A LOCAL View of the Polynomial Hierarchy*

Fabian Reiter ✉ ⦿

LIGM, Université Gustave Eiffel, Marne-la-Vallée, France

───── **Abstract** ─────

We extend classical methods of computational complexity to the setting of distributed computing, where they are sometimes more effective than in their original context. Our focus is on distributed decision in the LOCAL model, where multiple networked computers communicate via synchronous message-passing to collectively answer a question about their network topology. Rather unusually, we impose two orthogonal constraints on the running time of this model: the number of communication rounds is bounded by a constant, and the number of computation steps of each computer is polynomially bounded by the size of its local input and the messages it receives.

By letting two players take turns assigning certificates to all computers in the network, we obtain a generalization of the polynomial hierarchy (and hence of the complexity classes **P** and **NP**). We then extend some key results of complexity theory to this setting, in particular the Cook–Levin theorem (which identifies Boolean satisfiability as a complete problem for **NP**), and Fagin's theorem (which characterizes **NP** as the problems expressible in existential second-order logic). The original results can be recovered as the special case where the network consists of a single computer. But perhaps more surprisingly, the task of separating complexity classes becomes easier in the general case: we can show that our hierarchy is *infinite*, while it remains notoriously open whether the same is true in the case of a single computer. (By contrast, a collapse of our hierarchy would have implied a collapse of the polynomial hierarchy.) As an application, we propose quantifier alternation as a new approach to measuring the locality of problems in distributed computing.

arXiv:2305.09538v3 [cs.DC] 6 Nov 2023

───────

* Typeset with the `knowledge` package: technical terms and symbols are hyperlinked to their definitions.

Contents

1 Introduction

In this paper, we revisit classical computational complexity theory from the perspective of distributed network computing. As we will see, certain standard notions and techniques not only extend well to the distributed setting, but in some cases allow us to achieve more there than in the centralized setting. We begin by setting the context, and then present our approach and results.

1.1 Background

When solving a problem in a computer network using a distributed algorithm, a major concern is the issue of *locality*. At its core lies the question of how much information each computer needs to obtain about the rest of the network in order to solve the given problem. The less information needed, the more local the problem.

The LOCAL model. In the late 1980s, Linial [27] introduced an influential model of distributed computing that focuses entirely on locality, while abstracting away many other issues such as failures, asynchrony, and bandwidth limitations. In this model, which Peleg [33] later called the LOCAL model, a network consists of several computers that communicate with their neighbors by exchanging messages through a sequence of fault-free synchronous rounds. The computers, referred to as nodes, are all identical except for possessing globally unique identifiers. They have unlimited computational power to process their local input and the messages they receive in each round, and there is no limitation on the message sizes. The goal in this setting is for the nodes to collectively solve some graph problem related to the topology of their network. That is, the network serves both as the communication infrastructure and as the input graph. Typically, the problem is a construction task such as finding a (vertex or edge) coloring, a maximal matching, a maximal independent set, or a spanning tree. After a finite number of rounds, each node should produce a local output such as "my color is blue" or "I belong to the independent set", and the combined output of all nodes should yield a valid solution to the considered problem.

Since the LOCAL model imposes no constraints on computational power and message size, once the nodes have communicated for a number of rounds greater than the diameter of the network graph, they can in principle know the entire graph and thus solve any problem that can be solved by a single computer in the centralized setting. Therefore, if we equate the complexity of a problem with the number of rounds required to solve it, call this number the *round-time complexity*, and measure it as a function of the number of nodes, then all problems have a complexity that lies between constant round time (the purely local problems) and linear round time (the inherently global problems). From this perspective, investigating the locality of a problem amounts to determining its round-time complexity.

Constant round time. The role of constant round time in the LOCAL model is vaguely analogous to the role of polynomial time in centralized computing, in that it provides a first approximation of what constitutes an efficiently solvable problem. The rigorous study of the class of problems solvable in constant round time was initiated in the early 1990s by Naor and Stockmeyer [31]. To narrow down the area of investigation, they focused on construction problems for which the validity of a proposed solution can at least be *verified* in constant round time. For instance, a proposed vertex coloring can be easily verified in a single round of communication (each node compares its own color with those of its neighbors), whereas a proposed spanning tree cannot be verified locally (a sufficiently long cycle is

indistinguishable from a line). Given the analogy with centralized computing, construction problems verifiable in constant round time are sometimes referred to as the distributed analog of the complexity class **FNP**, i.e., the function problem variant of **NP** [37]. For technical reasons, Naor and Stockmeyer considered only the subclass of locally verifiable problems for which there are constant bounds on the maximum degree of the graphs and on the size of the local inputs and outputs. This subclass, which they called **LCL** (for *locally checkable labelings*), became the foundation of a fruitful research program on locality in distributed computing (see "Construction problems" at the end of Section 1.3).

Decision problems.   Although research in distributed computing has traditionally focused on construction problems, one of its newer branches, called *distributed decision* [9], takes more inspiration from classical complexity theory. In his PODC 2010 keynote talk [12], Fraigniaud suggested that decision problems, on which standard complexity theory is built, could also serve as the basis for a complexity theory of distributed computing. The rationale is that decision problems are easier to reduce to one another than construction problems, while still being general enough to express challenges that arise in a wide variety of models of distributed computing. Reductions between such problems could therefore reveal connections between different areas of distributed computing, or even connections to other fields.

To make joint decisions in a distributed setting, the simplest and most widely studied mechanism is acceptance by unanimity. This requires all computers to accept on yes-instances, and at least one computer to reject on no-instances. When viewed in this context, **LCL** can be reinterpreted as a class of decision problems on labeled graphs, which we will refer to as *graph properties*. The idea is that a graph is a yes-instance of a given graph property if its labeling represents a valid solution to the corresponding **LCL** problem (e.g., a valid coloring, or a maximal independent set). By generalizing this to arbitrary graphs with labels of arbitrary size, we arrive at the class of graph properties that are decidable in constant round time in the LOCAL model. This class was given the name **LD** (for *local decision*) in [14]. Following the above analogy with centralized computing, one can think of **LD** as a distributed analog of the complexity class **P**.

Nondeterminism.   Problems in **LD** are by definition purely local, and it is easy to come up with graph properties that lie outside this class. For instance, the nodes of a graph cannot locally decide whether the graph is a tree (again, because a sufficiently long cycle is indistinguishable from a line). However, a much larger class of properties can be verified if we take inspiration from the complexity class **NP** and allow some external entity to nondeterministically assign each node an additional label that acts as a certificate. In fact, if arbitrary certificates are allowed, then the nodes can verify any property decidable by a single computer in the centralized setting, because each certificate can in principle encode the entire graph along with the node's identifiers (see, e.g., [8, § 4.1]).

To obtain more interesting complexity classes, two types of restrictions on the certificates have been considered in the literature. The first is to require the certificates to be independent of the nodes' identifiers. Fraigniaud, Korman, and Peleg [14] explored this restriction in a model called *nondeterministic local decision* and showed that it strictly weakens expressiveness, as some properties dependent on the number of nodes in the graph can no longer be locally verified. (It was subsequently shown that the non-verifiable properties are precisely those that are not closed under lift [13].) The full power of unrestricted certificates can only be recovered in combination with additional resources such as randomization or an oracle providing the number of nodes.

The second type of restriction limits the size of certificates with respect to the total number of nodes. This idea was introduced by Korman, Kutten, and Peleg [23] in a model called *proof labeling schemes*, and further developed by Göös and Suomela [18] in a more general model called *locally checkable proofs*. As later argued by Feuilloley [8], certificate size provides, in a sense, an alternative measure of locality: purely local properties do not require any certificates, while inherently global properties require quadratic-size certificates (to encode an adjacency matrix of the entire graph). The property of 3-colorability is almost local, requiring only constant-size certificates, and interestingly, many natural properties such as non-2-colorability and Hamiltonicity require logarithmic-size certificates.

Alternation. Since nondeterminism with restrictions on the certificates provides additional power, but not enough to express all graph properties, a natural follow-up is to explore more computational resources from standard complexity theory and assess their impact on expressiveness. One such resource is quantifier alternation, the key concept underlying the polynomial hierarchy (a hierarchy of complexity classes that contains **P**, **NP**, and **coNP** at its lowest levels). Adapted to the LOCAL model, alternation can be thought of as a game between two players, Eve and Adam, who take turns assigning certificates to all nodes. Intuitively, Eve (the existential player) tries to prove that the input graph satisfies a given property, while Adam (the universal player) tries to disprove it. A constant-round-time distributed algorithm then serves as an arbiter to determine the winner based on the certificates provided.

This framework was investigated by Balliu, D'Angelo, Fraigniaud, and Olivetti [3] for identifier-independent certificates, and by Feuilloley, Fraigniaud, and Hirvonen [10] for logarithmic-size certificates. The two resulting alternation hierarchies turned out to be radically different. In the case of identifier-independent certificates, a single alternation between Adam and Eve already suffices to arbitrate any property decidable in the centralized setting. This means that the entire hierarchy collapses to its second level. On the other hand, for logarithmic-size certificates, there are graph properties that lie outside the corresponding hierarchy, and it remains open whether the hierarchy is infinite. The latter question is believed to be difficult, as it has been shown to be closely related to a long-standing open problem in communication complexity [11].

1.2 Contribution

This paper is motivated by the following question raised by Fraigniaud, Korman, and Peleg [14, § 5.1]: *What are the connections between classical computational complexity theory and local complexity theory?* Rather than viewing the classical theory as merely a source of inspiration, we aim to extend it directly to the setting of distributed decision. We approach this from the perspective that centralized computing corresponds to a special case of the LOCAL model, where the network consists of a single computer. More specifically, we introduce the class **LP** (for *local-polynomial time*), which consists of the graph properties that can be decided in a constant number of rounds in the LOCAL model under the following constraints: the number of computation steps of each computer in each round must be polynomially bounded by the size of its local input and the messages received, and the algorithm must work correctly even under identifier assignments that are only locally unique within a fixed radius. This class generalizes both the complexity class **P** (its restriction to graphs consisting of a single labeled node), and the class of **LCL** decision problems (its restriction to graphs of bounded maximum degree and constant label size).

Building on **LP**, we then define the *local-polynomial hierarchy* $\{\mathbf{\Sigma}_\ell^{\mathrm{LP}}, \mathbf{\Pi}_\ell^{\mathrm{LP}}\}_{\ell \in \mathbb{N}}$ analogously to the alternation hierarchies mentioned above, i.e., as a game between Eve and Adam

who alternately assign certificates to all nodes. The size of these certificates must be polynomially bounded with respect to a constant-radius neighborhood of the nodes. Hence, when restricted to single-node graphs, our hierarchy coincides exactly with the classical polynomial hierarchy. For example, the restriction of $\mathbf{\Sigma}_1^{\mathrm{LP}} = \mathbf{NLP}$ to single-node graphs coincides with the complexity class $\mathbf{\Sigma}_1^{\mathrm{P}} = \mathbf{NP}$. Aiming more at a conceptual than a technical contribution, we make three main points:

1. *Several key concepts and results from standard complexity theory generalize well to the distributed setting.* To illustrate this, we extend the notion of polynomial-time reductions to our model of computation, and then use it to establish a number of hardness and completeness results for the two lowest levels of the local-polynomial hierarchy. While some of these are meaningful only in the distributed setting, our results also include distributed generalizations of well-known classics, in particular the Cook–Levin theorem (which identifies Boolean satisfiability as a complete problem for $\mathbf{NP}$), and the fact that 3-colorability is $\mathbf{NP}$-complete. Similarly, we prove a distributed generalization of Fagin's theorem (which characterizes $\mathbf{NP}$ as the problems expressible in existential second-order logic). This gives us a logical, and thus machine-independent, characterization of the entire local-polynomial hierarchy, demonstrating the robustness of our definition. Moreover, whenever we generalize a classical result, the original version can be recovered by restricting networks to single computers.

2. *Sometimes standard techniques get us further in the distributed setting than they do in the centralized setting.* Specifically, we are able to show that the local-polynomial hierarchy is infinite, while it remains notoriously open whether this is also true when restricted to a single computer. (A collapse of our hierarchy would have implied a collapse of the classical polynomial hierarchy, but the converse does not hold.) As a consequence, our hardness and completeness results for the local-polynomial hierarchy immediately yield *unconditional* lower bounds on the complexity of the graph properties in question, i.e., lower bounds that do not rely on any complexity-theoretic assumptions. In addition, the constraints imposed by the distributed setting allow us to identify natural graph properties that lie outside our hierarchy.

3. *Descriptive complexity theory, the discipline of characterizing complexity classes in terms of equivalent logical formalisms, is particularly helpful in the distributed setting.*

   - *On the one hand, this approach gives us access to a large body of existing results in logic and automata theory.* In particular, our infiniteness result for the local-polynomial hierarchy leverages a corresponding result on monadic second-order logic established by Matz, Schweikardt, and Thomas [29], as well as a logical characterization of finite automata on pictures (so-called tiling systems) established by Giammarresi, Restivo, Seibert and Thomas [16]. Despite not being explicitly concerned with distributed computing, these results rely significantly on a form of locality. Moreover, to show that some graph properties lie outside our hierarchy, we make direct use of classical results from automata theory, namely the pumping lemma for regular languages and the Büchi-Elgot-Trakhtenbrot theorem (which provides a logical characterization of finite automata on words).

   - *On the other hand, descriptive complexity can offer a fresh perspective on distributed computing by imposing unconventional constraints that shed new light on familiar concepts.* For instance, the restriction to algorithms that work correctly under locally unique identifiers is necessary to prove our generalization of Fagin's theorem. But this restriction is also meaningful from a pure distributed computing point of view: in a sense, $\mathbf{LP}$ fully preserves the locality of $\mathbf{LCL}$ (which can be defined without identifiers),

whereas **LD** is somewhat less local due to its reliance on globally unique identifiers. Similarly, proving our generalization of Fagin's theorem requires a polynomial bound on the certificate sizes with respect to the constant-radius neighborhoods of the nodes. This again places a strong emphasis on locality, and contrasts sharply with previous approaches to distributed nondeterminism and alternation, which allow certificate sizes to depend on the whole graph. As a consequence, our approach may provide a new way to *measure the locality* of problems in distributed computing, a prospect we will discuss in Section 10.

1.3 Related work

Our base class **LP** generalizes the **LCL** problems of Naor and Stockmeyer [31] to arbitrary labeled graphs (when interpreting these problems as decision problems). However, it is less general than the class **LD** of Fraigniaud, Korman, and Peleg [14], since it imposes restrictions on the individual processing power of the nodes and requires correctness under locally unique identifiers. Hence, we have **LCL** $\subseteq$ **LP** $\subseteq$ **LD**, and it is easy to check that these inclusions are strict.

Alternation hierarchies. The work most closely related to this paper includes the different alternation hierarchies based on **LD**. The relationship is particularly clear for the previously mentioned identifier-independent hierarchy $\{\mathbf{ind\Sigma}_\ell^{\mathrm{LD}}, \mathbf{ind\Pi}_\ell^{\mathrm{LD}}\}_{\ell \in \mathbb{N}}$ of Balliu, D'Angelo, Fraigniaud, and Olivetti [3]. Since that hierarchy collapses to its second level $\mathbf{ind\Pi}_2^{\mathrm{LD}}$ and contains all decidable properties, it obviously subsumes our hierarchy $\{\mathbf{\Sigma}_\ell^{\mathrm{LP}}, \mathbf{\Pi}_\ell^{\mathrm{LP}}\}_{\ell \in \mathbb{N}}$, which is infinite and excludes some decidable properties. But even on the lower levels, it is easy to see that $\mathbf{\Sigma}_\ell^{\mathrm{LP}} \subseteq \mathbf{ind\Sigma}_\ell^{\mathrm{LD}}$ and $\mathbf{\Pi}_\ell^{\mathrm{LP}} \subseteq \mathbf{ind\Pi}_\ell^{\mathrm{LD}}$, essentially because the identifier-independent certificates chosen by the first player can be used to provide each node with a new, locally unique identifier whose validity can be verified in a constant number of communication rounds. The inclusion on the nondeterministic level is strict because NOT-ALL-SELECTED, the property of a labeled graph having at least one unselected node, lies in $\mathbf{ind\Sigma}_1^{\mathrm{LD}}$ but not in $\mathbf{\Sigma}_1^{\mathrm{LP}}$.

Recently, a polynomial-time version $\{\mathbf{ind\Sigma}_\ell^{\mathrm{LD/P}}, \mathbf{ind\Pi}_\ell^{\mathrm{LD/P}}\}_{\ell \in \mathbb{N}}$ of the identifier-independent hierarchy was investigated by Aldema Tshuva and Oshman [39]. Although at first glance their definition may seem similar to ours, it differs in a crucial point: the polynomial bound they impose on the processing time of the nodes is relative to the size of the entire input graph (including labels), rather than relative to the amount of information that the nodes receive locally. As a result, from its second level $\mathbf{ind\Pi}_2^{\mathrm{LD/P}}$ onward, their hierarchy is essentially equivalent to the centralized polynomial hierarchy (restricted to encodings of graphs), and thus it is unknown whether it collapses or not. Nevertheless, its relationship to our hierarchy mirrors that of the original identifier-independent hierarchy, i.e., $\mathbf{\Sigma}_\ell^{\mathrm{LP}} \subseteq \mathbf{ind\Sigma}_\ell^{\mathrm{LD/P}}$ and $\mathbf{\Pi}_\ell^{\mathrm{LP}} \subseteq \mathbf{ind\Pi}_\ell^{\mathrm{LD/P}}$ for all $\ell \in \mathbb{N}$, and the property NOT-ALL-SELECTED separates $\mathbf{ind\Sigma}_1^{\mathrm{LD/P}}$ from $\mathbf{\Sigma}_1^{\mathrm{LP}}$.

It is less obvious how exactly our hierarchy relates to the previously mentioned logarithmic-size hierarchy $\{\mathbf{log\Sigma}_\ell^{\mathrm{LD}}, \mathbf{log\Pi}_\ell^{\mathrm{LD}}\}_{\ell \in \mathbb{N}}$ of Feuilloley, Fraigniaud, and Hirvonen [10]. But at least when restricted to graphs of bounded maximum degree and constant label size, each level of the logarithmic-size hierarchy contains the corresponding level of our hierarchy, since our bound on the certificates sizes reduces to a constant bound for such graphs. That is, $\mathbf{\Sigma}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)} \subseteq \mathbf{log\Sigma}_\ell^{\mathrm{LD}}|_{\mathrm{GRAPH}(\Delta)}$ and $\mathbf{\Pi}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)} \subseteq \mathbf{log\Pi}_\ell^{\mathrm{LD}}|_{\mathrm{GRAPH}(\Delta)}$, where $\mathbf{C}|_{\mathrm{GRAPH}(\Delta)}$ denotes the aforementioned restriction of a class $\mathbf{C}$. Moreover, the nondeterministic classes are again separated by the property NOT-ALL-SELECTED, which lies in $\mathbf{log\Sigma}_1^{\mathrm{LD}}$ but not in $\mathbf{\Sigma}_1^{\mathrm{LP}}$.

