sal framework for more expressive graph neural networks. In Andreas Krause, Emma Brunskill,
829 Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *International*
830 *Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, vol-
831 *ume 202 of Proceedings of Machine Learning Research*, pp. 42742–42768. PMLR, 2023a. URL
832 <https://proceedings.mlr.press/v202/zhou23n.html>.
- 833 Junru Zhou, Jiarui Feng, Xiyuan Wang, and Muhan Zhang. Distance-restricted folklore
834 weisfeiler-leman gnns with provable cycle counting power. In Alice Oh, Tristan Nau-
835 mann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), *Advances*
836 *in Neural Information Processing Systems 36: Annual Conference on Neural Informa-*
837 *tion Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16,*
838 *2023*, 2023b. URL [http://papers.nips.cc/paper_files/paper/2023/hash/](http://papers.nips.cc/paper_files/paper/2023/hash/2e2e7c2e3c2e70fa2e9756dce728fcca-Abstract-Conference.html)
839 [2e2e7c2e3c2e70fa2e9756dce728fcca-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2023/hash/2e2e7c2e3c2e70fa2e9756dce728fcca-Abstract-Conference.html).
- 840 Wenhao Zhu, Tianyu Wen, Guojie Song, Liang Wang, and Bo Zheng. On structural expressive power
841 of graph transformers. In Ambuj K. Singh, Yizhou Sun, Leman Akoglu, Dimitrios Gunopulos,
842 Xifeng Yan, Ravi Kumar, Fatma Ozcan, and Jieping Ye (eds.), *Proceedings of the 29th ACM*
843 *SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2023, Long Beach, CA,*
844 *USA, August 6-10, 2023*, pp. 3628–3637. ACM, 2023. doi: 10.1145/3580305.3599451. URL
845 <https://doi.org/10.1145/3580305.3599451>.
- 846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

864		
865		
866		
867		
868	A	A Brief Discussion on the Relationship Between Distinct Features and Permutation Invariance or Equivariance
869		18
870		
871	B	A Brief Introduction to Communication Complexity
872		18
873		
874	C	A Brief Introduction to Boolean Circuits
875		18
876	D	Time Complexity of All-Pairs Shortest Paths in Unweighted Undirected Graphs
877		19
878	E	A Brief Introduction to Resistance Distance
879		19
880	F	Proof of Theorem 1
881		20
882	G	Proof of Theorem 2
883		21
884	H	Proof of Theorem 4
885		21
886	I	The Weisfeiler-Lehman Problem and Its Hardness (Theorem 5)
887		22
888		
889	J	Proof of Theorem 6
890		24
891	K	Proof of Theorem 7
892		25
893	L	Complexity of the PNF Model Checking Problem
894		26
895	M	Proof of Theorem 8
896		27
897		
898		
899		
900		
901		
902		
903		
904		
905		
906		
907		
908		
909		
910		
911		
912		
913		
914		
915		
916		
917		

A A BRIEF DISCUSSION ON THE RELATIONSHIP BETWEEN DISTINCT FEATURES AND PERMUTATION INVARIANCE OR EQUIVARIANCE

Permutation invariance and equivariance are two vital properties to discuss whether a GNN model possesses them. They are defined as follows:

Definition 2 (Permutation Invariant). *Let f be a graph neural network model on some graph G with n nodes, and let $\mathbf{P} \in \{0, 1\}^{n \times n}$ be a permutation matrix. If for any input feature matrix $\mathbf{X} \in \mathbb{R}^{n \times F}$,*

$$f(\mathbf{P}\mathbf{X}) = f(\mathbf{X}),$$

then f is said to be permutation invariant.

Definition 3 (Permutation Equivariant). *Let f be a graph neural network model on some graph G with n nodes, and let $\mathbf{P} \in \{0, 1\}^{n \times n}$ be a permutation matrix. If for any input feature matrix $\mathbf{X} \in \mathbb{R}^{n \times F}$,*

$$f(\mathbf{P}\mathbf{X}) = \mathbf{P}f(\mathbf{X}),$$

then f is said to be permutation equivariant.

In other words, a model is permutation invariant (equivariant) if permuting the input features yields identical (or correspondingly permuted) outputs, respectively.

It should be noted that permutation invariance or equivariance are properties determined by the GNN architecture f , or more specifically, by the update functions in each layer, rather than by the features \mathbf{X} used as input.

B A BRIEF INTRODUCTION TO COMMUNICATION COMPLEXITY

To prove lower bounds on the rounds of CONGEST algorithms, a key tool is the communication complexity which was first introduced by Yao (1979).

Two-party communication complexity involves two participants, Alice and Bob, who collaborate to compute a function $f : X \times Y \rightarrow Z$, where X and Y are their input domains, respectively. They agree on a strategy beforehand but are separated before receiving their inputs $(x, y) \in X \times Y$. They then exchange messages to compute $f(x, y)$, with the goal of minimizing the total number of bits exchanged.

In deterministic communication, the strategy is fixed, and the minimum number of bits required to compute f in this setting is known as the deterministic communication complexity, denoted by $D(f)$. Similarly, in randomized communication, where Alice and Bob can use random bits and a two-sided error of ϵ is allowed, the minimum number of bits required is the randomized communication complexity. If the randomness is private, it is denoted by $R_\epsilon^{\text{priv}}(f)$, and if it is public, it is denoted by $R_\epsilon^{\text{pub}}(f)$.

The Equality (EQ) problem between two n -bit strings, denoted by $\text{EQ}_n : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, is defined as

$$\text{EQ}_n(\mathbf{x}, \mathbf{y}) = \begin{cases} 1, & \mathbf{x} = \mathbf{y}, \\ 0, & \text{otherwise.} \end{cases}$$

It is arguably the most well-known problem in two-party communication complexity which has been extensively studied. We summarize its communication complexity under different settings in Table 1 below.

C A BRIEF INTRODUCTION TO BOOLEAN CIRCUITS

Boolean circuits are computational models used to represent Boolean function computations. A Boolean circuit with input size n can be described as a Directed Acyclic Graph (DAG), with n source nodes as inputs and one sink node as the output. The nodes represent gates, including NOT (\neg , with one input), AND (\wedge , with two or more inputs), OR (\vee , with two or more inputs), and threshold gates ($\text{Th}_{\mathbf{a}, \theta}$), which output 1 if and only if $\mathbf{a}^\top \mathbf{x} \geq \theta$, where \mathbf{a} and θ are independent of the input \mathbf{x} .