What most fundamentally distinguishes this work from all three **LD**-based hierarchies is that our hierarchy preserves some degree of locality. This is because we bound the size of a node's certificates with respect to its constant-radius neighborhood, so that each certificate can encode only a very limited amount of global information about the input graph. By contrast, in both identifier-independent hierarchies, "only the first few levels of alternation are needed to overcome the locality of a distributed algorithm", as noted by Aldema Tshuva and Oshman [39, § 1]. This is particularly evident in the original version of Balliu, D'Angelo, Fraigniaud, and Olivetti, where the second level $\mathbf{ind\Pi}_2^{\text{LD}}$ already contains all global properties. Similarly, in the logarithmic-size hierarchy of Feuilloley, Fraigniaud, and Hirvonen, the third level $\mathbf{log\Sigma}_3^{\text{LD}}$ is already powerful enough to express the existence of a nontrivial automorphism, a property that is inherently global when using certificate size as the measure of locality. (Göös and Suomela [18] have shown that it requires quadratic-size certificates, which is the highest possible complexity).

Descriptive complexity.   Another line of research closely related to this paper is the development of descriptive complexity in the setting of distributed computing. This was initiated by Hella et al. [20], who used several variants of modal logic to characterize synchronous constant-round-time algorithms for various models of distributed computing in anonymous networks. Their idea was later extended to arbitrary-round-time algorithms [24], asynchronous algorithms [34], and a stronger model with unique identifiers [4]. Our generalization of Fagin's theorem to the LOCAL model remains close in spirit to the work of Hella et al. The main difference is that we consider polynomial-time Turing machines instead of finite-state automata, and bounded first-order quantifiers instead of modal operators. Also, the result of Hella et al. already holds for deterministic models, whereas our result requires the presence of nondeterminism, or more generally, alternation. This parallels the situation in the centralized setting, where Fagin's theorem characterizes the class **NP**, but it remains a major open question whether the class **P** admits a similar characterization.

Construction problems.   More distantly related to this paper, recent years have also seen significant progress in the study of **LCL** problems as originally defined by Naor and Stockmeyer (i.e., as construction problems whose solutions can be verified locally). Much research has focused on classifying **LCL** problems according to the round time required to construct solutions for them in the LOCAL model. For over two decades, progress was slow as efforts were driven by individual problems rather than entire classes of problems. While there were known examples of **LCL** problems with constant, iterated logarithmic, and linear round-time complexities, it remained unclear whether problems with other complexities existed in the spectrum between constant and linear round time. However, this picture changed drastically in the mid 2010's, when a large number of positive and negative results were published within a few years, proving the existence of problems in some intermediate regions of the spectrum, and ruling out the existence of problems in other regions. Taken together, those results now provide a nearly complete classification of **LCL** problems, revealing essentially four complexity classes. In his SWAT 2020 keynote talk [37], Suomela interpreted those classes as follows: purely local problems, symmetry-breaking problems, inherently global problems, and an intriguing class of problems for which randomness provides a significant speedup. As the topic is well beyond the scope of this paper, and the publications are numerous, the reader is referred to (the transcript of) Suomela's talk, which summarizes recent progress and provides many references.

1.4 Organization

We begin with an informal overview of the paper in Section 2. The material covered there will be repeated later in much greater detail and formality. This is necessary because descriptive complexity involves mechanical translations between algorithms and logical formulas, forcing us to deal with the low-level aspects of both frameworks. In Section 3, we give some preliminaries on graphs and relational structures. Then, in Section 4, we introduce our model of computation, which extends standard Turing machines to the distributed setting, and define the local-polynomial hierarchy based on this model. Section 5 introduces the corresponding logical formalism, along with some examples of graph properties expressed as logical formulas. In Section 6, we provide a more flexible characterization of our complexity classes in order to simplify subsequent proofs. The actual results begin in Section 7, where we present Fagin's theorem and generalize it to the local-polynomial hierarchy. In Section 8, we introduce the notion of local-polynomial reductions, based on which we establish a number of hardness and completeness results, including a generalization of the Cook–Levin theorem and the **NLP**-completeness of 3-colorability. Section 9 constitutes the longest part of the paper, where we prove that the local-polynomial hierarchy is infinite. This involves a detour through tiling systems and monadic second-order logic on pictures. Finally, in Section 10, we discuss how the preceding results may be relevant to the study of locality in distributed computing.

2 Informal overview

In this paper, we study the computational complexity of graph properties in terms of a distributed model of computation. As is common in this type of setting, we always assume that graphs are finite, simple, undirected, and connected. In addition, our graphs are equipped with a labeling function that assigns a bit string to each node. The focus is exclusively on graph properties that are invariant under isomorphism. These properties typically depend on the graph's topology (e.g., 3-colorability, Eulerianness, or Hamiltonicity), but may also depend on its node labels (e.g., having all nodes labeled the same, or having the labeling form a valid 3-coloring).

2.1 Our complexity classes

To classify graph properties, we extend standard complexity classes from strings to graphs, treating strings as graphs consisting of a single labeled node.

Model of computation. We use distributed algorithms as decision procedures for graph properties. Given an input graph G and an assignment id of identifiers to the nodes of G, the goal is for the nodes to collectively decide whether G has a certain property L. To do so, they proceed in a sequence of synchronous communication rounds. In each round, each node first receives the messages sent by its neighbors in the previous round, then performs some local computations, and finally sends new messages to its neighbors. After a finite number of rounds, each node must have reached an individual verdict, and G is accepted if and only if the nodes unanimously agree on it. The collective decision must be independent of the particular identifier assignment id, as long as the latter satisfies a basic requirement of *local uniqueness*: id must assign different identifiers to any two nodes that lie within some fixed distance of each other. This can be seen as a precondition for the algorithm to work correctly.

In the following, we restrict our attention to distributed algorithms that are guaranteed to terminate in a constant number of communication rounds, and where the number of

computation steps of each node in each round is polynomially bounded by the size of its local input and the messages it receives. We call such algorithms *local-polynomial machines*, as we will formalize them using a model based on Turing machines (see Section 4).

The local-polynomial hierarchy.   Our complexity classes are based on a game between two players who intuitively argue whether a given graph G has some property L: Eve (the existential player) tries to prove that G has property L, and Adam (the universal player) tries to prove the opposite. Given some locally unique identifier assignment of G, the two players take turns choosing assignments of additional labels, called certificates, to the nodes of G. These certificates can be thought of as proofs (in Eve's case) and counterproofs (in Adam's case). They may depend on the provided identifiers, but their size must be polynomially bounded with respect to the amount of information contained in a node's constant-radius neighborhood (including all labels and identifiers therein). After a fixed number ℓ of moves, the winner is determined by a local-polynomial machine, which acts as an arbiter. Ultimately, the graph G has property L if and only if Eve has a winning strategy in this game, i.e., if she always wins when playing optimally. Depending on who makes the first move, L is classified as a $\mathbf{\Sigma}_\ell^{\mathrm{LP}}$-property (if Eve starts) or a $\mathbf{\Pi}_\ell^{\mathrm{LP}}$-property (if Adam starts).

To give an example, L belongs to $\mathbf{\Sigma}_3^{\mathrm{LP}}$ if it satisfies the following equivalence for every graph G and every admissible identifier assignment *id* of G:

$$G \in L \iff \exists \kappa_1 \, \forall \kappa_2 \, \exists \kappa_3 : M(G, \, id, \, \kappa_1 \cdot \kappa_2 \cdot \kappa_3) \equiv \textsc{accept},$$

where M is an appropriately chosen local-polynomial machine, and all quantifiers range over certificate assignments that satisfy the aforementioned polynomial bound. (The notation used here will be formally introduced in Section 4.)

We refer to the family of classes $\{\mathbf{\Sigma}_\ell^{\mathrm{LP}}, \mathbf{\Pi}_\ell^{\mathrm{LP}}\}_{\ell \in \mathbb{N}}$ as the *local-polynomial hierarchy*. Two classes at the lowest levels are of particular interest: $\mathbf{LP} = \mathbf{\Sigma}_0^{\mathrm{LP}}$ (for *local-polynomial time*) and $\mathbf{NLP} = \mathbf{\Sigma}_1^{\mathrm{LP}}$ (for *nondeterministic local-polynomial time*). Due to the asymmetric nature of acceptance by unanimity, classes on the same level of the local-polynomial hierarchy are neither complement classes of each other, nor are they closed under complementation (see Corollary 38 on page 63). Therefore, it makes sense to also consider the hierarchy of complement classes $\{\mathbf{co\Sigma}_\ell^{\mathrm{LP}}, \mathbf{co\Pi}_\ell^{\mathrm{LP}}\}_{\ell \in \mathbb{N}}$. The two hierarchies are illustrated in Figure 1, along with the inclusion and separation results shown in this paper.

Connection to standard complexity classes.   By restricting the classes $\mathbf{\Sigma}_\ell^{\mathrm{LP}}$ and $\mathbf{\Pi}_\ell^{\mathrm{LP}}$ to strings (i.e., labeled graphs consisting of a single node), we obtain the corresponding classes $\mathbf{\Sigma}_\ell^{\mathrm{P}}$ and $\mathbf{\Pi}_\ell^{\mathrm{P}}$ of the original polynomial hierarchy introduced by Meyer and Stockmeyer [30]. Since the same observation holds for the complement classes, this means that the local-polynomial hierarchy is identical to its complement hierarchy on strings. In particular, $\mathbf{P} = \mathbf{LP}|_{\textsc{node}} = \mathbf{coLP}|_{\textsc{node}}$ and $\mathbf{NP} = \mathbf{NLP}|_{\textsc{node}} = \mathbf{co\Pi}_1^{\mathrm{LP}}|_{\textsc{node}}$, where $\mathbf{C}|_{\textsc{node}}$ denotes the restriction of a complexity class $\mathbf{C}$ to single-node graphs. This means that any inclusion result for the local-polynomial hierarchy would imply the corresponding inclusion result for the polynomial hierarchy (e.g., $\mathbf{LP} = \mathbf{NLP}$ would imply $\mathbf{P} = \mathbf{NP}$), but not vice versa. Conversely, any separation result for the polynomial hierarchy would imply the corresponding separation result for the local-polynomial hierarchy (e.g., $\mathbf{P} \neq \mathbf{NP}$ would imply $\mathbf{LP} \neq \mathbf{NLP}$), but not vice versa. Thus, unfortunately, our infiniteness result for the local-polynomial hierarchy does not imply a corresponding result for the original polynomial hierarchy.

Figure 1 The local-polynomial hierarchy (left) and its complement hierarchy (right). The dotted classes on the far right are the same as those on the far left, repeated for the sake of readability. Only the lowest five levels are shown, but the pattern extends infinitely. (See Figure 11 on page 47 for an extended version of this figure with references to the corresponding proofs.) Each line (whether solid or dashed) indicates an inclusion of the lower class in the higher class. All inclusions represented by solid lines are proved to be strict, and classes on the same level (regardless of which hierarchy) are proved to be pairwise distinct, even when restricted to graphs of bounded maximum degree and constant label size. The inclusions represented by dashed lines are in fact equalities when restricted to the latter class of graphs, but this statement is unlikely to generalize to arbitrary graphs, where it holds if and only if $\mathbf{P} = \mathbf{coNP}$. This means that from a distributed computing perspective, the classes shown with thick borders are the most meaningful.

2.2 Extending classical reductions

Aiming to apply standard techniques of complexity theory to the distributed setting, we extend Karp's [22] notion of polynomial-time reduction to computer networks.

Local-polynomial reductions. In a nutshell, if there is a *local-polynomial reduction* from a property L to a property L', then this means that there exists a local-polynomial machine M_{red} that transforms an input graph G into a new graph G' such that G has property L if and only if G' has property L'. Hence, the existence of such a reduction implies that L' is at least as hard as L, since an efficient decider M' for L' could be converted into an efficient decider M for L, which would first run M_{red} and then simulate M' on the resulting graph.

To transform a graph G into a graph G' with a distributed algorithm, each node u of the input graph G computes a string that encodes a subgraph of the output graph G', including the labels of all nodes therein. We call this subgraph the *cluster* representing u in G'. Clusters of different nodes may not overlap, and edges between different clusters are only permitted if the clusters represent adjacent nodes in the original graph G. This setup allows the nodes of G to simulate a distributed algorithm running on G' by simulating the algorithm within their respective clusters and exchanging messages with their neighbors to simulate inter-cluster communication. (For a more formal presentation, see Section 8.)

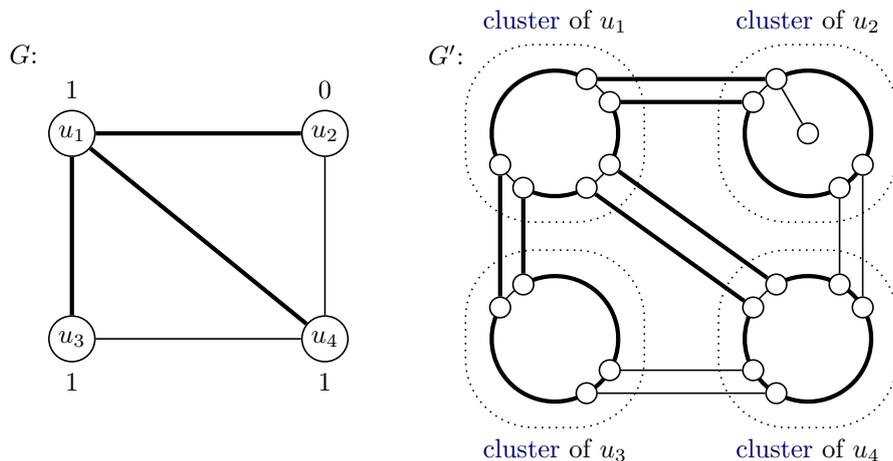**Hardness and completeness results.** Given the above notion of reduction, our definitions of hardness and completeness for different levels of the local-polynomial hierarchy should come as no surprise: a graph property L is hard for a complexity class $\mathbf{C}$ if there is a

local-polynomial reduction to L from every graph property in **C**, and L is complete for **C** if itself additionally lies in that class.

The basic approach to establishing reductions between specific problems in our setting is quite similar to that in the centralized setting, so conventional techniques continue to work well. Using fairly simple constructions, we can show that Eulerianness is **LP**-complete, while Hamiltonicity is both **LP**-hard and **coLP**-hard. Because of the incomparability of **LP** and **coLP** (see Figure 1), this immediately tells us that Hamiltonicity is a strictly harder problem than Eulerianness in our model of computation. What's more, we can sometimes even build directly on classical reductions by extending them to the distributed setting. In particular, we can generalize the Cook–Levin theorem from **NP** to **NLP** and, based on that, establish the **NLP**-completeness of 3-colorability. Again, this has direct implications: 3-colorability is neither in **LP** nor in **coNLP**, since both classes are separate from **NLP**. We now sketch two of the above reductions as examples. (More details are given in Section 8.)

An LP-hardness proof.   To show that Hamiltonicity is **LP**-hard, we provide a reduction to it from ALL-SELECTED, a trivially **LP**-complete graph property that requires all nodes to be labeled with the bit string 1. This reduction is illustrated in Figure 2.

Given an arbitrary graph G, we construct a graph G' that has a Hamiltonian cycle if and only if all nodes of G have label 1. The main idea is that a Hamiltonian cycle in G' represents a depth-first traversal of a spanning tree of G, using a method known as the Euler tour technique. For this purpose, each edge of G is represented by two edges in G', so that it can be traversed twice by a Hamiltonian cycle in G'. If all nodes of G are labeled with 1, then any spanning tree of G yields a Hamiltonian cycle of G'. However, if at least one node of G has a label different from 1 (such as node u_2 in Figure 2), then our construction includes an additional node of degree 1 to ensure that G' is not Hamiltonian. Note that the nodes of G can compute G' in a constant number of communication rounds and a number of computation steps polynomial in the size of their local input and the messages they receive. (For more details, see Proposition 16 on page 41.)

Figure 2 Example illustrating the reduction from ALL-SELECTED to HAMILTONIAN, used to show that the latter is **LP**-hard. The graph G has all node labels equal to 1 if and only if the graph G' has a Hamiltonian cycle. The thick edges in G form a spanning tree, which is replicated by the thick edges in G'. In this particular case, if node u_2 of G had label 1, then its cluster in G' would lack the "central" node, and thus the thick edges in G' would form a Hamiltonian cycle.

An NLP-completeness proof. 3-colorability clearly lies in **NLP**. To show that it is also **NLP**-hard, we build on the classical reduction from 3-SAT to the string-encoded version of 3-colorability, which gives us the desired result almost for free. Our extension of this construction to the distributed setting is illustrated in Figure 3. Here we generalize 3-SAT to graphs as follows to obtain an **NLP**-complete property: each node of the input graph is labeled with a Boolean formula, and the graph is said to be satisfiable if there exists an assignment of variable valuations to its nodes such that each valuation satisfies the formula of the corresponding node while being consistent with the valuations of all adjacent nodes. (Two adjacent nodes can have different variables, but any variables shared by both must be assigned the same values.)

Given an input graph G, we construct a graph G' that is 3-colorable if and only if G is satisfiable. For this purpose, each node of G is represented by a cluster that encodes its formula in such a way that a valid 3-coloring of the cluster represents a satisfying valuation of the formula. This is done by directly using the classical construction as it stands. In addition, to ensure that the variable valuations of adjacent nodes are consistent for all shared variables, the corresponding clusters are connected with auxiliary gadgets that force certain nodes to have the same color. Again, the nodes of G can compute G' in a constant number of communication rounds and a polynomial number of computation steps. (For more details, see Theorem 20 on page 45; for the classical reduction, see, e.g., [17, Prp. 2.27].)

Figure 3 Example illustrating the reduction from 3-SAT-GRAPH to 3-COLORABLE, used to show that the latter is **NLP**-complete. The Boolean graph G is satisfiable if and only if the graph G' is 3-colorable. The labels in G' serve explanatory purposes only and are not part of the graph.

2.3 A logical characterization

A central tool and recurring theme of this paper is Fagin's theorem [7]. In its original form (see Theorem 9 on page 29), it states that a formal language lies in **NP** if and only if it can be defined by a formula of existential second-order logic. Such formulas are of the form $\exists R_1 \ldots \exists R_n\, \varphi$, where $R_1, \ldots, R_n$ are second-order variables and φ is a first-order formula. In the context of this paper, these formulas are evaluated on bit strings represented as relational

structures. More precisely, the bits of a string are represented by a sequence of elements connected by a binary successor relation, and their values are determined by a unary relation. For instance, the string 010011 is represented by the structure

where the elements belonging to the unary relation are marked in black.

Extension to the distributed setting.   We show that Fagin's result can be generalized to obtain a similar logical characterization of the class **NLP** (see Theorem 11 on page 31). To evaluate logical formulas on labeled graphs, we use the structural representation illustrated in Figure 4. This representation contains an element for every node and every labeling bit of the graph. The nodes are connected symmetrically to their neighbors and asymmetrically to their labeling bits by two binary relations. In turn, the labeling bits of each node are interconnected by a successor relation and assigned a value by a unary relation, just as in the string representation described above.

Figure 4 A labeled graph G and its structural representation $\$G$. For later reference: the unary relation $\odot_1^{\$G}$ is represented by black elements, and the binary relations $\rightharpoonup_1^{\$G}$ and $\rightharpoonup_2^{\$G}$ by solid and dotted arrows, respectively.

Our generalization of Fagin's theorem states that a graph property lies in **NLP** if and only if it can be defined (on structural representations as above) by a formula of the following fragment of existential second-order logic: formulas are of the form $\exists R_1 \ldots \exists R_n \forall x\, \varphi$, where $R_1, \ldots, R_n$ are second-order variables, x is a first-order variable, and φ is a first-order formula in which all quantifiers are *bounded* to range only over locally accessible elements. For instance, existential quantification must be of the form $\exists z \rightleftharpoons y\, \psi$, which can be read as "there exists an element z connected to a known element y such that formula ψ is satisfied". This means that first-order quantification in φ is always relative to some element already fixed at an outer scope, and thus in effect that φ is *bounded* around the variable x. (For formal definitions and examples, see Section 5.)

Extension to higher levels of alternation.   Stockmeyer [36] showed that Fagin's theorem extends to the higher levels of the polynomial hierarchy. For example, the complexity class $\mathbf{\Pi}_2^{\mathrm{p}}$ is characterized by formulas of the form $\forall R_1 \ldots \forall R_m \exists S_1 \ldots \exists S_n\, \varphi$, consisting of a block of universal second-order quantifiers, followed by a block of existential second-order quantifiers, and then a first-order formula φ. We similarly extend our generalization of Fagin's theorem to the higher levels of the local-polynomial hierarchy (see Theorem 12 on page 35). For instance, the complexity class $\mathbf{\Pi}_2^{\mathrm{LP}}$ is characterized by formulas of the form $\forall R_1 \ldots \forall R_m \exists S_1 \ldots \exists S_n \forall x\, \varphi$, where the prefix of second-order quantifiers is as above, x is a

first-order variable, and φ is a first-order formula bounded around x. We refer to this logical characterization of the local-polynomial hierarchy as the *local second-order hierarchy*. All graph properties in this hierarchy can be defined by formulas consisting of alternating blocks of existential and universal second-order quantifiers, followed by a single universal first-order quantifier, and then a bounded first-order formula.

For each alternation level, we can recover Stockmeyer's result by restricting our corresponding statement to single-node graphs. Indeed, if the input graph consists of a single node, then all elements of its structural representation lie within distance 2 of each other, so the distinction between bounded and unbounded first-order quantification becomes irrelevant.

Proof outline. Unfortunately, there does not seem to be a straightforward way to use Fagin's original result as a black box to prove our generalization. So instead, we give a full proof from scratch, adapting the ideas of the original proof to the distributed setting. Here we sketch only the nondeterministic case, which is the easiest to present, but the proof extends to arbitrarily high levels of quantifier alternation.

The easy part is to translate a formula of the form $\exists R_1 \dots \exists R_n \, \forall x \, \varphi$ into a distributed Turing machine that verifies the same property. In essence, the certificates chosen by the prover (Eve) are used to encode the existentially quantified relations $R_1, \dots, R_n$, so that the nodes executing the machine just have to run a local algorithm to evaluate φ in their constant-radius neighborhood. They can do this in a polynomial number of computation steps by simply iterating over all possible interpretations of the first-order variables in φ.

The reverse translation, however, is more complicated. It involves encoding the space-time diagram of every Turing machine in the network by a collection of relations over the corresponding structural representation. The key insight that makes this possible is the same as in Fagin's original proof: since the number of computation steps of each machine is polynomially bounded by the size of its input, each cell of the corresponding space-time diagram can be represented by a tuple of nearby elements whose length depends on the degree of the bounding polynomial. What makes our generalized proof somewhat more cumbersome are the additional technicalities imposed by the distributed setting, in particular the assignment of locally unique identifiers (chosen small enough to be representable), and the exchange of messages between adjacent nodes. The latter requires that, for each pair of adjacent machines, our formula keeps track of the tape positions of the sent and received messages, so that the appropriate section of one machine's sending tape is copied to the appropriate section of the other machine's receiving tape.

Implications. Our generalization of Fagin's theorem serves several purposes in this paper:
1. It provides evidence that our definition of the local-polynomial hierarchy is robust, in the sense that the complexity classes defined do not inherently depend on technical details such as the chosen model of computation.
2. It gives us a convenient way to prove our generalization of the Cook–Levin theorem mentioned above. This is analogous to the centralized setting, where the Cook–Levin theorem can be obtained as a corollary of Fagin's theorem.
3. As we will see next, we make extensive use of the provided connection to logic to prove that the local-polynomial hierarchy is infinite.
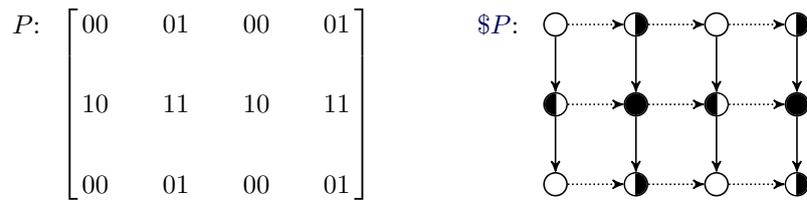
2.4 Infiniteness of our hierarchy

While in centralized computing the question of whether **P** equals **NP** remains a major open problem, the corresponding question in distributed computing—whether **LP** equals **NLP**—is

easily settled with an elementary argument: nondeterminism provides a means to break symmetry, which is impossible in a purely deterministic setting (see Proposition 21 on page 48). What seems less obvious, however, is how to separate complexity classes that lie higher in the local-polynomial hierarchy. This is where our generalization of Fagin's theorem proves particularly helpful, as it allows us to reformulate the problem in the well-studied framework of logic and automata theory. Our separation proof builds on two results from that area, both concerning monadic second-order logic on pictures. The main ideas are outlined below.

Logic on pictures.   *Monadic second-order logic* is the fragment of second-order logic that can only quantify over sets instead of arbitrary relations. This means, for instance, that the formulas of existential monadic second-order logic are of the form $\exists X_1 \ldots \exists X_n\, \varphi$, where the variables $X_1, \ldots, X_n$ represent sets of elements and φ is a first-order formula.

Meanwhile, *pictures* are matrices of fixed-length binary strings. To describe the properties of pictures using logical formulas, every picture is given a structural representation as shown in Figure 5. Specifically, the entries of the picture are represented by elements connected by a "vertical" and a "horizontal" successor relation, and the value of each bit is represented by a unary relation. (Formal definitions are given in Section 9.2.)

Figure 5 A 2-bit picture P and its structural representation $P.

Proof outline.   Our proof of the infiniteness of the local-polynomial hierarchy consists of two main parts, which remain mostly in the realm of logic:
1. First, we show that the local second-order hierarchy is infinite when restricted to pictures, and more precisely that all levels ending with a block of existential quantifiers are distinct (see Section 9.2.1). This is obtained by combining the following two results:
   a. We show that the local and the monadic second-order hierarchies on pictures are levelwise equivalent for all levels ending with a block of existential quantifiers. The main ingredient to prove this is an automata-theoretic characterization of existential monadic second-order logic on pictures due to Giammarresi, Restivo, Seibert, and Thomas [16] (see Theorem 29 on page 53). This characterization, which is itself based on a locality property of first-order logic, gives us a convenient way to establish the equivalence of the existential fragments of local and monadic second-order logic on pictures. For the higher levels of the hierarchies, the equivalence is then obtained by induction on the number of quantifier alternations.
   b. Matz, Schweikardt, and Thomas [29] have shown that the monadic second-order hierarchy on pictures is infinite (see Theorem 27 on page 51). Interestingly, one way to prove their result is based on the automata-theoretic characterization mentioned above, and thus ultimately on the same locality property of first-order logic.
2. Second, we transfer the previous infiniteness result for the local second-order hierarchy from pictures to graphs (see Section 9.2.2). We do this by encoding pictures as graphs

in such a way that formulas can be translated from one type of structure to the other without changing the alternation level of second-order quantifiers. By our generalization of Fagin's theorem, this implies that all levels of the local-polynomial hierarchy ending with a block of existential quantifiers are distinct (see Theorem 33 on page 59). We then complete this partial separation result with some additional arguments to arrive at the fuller separation result shown in Figure 1 (see Sections 9.1 and 9.3).

Implications. Besides the result itself, we derive two main benefits from the infiniteness of the local-polynomial hierarchy:

1. Hardness results give us unconditional lower bounds. That is, if we can show that a graph property is hard for some class of the hierarchy, then we immediately know that it does not lie in the classes below (see Corollaries 22, 25, and 26 starting on page 49).

2. Since alternation is the only way for nodes to obtain global information about their network, the infiniteness result also suggests that the local-polynomial hierarchy may provide a new way to measure the locality of problems in distributed computing (see the discussion in Section 10).

3   Preliminaries

We denote the empty set by $\emptyset$, the set of nonnegative integers by $\mathbb{N}$, the set of positive integers by $\mathbb{N}_{>0}$, and the set of integers by $\mathbb{Z}$. The absolute value of an integer n is denoted by $\mathrm{abs}(n)$. The cardinality of any set A is written as $\mathrm{card}(A)$, its power set as 2^A, and the set of finite strings over A as A^*. The length of a string s is denoted by $\mathrm{len}(s)$, and its i-th symbol by $s(i)$. By a slight abuse of notation, we sometimes lift functions from elements to sets, i.e., given $f\colon X \to Y$ and $A \subseteq X$, we write $f(A)$ for $\{f(a) \mid a \in A\}$. To denote integer intervals, we define $[m:n] = \{i \in \mathbb{Z} \mid m \le i \le n\}$ and $[n] = [0:n]$, for any $m, n \in \mathbb{Z}$. Angle brackets indicate excluded endpoints, e.g., $\langle m:n] = [m+1:n]$ and $[n\rangle = [0:n-1]$.

Throughout this paper, we assume some fixed but unspecified *encoding* of finite objects (e.g., integers, graphs, or tuples of finite objects) as binary strings. Sometimes we also implicitly identify such objects with their string representations.

Graphs. All graphs we consider are finite, simple, undirected, and connected. Formally, a *labeled graph*, or simply *graph*, is represented by a triple $G = (V^G, E^G, \lambda^G)$, where V^G is a finite nonempty set of *nodes*, E^G is a set of undirected *edges* consisting of 2-element subsets of V^G and containing, for every partition $\{V_0, V_1\}$ of V^G, at least one edge $\{u, v\}$ with $u \in V_0$ and $v \in V_1$, and $\lambda^G\colon V^G \to \{0,1\}^*$ is a *labeling* function that assigns a bit string to each node. We refer to the string $\lambda^G(u)$ as the *label* of node u and to the symbol $\lambda^G(u)(i)$ as the i-th *labeling bit* of u, for $i \in [1:\mathrm{len}(\lambda^G(u))]$. To simplify notation, we often write $u \in G$ instead of $u \in V^G$, and we define $\mathrm{card}(G)$, the *cardinality* of G, as $\mathrm{card}(V^G)$.

We denote by GRAPH the set of all labeled graphs and by NODE the set of *single-node graphs*, i.e., labeled graphs consisting of a single node. A *graph property* (sometimes called a "graph language") is a set $L \subseteq$ GRAPH that is closed under isomorphism. If a graph G belongs to L, then we also say that G has the property L.

We follow Diestel [6] for standard graph-theoretic terms such as *neighbor*, *degree*, *distance*, *diameter*, *induced subgraph*, and so on. The diameter of a graph G is denoted by $\mathrm{diam}(G)$. For $r \in \mathbb{N}$ and $u \in G$, the *r-neighborhood* $N_r^G(u)$ of u in G is the subgraph of G induced by all nodes at distance at most r from u. That is, $N_r^G(u)$ is the graph G' that consists of the

nodes at distance at most r from u and all edges connecting them, and whose labeling $\lambda^{G'}$ is the restriction of λ^G to $V^{G'}$.

Identifier assignments.   An *identifier assignment* of a graph G is a function $id\colon V^G \to \{0,1\}^*$ whose purpose is to distinguish between different nodes of G. We refer to $id(u)$ as the *identifier* of node u under id. Identifiers are ordered lexicographically, i.e., the *identifier order* is such that $id(u) < id(v)$ if either $id(u)$ is a proper prefix of $id(v)$, or $id(u)(i) < id(v)(i)$ at the first position i where the two strings differ.

   We say that id is r_{id}-*locally unique* for some $r_{\mathrm{id}} \in \mathbb{N}$ if it satisfies $id(u) \neq id(v)$ for all distinct nodes u and v that lie in the r_{id}-neighborhood of a common node w (or equivalently, in the $2r_{\mathrm{id}}$-neighborhood of each other). If $r_{\mathrm{id}} \geq \mathrm{diam}(G)/2$, we say that id is *globally unique.*

   An r_{id}-locally unique identifier assignment id of a graph G is called *small* (with respect to r_{id}) if for every node $u \in G$, the length of $id(u)$ is at most $\lceil \log_2 \mathrm{card}(N^G_{2r_{\mathrm{id}}}(u)) \rceil$, i.e., logarithmically bounded by the cardinality of u's $2r_{\mathrm{id}}$-neighborhood in G. When we want to emphasize that an r_{id}-locally unique identifier assignment is not necessarily small, we call it *arbitrary-sized.*

▶ Remark 1. For every graph G and integer $r_{\mathrm{id}} \in \mathbb{N}$, there exists an r_{id}-locally unique identifier assignment id of G that is small.

Proof. By definition, an identifier assignment id of G is r_{id}-locally unique if the identifier $id(u)$ of every node $u \in G$ is distinct from the identifiers of all other nodes in $N^G_{2r_{\mathrm{id}}}(u)$. Such an identifier assignment can be easily constructed if we may choose among at least $\mathrm{card}(N^G_{2r_{\mathrm{id}}}(u))$ possible values of $id(u)$. Hence, a bit string of length at most $\lceil \log_2 \mathrm{card}(N^G_{2r_{\mathrm{id}}}(u)) \rceil$ is sufficient.   ◁

Certificate assignments.   For any graph G and any identifier assignment id of G, a *certificate assignment* of (G, id) is a function $\kappa\colon V^G \to \{0,1\}^*$. We refer to $\kappa(u)$ as the *certificate* of node u under κ. Given $r \in \mathbb{N}$ and $p\colon \mathbb{N} \to \mathbb{N}$, we say that κ is (r,p)-*bounded* if for every node $u \in G$, the length of u's certificate is bounded by p with respect to the cardinality of u's r-neighborhood and the lengths of all labels and identifiers therein, i.e.,

$$\mathrm{len}(\kappa(u)) \leq p\Big(\sum_{v \in N^G_r(u)} 1 + \mathrm{len}(\lambda^G(v)) + \mathrm{len}(id(v)) \Big).$$

   We often represent several certificate assignments as a single function $\bar{\kappa}\colon V^G \to \{0,1,\#\}^*$, called a *certificate-list assignment*, where the symbol $\#$ is used to separate the individual certificates of each node. Given certificate assignments $\kappa_1, \kappa_2, \ldots, \kappa_\ell$, we write $\kappa_1 \cdot \kappa_2 \cdot \ldots \cdot \kappa_\ell$ for the certificate-list assignment $\bar{\kappa}$ such that $\bar{\kappa}(u) = \kappa_1(u) \# \kappa_2(u) \# \ldots \# \kappa_\ell(u)$ for all $u \in G$. We say that $\bar{\kappa}$ is (r,p)-*bounded* if κ_i is (r,p)-bounded for every $i \in [1:\ell]$.

Structural representations.   We will evaluate logical formulas on relational *structures* of the form $S = (D^S, \odot^S_1, \ldots, \odot^S_m, \rightharpoonup^S_1, \ldots, \rightharpoonup^S_n)$, where D^S is a finite nonempty set of *elements*, called the *domain* of S, each $\odot^S_i$ is a subset of D^S, for $i \in [1:m]$, and each $\rightharpoonup^S_i$ is a binary relation over D^S, for $i \in [1:n]$. We refer to (m,n) as the *signature* of S. To simplify notation, we often write $a \in S$ instead of $a \in D^S$, and we define $\mathrm{card}(S)$, the *cardinality* of S, as $\mathrm{card}(D^S)$. We also write $a \rightleftharpoons^S b$ to indicate that $a \rightharpoonup^S_i b$ or $b \rightharpoonup^S_i a$ for some $i \in [1:n]$.

   To evaluate logical formulas on graphs, we identify each graph G with a structure $\$G = (D^{\$G}, \odot^{\$G}_1, \rightharpoonup^{\$G}_1, \rightharpoonup^{\$G}_2)$ of signature $(1,2)$, called the *structural representation* of G. This

structure contains one element u for each node $u \in G$ and one element (u, i) for each of u's labeling bits, i.e.,

$$D^{\$G} = V^G \cup \{(u, i) \mid u \in V^G, i \in [1 : \text{len}(\lambda^G(u))]\}.$$

The set $\odot_1^{\$G}$ corresponds to the labeling bits whose value is 1, i.e., $(u, i) \in \odot_1^{\$G}$ if and only if $\lambda^G(u)(i) = 1$. The relation $\rightharpoonup_1^{\$G}$ represents the edges in E^G and the successor relation of the labeling bits, i.e., $u \rightharpoonup_1^{\$G} v$ if and only if $\{u, v\} \in E^G$, and $(u, i) \rightharpoonup_1^{\$G} (v, j)$ if and only if $u = v$ and $j = i + 1$. Finally, the relation $\rightharpoonup_2^{\$G}$ determines which node owns which labeling bits, i.e., $u \rightharpoonup_2^{\$G} (v, i)$ if and only if $u = v$. An example is provided in Figure 4 on page 12.

For $r \in \mathbb{N}$ and $u \in G$, the structural representation of u's r-neighborhood $N_r^G(u)$ is denoted by $N_r^{\$G}(u)$. For instance, if u is the upper right node of the graph G depicted in Figure 4, then $\text{card}(N_0^{\$G}(u)) = 4$, $\text{card}(N_1^{\$G}(u)) = 8$, and $N_2^{\$G}(u) = \$G$.

A *structure property* is a set L of structures that is closed under isomorphism. In particular, since we identify graphs with their structural representations, every graph property is also a structure property. We will often restrict a given class **C** of structure properties (e.g., **NLP**) to structures that have some presupposed property K (e.g., NODE). In such cases, we write $\mathbf{C}|_K$ for the restriction of **C** to K, i.e., $\mathbf{C}|_K = \{L \cap K \mid L \in \mathbf{C}\}$.

4 Distributed Turing machines

We formalize synchronous distributed algorithms using the notion of distributed Turing machines. As illustrated in Figure 6, such machines are equipped with three one-way infinite *tapes*: a *receiving* tape to store incoming messages, an *internal* tape to store the machine's internal state and perform local computations, and a *sending* tape to store outgoing messages.

Figure 6 Local copy of a distributed Turing machine being executed by a node.

Formal representation. A *distributed Turing machine* is represented by a tuple $M = (Q, \delta)$ consisting of a finite set Q of *states* and a *transition function* $\delta \colon Q \times \Sigma^3 \to Q \times \Sigma^3 \times \{-1, 0, 1\}^3$. Here, Σ is the tape alphabet $\{\vdash, \square, \#, 0, 1\}$ with the *left-end marker* $\vdash$, the *blank symbol* $\square$, and the *separator* $\#$. We assume that Q always contains the designated states q_{start}, q_{pause}, and q_{stop}.

When we refer to the *content* of a tape, we mean the sequence of symbols written on the tape ignoring any leading or trailing occurrences of $\vdash$ and $\square$. In particular, if the first cell of the tape contains $\vdash$ and the remaining cells contain $\square$, we consider the tape to be empty.