Table 1: The communication complexity of the EQ_n function under different settings.

Function	Deterministic $D(\cdot)$	Randomized			
		Private Coin		Public Coin	
		$R_0^{\text{prv}}(\cdot)$	$R_{1/3}^{\text{prv}}(\cdot)$	$R_0^{\text{pub}}(\cdot)$	$R_{1/3}^{\text{pub}}(\cdot)$
EQ_n	$\Theta(n)^\dagger$	$\Theta(n)^\dagger$	$\Theta(\log n)^\dagger$	$\Theta(n)^*$	$\Theta(1)^\dagger$

† The proofs can be found in (Kushilevitz & Nisan, 1997).

* Since $R_0^{\text{prv}}(f) = O(R_0^{\text{pub}}(f) + \log n)$, as per Exercise 3.15 in (Kushilevitz & Nisan, 1997).

A circuit family $(C_n)_{n \geq 1}$ is a sequence of circuits with input size growing from 1 to infinity. A family is L-uniform if a Turing machine can construct C_n in $O(\log n)$ space, given n in unary. Circuit families are assumed uniform unless otherwise stated. There are three major circuit complexity classes:

- AC^k consists of problems solvable by circuits with \neg , \wedge , and \vee gates, polynomial size, depth $O(\log^k n)$, and unbounded fan-in.
- NC^k includes problems solvable by circuits with \neg , \wedge , and \vee gates, polynomial size, depth $O(\log^k n)$, and fan-in 2.
- TC^k comprises problems solvable by circuits with polynomial size, depth $O(\log^k n)$, unbounded fan-in gates.

D TIME COMPLEXITY OF ALL-PAIRS SHORTEST PATHS IN UNWEIGHTED UNDIRECTED GRAPHS

The shortest path problem is one of the fundamental problems in graph theory. The All-Pairs Shortest Path (APSP) problem seeks to determine the shortest path distance between all pairs of nodes in a given graph G . To the best of our knowledge, the fastest algorithm for APSP on unweighted and undirected graphs can be formally stated as follows:

Lemma 1 (Folklore; (Seidel, 1995)). *The computation of APSP for an unweighted, undirected graph with n nodes and m edges can be achieved with a time complexity of $\tilde{O}(\min(nm, n^\omega))$, where $\omega < 2.372$ is the matrix multiplication exponent.*

E A BRIEF INTRODUCTION TO RESISTANCE DISTANCE

In this section, we introduce the concept of Resistance Distance (RD), covering its definition and the time complexity of approximately computing All-Pairs Resistance Distances (APRD). We begin with the definition of resistance distance:

Definition 4 (Resistance Distance). *Given an undirected graph G and a pair of nodes s and t , the resistance distance between s and t , denoted by $R(s, t)$, is defined as:*

$$R(s, t) = (e_s - e_t)^\top \mathbf{L}^\dagger (e_s - e_t) = \mathbf{L}_{ss}^\dagger - \mathbf{L}_{st}^\dagger - \mathbf{L}_{ts}^\dagger + \mathbf{L}_{tt}^\dagger, \quad (7)$$

where e_s is a one-hot vector with a 1 in the s -th position, and \mathbf{L}^\dagger is the Moore-Penrose pseudo-inverse of the graph Laplacian matrix $\mathbf{L} := \mathbf{D} - \mathbf{A}$, satisfying $\mathbf{L}\mathbf{L}^\dagger = \mathbf{\Pi}$ and $\text{span}(\mathbf{L}^\dagger) = \text{span}(\mathbf{L}) = \{\mathbf{v} \in \mathbb{R}^n : \mathbf{v}^\top \mathbf{1} = 0\}$. Here, $\mathbf{\Pi} = \mathbf{I}_n - \frac{1}{n}\mathbf{1}\mathbf{1}^\top$ is the projection matrix onto $\text{span}(\mathbf{L})$.

As shown by Klein & Randić (1993), $R(s, t)$ is a valid distance metric on graphs. Additionally, we present the following lemma, which connects resistance distance to spanning trees:

Lemma 2 ((Lovász, 1993; Hayashi et al., 2016)). *Given an edge (s, t) in an unweighted undirected graph G , we have*

$$R(s, t) = \Pr_{T \sim \mu_G} (\mathbb{I}[(s, t) \in E(T)]),$$

where T is a spanning tree sampled from the uniform distribution of spanning trees of G , denoted by μ_G , and $\mathbb{I}[\cdot]$ is the indicator function.

Next, we define the approximate computation of APRD:

Definition 5 (Approximate Computation of APRD). *Given an undirected, unweighted graph $G = (V, E)$, an error threshold $\epsilon > 0$, and a failure probability $0 \leq p_f \leq 1$, compute a matrix $\mathbf{R} \in \mathbb{R}^{n \times n}$ such that for any node pair u, v ,*

$$\Pr(|\mathbf{R}_{uv} - R(u, v)| > \epsilon R(u, v)) \leq p_f.$$

To the best of our knowledge, the fastest algorithm for approximating APRD can be formally stated as follows:

Lemma 3 ((Dwaraknath et al., 2023)). *The approximate computation of APRD for a graph with n nodes and m edges can be achieved with a time complexity of*

$$\tilde{O}\left(\min\left(nm, n^\omega, \frac{m}{\epsilon} \kappa(\mathbf{D}^{-1/2} \mathbf{L} \mathbf{D}^{-1/2}) + n^2\right)\right),$$

where $\omega < 2.372$ is the matrix multiplication exponent, and κ denotes the condition number of the matrix.

Note that the $\tilde{O}(nm)$ time complexity is achieved using near-linear time Laplacian solvers, as proposed in a series of works (Spielman & Teng, 2004; Koutis et al., 2010; 2011; Cohen et al., 2014; Jambulapati & Sidford, 2021), while the $\tilde{O}(n^\omega)$ complexity comes from fast matrix multiplication techniques. However, under the $\Theta\left(\frac{1}{n}\right)$ error requirement, as discussed in Section 3, the time complexity degenerates to $\tilde{O}(\min(nm, n^\omega))$, which matches the time complexity of APSP.