Execution. A distributed Turing machine $M = (Q, \delta)$ can be executed on any graph G, under any identifier assignment id of G and any certificate-list assignment $\bar{\kappa}$ of (G, id), provided

that *id* is at least 1-locally unique. An *execution* consists of a sequence of synchronous *communication rounds*, where all nodes start at the same time and run their own copy of M. In every round, each node $u \in G$ goes through three *phases*: (1) it receives messages from its neighbors, (2) it performs local computations, and (3) it sends messages to its neighbors. We now describe these phases in detail.

1. In the first phase, the messages $\mu_1, \ldots, \mu_d \in \{0, 1\}^*$ that u receives from its neighbors $v_1, \ldots, v_d$ are concatenated using the separator $\#$ (including a trailing $\#$) and written on u's receiving tape. Any previous content is discarded so that the new content of the receiving tape is the string $\mu_1 \# \cdots \# \mu_d \#$. In particular, if we are in the first round, the content is $\#^d$, which indicates that u has not yet received any (nonempty) messages. In later rounds, $\mu_1, \ldots, \mu_d$ correspond to the messages that were sent by the neighbors in the previous round, sorted in ascending identifier order. That is, we assume $id(v_1) < \ldots < id(v_d)$.

2. In the second phase, u's copy of M behaves like a standard Turing machine with three tapes. The receiving tape is initialized as stated above, while the sending tape is initially empty, meaning that any content from the previous round is erased. In case we are in the first round, the internal tape is initialized to the string $\lambda^G(u) \# id(u) \# \bar{\kappa}(u)$, i.e., the node gets a copy of its label, identifier, and certificates. Otherwise, the content of the internal tape remains the same as at the end of the previous round. Now, if the machine ended up in state q_{stop} in the previous round, then it remains in that state and immediately goes to phase 3. Otherwise, it starts its local computation in state q_{start} with all three tape heads on the leftmost cell of the tapes, and then goes through a sequence of *computation steps*. In each step, depending on the current state and the symbols currently scanned on the three tapes, the transition function δ tells M how to update its state and the symbols on the tapes, and also in which directions to move the three tape heads. The local computation halts as soon as the machine reaches one of the states q_{pause} or q_{stop}.

3. In the third phase, the messages $\mu'_1, \ldots, \mu'_d \in \{0, 1\}^*$ sent to the neighbors $v_1, \ldots, v_d$ correspond to the first d bit strings stored on the sending tape, using the symbol $\#$ as a separator and ignoring any $\square$'s. The order of the neighbors is the same as in phase 1, i.e., the ascending order of identifiers. In case there are not enough messages on the sending tape, the missing ones default to the empty string. In particular, if u has already reached state q_{stop} in the previous round, then its sending tape remains empty, so all neighbors receive an empty message from u.

The execution *terminates* in $r \in \mathbb{N}_{>0}$ rounds if all nodes have reached state q_{stop} by the end of the r-th round. Note that this implies that the local computations of all nodes halt in every round. Throughout this paper, we will restrict ourselves to distributed Turing machines whose executions terminate on every graph under all identifier and certificate-list assignments.

Result and decision.   The *result $M(G, id, \bar{\kappa})$* computed by M on graph G under identifier assignment *id* and certificate-list assignment $\bar{\kappa}$ is the graph G' whose nodes and edges are the same as those of G, and whose labeling function $\lambda^{G'}$ assigns to each node u the bit string written on u's internal tape after M's execution has terminated. To guarantee that this is indeed a bit string, any symbols other than 0 and 1 are ignored. In case we do not need any certificate assignments, we simply write $M(G, id)$ to denote the result computed by M on G under *id* and the trivial certificate-list assignment that assigns the empty string to every node of G.

A distributed Turing machine can act as a consensus-based decision procedure, where all nodes must agree in order for a given input to be accepted. More precisely, when executing M on G under id and $\bar{\kappa}$, the individual *verdict* of node $u \in G$ is the string s with which u is labeled in the result $M(G, id, \bar{\kappa})$. We say that u *accepts* in $M(G, id, \bar{\kappa})$ if $s = 1$, and that u *rejects* in $M(G, id, \bar{\kappa})$ otherwise. Based on that, G is *accepted* by M under id and $\bar{\kappa}$, written $M(G, id, \bar{\kappa}) \equiv \text{ACCEPT}$, if every node accepts in $M(G, id, \bar{\kappa})$. Conversely, G is *rejected* by M under id and $\bar{\kappa}$ if at least one node rejects in $M(G, id, \bar{\kappa})$.

Running time. In order to measure the running time of distributed Turing machines, we will use two different metrics: round time, which corresponds to the number of communication rounds in an execution, and step time, which gives the number of computation steps made by a single node in one round.

More precisely, for any execution of a distributed Turing machine M, the *round running time* is the number of rounds until all nodes have reached state q_{stop}. Given some function $f \colon \mathbb{N} \to \mathbb{N}$, we say that M runs in round time f if the round running time is bounded by f with respect to the cardinality of the graph on which M is executed. This means that for every graph G, every identifier assignment id of G, and every certificate-list assignment $\bar{\kappa}$ of (G, id), all nodes of G reach state q_{stop} after at most $f(\text{card}(G))$ rounds in the corresponding execution of M. Accordingly, M runs in *constant round time* if this holds for some constant function f.

On the other hand, the *step running time* of node u in round $i \in \mathbb{N}_{>0}$ of an execution of M is the number of local computation steps that u makes during (phase 2 of) round i. For $f \colon \mathbb{N} \to \mathbb{N}$, we say that M runs in step time f if in every execution, the step running time of every node u in every round i is bounded by f with respect to the length of u's initial tape contents in round i. This means that if M starts in state q_{start} with some arbitrary strings $s \in \{0, 1, \#\}^*$ and $t \in \{0, 1, \#, \square\}^*$ written on its receiving and internal tapes, then M reaches q_{pause} or q_{stop} after at most $f(\text{len}(s) + \text{len}(t))$ steps. Accordingly, M runs in *polynomial step time* if this holds for some polynomial function f.

A *local-polynomial machine* is a distributed Turing machine that runs in constant round time and polynomial step time.

Arbiters and the local-polynomial hierarchy. As explained in Section 2.1, each graph property L in the local-polynomial hierarchy corresponds to a game between two players: *Eve*, who tries to prove that a given graph G has property L, and *Adam*, who tries to prove the opposite. The players take turns labeling the nodes of G with certificates, which serve as proofs (in Eve's case) and counterproofs (in Adam's case). After a fixed number ℓ of moves, the winner is determined by a local-polynomial machine M, and the graph G has property L if and only if Eve has a winning strategy in this game. Depending on who makes the first move, L is classified as a $\mathbf{\Sigma}_{\ell}^{\text{LP}}$-property (if Eve starts) or a $\mathbf{\Pi}_{\ell}^{\text{LP}}$-property (if Adam starts).

Formally, we represent Eve's and Adam's choices by quantifying existentially and universally, respectively, over the certificate assignments chosen by each player. More precisely, for $\ell \in \mathbb{N}$, a graph property L belongs to $\mathbf{\Sigma}_{\ell}^{\text{LP}}$ if there exists a local-polynomial machine M, constants $r_{\text{id}}, r \in \mathbb{N}_{>0}$, and a polynomial function p such that the following equivalence holds for every graph G and every r_{id}-locally unique identifier assignment id of G:

$$G \in L \iff \exists \kappa_1 \, \forall \kappa_2 \dots \mathcal{Q}\kappa_\ell : M(G, id, \kappa_1 \cdot \kappa_2 \cdot \dots \cdot \kappa_\ell) \equiv \text{ACCEPT},$$

where $\mathcal{Q}$ is $\forall$ if ℓ is even and $\exists$ otherwise, and all quantifiers range over (r, p)-bounded certificate assignments of (G, id). We say that M *arbitrates* L with respect to $\mathbf{\Sigma}_{\ell}^{\text{LP}}$ and call it a $\mathbf{\Sigma}_{\ell}^{\text{LP}}$-*arbiter* for L under r_{id}-locally unique identifiers and (r, p)-bounded certificates.

The class $\mathbf{\Pi}_\ell^{\mathrm{LP}}$ and the notion of $\mathbf{\Pi}_\ell^{\mathrm{LP}}$-*arbiters* are defined analogously, with the only difference that quantifier alternation starts with a universal quantifier instead of an existential one. That is, for $\mathbf{\Pi}_\ell^{\mathrm{LP}}$, we modify the above equivalence to read "$\forall \kappa_1 \exists \kappa_2 \ldots \mathbb{Q}\kappa_\ell$", where $\mathbb{Q}$ is $\exists$ if ℓ is even and $\forall$ otherwise. We refer to the family of classes $\{\mathbf{\Sigma}_\ell^{\mathrm{LP}}, \mathbf{\Pi}_\ell^{\mathrm{LP}}\}_{\ell \in \mathbb{N}}$ as the *local-polynomial hierarchy*.

Note that the certificate assignments $\kappa_1, \ldots, \kappa_\ell$ may depend on the identifier assignment *id*. Moreover, the individual verdict of a single node may vary depending on the identifiers and certificates in its neighborhood. However, G's membership in L (and thus whether Eve has a winning strategy) must be independent of the particular identifier assignment.

Two classes at the lowest levels of the hierarchy are of particular interest: $\mathbf{LP} = \mathbf{\Sigma}_0^{\mathrm{LP}}$ (for *local-polynomial time*) is the class of graph properties that can be *decided* by a local-polynomial machine, and $\mathbf{NLP} = \mathbf{\Sigma}_1^{\mathrm{LP}}$ (for *nondeterministic local-polynomial time*) is the class of graph properties that can be *verified* by a local-polynomial machine. Accordingly, $\mathbf{\Sigma}_0^{\mathrm{LP}}$-arbiters and $\mathbf{\Sigma}_1^{\mathrm{LP}}$-arbiters are also called $\mathbf{LP}$-*deciders* and $\mathbf{NLP}$-*verifiers*, respectively.

Complement hierarchy.   The *complement class* of a class $\mathbf{C}$ of graph properties is the class $\{\bar{L} \mid L \in \mathbf{C}\}$, where $\bar{L}$ denotes the *complement* of a graph property L, i.e., $\bar{L} = \mathrm{GRAPH} \setminus L$. For $\ell \in \mathbb{N}$, we denote the complement classes of $\mathbf{\Sigma}_\ell^{\mathrm{LP}}$ and $\mathbf{\Pi}_\ell^{\mathrm{LP}}$ by $\mathbf{co\Sigma}_\ell^{\mathrm{LP}}$ and $\mathbf{co\Pi}_\ell^{\mathrm{LP}}$, and also often denote $\mathbf{co\Sigma}_0^{\mathrm{LP}}$ and $\mathbf{co\Sigma}_1^{\mathrm{LP}}$ by $\mathbf{coLP}$ and $\mathbf{coNLP}$, respectively. As we shall see in Corollary 38, classes on the same level of the local-polynomial hierarchy are neither complement classes of each other, nor are they closed under complementation, so it makes sense to consider their complement classes in their own right. We will refer to the family of classes $\{\mathbf{co\Sigma}_\ell^{\mathrm{LP}}, \mathbf{co\Pi}_\ell^{\mathrm{LP}}\}_{\ell \in \mathbb{N}}$ as the *complement hierarchy* of the local-polynomial hierarchy.

Connection to standard complexity classes.   On single-node graphs, distributed Turing machines are equivalent to standard Turing machines that take as input the label and certificates of the unique node. The node's identifier is irrelevant and can therefore be assumed empty, so the condition of the certificates being (r, p)-bounded reduces to them being polynomially bounded in the length of the label. Hence, by restricting the classes $\mathbf{\Sigma}_\ell^{\mathrm{LP}}$ and $\mathbf{\Pi}_\ell^{\mathrm{LP}}$ to NODE and identifying single-node graphs with strings, we obtain the corresponding classes $\mathbf{\Sigma}_\ell^{\mathrm{P}}$ and $\mathbf{\Pi}_\ell^{\mathrm{P}}$ of the classical *polynomial hierarchy* introduced by Meyer and Stockmeyer [30] (see, e.g., [2, § 5.2]). In particular, $\mathbf{P} = \mathbf{LP}|_{\mathrm{NODE}} = \mathbf{coLP}|_{\mathrm{NODE}}$ and $\mathbf{NP} = \mathbf{NLP}|_{\mathrm{NODE}} = \mathbf{co\Pi}_1^{\mathrm{LP}}|_{\mathrm{NODE}}$.

5   Logic with bounded quantifiers

We now introduce a logical formalism that will provide a purely syntactic characterization of most of the complexity classes defined in the previous section. This characterization will be presented in Section 7.

5.1   Definitions

We begin with the necessary formal definitions, and then illustrate them with a series of examples in Section 5.2, using standard graph properties such as 3-colorability and Hamiltonicity. The reader may wish to skip ahead to the examples and refer back to this subsection as needed.

Variables and interpretations.   Let $\mathcal{V}_{\mathrm{FO}}$ be an infinite supply of *first-order variables* and $\mathcal{V}_{\mathrm{SO}} = \bigcup_{k \geq 1} \mathcal{V}_{\mathrm{SO}(k)}$ be an infinite supply of *second-order variables*, also called *relation variables*,

where $\mathcal{V}_{\mathrm{so}(k)}$ contains the second-order variables of *arity* k and $\mathcal{V}_{\mathrm{so}(k)} \cap \mathcal{V}_{\mathrm{so}(k')} = \emptyset$ for $k \neq k'$. We collectively refer to the elements of $\mathcal{V}_{\mathrm{FO}}$ and $\mathcal{V}_{\mathrm{SO}}$ as *variables*.

A *variable assignment* σ of a set of variables $\mathcal{V} \subseteq \mathcal{V}_{\mathrm{FO}} \cup \mathcal{V}_{\mathrm{SO}}$ on a structure S is a function that maps each first-order variable of $\mathcal{V}$ to an element of D^S and each second-order variable of $\mathcal{V}$ to a relation of matching arity over D^S. The value $\sigma(R)$ assigned to a variable R is called the *interpretation* of R under σ. We sometimes write $\sigma[R \mapsto A]$ to denote the variable assignment that is identical to σ except for mapping R to A. Moreover, if σ is irrelevant or clear from context, we may also omit it to simplify the exposition, and refer directly to R as an element or a relation when we really mean $\sigma(R)$.

Syntax and semantics. To avoid repetitions, we first define the syntax and semantics of a generalized class of logical formulas, and then specify which particular subclasses we are interested in.

Table 1 shows how *logical formulas*, or simply *formulas*, are built up inductively (in the first column), and what they mean (in the third column). It also specifies the set $\mathrm{free}(\varphi)$ of variables that occur *freely* in a given formula φ, i.e., outside the scope of any quantifier. When we need to distinguish between first-order and second-order variables, we use the notations $\mathrm{free}_{\mathrm{FO}}(\varphi) = \mathrm{free}(\varphi) \cap \mathcal{V}_{\mathrm{FO}}$ and $\mathrm{free}_{\mathrm{SO}}(\varphi) = \mathrm{free}(\varphi) \cap \mathcal{V}_{\mathrm{SO}}$. If $\mathrm{free}(\varphi) = \emptyset$, then φ is called a *sentence*.

■ **Table 1** Syntax and semantics of all logics considered in this paper.

	Syntax Formula ψ	*Free variables* Set $\mathrm{free}(\psi)$	*Semantics* Necessary and sufficient condition for $S, \sigma \models \psi$
1.	$\odot_i x$	$\{x\}$	$\sigma(x) \in \odot_i^S$
2.	$x \rightharpoonup_i y$	$\{x, y\}$	$\sigma(x) \rightharpoonup_i^S \sigma(y)$
3.	$x \doteq y$	$\{x, y\}$	$\sigma(x) = \sigma(y)$
4.	$R(x_1, \ldots, x_k)$	$\{R, x_1, \ldots, x_k\}$	$\big(\sigma(x_1), \ldots, \sigma(x_k)\big) \in \sigma(R)$
5.	$\neg\varphi$	$\mathrm{free}(\varphi)$	not $S, \sigma \models \varphi$
6.	$\varphi_1 \vee \varphi_2$	$\mathrm{free}(\varphi_1) \cup \mathrm{free}(\varphi_2)$	$S, \sigma \models \varphi_1$ or $S, \sigma \models \varphi_2$
7.	$\exists x\, \varphi$	$\mathrm{free}(\varphi) \setminus \{x\}$	$S, \sigma[x \mapsto a] \models \varphi$ for some $a \in D^S$
8.	$\underbrace{\exists x \rightleftharpoons y\, \varphi}_{\text{where } x \neq y}$	$\{y\} \cup \mathrm{free}(\varphi) \setminus \{x\}$	$S, \sigma[x \mapsto a] \models \varphi$ for some $a \in D^S$ s.t. $a \rightleftharpoons^S \sigma(y)$
9.	$\exists R\, \varphi$	$\mathrm{free}(\varphi) \setminus \{R\}$	$S, \sigma[R \mapsto A] \models \varphi$ for some $A \subseteq (D^S)^k$

Here, $i, k \in \mathbb{N}_{>0}$, $x, x_1, \ldots, x_k, y \in \mathcal{V}_{\mathrm{FO}}$, $R \in \mathcal{V}_{\mathrm{so}(k)}$, and $\varphi, \varphi_1, \varphi_2$ are formulas.

The truth of a formula φ can be evaluated on a structure S of signature (m, n) under a variable assignment σ of $\mathrm{free}(\varphi)$ on S, provided that φ does not contain any expressions of the form $\odot_i x$ or $x \rightharpoonup_j y$ for $i > m$ and $j > n$. Assuming this basic requirement is met, the third column of Table 1 specifies in which cases S *satisfies* φ under σ, written $S, \sigma \models \varphi$. If φ is a sentence, σ is irrelevant, so we simply say that S *satisfies* φ and write $S \models \varphi$. The property *defined* by a sentence φ on a class of structures K is the set $\{S \in K \mid S \models \varphi\}$.

Lines 1 to 4 of Table 1 correspond to *atomic* formulas. An atomic formula of the form $\odot_i x$ or $x \rightharpoonup_i y$ refers to the corresponding set $\odot_i^S$ or binary relation $\rightharpoonup_i^S$ given by the structure S,

while an atomic formula of the form $R(x_1, \ldots, x_k)$ refers to an additional relation $\sigma(R)$ given by the variable assignment σ. Lines 5 and 6 describe the usual *Boolean connectives*, and the remaining lines correspond to *quantifiers* over different scopes: *first-order quantification* on lines 7 and 8 ranges over elements, and *second-order quantification* on line 9 ranges over relations.

Of particular interest for this paper is the *bounded* version of first-order quantification shown on line 8. Intuitively, $\exists x \rightleftharpoons y \, \varphi$ can be read as "there exists an element x connected to y such that φ is satisfied". Here, "connected" means that the elements of S represented by x and y are related by some relation $\rightharpoonup_i^S$ or its inverse. Thus, bounded first-order quantification is relative to an already fixed element, represented here by the free variable y.

Syntactic sugar.   By nesting bounded first-order quantifiers, we can quantify over elements that lie within a given distance $r \in \mathbb{N}$ from the fixed element. To simplify this, we introduce the shorthand notation $\exists x \overset{\leq r}{\rightleftharpoons} y \, \varphi$, which is defined inductively as follows, for any $x, y \in \mathcal{V}_{\mathrm{FO}}$ and formula φ:

$$\exists x \overset{\leq 0}{\rightleftharpoons} y \, \varphi \quad \text{is equivalent to} \quad \varphi[x \mapsto y], \quad \text{and}$$

$$\exists x \overset{\leq r+1}{\rightleftharpoons} y \, \varphi \quad \text{is equivalent to} \quad \exists x \overset{\leq r}{\rightleftharpoons} y \, \big(\varphi \vee \exists x' \rightleftharpoons x \, \varphi[x \mapsto x']\big),$$

where $\varphi[x \mapsto y]$ denotes the formula obtained from φ by substituting every free occurrence of x by y, and x' is a fresh first-order variable that does not occur in φ.

For additional convenience, we will make liberal use of truth constants (i.e., $\bot$, $\top$) and the remaining operators of predicate logic (i.e., $\wedge$, $\rightarrow$, $\leftrightarrow$, $\forall$), use shorthand notations such as $x \neq y$, and we may leave out some parentheses, assuming that $\vee$ and $\wedge$ take precedence over $\rightarrow$ and $\leftrightarrow$. Moreover, a sequence consisting solely of *existential* ($\exists$) or solely of *universal* ($\forall$) quantifiers may be combined into a single quantifier that binds a tuple of variables. For instance, we may write $\forall \bar{R} \, \varphi$ instead of $\forall R_1 \ldots \forall R_n \, \varphi$, where $\bar{R} = (R_1, \ldots, R_n)$.

Formulas expressing relations.   Given a formula φ with $\mathrm{free}_{\mathrm{FO}}(\varphi) = \{x_1, \ldots, x_n\}$, we often write $\varphi(x_1, \ldots, x_n)$ instead of simply φ to convey the intention that φ expresses some relation between the elements represented by $x_1, \ldots, x_n$. If we then want to express that the same relation holds between some other variables $y_1, \ldots, y_n$ that do not occur in the scope of any quantifier in φ, we write $\varphi(y_1, \ldots, y_n)$ to denote the formula obtained from φ by simultaneously replacing all free occurrences of $x_1, \ldots, x_n$ by $y_1, \ldots, y_n$, respectively.

Fragments of first-order logic.   For our purposes, the class FO of formulas of *first-order logic* is generated by the grammar

$$\varphi ::= \odot_i x \mid x \rightharpoonup_i y \mid x \doteq y \mid R(x_1, \ldots, x_k) \mid \neg \varphi \mid \varphi \vee \varphi \mid \exists x \, \varphi, \tag{FO}$$

where $i, k \in \mathbb{N}_{>0}$, $x, x_1, \ldots, x_k, y \in \mathcal{V}_{\mathrm{FO}}$, and $R \in \mathcal{V}_{\mathrm{SO}(k)}$.

The class BF of formulas of the *bounded fragment* of first-order logic is defined by a similar grammar, the only difference being that first-order quantification is bounded:

$$\varphi ::= \odot_i x \mid x \rightharpoonup_i y \mid x \doteq y \mid R(x_1, \ldots, x_k) \mid \neg \varphi \mid \varphi \vee \varphi \mid \underbrace{\exists x \rightleftharpoons y \, \varphi}_{\text{where } x \neq y} \tag{BF}$$

To give some basic examples, when evaluated on (the structural representation of) a graph, the following BF-formulas state that the element represented by the first-order variable x

corresponds to a node, to a labeling bit of value 0, and to a labeling bit of value 1, respectively:

$$IsNode(x) = \neg \exists y \rightleftharpoons x \, (y \rightharpoonup_2 x) \qquad IsBit_0(x) = \neg IsNode(x) \wedge \neg \odot_1 x$$

$$IsBit_1(x) = \neg IsNode(x) \wedge \odot_1 x$$

The first formula is particularly useful, as we will often restrict quantification to nodes. To simplify this, we introduce the notation $\exists^\circ x \, \varphi$ to abbreviate $\exists x \, (IsNode(x) \wedge \varphi)$, and $\exists^\circ x \overset{\leq r}{\rightleftharpoons} y \, \varphi$ to abbreviate $\exists x \overset{\leq r}{\rightleftharpoons} y \, (IsNode(x) \wedge \varphi)$, and similarly for universal quantifiers.

Since every BF-formula contains at least one free first-order variable, evaluating such a formula always requires a variable assignment that provides an element as a "starting point". To remedy this, we introduce LFO, the class of formulas of *local first-order logic*, which are BF-formulas prefixed by a single universal first-order quantifier. That is, LFO consists of formulas of the form $\forall x \, \varphi$, where $x \in \mathcal{V}_{\mathrm{FO}}$ and $\varphi \in \mathrm{BF}$.

Second-order hierarchies. FO and LFO form the basis of two hierarchies of alternating second-order quantifiers. The first, called the *second-order hierarchy*, starts with the base class $\Sigma_0^{\mathrm{FO}} = \Pi_0^{\mathrm{FO}} = \mathrm{FO}$, and continues for $\ell > 0$ with the classes $\Sigma_\ell^{\mathrm{FO}}$ and Π_ℓ^{FO} that are obtained by prepending blocks of existential and universal second-order quantifiers to formulas of $\Pi_{\ell-1}^{\mathrm{FO}}$ and $\Sigma_{\ell-1}^{\mathrm{FO}}$, respectively. That is, $\Sigma_\ell^{\mathrm{FO}}$ consists of formulas of the form $\exists R_1 \ldots \exists R_n \, \varphi$, where $R_1, \ldots, R_n \in \mathcal{V}_{\mathrm{SO}}$ and $\varphi \in \Pi_{\ell-1}^{\mathrm{FO}}$, whereas Π_ℓ^{FO} consists of formulas of the form $\forall R_1 \ldots \forall R_n \, \varphi$, where $\varphi \in \Sigma_{\ell-1}^{\mathrm{FO}}$.

The other hierarchy, called the *local second-order hierarchy*, is defined the same way, except that it starts with the base class LFO instead of FO. That is, $\Sigma_0^{\mathrm{LFO}} = \Pi_0^{\mathrm{LFO}} = \mathrm{LFO}$, and for $\ell > 0$, the classes $\Sigma_\ell^{\mathrm{LFO}}$ and Π_ℓ^{LFO} are obtained by prepending blocks of existential and universal second-order quantifiers to formulas of $\Pi_{\ell-1}^{\mathrm{LFO}}$ and $\Sigma_{\ell-1}^{\mathrm{LFO}}$, respectively. Notice that it would *not* be equivalent to define Π_ℓ^{LFO} as the set of negations of formulas in $\Sigma_\ell^{\mathrm{LFO}}$ because LFO is not closed under negation.

As with the local-polynomial hierarchy, it can be helpful to think of formulas of the (local) second-order hierarchy as a two-player game between *Eve* and *Adam*, who choose the existentially and universally quantified relations, respectively. From this point of view, the referee of the game corresponds to the FO- or LFO-subformula nested inside the second-order quantifications, and the whole formula is satisfied by the input structure precisely if Eve has a winning strategy.

Second-order logic is the union of all classes of the second-order hierarchy, and similarly *local second-order logic* is the union of all classes of the local second-order hierarchy. The classes Σ_1^{FO} and Σ_1^{LFO} will be referred to as the the *existential fragments* of second-order logic and local second-order logic, respectively.

Classes of definable properties. For any class of formulas $\mathrm{C} \in \{\Sigma_\ell^{\mathrm{FO}}, \Pi_\ell^{\mathrm{FO}}, \Sigma_\ell^{\mathrm{LFO}}, \Pi_\ell^{\mathrm{LFO}}\}$ with $\ell \in \mathbb{N}$, we use the corresponding boldface notation $\mathbf{C} \in \{\mathbf{\Sigma}_\ell^{\mathrm{FO}}, \mathbf{\Pi}_\ell^{\mathrm{FO}}, \mathbf{\Sigma}_\ell^{\mathrm{LFO}}, \mathbf{\Pi}_\ell^{\mathrm{LFO}}\}$ to denote the class of structure properties that can be defined by a formula of C. It is worth noting that $\mathbf{\Sigma}_\ell^{\mathrm{FO}}|_{\mathrm{NODE}} = \mathbf{\Sigma}_\ell^{\mathrm{LFO}}|_{\mathrm{NODE}}$ and $\mathbf{\Pi}_\ell^{\mathrm{FO}}|_{\mathrm{NODE}} = \mathbf{\Pi}_\ell^{\mathrm{LFO}}|_{\mathrm{NODE}}$, since the distinction between bounded and unbounded quantification is irrelevant on single-node graphs.

5.2 Example formulas

We now show how to express a number of graph properties in local second-order logic, starting with a very simple property that the nodes can check locally: ALL-SELECTED, the

set of labeled graphs in which all nodes are assigned label 1 (i.e., they are all "selected").

▶ **Example 2.** We can easily define ALL-SELECTED on GRAPH with the LFO-formula $\forall°x\, IsSelected(x)$, where x is a first-order variable, and

$$IsSelected(x) \ = \ \exists y \rightleftharpoons x \left(IsBit_1(y) \ \wedge \ \neg \exists z \rightleftharpoons y \left(z \rightharpoonup_1 y \ \vee \ y \rightharpoonup_1 z \right) \right)$$

is a BF-formula that states that the node represented by x is labeled with the string 1. Here, the first-order variable y is used to represent x's unique labeling bit.   ⌟

Next, we consider the property of being 3-*colorable*. For $k \in \mathbb{N}_{>0}$, the set k-COLORABLE consists of the graphs G for which there exists a function $f \colon V^G \to [k\rangle$ such that $f(u) \neq f(v)$ for all $\{u, v\} \in E^G$.

▶ **Example 3.** We can define 3-COLORABLE on GRAPH with the Σ_1^{LFO}-formula

$$\exists C_0, C_1, C_2 \ \forall°x\, WellColored(x),$$

where C_0, C_1 and C_2 are unary second-order variables intended to represent the sets of nodes colored with 0, 1 and 2, respectively, x is a first-order variable, and

$$WellColored(x) = \left(\bigvee_{i \in [3\rangle} C_i(x) \right) \wedge \left(\bigwedge_{i,j \in [3\rangle \,:\, i \neq j} \neg \big(C_i(x) \wedge C_j(x) \big) \right) \wedge \forall°y \rightleftharpoons x \left(\bigwedge_{i \in [3\rangle} \neg \big(C_i(x) \wedge C_i(y) \big) \right)$$

states that the node represented by x is correctly colored. More precisely, the first two conjuncts express that x is assigned one color and one color only, while the third conjunct expresses that x's color is different from its neighbors' colors.   ⌟

To make things a little more challenging, let us now consider the complement of ALL-SELECTED, which we denote by NOT-ALL-SELECTED. This property is more difficult to express in local second-order logic. In fact, as we will see in the proof of Proposition 23, it is not Σ_1^{LFO}-definable.

▶ **Example 4.** A straightforward way to define NOT-ALL-SELECTED on GRAPH would be to negate the formula from Example 2, yielding the FO-formula $\exists°x\, \neg IsSelected(x)$. But this formula does not belong to local second-order logic because of the unbounded existential first-order quantification over x. To remedy this, we can rewrite it as an equivalent Σ_3^{LFO}-formula *ExistsUnselectedNode*, which intuitively describes the following game: First, Eve tries to cover the input graph with a spanning forest whose roots include only unselected nodes. She represents this forest by a binary relation variable P, where $P(x, y)$ is intended to mean "the parent of x is y". If she succeeds, one should thus always reach an unselected node by following parent pointers. Then, Adam tries to disprove Eve's claim that P represents a forest by showing that the relation contains a directed cycle. To do so, he chooses a subset X of nodes, and then asks Eve to assign a charge (positive or negative) to each node such that roots are positive, children outside X have the same charge as their parent, and children in X have the opposite charge of their parent. Now, if P is cycle-free, then Eve can charge the nodes as requested by simply traversing the paths of each tree from top to bottom, starting with a positive charge at the root, and inverting the charge every time she encounters a node in X. However, if P contains a cycle, then Adam can choose X to be a singleton set containing exactly one node of the cycle. By doing so, he prevents Eve from winning because she will either have to charge the node in X like its parent, or charge another node of the cycle differently than its parent.

Formally, we represent the positive and negative charges by a unary relation variable Y (interpreted as the set of positive nodes), and write

$$ExistsUnselectedNode \;=\; \exists P \,\forall X \,\exists Y \,\forall^\circ x \,\big(PointsTo[\neg IsSelected](x)\big).$$

Here, the subformula $PointsTo[\neg IsSelected](x)$ basically states that x's parent pointer points in the direction of an unselected node, assuming that both players play optimally and that Eve wins the game described above. Since the same idea will be useful later for conditions other than $\neg IsSelected(x)$, we present this subformula as a formula schema that can be instantiated with any BF-formula $\vartheta(x)$:

$$PointsTo[\vartheta](x) \;=\; UniqueParent(x) \,\wedge\, RootCase[\vartheta](x) \,\wedge\, ChildCase(x),$$

where

$$UniqueParent(x) \;=\; \exists^\circ y \overset{\leq 1}{\rightleftharpoons} x \Big(P(x,y) \,\wedge\, \forall^\circ z \overset{\leq 1}{\rightleftharpoons} x \big(P(x,z) \to z \doteq y\big)\Big)$$

states that x has exactly one parent (possibly itself, in which case it is a root),

$$RootCase[\vartheta](x) \;=\; P(x,x) \,\to\, \big(\vartheta(x) \wedge Y(x)\big)$$

states that if x is a root, then it satisfies the target condition ϑ and is positively charged, and

$$ChildCase(x) \;=\; \neg P(x,x) \,\to\, \exists^\circ y \rightleftharpoons x \Big(P(x,y) \,\wedge\, \big(Y(x) \leftrightarrow \neg(Y(y) \leftrightarrow X(x))\big)\Big)$$

states that if x is a child, then it has the same charge as its parent if it lies outside X, and the opposite charge of its parent if it belongs to X. ⌟

The spanning-forest construction described in Example 4 can be generalized to express the complement of any graph property that is definable in local second-order logic. We now illustrate this using the complement of 3-COLORABLE, which we denote by NON-3-COLORABLE. As we will show in Corollary 25, this property is not Σ_1^{LFO}-definable either.

▶ **Example 5.** To define NON-3-COLORABLE, we could simply negate the formula from Example 3, yielding the Π_1^{FO}-formula $\forall C_0, C_1, C_2 \,\exists^\circ x \,\neg WellColored(x)$. But just as in Example 4, this formula does not belong to local second-order logic because of the unbounded existential first-order quantification over x. Fortunately, the solution is also very similar: we can rewrite our initial attempt as the equivalent Π_4^{LFO}-formula $\forall C_0, C_1, C_2 \,ExistsBadNode$, using the subformula

$$ExistsBadNode \;=\; \exists P \,\forall X \,\exists Y \,\forall^\circ x \,\big(PointsTo[\neg WellColored](x)\big),$$

where P is a binary relation, X and Y are sets, and the subformula $PointsTo[\neg WellColored](x)$ is an instantiation of the formula schema from Example 4. ⌟

Next, we turn to HAMILTONIAN, the property of graphs that contain a *Hamiltonian cycle*, i.e., a cycle that goes through each node exactly once. Again, the spanning-forest construction from Example 4 proves useful to express this property in local second-order logic.

▶ **Example 6.** We present a Σ_5^{LFO}-formula that defines HAMILTONIAN on GRAPH based on the following characterization: a graph is Hamiltonian if and only if it contains a spanning subgraph (i.e., a subgraph containing all nodes) that is 2-regular (i.e., all nodes have degree 2) and connected (i.e., any two nodes are linked by a path).

Intuitively, this property can be tested through the following game: First, Eve chooses a 2-regular spanning subgraph, which she represents by a binary relation H. The intended meaning of $H(x, y)$ is "the edge $\{x, y\}$ belongs to the subgraph". She claims that the chosen subgraph is a Hamiltonian cycle. Next, Adam tries to disprove this claim by showing that Eve's subgraph is disconnected, i.e., that it consists of multiple disjoint cycles. He does this by choosing a nontrivial subset S of nodes that he claims form such a cycle. Then, Eve tries to point out a mistake in Adam's counterproof. Assuming that she did indeed choose a Hamiltonian cycle in the first step, there are only two possibilities: either Adam chose a trivial subset (i.e., the empty set or the set of all nodes), or he partitioned the Hamiltonian cycle into two nonempty sets of nodes. To tell the nodes which of the two cases applies, Eve assigns a bit to each of them, represented by a unary relation C. In the first case, represented by $\neg C(x)$, the game is over, and she wins if either all or none of the nodes belong to S. In the second case, represented by $C(x)$, she must show that there is a discontinuity in the Hamiltonian cycle, i.e., two adjacent nodes on the cycle that do not agree about their membership in S. She does this using the technique from Example 4, i.e., by constructing a spanning forest P whose roots witness a discontinuity. Proving the correctness of her forest adds two more steps to the game, where Adam challenges her with a set X and she responds with a set Y (see Example 4).

Formally, we define HAMILTONIAN with the Σ_5^{LFO}-formula

$$\exists H \, \forall S \, \exists C, P \, \forall X \, \exists Y \, \forall^\circ x \, \big(DegreeTwo(x) \, \wedge \, ConnectivityTest(x)\big),$$

where

$$DegreeTwo(x) \; = \; \exists^\circ y_1, y_2 \rightleftharpoons x \left(\begin{array}{l} y_1 \not\doteq y_2 \, \wedge \, \bigwedge_{i \in \{1,2\}} \big(H(x, y_i) \wedge H(y_i, x)\big) \, \wedge \\ \forall^\circ z \rightleftharpoons x \, \big(H(x, z) \vee H(z, x) \, \rightarrow \, \bigvee_{i \in \{1,2\}} (z \doteq y_i)\big) \end{array} \right)$$

states that x has exactly two neighbors in the spanning subgraph represented by the relation H (which must be symmetric), and

$$ConnectivityTest(x) \; = \; InAgreementOn[C](x) \, \wedge \, TrivialCase(x) \, \wedge \, PartitionedCase(x)$$

states that, as far as x can tell, Eve has correctly pointed out a mistake in Adam's counterproof, so x "believes" Eve's claim that her subgraph is connected. This belief is correct if it is shared by all nodes. More precisely, the subformula

$$InAgreementOn[C](x) \; = \; \forall^\circ y \rightleftharpoons x \, \big(C(x) \leftrightarrow C(y)\big)$$

ensures that all nodes agree on the type of mistake Adam has made,

$$TrivialCase(x) \; = \; \neg C(x) \, \rightarrow \, InAgreementOn[S](x)$$

covers the case where he has chosen a trivial partition (meaning that all nodes must agree on whether they belong to S), and

$$PartitionedCase(x) \; = \; C(x) \, \rightarrow \, PointsTo[DiscontinuityAt](x)$$

covers the case where Adam's partition creates a discontinuity in the Hamiltonian cycle. The subformula $PointsTo[DiscontinuityAt](x)$ ensures that x's parent pointer points in the direction of a discontinuity. It is an instantiation of the formula schema from Example 4 with the BF-formula

$$DiscontinuityAt(x) \; = \; \exists^\circ y \rightleftharpoons x \, \big(H(x, y) \wedge (S(x) \leftrightarrow \neg S(y))\big),$$

which states that x and one of its cycle neighbors are on opposite sides of the partition. ⌟

Finally, let us adapt Example 6 to define NON-HAMILTONIAN, the complement property of HAMILTONIAN.