F PROOF OF THEOREM 1

Theorem 1. *For any integer $n \geq 3$, there exist two graphs with n nodes, P_n and C_n , such that for any two adjacent nodes in P_n , the resistance distance is 1, while in C_n , it is $1 - 1/n$. Furthermore, P_n is neither vertex- nor edge-biconnected, whereas C_n is both vertex- and edge-biconnected.*

Proof. Let P_n be the path graph with n nodes, and C_n be the cycle graph with n nodes.



Figure 4: P_6 and C_6 .

We observe that P_n has exactly one spanning tree, which is the graph itself. In contrast, C_n has n spanning trees, each formed by removing a single edge from $E(C_n)$. Therefore, for each edge $(u, v) \in E(P_n)$, we have

$$\Pr_{T \sim \mu_{P_n}}(\mathbb{I}[(u, v) \in E(T)]) = 1,$$

and for each edge $(u, v) \in E(C_n)$,

$$\Pr_{T \sim \mu_{C_n}}(\mathbb{I}[(u, v) \in E(T)]) = \frac{n-1}{n} = 1 - \frac{1}{n}.$$

Thus, by applying Lemma 2, we arrive at the desired conclusion: for $(u, v) \in E(P_n)$, $R(u, v) = 1$, while for $(u, v) \in E(C_n)$, $R(u, v) = 1 - \frac{1}{n}$. The biconnectivity of these two graph types is evident. \square

1080 G PROOF OF THEOREM 2

1081
1082 **Theorem 2.** *Given any edge $(u, v) \in E$ in an undirected, unweighted graph G , (u, v) is a cut edge*
1083 *if and only if the resistance distance $R(u, v) = 1$.*

1084
1085 *Proof.* It is straightforward to show that an edge (u, v) is a cut edge if and only if it is included in
1086 every spanning tree of G . Therefore, by Lemma 2, we have

$$1087 R(u, v) = \Pr_{T \sim \mu_G} (\mathbb{I}[(u, v) \in E(T)]) = 1.$$

1088
1089 □

1091 H PROOF OF THEOREM 4

1092
1093 Before proving Theorem 4, we present some basic facts about the CONGEST model. First, we show
1094 that a spanning tree rooted at a node u can be constructed using the FLOOD algorithm.

1095 **Lemma 4** ((Peleg, 2000), FLOOD Algorithm). *There exists a CONGEST algorithm in which a*
1096 *designated node $u \in V$ can construct a spanning tree T rooted at u with depth $\text{depth}(T) =$*
1097 *$\max_v d_G(u, v)$ in $\max_v d_G(u, v) = O(D)$ rounds, where D is the diameter of the graph.*

1098 The idea behind the FLOOD algorithm is straightforward: Initially, the source node u sends a special
1099 token to all its neighbors. Each node, upon receiving the token for the first time, stores it and
1100 forwards it to its neighbors. If a node receives the token again, it discards it and does nothing.

1101
1102 Additionally, we include the following lemmas, which describe the ability to broadcast and collect
1103 messages to and from a designated node.

1104 **Lemma 5** ((Peleg, 2000), DOWNCAST Algorithm). *There exists a CONGEST algorithm in which,*
1105 *given M messages (of $\Theta(\log n)$ bit) stored at a designated node $u \in V$, and a spanning tree T*
1106 *rooted at u , the messages can be broadcast to other nodes in $O(\text{depth}(T) + M)$ rounds.*

1107 **Lemma 6** ((Peleg, 2000), UPCAST Algorithm). *There exists a CONGEST algorithm in which, given*
1108 *M messages stored at different nodes and a spanning tree T rooted at u , the messages can be*
1109 *collected at node u in $O(\text{depth}(T) + M)$ rounds.*

1110 It is important to note that the conclusions for the DOWNCAST and UPCAST algorithms above are
1111 derived under the standard CONGEST model, where each edge can transmit only $O(1)$ messages of
1112 size $\Theta(\log n)$ bits per communication round. If we relax this restriction to allow the transmission of
1113 w messages of size $\Theta(\log n)$ bits per round, the round complexities of the two algorithms reduce to
1114 $O\left(\text{depth}(T) + \frac{M}{w}\right)$ by grouping messages together.

1115
1116 With these tools in hand, we are now ready to prove the theorem.

1117 **Theorem 4.** *If we allow a single node to have unbounded computational power to solve any com-*
1118 *putable problem, then every NP-Complete decision problems on undirected unweighted graphs can*
1119 *be solved by the CONGEST model in $O(m)$ rounds.*

1120
1121 *Proof.* Given an NP-Complete problem on an undirected, unweighted graph, such as deciding
1122 whether there is a k -clique in a graph G , we proceed as follows. Since each node is assigned a
1123 unique ID in $[n]$, we designate node 0 as the leader without loss of generality.

1124 First, in one round, each node u collects the IDs of its neighbors and forms $d(u)$ messages of the
1125 form $(\text{ID}(u), \text{ID}(v))$ for each $v \in N(u)$. Next, we invoke the FLOOD algorithm to construct a
1126 spanning tree rooted at node 0 in $O(D)$ rounds. Afterward, we apply the UPCAST algorithm to
1127 gather a total of $\sum_{u \in V} d(u) = \Theta(m)$ messages in $O(D + m)$ rounds. At this point, node 0 has
1128 complete knowledge of the graph's topology. Since node 0 can solve any computable problem, it
1129 can solve the NP-Complete problem locally in one round. Finally, node 0 uses the DOWNCAST
1130 algorithm to broadcast the result to all other nodes in $O(D)$ rounds.

1131 Thus, the total number of communication rounds is

$$1132 1 + O(D) + O(D + m) + O(D) = O(m).$$

1133 □

I THE WEISFEILER-LEHMAN PROBLEM AND ITS HARDNESS (THEOREM 5)

In this section, we formally define the meaning of one iteration of the Weisfeiler-Lehman (WL) test and prove its computational hardness. To facilitate our presentation, we first define the WL relation.

Definition 6 (Weisfeiler-Lehman Relation). *Given an unweighted, undirected graph G , the Weisfeiler-Lehman relation on G , denoted by $WL(G)$, is the set of color vector pairs $(\mathbf{x}, \mathbf{y}) \in [n]^n \times [n]^n$ satisfying:*

$$\forall u \in V, v \in V, \mathbf{y}_u = \mathbf{y}_v \Leftrightarrow \mathbf{x}_u = \mathbf{x}_v \wedge \{\{\mathbf{x}_z : z \in N(u)\}\} = \{\{\mathbf{x}_z : z \in N(v)\}\}.$$

The formal definition of one iteration of the WL test is captured by the deterministic, zero-error randomized, and bounded-error randomized WL problems, as described below.

Definition 7 (Weisfeiler-Lehman Problem). *Given an unweighted, undirected graph G with n nodes, where each node v is initially assigned a color $\mathbf{x}_v \in [n]$, the goal is to assign a new color $\mathbf{y}_v \in [n]$ for each node such that:*

- **Deterministic:** $(\mathbf{x}, \mathbf{y}) \in WL(G)$, or
- **Randomized, Zero-Error:** $\Pr((\mathbf{x}, \mathbf{y}) \in WL(G)) = 1$, or
- **Randomized, Bounded-Error:** $\Pr((\mathbf{x}, \mathbf{y}) \in WL(G)) \geq 1 - \epsilon$.

We now formally state and prove theorems regarding the hardness of one iteration of the WL test.

Theorem 5. *If an RL-CONGEST model can simulate **one** iteration of the WL test without preprocessing, either deterministically or randomly with zero error, regardless of the computational power available to each node, the model’s width w and depth d must satisfy $d = \Omega\left(D + \frac{m}{w \log n}\right)$, given that $w = o\left(\frac{n}{\log n}\right)$.*

The proof of the theorem relies on tools from communication complexity, so we recommend that readers refer to Appendix B for a basic understanding of these concepts.

Proof. We will prove that for any positive integer n and any positive integer m such that $m \in [n, n^2]$, there exists a hard-case graph with $\Theta(n)$ nodes and $\Theta(m)$ edges. Given n and m , we first construct an incomplete “basic” graph $G_{(n,m)}$ with $\Theta(n)$ nodes and $\Theta(m)$ edges, partitioned between Alice (A) and Bob (B), as follows:

- Alice and Bob each have nodes $x^{(A)}$ and $x^{(B)}$, connected by an edge;
- They also hold nodes $w_i^{(A)}$ and $w_i^{(B)}$ ($i = 1, 2, \dots, \lceil \frac{m}{n} \rceil$), connected to $x^{(A)}$ and $x^{(B)}$ respectively;
- Additionally, they possess nodes $u_i^{(A)}, u_i^{(B)}, v_i^{(A)},$ and $v_i^{(B)}$ ($i = 1, 2, \dots, n$);
- Add edges $\left\{ \left(u_n^{(A)}, v_n^{(A)} \right) \right\} \cup \left\{ \left(u_i^{(A)}, u_{i+1}^{(A)} \right) : i \in \{1, 2, \dots, n-1\} \right\} \cup \left\{ \left(v_i^{(A)}, v_{i+1}^{(A)} \right) : i \in \{1, 2, \dots, n-1\} \right\}$ to form a path with nodes $u_i^{(A)}$, and $v_i^{(A)}$. Repeat similarly for nodes $u_i^{(B)}$, and $v_i^{(B)}$, to form another path.


Then, we assign each node’s color \mathbf{x}_u as follows:

- $\mathbf{x}_{x^{(A)}} = \mathbf{x}_{x^{(B)}} = 0$;

⁶Our theorems also hold for any bounded color spaces, such as $[p(n)]$ for any polynomial p . However, since n colors are always sufficient, we can simply set the color space to $[n]$.

- 1188 • For each $i \in \{1, 2, \dots, n\}$, $\mathbf{x}_{u_i^{(A)}} = \mathbf{x}_{u_i^{(B)}} = i$;
- 1189
- 1190 • For each $i \in \{1, 2, \dots, n\}$, $\mathbf{x}_{v_i^{(A)}} = \mathbf{x}_{v_i^{(B)}} = n + i$;
- 1191
- 1192 • For each $i \in \left\{1, 2, \dots, \left\lceil \frac{m}{n} \right\rceil\right\}$, $\mathbf{x}_{w_i^{(A)}} = \mathbf{x}_{w_i^{(B)}} = 2n + i$;
- 1193
- 1194
- 1195

1196 The constructed basic graph is illustrated in Figure 5.



1222 Figure 5: The constructed basic graph $G_{(n,m)}$. Nodes are colored according to \mathbf{x} .

1223 Alice and Bob also fix a bijection c in advance between the set of index pairs $\left\{1, 2, \dots, \left\lceil \frac{m}{n} \right\rceil\right\} \times$
 1224 $\{1, 2, \dots, n\}$ and the set $\left\{1, 2, \dots, n \cdot \left\lceil \frac{m}{n} \right\rceil\right\}$ ⁷. For example, define $c((i, j)) = (i - 1)n + j$ and
 1225 $c^{-1}(i) = \left(\left\lceil \frac{i}{n} \right\rceil, (i - 1) \bmod n + 1\right)$.

1226 Given an instance $(\mathbf{a}, \mathbf{b}) \in \{0, 1\}^m \times \{0, 1\}^m$ of EQ_m , Alice receives $\mathbf{a} = (\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_m)$ and
 1227 Bob receives $\mathbf{b} = (\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_m)$, they complete $G_{(n,m)}$ to $G_{(n,m);(\mathbf{a},\mathbf{b})}$ with $\Theta(n)$ nodes and
 1228 $\Theta(m)$ edges as follows:

- 1229
- 1230 • For each $k \in \{1, 2, \dots, m\}$, let $(i, j) = c^{-1}(k)$. If $\mathbf{a}_k = 0$, Alice adds edge $(w_i^{(A)}, u_j^{(A)})$;
- 1231 otherwise, Alice adds edge $(w_i^{(A)}, v_j^{(A)})$;
- 1232
- 1233 • For each $k \in \{1, 2, \dots, m\}$, let $(i, j) = c^{-1}(k)$. If $\mathbf{b}_k = 0$, Bob adds edge $(w_i^{(B)}, u_j^{(B)})$;
- 1234 otherwise, Bob adds edge $(w_i^{(B)}, v_j^{(B)})$;
- 1235
- 1236
- 1237
- 1238
- 1239
- 1240

1241 ⁷Since $\frac{m}{n} \leq \left\lceil \frac{m}{n} \right\rceil < \frac{m}{n} + 1$, we have $m \leq n \left\lceil \frac{m}{n} \right\rceil < m + n$.