▶ **Example 7.** We essentially reverse the roles of Eve and Adam in the game from Example 6, but the asymmetric nature of local second-order logic allows us to save one alternation. This time, Adam starts by proposing a supposed Hamiltonian cycle, and then Eve tries to prove that his proposition is incorrect. If she is right, there are two possible cases: either Adam did not propose a 2-regular spanning subgraph, or his subgraph is not connected. Again, Eve assigns a bit to each node to indicate which of the two cases applies (giving priority to the first if both apply). In the first case, she constructs a spanning forest whose roots are nodes at which the 2-regularity condition is violated. In the second case, she chooses a set S containing exactly one component of Adam's subgraph, and then constructs a spanning forest whose roots witness that S does indeed divide the subgraph into two nonempty parts. (There must be at least one witness, since we require by definition that the input graph is connected.)

Reusing the formula schema from Example 4 and some of the subformulas from Example 6, we can formalize this game as the Π_4^{LFO}-formula

$$\forall H \, \exists C, S, P \, \forall X \, \exists Y \, \forall^\circ x \, \big(InAgreementOn[C](x) \, \wedge \, InvalidCase(x) \, \wedge \, DisjointCase(x) \big),$$

where

$$InvalidCase(x) \; = \; \neg C(x) \; \rightarrow \; PointsTo[\neg DegreeTwo](x)$$

covers the case where Adam violated the 2-regularity condition, and

$$DisjointCase(x) \; = \; C(x) \; \rightarrow \; \neg DiscontinuityAt(x) \, \wedge \, PointsTo[DivisionAt](x)$$

covers the case where his subgraph consists of multiple components. In the latter formula, the first conjunct ensures that Eve's partition does not divide any component, while the second conjunct ensures that her partition is nontrivial, using the subformula

$$DivisionAt(x) \; = \; \neg InAgreementOn[S](x),$$

to state that x sees nodes both inside and outside of S. ⌟

6 Restrictive arbiters

The notion of arbiters defined in Section 4 was kept simple for the sake of presentation, but it can be cumbersome when constructing arbiters for specific graph properties. In this section, we provide a more flexible definition that allows us to impose additional restrictions on the input graphs and certificates. We then prove its equivalence to the original definition.

Certificate restrictors. Let r_{id} and r be positive integers, and p be a polynomial function. A *certificate restrictor* for (r, p)-bounded certificates under r_{id}-locally unique identifiers is a local-polynomial machine M that satisfies the following property for every graph G, every r_{id}-locally unique identifier assignment id of G, every (r, p)-bounded certificate-list assignment $\bar{\kappa}$ of (G, id), and every (r, p)-bounded certificate assignment κ of (G, id): if some node $u \in G$ rejects in $M(G, id, \bar{\kappa} \cdot \kappa)$, then there exists an (r, p)-bounded certificate assignment κ' differing from κ only in the certificate assigned to u such that u accepts in $M(G, id, \bar{\kappa} \cdot \kappa')$ while the verdict of all other nodes remains the same as in $M(G, id, \bar{\kappa} \cdot \kappa)$. We refer to this property as *local repairability*. Moreover, we say that M is *trivial* if $M(G, id, \bar{\kappa} \cdot \kappa) \equiv$ ACCEPT for all choices of G, id, $\bar{\kappa}$, κ.

Restrictive arbiters.   Let ℓ be a nonnegative integer, r_{id} and r be positive integers, p be a polynomial function, K be an **LP**-property, and $M_1, \ldots, M_\ell$ be certificate restrictors for (r, p)-bounded certificates under r_{id}-locally unique identifiers. A *restrictive $\mathbf{\Sigma}_\ell^{\mathrm{LP}}$-arbiter* for a graph property L on K under r_{id}-locally unique identifiers and (r, p)-bounded certificates restricted by $M_1, \ldots, M_\ell$ is a local-polynomial machine M that satisfies the following equivalence for every graph $G \in K$ and every r_{id}-locally unique identifier assignment *id* of G:

$$G \in L \iff \exists \kappa_1 \, \forall \kappa_2 \ldots \mathcal{Q} \kappa_\ell : M(G, \, id, \, \kappa_1 \cdot \kappa_2 \cdot \ldots \cdot \kappa_\ell) \equiv \text{ACCEPT},$$

where $\mathcal{Q}$ is $\forall$ if ℓ is even and $\exists$ otherwise, and all quantifiers range over (r, p)-bounded certificate assignments of (G, id) with the additional restriction that $M_i(G, \, id, \, \kappa_1 \cdot \ldots \cdot \kappa_i) \equiv \text{ACCEPT}$ for all $i \in [1 : \ell]$. If all certificate restrictors are trivial, we say that M operates under *unrestricted* (r, p)-bounded certificates. We analogously define *restrictive $\mathbf{\Pi}_\ell^{\mathrm{LP}}$-arbiters*.

Notice that the notion of $\mathbf{\Sigma}_\ell^{\mathrm{LP}}$- and $\mathbf{\Pi}_\ell^{\mathrm{LP}}$-arbiters for L introduced on page 19 coincides with the notion of restrictive $\mathbf{\Sigma}_\ell^{\mathrm{LP}}$- and $\mathbf{\Pi}_\ell^{\mathrm{LP}}$-arbiters for L on GRAPH under unrestricted certificates. We will refer to such arbiters as *permissive arbiters* when we want to emphasize the distinction from other restrictive arbiters. Although not every restrictive arbiter is permissive, we can prove the following lemma, which allows us to use arbitrary restrictive arbiters whenever it is more convenient.

▶ **Lemma 8.** *Let $\ell \in \mathbb{N}$ and $K, L \subseteq$ GRAPH with $K \in$ **LP**. The graph property $L \cap K$ belongs to $\mathbf{\Sigma}_\ell^{\mathrm{LP}}|_K$ if and only if L has a restrictive $\mathbf{\Sigma}_\ell^{\mathrm{LP}}$-arbiter on K. The analogous statement holds for $\mathbf{\Pi}_\ell^{\mathrm{LP}}|_K$.*

Proof. We prove only the first statement, since the proof for $\mathbf{\Pi}_\ell^{\mathrm{LP}}|_K$ is completely analogous. By definition, if $L \cap K$ belongs to $\mathbf{\Sigma}_\ell^{\mathrm{LP}}|_K$, then there exists a permissive $\mathbf{\Sigma}_\ell^{\mathrm{LP}}$-arbiter M for a graph property L' such that $L' \cap K = L \cap K$, and thus M is also a restrictive $\mathbf{\Sigma}_\ell^{\mathrm{LP}}$-arbiter for L on K.

For the converse, we have to convert an arbitrary restrictive arbiter for L on K into a permissive arbiter for some graph property L' such that $L' \cap K = L \cap K$. We do this for $L' = L \cap K$, proceeding in two steps by first removing the restrictions on the input graphs and then on the certificates.

Let M^a be a restrictive $\mathbf{\Sigma}_\ell^{\mathrm{LP}}$-arbiter for L on K under r_{id}^a-locally unique identifiers and (r^a, p^a)-bounded certificates restricted by $M_1^a, \ldots, M_\ell^a$.

1. We start by converting M^a into a restrictive $\mathbf{\Sigma}_\ell^{\mathrm{LP}}$-arbiter M^b for $L \cap K$ on arbitrary graphs. Since K is in **LP**, there exists an **LP**-decider M^K for that property and an integer $r_{\mathrm{id}}^K \in \mathbb{N}_{>0}$ such that M^K operates under r_{id}^K-locally unique identifiers. When executing M^b on a graph G under an identifier assignment *id* and a certificate-list assignment $\bar{\kappa}$, the nodes first simulate M^K to check whether G belongs to K. Any node that rejects in the simulation also immediately rejects in $M^b(G, id, \bar{\kappa})$, so G can only be accepted if it belongs to K. Then, the nodes that have not rejected simulate M^a and return the verdict reached in that second simulation (unless they learn about some node that has previously rejected, in which case they also reject). The machine M^b obtained this way operates on arbitrary graphs under r_{id}^b-locally unique identifiers and (r^b, p^b)-bounded certificates restricted by $M_1^b, \ldots, M_\ell^b$, where $r_{\mathrm{id}}^b = \max\{r_{\mathrm{id}}^a, r_{\mathrm{id}}^K\}$, $r^b = r^a$, $p^b = p^a$, and $M_i^b = M_i^a$ for $i \in [1 : \ell]$.

2. Now we convert M^b into a permissive $\mathbf{\Sigma}_\ell^{\mathrm{LP}}$-arbiter M^c for $L \cap K$. By definition, for every graph G and every r_{id}^b-locally unique identifier assignment *id* of G,

$$G \in L \cap K \iff \exists \kappa_1 \, \forall \kappa_2 \ldots \mathcal{Q} \kappa_\ell : M^b(G, \, id, \, \kappa_1 \cdot \kappa_2 \cdot \ldots \cdot \kappa_\ell) \equiv \text{ACCEPT},$$

where all quantifiers range over (r^b, p^b)-bounded certificate assignments of (G, id) with the additional restriction that $M_i^b(G, id, \kappa_1 \cdot \ldots \cdot \kappa_i) \equiv \text{ACCEPT}$ for all $i \in [1 : \ell]$. The new machine M^c has to satisfy the analogous equivalence without the additional restriction on the certificate assignments.

When executing M^c on G under id and $\kappa_1 \cdot \ldots \cdot \kappa_\ell$, the nodes first simulate $M_1^b, \ldots, M_\ell^b$ to check if the given certificates satisfy the imposed restrictions. As a result of this first phase, each node u stores a flag ok_i for each $i \in [1 : \ell]$ to indicate whether u accepts in $M_i^b(G, id, \kappa_1 \cdot \ldots \cdot \kappa_i)$. Then, the nodes simulate M^b while simultaneously updating their flags to propagate errors. That is, if a node sees that the ok_i flag of one of its neighbors is false, then it also sets its own ok_i flag to false. Once the simulation of M^b has terminated, each node u goes *sequentially* through its flags $ok_1, \ldots, ok_\ell$ to verify that they are all true. If it encounters an ok_i flag that is false, u aborts, writes a verdict on its internal tape, and enters state q_{stop}. The verdict depends on whether the certificate assignment κ_i is quantified existentially or universally: in the first case, the verdict is 0 (reject), whereas in the second it is 1 (accept). Finally, if it did not stop before, u returns the same verdict it would have returned when executing M^b.

Note that the sequential verification and early termination described above ensures that quantifications are relativized in the same way as for M^b. More precisely, if the first certificate assignment violating the restrictions is quantified existentially, then the input graph is rejected because all nodes that know about the violation return 0. If instead the first certificate assignment κ_i violating the restrictions of the corresponding machine M_i^b is quantified universally, then there are two possibilities: either the input graph is accepted (the desired outcome), or it is rejected because of some node u that does not know about the violation. However, in the latter case, there exists an (r^b, p^b)-bounded certificate assignment κ_i' that does not violate the restrictions of M_i^b but for which u still rejects. This is because M_i^b satisfies local repairability, so all defects in κ_i can be fixed without affecting u. Intuitively speaking, u's verdict is legitimate since it is independent of the violation. ◄

7 A generalization of Fagin's theorem

The founding result of descriptive complexity theory is Fagin's theorem, which provides a logical, and thus machine-independent, characterization of the complexity class **NP** (see, e.g., [19, Thm. 3.2.4]). In the context of this paper, we can state it as follows.[1]

▶ **Theorem 9** (Fagin [7]). *On single-node graphs, a property can be verified by a local-polynomial machine if and only if it can be defined by a formula of the existential fragment of local second-order logic. In symbols, $\mathbf{NLP}|_{\text{NODE}} = \mathbf{\Sigma}_1^{\text{LFO}}|_{\text{NODE}}$, or equivalently, $\mathbf{NP} = \mathbf{\Sigma}_1^{\text{FO}}|_{\text{NODE}}$.*

The inclusion from right to left is straightforward because any Σ_1^{FO}-formula $\exists R_1 \ldots \exists R_n \, \varphi$ can be evaluated in polynomial time by a Turing machine that is given some interpretation of $R_1, \ldots, R_n$ (encoded in the certificate chosen by Eve). The machine can use brute force

---

[1] In the literature, Fagin's theorem is usually stated in terms of arbitrary graphs (or even arbitrary structures) instead of labeled single-node graphs. More specifically, a graph property can be verified by a (centralized) Turing machine operating in polynomial time on encodings of graphs if and only if it can be defined by a formula of the existential fragment of second-order logic. In symbols, $\mathbf{NP}|_{enc(\text{GRAPH})} = \{enc(L) \mid L \in \mathbf{\Sigma}_1^{\text{FO}}|_{\text{GRAPH}}\}$, where $enc \colon \text{GRAPH} \to \text{NODE}$ is some encoding of graphs as binary strings. However, the statement presented here is equivalent, since it is immaterial whether we encode graphs as strings or vice versa (see, e.g., Problem 8.4.12 in Papadimitriou's book [32]).

to check whether the first-order formula φ is satisfied under the given interpretation of $R_1, \ldots, R_n$, by simply iterating over all possible interpretations of the first-order variables in φ. The reverse inclusion, however, is more intricate, as it involves encoding the space-time diagram of a Turing machine by a collection of relations over the input structure. The key insight that makes this possible is the following: since the machine's running time is polynomially bounded by the structure's cardinality, each cell of the space-time diagram can be represented by a tuple of elements whose length depends on the degree of the bounding polynomial.

In this section, we generalize Theorem 9 from single-node graphs to arbitrary graphs, thereby providing a logical characterization of the complexity class **NLP**. We then further generalize this result to obtain similar characterizations of the higher levels of the local-polynomial hierarchy.

Our proofs make use of the following lemma, which basically states that in the execution of a local-polynomial machine, the space-time diagram of each node in each round is polynomially bounded by the cardinality of a constant-radius neighborhood of the node. For a given execution of a machine M, the *space usage* of node u in round i is the maximum number of tape cells that u occupies in round i. More precisely, if we denote by t the step running time of u in round i and by ℓ_j the total length of u's tape contents after its j-th computation step, then u's space usage in round i is $\max\{\ell_j \mid 0 \le j \le t\}$.

▶ **Lemma 10.** *Let ℓ and r be positive integers, p be a polynomial function, and M be a local-polynomial machine running in round time r and step time p. There exists a polynomial function f such that the following holds for every labeled graph G, every small r-locally unique identifier assignment id of G, and all (r, p)-bounded certificate assignments $\kappa_1, \ldots, \kappa_\ell$ of (G, id): in the execution of M on G under id and $\kappa_1 \cdot \ldots \cdot \kappa_\ell$, the step running time and space usage of each node $u \in G$ in each round $i \in [1\!:\!r]$ are at most $f\big(\mathrm{card}(N_{4r}^{\$G}(u))\big)$, i.e., f applied to the number of nodes and labeling bits of u's $4r$-neighborhood in G.*

Proof. By definition, if id is a small r-locally unique identifier assignment id of G, then $\mathrm{len}(id(u)) \le \lceil \log_2 \mathrm{card}(N_{2r}^G(u)) \rceil \le \lceil \log_2 \mathrm{card}(N_{2r}^{\$G}(u)) \rceil$ for every node $u \in G$. Consider arbitrary (r, p)-bounded certificate assignments $\kappa_1, \ldots, \kappa_\ell$ of (G, id), let $\bar{\kappa} = \kappa_1 \cdot \ldots \cdot \kappa_\ell$, and let us denote by $b_i(u)$ the maximum of u's step running time and space usage in round $i \in [1\!:\!r]$ of M's execution on G under id and $\bar{\kappa}$. Furthermore, for $k \in [3r\!:\!4r-1]$, let us write $n_k(u)$ as a shorthand for $\mathrm{card}(N_k^{\$G}(u))$. We now show by induction that for each round $i \in [1\!:\!r]$, there is a polynomial f_i (independent of G, id, and $\bar{\kappa}$) such that $b_i(u) \le f_i(n_{3r+i-1}(u))$.

Assuming node u has d neighbors $v_1, \ldots, v_d$, its initial tape contents in round 1 consist of the string $\#^d$ on its receiving tape and the string $\lambda^G(u) \# id(u) \# \bar{\kappa}(u)$ on its internal tape. Since M runs in step time p, node u's step running time in round 1 cannot exceed $p\big(d + \mathrm{len}(\lambda^G(u) \# id(u) \# \bar{\kappa}(u))\big)$. This also implies that u's space usage in round 1, and thus $b_1(u)$, cannot exceed $3 \cdot p\big(d + \mathrm{len}(\lambda^G(u) \# id(u) \# \bar{\kappa}(u))\big)$, since all three tape heads start on the leftmost cell of their respective tape and can advance by at most one cell in each computation step. Moreover, we know that $d \le n_1(u)$, $\mathrm{len}(\lambda^G(u)) \le n_0(u)$, $\mathrm{len}(id(u)) \le \lceil \log_2 n_{2r}(u) \rceil$, and (since each κ_i is (r, p)-bounded),

$$\mathrm{len}(\bar{\kappa}(u)) \le \ell + \ell \cdot p\Big(\sum_{v \in N_r^G(u)} \mathrm{len}(\lambda^G(v) \# id(v))\Big) \le p'(n_{3r}(u)),$$

where p' is a polynomial that depends only on ℓ and p. The last inequality stems from the fact that the $2r$-neighborhood of any node $v \in N_r^G(u)$ is included in u's $3r$-neighborhood.

We can therefore conclude that $b_1(u) \leq f_1(n_{3r}(u))$ for some polynomial f_1 that can be easily derived from p'.

Now, let us assume by induction that there exists a polynomial f_i such that $b_i(u) \leq f_i(n_{3r+i-1}(u))$ for every node $u \in G$. At the beginning of round $i+1$, u's internal tape contains a string s of length less than $b_i(u)$ and its receiving tape contains a string of the form $\mu_1 \# \ldots \# \mu_d \#$, where μ_j is a message of length less than $b_i(v_j)$ that was sent by neighbor v_j in round i. Again, since we know that M runs in step time p, this gives us an upper bound on u's step running time and space usage: $b_{i+1}(u) \leq 3 \cdot p\big(\text{len}(s) + \text{len}(\mu_1 \# \ldots \# \mu_d \#)\big)$. From the induction hypothesis we obtain that $\text{len}(s) + \text{len}(\mu_1 \# \ldots \# \mu_d \#) \leq d + \sum_{w \in \{u,v_1,\ldots v_d\}} f_i(n_{3r+i-1}(w))$, which cannot exceed $n_1(u) \cdot \big(f_i(n_{3r+i}(u)) + 1\big)$, given that the $(3r + i - 1)$-neighborhood of every neighbor v_j of u is included in u's $(3r+i)$-neighborhood. We thus obtain the bound $b_{i+1}(u) \leq f_{i+1}(n_{3r+i}(u))$, where f_{i+1} is a polynomial that can be easily derived from p and f_i.

Ultimately, we have $b_i(u) \leq f_r(n_{4r-1}(u))$ for all $i \in [1:r]$, which implies our claim. ◄

We now generalize Fagin's theorem from **NP** to **NLP**, resulting in the following statement. Notice that the original result (Theorem 9) can be recovered by restricting both sides of the equivalence to single-node graphs.

▶ **Theorem 11.** *On arbitrary graphs, a property can be verified by a local-polynomial machine if and only if it can be defined by a formula of the existential fragment of local second-order logic. In symbols,* $\mathbf{NLP} = \mathbf{\Sigma}_1^{\text{LFO}}|_{\text{GRAPH}}$.