We claim that given $G_{(n,m);(a,b)}$, for any new color vector \mathbf{y} such that $(\mathbf{x}, \mathbf{y}) \in \text{WL}(G_{(n,m);(a,b)})$, we have:

$$\mathbf{a} = \mathbf{b} \iff \forall i \in \left\{1, 2, \dots, \left\lceil \frac{m}{n} \right\rceil\right\}, \mathbf{y}_{w_i^{(A)}} = \mathbf{y}_{w_i^{(B)}}.$$

- “ \Rightarrow ” is straightforward. Since the construction is symmetric, for any $i \in \{1, \dots, \lceil \frac{m}{n} \rceil\}$ we have $\left\{\left\{\mathbf{x}_k : k \in N(w_i^{(A)})\right\}\right\} = \left\{\left\{\mathbf{x}_k : k \in N(w_i^{(B)})\right\}\right\}$ and thus

$$\left(\mathbf{x}_{w_i^{(A)}}, \left\{\left\{\mathbf{x}_k : k \in N(w_i^{(A)})\right\}\right\}\right) = \left(\mathbf{x}_{w_i^{(B)}}, \left\{\left\{\mathbf{x}_k : k \in N(w_i^{(B)})\right\}\right\}\right).$$

Therefore, $\forall i \in \left\{1, \dots, \left\lceil \frac{m}{n} \right\rceil\right\}, \mathbf{y}_{w_i^{(A)}} = \mathbf{y}_{w_i^{(B)}}.$

- To prove the other direction, we show its contrapositive: $\mathbf{a} \neq \mathbf{b} \Rightarrow \exists i, \mathbf{y}_{w_i^{(A)}} \neq \mathbf{y}_{w_i^{(B)}}.$ Since $\mathbf{a} \neq \mathbf{b}$, there exists $1 \leq k \leq m$ such that $\mathbf{a}_k \neq \mathbf{b}_k.$ Without loss of generality, assume $\mathbf{a}_k = 0$ and $\mathbf{b}_k = 1,$ and let $(i, j) = c^{-1}(k).$ According to our construction, for Alice, $j \in \left\{\left\{\mathbf{x}_v : v \in N(w_i^{(A)})\right\}\right\}$ and $n + j \notin \left\{\left\{\mathbf{x}_v : v \in N(w_i^{(A)})\right\}\right\},$ while for Bob, $j \notin \left\{\left\{\mathbf{x}_v : v \in N(w_i^{(B)})\right\}\right\}$ and $n + j \in \left\{\left\{\mathbf{x}_v : v \in N(w_i^{(B)})\right\}\right\}.$ Therefore, $\mathbf{y}_{w_i^{(A)}} \neq \mathbf{y}_{w_i^{(B)}}.$

Now, assume that $\text{WL}(G_{(n,m);(a,b)})$ can be solved by an RL-CONGEST model on $G_{(n,m);(a,b)}$ deterministically, or randomly with zero error, in d rounds. We can construct an algorithm that solves EQ_m in no more than $d + 1 + \left\lceil \frac{m}{n} \right\rceil$ rounds. This can be done by first using d rounds to compute $\mathbf{y},$ then using 1 round for $x^{(A)}$ and $x^{(B)}$ to collect the colors of their neighbors, and finally using $\left\lceil \frac{m}{n} \right\rceil$ rounds to compare the colors between $w_i^{(A)}$ and $w_i^{(B)}.$

According to the results in communication complexity, we have $D(\text{EQ}_m) = R_0(\text{EQ}_m) = \Theta(m),$ which implies that the total number of communicated bits satisfies

$$\left(d + 1 + \left\lceil \frac{m}{n} \right\rceil\right) w \log |V(G_{(n,m);(a,b)})| = \Omega(m).$$

This is equivalent to

$$d + \Theta\left(\frac{m}{n}\right) = \Omega\left(\frac{m}{w \log n}\right).$$

Therefore, when $w = o\left(\frac{n}{\log n}\right),$ we have $d = \Omega\left(\frac{m}{w \log n}\right).$

The $\Omega(D)$ component is straightforward. By adding a path of length $\Theta(D)$ between $x^{(A)}$ and $x^{(B)},$ any message exchange between the two parties will require $\Omega(D)$ rounds. \square

J PROOF OF THEOREM 6

Theorem 6. *There exists a deterministic RL-CONGEST algorithm that can simulate **one** iteration of the WL test without preprocessing, with width w and depth d satisfying $d = O\left(D + \frac{m}{w}\right).$ Additionally, it is sufficient to set the nodes’ computational resource class to $\mathcal{C} = \text{DTIME}(n^2 \log n).$*

Proof. We present the framework of our algorithm and analyze the round complexity for each step:

1. Each node u sends a message (u, \mathbf{x}_u) to its neighbors and receives messages from them to form the set $S_u = \{(u, v, \mathbf{x}_v) : v \in N(u) \cup u\}.$ This process takes $O(1)$ rounds.
2. Node 0 initiates the FLOOD algorithm to construct a BFS spanning tree rooted at node 0. This process takes $O(D)$ rounds.

- 1296 3. The UPCASt algorithm is used to collect all sets S_u at the root node 0 along the spanning
 1297 tree. This process takes $O\left(D + \frac{m}{w}\right)$ rounds, as there are $\sum_{u \in V} O(d_u) = O(m)$ messages
 1298 to gather, and each edge can transmit w messages per round.
 1299
- 1300 4. Node 0 merges all S_u to form the set $K = \{(u, (\mathbf{x}_u, \{\mathbf{x}_v : v \in N(u)\})) : u \in V\}$. This
 1301 step can be completed in one round and requires $O(m)$ time⁸, since there are $O(m)$ tuples
 1302 in $\bigcup_u S_u$.
 1303
- 1304 5. Node 0 sorts K by $(\mathbf{x}_u, \{\mathbf{x}_v : v \in N(u)\})$ to create the ordered set $K' =$
 1305 $\{((\mathbf{x}_u, \{\mathbf{x}_v : v \in N(u)\}), u) : u \in V\}$. It then determines new colors for each node
 1306 by the rank of $(\mathbf{x}_u, \{\mathbf{x}_v : v \in N(u)\})$. The computed color mapping is represented as
 1307 $\{(u, \mathbf{y}_u) : u \in V\}$. This process can be done in one round and requires $O(n \log n)$ com-
 1308 parisons, with each comparison taking $O(\Delta)$ time, resulting in $O(n\Delta \log n)$ time.
 1309
- 1310 6. The DOWNCAST algorithm is used to send the results back to each node along the spanning
 1311 tree. This process takes $O\left(D + \frac{n}{w}\right)$ rounds, as there are $O(n)$ messages to transmit.
 1312