As this result will be further generalized below, we do not explicitly prove the backward direction, which is a simple special case of the backward direction of Theorem 12 on page 35. However, we do explicitly prove the forward direction to provide a more accessible introduction to the general case presented on page 36. The key idea of encoding a space-time diagram as a collection of relations remains the same as in Fagin's original proof, but we have to deal with additional issues such as locally unique identifiers and the exchange of messages between adjacent nodes.

Proof of Theorem 11 − Forward direction. Let L be a graph property in **NLP**, and let $M = (Q, \delta)$ be an **NLP**-verifier for L that operates under r_{id}-locally unique identifiers and (r_1, p_1)-bounded certificates, and runs in round time r_2 and step time p_2. Moreover, let $r = \max\{r_{\text{id}}, r_1, r_2\}$, and let p be a polynomial that bounds both p_1 and p_2. By the proof of Lemma 8, we may assume without loss of generality that M rejects under any certificate assignment violating the (r_1, p_1)-boundedness condition, so it does not matter if Eve chooses certificates that are too large. Now, we fix $k \in \mathbb{N}$ such that the polynomial f described in Lemma 10 (for $\ell = 1$, and our choices of r, p, and M) satisfies $f(n) < n^k$ for $n > 1$. By Lemma 10, for every graph G whose structural representation has at least two elements,[2] every small r-locally unique identifier assignment id of G, and every (r, p)-bounded certificate assignment κ of (G, id), the step running time and space usage of each node $u \in G$ are bounded by $\big(\text{card}(N_{4r}^{\$G}(u))\big)^k$ in each round $i \in [1:r]$ of the corresponding execution of M. Intuitively, this gives us a bound on the amount of information required to describe Eve's choice of certificates and the subsequent execution of the verifier M, assuming that the nodes of the input graph are assigned small identifiers. Although our description below does not

---

[2] For the single-node graph G whose node u is labeled with the empty string, we have $(\text{card}(N_{4r}^{\$G}(u)))^k = 1$ for all $k \in \mathbb{N}$, since the structural representation $\$G$ has only one element. If G belongs to L, we can easily treat it as a special case in the formula φ_M described here.

explicitly state this assumption, it may run out of "space" if the provided identifiers are too large. However, this is not a problem because incomplete executions are simply ignored.

We convert M into a Σ_1^{LFO}-sentence defining L that is of the form $\varphi_M = \exists \bar{R} \, \forall^{\circ} x \, \psi(x)$, where $\bar{R} = (Z, C, \tilde{C}, I_0, I_1, \dots)$ is a collection of second-order variables intended to represent an accepting execution of M, and $\psi(x)$ is a BF-formula stating that, from the point of view of node x, the variables in $\bar{R}$ do indeed represent a valid execution of M in which x accepts. We start by introducing the relation variables in $\bar{R}$, together with their intended interpretations.

Z: a $(2k+1)$-ary relation that associates with each node x a linear order on the k-tuples of elements in x's $4r$-neighborhood. The intended meaning of $Z(x, \bar{p}_1, \bar{p}_2)$ is: "from x's point of view, tuple $\bar{p}_1$ is strictly smaller than tuple $\bar{p}_2$".

C: a $(4k+2)$-ary relation that establishes a correspondence between the linear orders of two nodes x_1 and x_2 that lie at distance at most $2r$ of each other. The intended meaning of $C(x_1, \bar{p}_1, \bar{p}_1', x_2, \bar{p}_2, \bar{p}_2')$ is: "The number of steps from $\bar{p}_1$ to $\bar{p}_1'$ in the linear order of x_1 is the same as the number of steps from $\bar{p}_2$ to $\bar{p}_2'$ in the linear order of x_2". Or, to put it more loosely, "$\bar{p}_1' - \bar{p}_1$ for x_1 is the same as $\bar{p}_2' - \bar{p}_2$ for x_2".

$\tilde{C}$: a $(k+2)$-ary relation that establishes a correspondence between the linear order of node x (defined by Z) and the order of the bits in x's label (represented by elements of the input structure). The intended meaning of $\tilde{C}(x, z, \bar{p})$ is: "the position of bit z in x's label corresponds to the position of $\bar{p}$ in x's linear order".

I_α: a family of $(k+1)$-ary relations, for $\alpha \in \{0, 1\}$, whose purpose is to represent the r-locally unique identifier of each node x. The intended meaning of $I_\alpha(x, \bar{p})$ is: "the $\bar{p}$-th bit of x's identifier is an α".

O: a binary order relation that compares, with respect to their identifiers, two nodes x_1 and x_2 that have some common neighbor. The intended meaning of $O(x_1, x_2)$ is: "x_1's identifier is smaller than x_2's identifier".

K_α: a family of $(k+1)$-ary relations, for $\alpha \in \{0, 1\}$, whose purpose is to represent the (r, p)-bounded certificate of each node x. The intended meaning of $K_\alpha(x, \bar{p})$ is: "the $\bar{p}$-th bit of x's certificate is an α".

S_q^i: a family of $(k+1)$-ary relations, for $i \in [1:r]$ and $q \in Q$, that indicate the state of each node in every communication round and computation step. The intended meaning of $S_q^i(x, \bar{t})$ is: "in round i, at step $\bar{t}$, node x is in state q".

H_β^i: a family of $(2k+1)$-ary relations, for $i \in [1:r]$ and $\beta \in \{rcv, int, snd\}$, that indicate the positions of the three tape heads of each node in every communication round and computation step. The intended meaning of $H_\beta^i(x, \bar{t}, \bar{p})$ is: "in round i, at step $\bar{t}$, node x's head on tape β is at position $\bar{p}$". Here, we adopt the convention that rcv, int, and snd refer to the receiving, internal, and sending tapes, respectively.

$T_{\beta, \alpha}^i$: a family of $(2k+1)$-ary relations, for $i \in [1:r]$, $\beta \in \{rcv, int, snd\}$, and $\alpha \in \{\vdash, \square, \#, 0, 1\}$, that indicate the tape contents of each node in every communication round and computation step. The intended meaning of $T_{\beta, \alpha}^i(x, \bar{t}, \bar{p})$ is: "in round i, at step $\bar{t}$, node x's tape β contains an α at position $\bar{p}$".

X_β^i: a family of $(k+2)$-ary relations, for $i \in [1:r]$ and $\beta \in \{rcv, snd\}$, that indicate the tape positions of the incoming and outgoing messages that each node exchanges with its neighbors in every communication round. The intended meaning of $X_\beta^i(x, y, \bar{p})$ is: "in round i, the message exchanged between x and y is written immediately after position $\bar{p}$ on x's tape β" (which implies that the symbol at position $\bar{p}$ is either a $\vdash$ or a $\#$). For $\beta = rcv$, this corresponds to the incoming message that x receives from y at the beginning of round i, whereas for $\beta = snd$, this corresponds to the outgoing message that x sends to y at the end of round i.

It remains to specify the BF-formula $\psi(x)$, which is evaluated from the point of view of the node represented by the first-order variable x. We define ψ as the conjunction of the following BF-formulas, most of which are described only on an intuitive level to keep the exposition readable.

LinearTupleOrder(x) states that the relation Z does indeed yield a linear order over the k-tuples in x's $4r$-neighborhood, as described above. We can write this formula as follows:

$$\forall \bar{p}_1, \bar{p}_2, \bar{p}_3 \stackrel{\leq 4r+1}{\Longleftarrow} x \left(\begin{array}{l} \neg Z(x, \bar{p}_1, \bar{p}_1) \ \wedge \ \left(\bar{p}_1 \neq \bar{p}_2 \ \rightarrow \ Z(x, \bar{p}_1, \bar{p}_2) \vee Z(x, \bar{p}_2, \bar{p}_1) \right) \\ \wedge \ \left(Z(x, \bar{p}_1, \bar{p}_2) \wedge Z(x, \bar{p}_2, \bar{p}_3) \ \rightarrow \ Z(x, \bar{p}_1, \bar{p}_3) \right) \end{array} \right)$$

The formula "looks" up to distance $4r+1$ because the labeling bits of a node at distance $4r$ lie $4r + 1$ steps away in the structural representation of the input graph.

Correspondences(x) enforces that the relation C establishes the desired correspondence between the linear order of x and the linear order of every other node in x's $2r$-neighborhood, and that the relation $\tilde{C}$ establishes the desired correspondence between the label and linear order of x. Both properties can be easily specified inductively. For C, the base case states that $C(x, \bar{p}_1, \bar{p}_1, y, \bar{p}_2, \bar{p}_2)$ must hold for every node y in x's $2r$-neighborhood and all k-tuples $\bar{p}_1$ and $\bar{p}_2$ in the $4r$-neighborhoods of x and y, respectively. The induction step then states that $C(x, \bar{p}_1, \bar{p}_1', y, \bar{p}_2, \bar{p}_2')$ implies $C(x, \bar{p}_1, "\bar{p}_1' + 1", y, \bar{p}_2, "\bar{p}_2' + 1")$, where "$\bar{p}_1' + 1$" and "$\bar{p}_2' + 1$" represent the direct successors of $\bar{p}_1'$ and $\bar{p}_2'$ with respect to x and y, respectively. For $\tilde{C}$, the specification is very similar.

UniqueIdentifier(x) ensures that the relations I_0 and I_1 represent a binary string for x and that this string constitutes an r-locally unique identifier. In detail, this means that, with respect to x, each position $\bar{p}$ can be either unlabeled, labeled with 0, or labeled with 1 (but not both). If $\bar{p}$ is labeled, then so must be its predecessor "$\bar{p}-1$" in the order defined by Z. Furthermore, for every node y in x's $2r$-neighborhood, there exists a position $\bar{p}$ at which the labelings of x and y differ, entailing that x's identifier is r-locally unique. In order to refer to "the labeling of y at position $\bar{p}$", we use C to relate $\bar{p}$ to another k-tuple $\bar{p}'$ that represents the same position as $\bar{p}$ from the point of view of node y.

NeighborOrder(x) states that in x's 1-neighborhood, the relation O agrees with the identifier order of the nodes. That is, for all neighbors y and z of x, we have $O(y, z)$ precisely if the identifier of y is smaller than the identifier of z with respect to the identifier order. This is the case if at the first position $\bar{p}$ where the identifiers of y and z differ, either y's bit is smaller than z's bit or we have reached the end of y's identifier while z's identifier still goes on. Again, we make use of the correspondence relation C to relate matching positions of y and z.

Certificate(x) is very similar to *UniqueIdentifier*(x). It ensures that the relations K_0 and K_1 represent a certificate of x consisting of 0's and 1's.

ExecGroundRules(x) formalizes some basic properties that any execution of M must satisfy at node x. In particular, in each round $i \in [1:r]$, at every step $\bar{t}$, the machine must be in exactly one state $q \in Q$, there must be exactly one symbol $\alpha \in \{\vdash, \square, \#, 0, 1\}$ written at each position $\bar{p}$ of each of the three tapes (the symbol of the first position always being $\vdash$), and each tape head must be located at exactly one position. Moreover, in each round there must be some step $\bar{t}$ at which the machine halts by reaching one of the states q_{pause} or q_{stop}. Intuitively speaking, x must respect the basic "mechanics" of Turing machines and may not run out of space or time (both of which are bounded by the number of k-tuples in x's $4r$-neighborhood).

OwnInput(x) ensures that at step 0 of round 1, node x's internal tape contains the string $\lambda^G(x) \# id(x) \# \kappa(x)$, where $\lambda^G(x)$ is the label of x that is represented by the unary relation $\odot_1^{\$G}$ of the input structure $\$G$, $id(x)$ is the identifier of x that is represented by the relations I_0 and I_1, and $\kappa(x)$ is the certificate of x that is represented by the relations K_0 and K_1. To check that $\lambda^G(x)$ is written at the beginning of the internal tape, i.e., just after the left-end marker $\vdash$, we make use of the relation $\tilde{C}$ as follows: for every labeling bit z and every position $\bar{p}$, if $\tilde{C}(x, z, \bar{p})$ and $IsBit_\alpha(z)$, for $\alpha \in \{0, 1\}$, then we must have $T^1_{int,\alpha}(x, \bar{t}_0, "\bar{p}+1")$, where $\bar{t}_0$ is the tuple representing computation step 0 and "$\bar{p}+1$" denotes the direct successor of $\bar{p}$ (according to the order defined by Z with respect to x). Based on that, the position of the first separator $\#$ must be x's smallest position $\bar{s}$ such that $\tilde{C}(x, z, "\bar{s}-1")$ does not hold for any labeling bit z. Next, we check that the first separator is followed by $id(x)$, using the relation C to express that $id(x)$ is shifted by "$\bar{s}+1$" positions to the right of the initial position $\bar{p}_0$: for all positions $\bar{p}$ and $\bar{p}'$, if $C(x, \bar{p}_0, \bar{p}, x, "\bar{s}+1", \bar{p}')$ and $I_\alpha(x, \bar{p})$, for $\alpha \in \{0, 1\}$, then we must have $T^1_{int,\alpha}(x, \bar{t}_0, \bar{p}')$. Finally, we check that the second separator $\#$ is followed by $\kappa(x)$, proceeding completely analogously with K_α instead of I_α.

ReceiveMessages$_i(x)$, for $i \in [1:r]$, states that the messages received by x at the beginning of round i are written on x's receiving tape, each followed by the separator $\#$, and sorted according to the identifier order of the senders. To accomplish this, the formula also guarantees by induction that the relation X^i_{rcv} correctly represents the starting positions of the messages that x receives from each neighbor.

- The base case of the definition of X^i_{rcv} is straightforward, since the message from x's first neighbor y_1 (with respect to the order relation O) must start immediately after the left-end marker $\vdash$ on x's receiving tape. That is, $X^i_{rcv}(x, y_1, \bar{p}_1)$, where $\bar{p}_1$ is x's first position (with respect to the relation Z).

- Now, we need to distinguish two cases. First, suppose that for a given neighbor y, we have $X^{i-1}_{snd}(y, x, \bar{p}')$ and $X^i_{rcv}(x, y, \bar{p})$, i.e., the message that y sends to x at the end of round $i-1$ is stored right after position $\bar{p}'$ on y's sending tape, and the message that x receives from y at the beginning of round i is stored right after position $\bar{p}$ on x's receiving tape. Based on this information, the formula *ReceiveMessages*$_i(x)$ ensures that the two messages are indeed the same by stating that on y's sending tape, at the end of round $i-1$, every position $\bar{s}'$ located between $\bar{p}'$ and the next occurrence of $\#$ (including the latter position) contains the same symbol as the corresponding position $\bar{s}$ on x's receiving tape at the beginning of round i. (Without loss of generality, we may assume that the messages on y's sending tape are always followed by a $\#$ and do not contain any useless $\square$'s.) The fact that position $\bar{s}$ corresponds to position $\bar{s}'$ is expressed by the relation $C(x, \bar{p}, \bar{s}, y, \bar{p}', \bar{s}')$, which states that the distance from $\bar{p}$ to $\bar{s}$ on x's receiving tape is the same as the distance from $\bar{p}'$ to $\bar{s}'$ on y's sending tape.
  The second case is when we have $X^i_{rcv}(x, y, \bar{p})$ but there is no $\bar{p}'$ such that $X^{i-1}_{snd}(y, x, \bar{p}')$. This means that y has not written any message for x on its sending tape at the end of round $i-1$, and therefore that x receives the empty string from y in round i. (Note that this happens in particular for $i = 1$.) In this case, our formula simply states that at the beginning of round i, the receiving tape of x contains the separator $\#$ at position "$\bar{p}+1$" (which, as before, represents the direct successor of $\bar{p}$ with respect to Z).

- Finally, to complete the inductive definition of X^i_{rcv}, we state that $X^i_{rcv}(x, y, \bar{p})$ implies $X^i_{rcv}(x, z, \bar{s})$ if z is the smallest neighbor of x strictly greater than y (with respect to O) and $\bar{s}$ is the smallest position of x that is strictly greater than $\bar{p}$ (with respect

to Z) and contains the symbol $\#$. To make sure that the relation X^i_{rcv} is minimal, we also require that for every neighbor y of x, there is only one position $\bar{p}$ such that $X^i_{rcv}(x, y, \bar{p})$.

InitLocalConfig$_i$ (x), for $i \in [1 : r]$, provides the missing parts of the description of x's local configuration at time 0 in round i: First, if $i \geq 2$, the formula stipulates that the node's internal tape contains the same string as at the end of round $i - 1$; for $i = 1$, the initial content of the internal tape in round 1 has already been specified by the above formula *OwnInput*(x). Second, the sending tape must initially be completely empty. Third, the machine's state must be reset to q_{start}, unless it has reached q_{stop} in round $i - 1$, in which case the state remains unchanged.

ComputeLocally$_i$ (x), for $i \in [1 : r]$, describes how M performs a single computation step at node x in round i. The formula essentially states the following: On the one hand, for the receiving tape, the cell contents remain the same at all steps $\bar{t}$, while for the internal and sending tapes, the symbol at position $\bar{p}$ remains unchanged between steps $\bar{t}$ and "$\bar{t} + 1$" if the corresponding tape head is not located at position $\bar{p}$ at step $\bar{t}$. On the other hand, if at step $\bar{t}$ the machine is in state q and reads the symbols α_1, α_2, α_3 at positions $\bar{p}_1$, $\bar{p}_2$, $\bar{p}_3$ of the receiving, internal, and sending tapes, respectively, then the configuration at time "$\bar{t} + 1$" must take into account the updates specified by the transition function δ. That is, if $\delta(q, \alpha_1, \alpha_2, \alpha_3) = (q', \alpha'_2, \alpha'_3, m_1, m_2, m_3)$, then at time "$\bar{t} + 1$", the machine is in state q', position $\bar{p}_2$ of the internal tape contains the symbol α'_2, position $\bar{p}_3$ of the sending tape contains the symbol α'_3, and the three tape heads are located at positions "$\bar{p}_1 + m_1$", "$\bar{p}_2 + m_2$", and "$\bar{p}_3 + m_3$".

SendMessages$_i$ (x), for $i \in [1 : r]$, guarantees that the relation X^i_{snd} correctly represents the starting positions of the messages that x sends to each neighbor at the end of round i. Since the inductive definition of X^i_{snd} is very similar to that of X^i_{rcv} given in *ReceiveMessages$_i$*(x), we do not further elaborate on it. Let us only note that in contrast to X^i_{rcv}, for some neighbors y there might be no $\bar{p}$ such that $X^i_{snd}(x, y, \bar{p})$. This happens when x does not write enough messages on its sending tape for all its neighbors (in which case the missing messages default to the empty string).

Accept(x) states that x must eventually reach state q_{stop} in some round $i \in [1 : r]$, with the string 1 written on its internal tape. ◀

Fagin's theorem was extended by Stockmeyer [36] to the higher levels of the polynomial hierarchy, thus establishing a levelwise correspondence with the second-order hierarchy (see, e.g., [26, Cor. 9.9]). In the next theorem, we show that this extension also carries over to the distributed setting. Again, the original result can be recovered by restricting both sides of the equivalences to single-node graphs.