1312 Thus, the WL problem can be computed by the RL-CONGEST model in
 1313

$$1314 \quad d = O\left(1 + D + D + \frac{m}{w} + D + \frac{n}{w}\right) = O\left(D + \frac{m}{w}\right)$$

1315 rounds. The computational resource class of each node is determined by steps 4 and 5. \square
 1316
 1317

1318 K PROOF OF THEOREM 7

1320 **Theorem 7.** *There exists a deterministic RL-CONGEST algorithm that can simulate **one** iteration*
 1321 *of the WL test by adding a virtual node, which connects to other nodes, as preprocessing. The*
 1322 *algorithm operates with width w and depth d satisfying $dw = O(\Delta)$, where Δ is the maximum*
 1323 *degree of the graph before the addition of the virtual node. Additionally, it is sufficient to set the*
 1324 *nodes' computational resource class to $\mathcal{C} = \text{DTIME}(n^2 \log n)$.*
 1325

1326 *Proof.* Without loss of generality, we denote the added virtual node as node n , which is known to all
 1327 nodes in the original graph. We then outline the framework of our algorithm and analyze the round
 1328 complexity for each step:

- 1329
- 1330 1. Each node u , except the virtual node, sends a message (u, \mathbf{x}_u) to its neighbors and receives
 1331 messages from them, forming the set $S_u = \{(u, v, \mathbf{x}_v) : v \in N(u) \cup u\}$. This process
 1332 takes $O(1)$ rounds.
 - 1333 2. The virtual node n sends a token to each node to notify them of the edge along which the
 1334 virtual node can be reached. This process takes 1 round.
 - 1335 3. Each node u , except the virtual node, sends its set S_u to the virtual node n along the
 1336 edge connecting to it. This process takes $O\left(\frac{\Delta}{w}\right)$ rounds, as each node has at most $\Theta(\Delta)$
 1337 messages to send, excluding the virtual node.
 1338
 - 1339 4. The virtual node n merges all sets S_u to form the set $K =$
 1340 $\{(u, (\mathbf{x}_u, \{\mathbf{x}_v : v \in N(u)\})) : u \in V\}$. This step can be completed in 1 round and
 1341 requires $O(m)$ time⁹, as there are $O(m)$ tuples in $\bigcup_u S_u$.
 1342
 - 1343 5. The virtual node n sorts K by $(\mathbf{x}_u, \{\mathbf{x}_v : v \in N(u)\})$ to create the ordered set $K' =$
 1344 $\{((\mathbf{x}_u, \{\mathbf{x}_v : v \in N(u)\}), u) : u \in V\}$. It then assigns new colors to each node based on
 1345 the rank of $(\mathbf{x}_u, \{\mathbf{x}_v : v \in N(u)\})$. The color mapping is represented as $\{(u, \mathbf{y}_u) : u \in$
 1346 $V\}$. This process can be completed in 1 round and requires $O(n \log n)$ comparisons, with
 1347 each comparison taking $O(\Delta)$ time, resulting in $O(n\Delta \log n)$ total time.
 1348

1349 ⁸Assuming a WordRAM machine where each word consists of $\Theta(\log n)$ bits.

⁹Assuming a WordRAM model where each word consists of $\Theta(\log n)$ bits.

- 1350 6. The virtual node n sends the new colors \mathbf{y}_u to their corresponding nodes u along the con-
 1351 necting edges. This process takes 1 round, as each edge transmits only one message.
 1352

1353 Thus, the WL problem can be computed in the RL-CONGEST model with the addition of a virtual
 1354 node in

$$1355 d = O\left(1 + 1 + \frac{\Delta}{w} + 1\right) = O\left(\frac{\Delta}{w}\right)$$

1356 rounds, or equivalently, $dw = O(\Delta)$. The computational resource class of each node is determined
 1357 by steps 4 and 5. \square

1361 L COMPLEXITY OF THE PNF MODEL CHECKING PROBLEM

1362
 1363 As mentioned in Section 4.3, the \mathcal{L}^k model checking problem is crucial in various areas of computer
 1364 science. We will introduce some key results that characterize the complexity of its special case, the
 1365 PNF \mathcal{L}^k model checking problem.

1366 On one hand, the first two results establish the upper bound of the time complexity for the PNF \mathcal{L}^k
 1367 model checking problem **on graphs**, where the predicates are limited to $E(x, y)$ and $=(x, y)$.
 1368

1369 **Theorem 9** ((Eisenbrand & Grandoni, 2004; Williams, 2014)). *For $k \geq 3$, every k -quantifier (PNF)
 1370 first-order sentence on n -node graphs can be decided in $\tilde{O}(n^{k-3+\omega})$ time.*

1371 **Theorem 10** ((Eisenbrand & Grandoni, 2004; Williams, 2014)). *For $k \geq 9$, every k -quantifier
 1372 (PNF) first-order sentence on n -node graphs can be decided in $n^{k-1+o(1)}$ time.*

1373
 1374 On the other hand, the following result establishes a lower bound for the PNF \mathcal{L}^k model checking
 1375 problem on graphs, conditioned on a well-known hypothesis in theoretical computer science:

1376 **Theorem 11** ((Puatracscu & Williams, 2010; Williams, 2014)). *For $k \geq 4$, if the model checking
 1377 problem for k -quantifier (PNF) first-order sentences over graphs can be solved in $O(n^{k-1-\epsilon})$ time
 1378 for some $\epsilon > 0$, then the Strong Exponential Time Hypothesis (SETH) is false.*

1379
 1380 Note that the algorithms in Theorem 9 and 10 leverage fast matrix multiplication techniques to im-
 1381 prove the brute-force algorithm by a factor of n , in contrast to the combinatorial operations typically
 1382 employed in GNNs. Motivated by this, many works have attempted to find a faster combinatorial al-
 1383 gorithm that does not rely on fast matrix multiplication and can outperform the brute-force approach
 1384 by a polynomial factor of n^ϵ for some $\epsilon > 0$, but to date, all such attempts have failed. This has led
 1385 to the formulation of the combinatorial k -clique conjecture, which is stated as follows:

1386 **Conjecture 1** (Combinatorial k -Clique Conjecture, (Abboud et al., 2017; Bringmann et al., 2017;
 1387 Abboud et al., 2024)). *For any $k \geq 3$ and any $\epsilon > 0$, no combinatorial algorithm can determine
 1388 whether a graph contains a k -clique in $O(n^{k-\epsilon})$ time.*

1389 Since the k -clique problem is a special case of the PNF \mathcal{L}^k model checking problem, it follows
 1390 that no combinatorial algorithm can solve the PNF \mathcal{L}^k model checking problem in $O(n^{k-\epsilon})$ time
 1391 for $k \geq 3$. Therefore, when focusing exclusively on combinatorial algorithms, it is reasonable to
 1392 classify the PNF \mathcal{L}^k model checking problem in the $\tilde{\Theta}(n^k)$ complexity class.

1393 For more general cases, such as PNF \mathcal{L}^k model checking not limited to graphs, there are also results
 1394 concerning both upper and lower bounds. First, Gao et al. (2017) designed an algorithm to solve the
 1395 PNF \mathcal{L}^k model checking problem in $m^{k-1-o(1)}$ time, where m is the size of the input structure, i.e.,
 1396 the number of all tuples in the relations, which is equivalent to the number of edges when the model
 1397 is a graph.

1398 **Theorem 12** ((Gao et al., 2017)). *There exists an algorithm that solves the (PNF) \mathcal{L}^k model checking
 1399 problem in time $m^{k-1}/2^{\Theta(\sqrt{\log m})}$.*

1400 They also prove the near-optimality of their algorithm under SETH:

1401 **Theorem 13** ((Gao et al., 2017)). *Assuming SETH, no algorithm can solve the (PNF) \mathcal{L}^k model
 1402 checking problem in $O(m^{k-1-\epsilon})$ time for any $\epsilon > 0$.*

Therefore, it is reasonable to classify the PNF \mathcal{L}^k model checking problem in the $\tilde{\Theta}(m^{k-1})$ complexity class. By combining the two results, it is reasonable to classify the PNF \mathcal{L}^k model checking problem in the $\Theta(\min\{n^k, m^{k-1}\})$ complexity class.

M PROOF OF THEOREM 8

Theorem 8. *If constructing the k -WL graph and additional features in $O(k \cdot \text{poly}|\phi| \cdot n^{k+1})$ time is allowed, where $|\phi|$ represents the length of the quantifier-free formula, the RL-CONGEST model can solve the PNF \mathcal{C}^k model checking problem in $O(k^2)$ rounds. Additionally, the computational resources required by each node are $\mathcal{C} = \text{DTIME}(k^2 n)$.*

Proof. We prove this theorem by directly constructing a solution under the RL-CONGEST framework to solve the PNF \mathcal{C}^k model checking problem. This involves constructing the k -WL graph and computing features as part of the preprocessing, followed by a message-passing process.

Given an input graph G and a PNF \mathcal{C}^k sentence $\varphi := (Q_1 u_1)(Q_2 u_2) \cdots (Q_k u_k) \phi(u_1, u_2, \dots, u_k)$, where Q_i are quantifiers from $\{\forall, \exists, \exists^{\leq i}, \exists^{=i}\}$, and $\phi(u_1, \dots, u_k)$ is a quantifier-free formula¹⁰ with predicates limited to $E(\cdot, \cdot)$ and $=(\cdot, \cdot)$, the overall framework using the RL-CONGEST model to decide whether $G \models \varphi$ proceeds as follows:

1. Preprocessing:

- **Construct the k -WL graph G' :** Enumerate all k -tuples $\mathbf{v} = (v_1, \dots, v_k) \in V^k$ to form the new node set V' for the k -WL graph. Each node has a unique ID ($\text{ID}(v_1), \dots, \text{ID}(v_k)$) in $[n]^k$. For simplicity, we use v_i and $\text{ID}(v_i)$ interchangeably. Then, connect nodes with Hamming distance one, that is, $E' = \{(\mathbf{u}, \mathbf{v}) \in V' \times V' : d_H(\mathbf{u}, \mathbf{v}) = 1\}$.
- **Construct a circuit for $\phi(u_1, \dots, u_k)$:** We define a bijection f between $\{1, 2, \dots, k\} \times \{1, 2, \dots, k\}$ and $\{1, 2, \dots, k^2\}$ such that $f(i, j) = (i-1)k + j$. Next, we introduce $2k^2$ auxiliary variables corresponding to $E(u_i, u_j)$ and $=(u_i, u_j)$:
 - For each pair $(i, j) \in \{1, \dots, k\} \times \{1, \dots, k\}$, let variable $x_{f(i,j)}$ represent $E(u_i, u_j)$.
 - For each pair $(i, j) \in \{1, \dots, k\} \times \{1, \dots, k\}$, let variable $x_{k^2+f(i,j)}$ represent $=(u_i, u_j)$.

A Boolean circuit $C_\phi(x_1, \dots, x_{k^2}, \dots, x_{2k^2})$ is built, taking $2k^2$ inputs. This is done by replacing atomic formulas $E(u_i, u_j)$ and $=(u_i, u_j)$ in ϕ with $x_{f(i,j)}$ and $x_{k^2+f(i,j)}$, respectively, and adding logic gates to form C_ϕ .

- **Compute features $\phi(\mathbf{v})$:** For each node $\mathbf{v} = (v_1, \dots, v_k)$, representing an assignment of the variables (u_1, \dots, u_k) , we compute a feature vector $\mathbf{x}(\mathbf{v}) = (x_1(\mathbf{v}), \dots, x_{2k^2}(\mathbf{v}))$, where:
 - For each pair (i, j) , $x_{f(i,j)}(\mathbf{v}) = \mathbb{I}[E(v_i, v_j)] \in \{0, 1\}$, indicating whether v_i and v_j are connected.
 - Similarly, $x_{k^2+f(i,j)}(\mathbf{v}) = \mathbb{I}[=(v_i, v_j)] \in \{0, 1\}$, indicating whether v_i and v_j are the same node in G .

By applying the C_ϕ circuit with input $\mathbf{x}(\mathbf{v})$ to each vertex \mathbf{v} , we obtain a one-bit feature for each node. It can be easily verified that $C_\phi(\mathbf{x}(\mathbf{v})) = \phi(\mathbf{v})$ for any node $\mathbf{v} \in V'$.

We illustrate this preprocessing process with an example. Let G be the graph in Figure 6, $k = 3$, and $\phi(u_1, u_2, u_3) := E(u_1, u_2) \wedge E(u_2, u_3) \wedge E(u_3, u_1)$. The auxiliary variables x_i and their corresponding atomic formulas are listed in Table 2. The circuit C_ϕ is given by $C_\phi(x_1, \dots, x_{18}) = x_2 \wedge x_6 \wedge x_7$. For node $\mathbf{v} = (0, 1, 3) \in V'$, the values of each auxiliary variable on \mathbf{v} are shown in the third column of Table 2. Applying the circuit to $\mathbf{x}((0, 1, 3))$, we obtain an output of 1.

We then analyze the time complexity of the preprocessing phase. For the construction of the k -WL graph, since there are $\Theta(n^k)$ nodes and $\Theta(kn^{k+1})$ edges, the construction takes $\Theta(kn^{k+1})$

¹⁰We assume that the logical connectives are limited to \wedge , \vee , and \neg . If other connectives are present, they can be transformed into an equivalent formula using only these connectives in $\text{poly}|\phi|$ time, where $|\phi|$ is the length of the formula.




Figure 6: An example graph illustrating the preprocessing process.

Table 2: Auxiliary variables, their corresponding atomic formulas, and their values on node $(0, 1, 3)$.

Auxiliary Variables	Corresponding Atomic Formulas	Values on Node $(0, 1, 3)$
x_1	$E(u_1, u_1)$	$= [E(0, 0)] = 0$
x_2	$E(u_1, u_2)$	$= [E(0, 1)] = 1$
x_3	$E(u_1, u_3)$	$= [E(0, 3)] = 1$
x_4	$E(u_2, u_1)$	$= [E(1, 0)] = 1$
x_5	$E(u_2, u_2)$	$= [E(1, 1)] = 0$
x_6	$E(u_2, u_3)$	$= [E(1, 3)] = 1$
x_7	$E(u_3, u_1)$	$= [E(3, 0)] = 1$
x_8	$E(u_3, u_2)$	$= [E(3, 1)] = 1$
x_9	$E(u_3, u_3)$	$= [E(3, 3)] = 0$
x_{10}	$=(u_1, u_1)$	$= [= (0, 0)] = 1$
x_{11}	$=(u_1, u_2)$	$= [= (0, 1)] = 0$
x_{12}	$=(u_1, u_3)$	$= [= (0, 3)] = 0$
x_{13}	$=(u_2, u_1)$	$= [= (1, 0)] = 0$
x_{14}	$=(u_2, u_2)$	$= [= (1, 1)] = 1$
x_{15}	$=(u_2, u_3)$	$= [= (1, 3)] = 0$
x_{16}	$=(u_3, u_1)$	$= [= (3, 0)] = 0$
x_{17}	$=(u_3, u_2)$	$= [= (3, 1)] = 0$
x_{18}	$=(u_3, u_3)$	$= [= (3, 3)] = 1$

time. The circuit construction involves simple operations, such as substituting variables, which takes $O(\text{poly}|\phi|)$ time. For the feature computation, we compute $\Theta(k^2)$ features for each node $\mathbf{v} \in V^k$ and pass it through the circuit C_ϕ to obtain $\phi(\mathbf{v})$. This process takes $O(\text{poly}|\phi| \cdot n^k)$ time. Thus, the total time complexity for the preprocessing phase is $O(\text{poly}|\phi| \cdot n^k + kn^{k+1}) = O(k \cdot \text{poly}|\phi| \cdot n^{k+1})$.

2. **Message-Passing with Limited Computational Resources:** We present an $O(k^2)$ -round message-passing algorithm on G' , where each node $\mathbf{v} \in V'$ has a feature $\phi(\mathbf{v})$ obtained during the preprocessing phase. The key idea is to treat the k -WL graph as an implicit n -ary prefix tree and perform dynamic programming to eliminate quantifiers over this tree. The algorithm is described as follows:

(a) **Initial Setup:**

- Each node $\mathbf{v} \in V'$ has a one-bit feature $\phi(\mathbf{v})$ from the preprocessing phase.
- For each node $\mathbf{v} \in V'$:
 - Sends its ID along with its feature to its neighbors¹¹.
 - Receives messages from neighbors.

(b) **Main Iteration ($\ell = k$ to 1):**

- We define the active nodes in the ℓ -th iteration as those whose IDs take the form $(v_1, \dots, v_{\ell-1}, 0, 0, \dots)$, meaning the last $k - \ell + 1$ coordinates are zeros.
- For each active node $(v_1, \dots, v_{\ell-1}, 0, \dots)$:
 - Collects all received features from nodes whose IDs match the pattern $(v_1, \dots, v_{\ell-1}, 0, \dots)$, where $v_\ell \in V$. This forms a multiset of n bits.
 - Updates its feature based on the quantifier Q_ℓ :
 - * If Q_ℓ is \forall , set the node's feature to 1 if all n bits in the multiset are 1; otherwise, set it to 0.

¹¹This message passing takes $O(k)$ communication rounds.

- 1512 * If Q_ℓ is \exists , set the feature to 1 if at least one 1 exists in the multiset; otherwise, set it
 1513 to 0.
 1514 * Similarly, handle other quantifiers ($\exists^{\geq i}$, $\exists^{\leq i}$, etc.).
 1515 – Sends its updated feature and ID to its neighbors.
 1516 – Receives messages from neighbors.
 1517 • Non-active nodes change their feature to \perp and remain silent for the rest of the process.
 1518 (c) **Final Result:** The final result bit is found at the node $(0, 0, \dots, 0) \in V'$. To propagate the
 1519 result to all nodes in V' , we use an additional k rounds.
 1520

1521 Since there are k iterations, and each iteration requires $O(k)$ communication rounds to send
 1522 messages, the total number of communication rounds is $k \cdot O(k) + k = O(k^2)$. In each iteration,
 1523 each node examines $O(kn)$ messages, each containing $O(k)$ words, resulting in $O(k^2n)$ time
 1524 complexity. Therefore, the computational complexity class for this RL-CONGEST algorithm is
 1525 $\mathbf{C} = \text{DTIME}(k^2n)$.

□

1526
 1527
 1528
 1529
 1530
 1531
 1532
 1533
 1534
 1535
 1536
 1537
 1538
 1539
 1540
 1541
 1542
 1543
 1544
 1545
 1546
 1547
 1548
 1549
 1550
 1551
 1552
 1553
 1554
 1555
 1556
 1557
 1558
 1559
 1560
 1561
 1562
 1563
 1564
 1565