▶ **Theorem 12.** *The local-polynomial hierarchy and the local second-order hierarchy on graphs are levelwise equivalent from level 1 onwards. More precisely, $\boldsymbol{\Sigma}^{\mathrm{LP}}_\ell = \boldsymbol{\Sigma}^{\mathrm{LFO}}_\ell|_{\mathrm{GRAPH}}$ and $\boldsymbol{\Pi}^{\mathrm{LP}}_\ell = \boldsymbol{\Pi}^{\mathrm{LFO}}_\ell|_{\mathrm{GRAPH}}$ for all $\ell \in \mathbb{N}_{>0}$.*

Proof of Theorem 12 – Backward direction. Let us first assume that we are given a $\Sigma^{\mathrm{LFO}}_\ell$-sentence $\varphi = \exists \bar{R}_1 \forall \bar{R}_2 \ldots \mathcal{Q} \bar{R}_\ell \, \forall x \, \psi(x)$, where each $\bar{R}_i$ is a tuple of second-order variables, and $\mathcal{Q}$ is $\forall$ if ℓ is even and $\exists$ otherwise. Let L be the property defined by φ on graphs, and let r be the maximum nesting depth of bounded first-order quantifiers in the BF-formula ψ (intuitively, the distance up to which ψ can "see"). We construct a restrictive $\boldsymbol{\Sigma}^{\mathrm{LP}}_\ell$-arbiter M_φ for L under r-locally unique identifiers and (r, p)-bounded certificates restricted by machines $M_1, \ldots, M_\ell$. Here, p is a polynomial chosen based on the number and arities of the variables in $\bar{R}_1, \ldots, \bar{R}_\ell$ such that the certificates described below satisfy the (r, p)-boundedness condition. The

intention is that for each $i \in [1:\ell]$ and each node u of the input graph G, the certificate $\kappa_i(u)$ encodes part of a variable assignment on $\$G$ that assigns interpretations to the relation variables in $\bar{R}_i$. More precisely, each machine M_i restricts quantification over κ_i such that for each relation variable R in $\bar{R}_i$ of arity k, the certificate $\kappa_i(u)$ must encode a set of k-tuples whose first element represents either u or one of its labeling bits, and whose remaining elements all represent nodes or labeling bits that lie in the $2r$-neighborhood of u. In order to refer to the elements corresponding to a particular node v, the certificate makes use of v's locally unique identifier $id(v)$. Since these additional restrictions on the certificates clearly satisfy local repairability, Lemma 8 allows us to subsequently convert M_φ into an equivalent $\mathbf{\Sigma}_\ell^{\mathrm{LP}}$-arbiter without restrictions.

The machine M_φ proceeds in $r + 1$ rounds. In the first r rounds, each node u collects information about its r-neighborhood, which allows it to reconstruct $N_r^G(u)$ and all the identifiers and certificates in that subgraph. Then, in the last round, each node u evaluates ψ locally on $N_r^{\$G}(u)$ and accepts if and only if ψ is satisfied at the elements representing u and its labeling bits. Since ψ only makes use of first-order quantification, this can be done in polynomial time (simply by exhaustive search).

Note that the certificate assignments $\kappa_1, \ldots, \kappa_\ell$ encode relations $\bar{R}_1, \ldots, \bar{R}_\ell$ on $\$G$ that relate only elements whose associated nodes lie at distance at most $2r$ from each other. However, this does not entail any loss of generality because the formula ψ (which belongs to BF) can only make statements about elements that lie this close together anyway.

The construction is completely analogous if we are given a Π_ℓ^{LFO}-sentence instead of a $\Sigma_\ell^{\mathrm{LFO}}$-sentence. ◀

Proof of Theorem 12 – Forward direction. We start by showing that $\mathbf{\Sigma}_\ell^{\mathrm{LP}} \subseteq \mathbf{\Sigma}_\ell^{\mathrm{LFO}}|_{\mathrm{GRAPH}}$. Let L be a graph property in $\mathbf{\Sigma}_\ell^{\mathrm{LP}}$, and let M be a $\mathbf{\Sigma}_\ell^{\mathrm{LP}}$-arbiter for L that operates under r_{id}-locally unique identifiers and (r_1, p_1)-bounded certificates, and runs in round time r_2 and step time p_2. Moreover, let $r = \max\{r_{\mathrm{id}}, r_1, r_2\}$, and let p be a polynomial that bounds both p_1 and p_2. By the proof of Lemma 8, we may assume without loss of generality that M relativizes quantification to certificate assignments that satisfy the (r_1, p_1)-boundedness condition, so it does not matter if Eve and Adam choose certificates that are too large. Now, we fix $k \in \mathbb{N}$ such that the polynomial f described in Lemma 10 (for our choices of ℓ, r, p, and M) satisfies $f(n) < n^k$ for $n > 1$. By Lemma 10, for every graph G whose structural representation has at least two elements,[3] every small r-locally unique identifier assignment id of G, and all (r, p)-bounded certificate assignments $\kappa_1, \ldots, \kappa_\ell$ of (G, id), the step running time and space usage of each node $u \in G$ are bounded by $\left(\mathrm{card}(N_{4r}^{\$G}(u))\right)^k$ in each round $i \in [1:r]$ of the corresponding execution of M. Intuitively, this gives us a bound on the amount of information required to describe a game between Eve and Adam and the subsequent execution of the arbiter M, assuming that the nodes of the input graph are assigned small identifiers.

To convert M into a $\Sigma_\ell^{\mathrm{LFO}}$-sentence φ_M defining L, we use the same relation variables as in the proof of Theorem 11 (with the same intended interpretations), except that instead of (K_0, K_1) we now use $(K_0^j, K_1^j)_{j \in [1:\ell]}$ to represent the certificate assignments $\kappa_1, \ldots, \kappa_\ell$. More precisely, our formula φ_M is of the form

$$\exists \bar{R}_{aux} \, \exists \bar{R}_1 \, \forall \bar{R}_2 \ldots \wp \bar{R}_\ell \, \wp \bar{R}_{exe} \, \forall^\circ x \, \psi(x),$$

---

[3] See Footnote 2 in the proof of Theorem 11 (forward direction) on page 31.

where $Ǒ$ is $\forall$ if ℓ is even and $\exists$ otherwise, $\bar{R}_{aux} = (Z, C, \tilde{C}, I_0, I_1, O)$ is a collection of auxiliary relation variables that help us specify the remaining relations, I_0 and I_1 are intended to represent an r-locally unique identifier assignment id, each $\bar{R}_j$ is a pair (K_0^j, K_1^j) of variables intended to represent a certificate assignment κ_j (for $j \in [1:\ell]$), and $\bar{R}_{exe} = (\dots, S_q^i, \dots, H_\beta^i, \dots, T_{\beta,\alpha}^i, \dots, X_\beta^i, \dots)$ is a collection of variables intended to represent the execution of M on the input graph under id and $\kappa_1 \cdot \dots \cdot \kappa_\ell$.

Similarly to the proof of Theorem 11, we would like to state in the BF-formula $\psi(x)$ that from the point of view of node x, the relations in $\bar{R}_{aux}, \bar{R}_1, \dots, \bar{R}_{exe}$ are valid in the sense that they correspond to their intended interpretations. However, in order to relativize all quantifications to valid relations, we now have to take into account whether a relation is quantified existentially or universally: $\psi(x)$ must be false if the first invalid relation known to x is chosen existentially, but true if it is chosen universally. Intuitively speaking, a relation is invalid from x's point of view if it does not correctly represent the part of the execution that affects x (in particular, x's local computations and the information x receives from its neighbors). Unfortunately, this also means that for relations chosen universally, x cannot simply rely on the fact that other nodes will detect invalidities; if x is affected by an invalidity within its neighborhood, then it must be "aware" of this itself to ensure that $\psi(x)$ holds true. Therefore, unlike in the proof of Theorem 11, it is not sufficient for $\psi(x)$ to only check x's local computations and message exchanges; it must also check those of the nodes that have a direct or indirect influence on x. This is very similar to the quantifier relativization algorithm described in the second part of the proof of Lemma 8.

We now give a formal definition of $\psi(x)$ using the helper formulas $LinearTupleOrder(x)$, $Correspondences(x)$, and so on, introduced in the proof of Theorem 11. The structure of $\psi(x)$ depends on the prefix of second-order quantifiers of φ_M. At the outermost level, $\psi(x)$ checks that the relations in $\bar{R}_{aux}$ are valid from the point of view of all nodes in the r-neighborhood of x, using the subformula

$$\psi_{aux}(x) = \forall^\circ y \stackrel{\leq r}{=\!=} x \Big(LinearTupleOrder(y) \wedge Correspondences(y) \wedge$$
$$UniqueIdentifier(y) \wedge NeighborOrder(y) \Big).$$

Since $\bar{R}_{aux}$ is quantified existentially, $\psi_{aux}(x)$ is used as a conjunct in $\psi(x)$, i.e.,

$$\psi(x) = \psi_{aux}(x) \wedge \psi_1(x).$$

The second conjunct $\psi_1(x)$ checks the validity of the remaining relations, starting with those in $\bar{R}_1$. Depending on whether a relation is quantified existentially or universally, its invalidity makes the remainder of the formula hold false or true, respectively. This leads to the inductive definition

$$\psi_j(x) = \begin{cases} \big(\forall^\circ y \stackrel{\leq r}{=\!=} x \, Certificate_j(y)\big) \to \psi_{j+1}(x) & \text{if } j \text{ is even,} \\ \big(\forall^\circ y \stackrel{\leq r}{=\!=} x \, Certificate_j(y)\big) \wedge \psi_{j+1}(x) & \text{if } j \text{ is odd,} \end{cases}$$

for $j \in [1:\ell]$, where the helper formula $Certificate_j(y)$ is a variant of $Certificate(y)$ that ensures that the relations K_0^j and K_1^j correctly represent a certificate of y consisting of 0's and 1's. For the last formula $\psi_\ell(x)$, the subformula $\psi_{\ell+1}(x)$ is defined below. It checks that $\bar{R}_{exe}$ is valid from x's point of view and that x eventually accepts:

$$\psi_{\ell+1}(x) = \begin{cases} \psi_{exe}(x) \to Accept(x) & \text{if } \ell \text{ is even,} \\ \psi_{exe}(x) \wedge Accept(x) & \text{if } \ell \text{ is odd,} \end{cases}$$

where the formula

$$\psi_{exe}(x) = \forall^\circ y \overset{\leq r}{=\!=\!=} x \Big(ExecGroundRules(y) \wedge OwnInput'(y) \Big) \wedge$$

$$\bigwedge_{i \in [1:r]} \forall^\circ y \overset{\leq r-i}{=\!=\!=} x \begin{pmatrix} ReceiveMessages_i(y) \wedge InitLocalConfig_i(y) \wedge \\ ComputeLocally_i(y) \wedge SendMessages_i(y) \end{pmatrix}$$

states that the part of the execution that has an influence on x is correctly represented by the appropriate relations. (The helper formula $OwnInput'(y)$ is a variant of $OwnInput(y)$ that takes into account the certificate assignments $\kappa_1, \ldots, \kappa_\ell$.) Since x is not influenced by the local computations that nodes at distance i make after round $r - i$, we only check those they make between rounds 1 and $r - i$.

To show that $\mathbf{\Pi}_\ell^{\mathrm{LP}} \subseteq \mathbf{\Pi}_\ell^{\mathrm{LFO}}|_{\mathrm{GRAPH}}$, the construction is almost the same, but the formula φ_M is now of the form

$$\forall \bar{R}_{aux} \forall \bar{R}_1 \exists \bar{R}_2 \ldots \mathbb{Q} \bar{R}_\ell \mathbb{Q} \bar{R}_{exe} \forall^\circ x \, \psi(x),$$

where $\mathbb{Q}$ is $\exists$ if ℓ is even and $\forall$ otherwise. Since the auxiliary relation variables in $\bar{R}_{aux}$ are now quantified universally, the subformula $\psi(x)$ becomes

$$\psi(x) = \psi_{aux}(x) \to \psi_1(x).$$

Similarly, the roles of even and odd indices are swapped in the definitions of $\psi_j(x)$, for $j \in [1 : \ell + 1]$, to account for the fact that the quantifier alternation now starts with a block of universal quantifiers. ◄

8 Hardness and completeness results

Reductions play a fundamental role in classical complexity theory, allowing us to compare the computational complexities of different problems, even when we are unable to determine any of them absolutely. Fraigniaud, Korman, and Peleg [14] transferred this idea to the setting of distributed decision, where they introduced the notion of *local reductions*. A local reduction from a property L to a property L' is performed by a distributed algorithm that modifies the labeling of the input graph in constant round time such that the original graph has property L if and only if the relabeled graph has property L'. While this unconstrained definition turned out to be too strong for the types of certificates considered, Balliu, D'Angelo, Fraigniaud, and Olivetti [3] later refined the concept and used it to establish completeness results for two classes of their identifier-independent hierarchy. However, they also noted that the local reductions involved were "very much time consuming at each node" and left as an open problem "whether non-trivial hardness results can be established under polynomial-time local reductions".

In this section, we extend Karp's [22] notion of polynomial-time reductions to computer networks. Our definition can also be seen as a further refinement of the aforementioned local reductions, where we impose bounds on the step running time of the algorithms performing the reductions and additionally require them to work under locally unique identifiers. As a presentation choice, we consider graph transformations that are slightly more general than just relabelings. These transformations may implicitly encode parts of the topology of the new graph in the output computed at each node of the original graph, thus allowing reductions to more natural graph properties. We then establish several hardness and completeness results for the two lowest levels of the local-polynomial hierarchy. We start with relatively

simple reductions showing that Eulerianness is **LP**-complete, while Hamiltonicity is both **LP**-hard and **coLP**-hard. Then, building on Theorem 11, we generalize the Cook–Levin theorem from **NP** to **NLP**, which gives us a first **NLP**-complete graph property. Finally, using standard techniques from complexity theory, we use this property to establish the **NLP**-completeness of 3-colorability.

In a way, the framework of reductions proves even more beneficial in the distributed setting than in the centralized setting. Indeed, since we will prove the infiniteness of the local-polynomial hierarchy in Section 9, all the hardness and completeness results presented below immediately yield *unconditional* lower bounds on the complexity of the properties in question, i.e., lower bounds that do not rely on any complexity-theoretic assumptions (see Corollaries 22, 25 and 26 in Section 9.1).

Clusters and implementable functions. Before we can define an appropriate notion of reduction, we need to generalize the idea of computable functions to our model of computation. Intuitively, the result $M(G, id)$ computed by a distributed Turing machine M on a graph G under an identifier assignment id can be interpreted as the encoding of a new graph G'. More precisely, for each node u of the original graph G, the output label computed by u is taken to encode a subgraph of G', which we refer to as u's cluster. Since G' may depend on the identifiers provided by id, and these are not considered part of the input, we shall regard M as implementing a "nondeterministic" function $f \colon \text{GRAPH} \to 2^{\text{GRAPH}}$ instead of a function $f \colon \text{GRAPH} \to \text{GRAPH}$.

Formally, a *cluster map* from a graph G' to a graph G is a function $g \colon V^{G'} \to V^G$ such that $\{u, v\} \in E^{G'}$ implies $g(u) = g(v)$ or $\{g(u), g(v)\} \in E^G$. With respect to g, the *cluster* of any node $u \in G$ is the induced subgraph of G' whose nodes are mapped to u, including the labels of those nodes.

Let $r_{\text{id}} \in \mathbb{N}_{>0}$. A function $f \colon \text{GRAPH} \to 2^{\text{GRAPH}}$ is *implemented* by a distributed Turing machine M under r_{id}-locally unique identifiers if for every $G \in \text{GRAPH}$ and every r_{id}-locally unique identifier assignment id of G, the result $M(G, id)$ computed by M on G under id represents a graph $G' \in f(G)$ in the following sense: there is a cluster map g from G' to G such that the label $\lambda^{M(G, id)}(u)$ computed at each node $u \in G$ encodes u's cluster with respect to g and all edges to u's neighbors' clusters (i.e., all edges between nodes of u's cluster and nodes of u's neighbors' clusters). Note that since the specific graph $G' \in f(G)$ computed by M may depend on id, its labels may refer to nodes of G by their identifiers.

A function $f \colon \text{GRAPH} \to 2^{\text{GRAPH}}$ is called *topology-preserving* if any two graphs $G \in \text{GRAPH}$ and $G' \in f(G)$ are identical except for their labeling, i.e., $V^G = V^{G'}$ and $E^G = E^{G'}$.

Reductions, hardness, and completeness. Basically, a local-polynomial reduction from a property L to a property L' is a graph transformation, implementable by a local-polynomial machine, that turns an input graph G into a new graph G' such that G has property L if and only if G' has property L'. The existence of such a reduction implies that L' is at least as hard as L, since it allows us to convert a hypothetical decider M' for L' into a decider M for L, which would work as follows: First, M would simulate a machine M_{red} that performs the reduction, thereby transforming G into G'. Then it would simulate M' on G', and finally each node of G would accept precisely if all nodes of its cluster did so in the simulation.

More formally, let $K, L, L' \subseteq \text{GRAPH}$. A *local-polynomial reduction* from L to L' on K is a function $f \colon \text{GRAPH} \to 2^{\text{GRAPH}}$ implementable by a local-polynomial machine such that for all graphs $G \in K$ and $G' \in f(G)$, we have $G \in L$ if and only if $G' \in L'$. If such a function exists, we denote this fact by $L \leq^{\text{LP}}_K L'$. Given a class **C** of graph properties, we

say that L' is **C**-*hard* on K under local-polynomial reductions if $L \leq_K^{\text{LP}} L'$ for all $L \in \mathbf{C}$, and we say that L' is **C**-*complete* on K under local-polynomial reductions if additionally $L' \in \mathbf{C}$. Since we will not consider other types of reductions, we usually omit mentioning "under local-polynomial reductions", and to avoid specifying K every time, we stipulate that "**C**-hard" and "**C**-complete" imply $K = \text{GRAPH}$ if $\mathbf{C} \in \{\boldsymbol{\Sigma}_\ell^{\text{LP}}, \boldsymbol{\Pi}_\ell^{\text{LP}}, \mathbf{co}\boldsymbol{\Sigma}_\ell^{\text{LP}}, \mathbf{co}\boldsymbol{\Pi}_\ell^{\text{LP}}\}_{\ell \in \mathbb{N}}$, and $K = \text{NODE}$ if $\mathbf{C} \in \{\boldsymbol{\Sigma}_\ell^{\text{P}}, \boldsymbol{\Pi}_\ell^{\text{P}}\}_{\ell \in \mathbb{N}}$.

▶ Remark 13. If a property L is **NLP**-hard, then it is also **NP**-hard, since NODE ⊆ GRAPH and $\mathbf{NP} = \mathbf{NLP}|_{\text{NODE}}$. Moreover, if L is **NLP**-complete under local-polynomial reductions that are also topology-preserving, then $L \cap \text{NODE}$ is **NP**-complete. This observation generalizes to all other levels of the local-polynomial and polynomial hierarchies.

In the remainder of this section, we establish a series of hardness and completeness results for the two lowest levels of the local-polynomial hierarchy. The first is trivial, but will be useful below. It concerns the property ALL-SELECTED introduced in Section 5.2.

▶ Remark 14. ALL-SELECTED is **LP**-complete. This holds even if we impose that all local-polynomial reductions must be topology-preserving.

Proof. The property obviously lies in **LP**, and its **LP**-hardness under topology-preserving reductions is established by the basic observation that any graph property L decided by a local-polynomial machine M can be reduced to ALL-SELECTED simply by executing M.   ◁
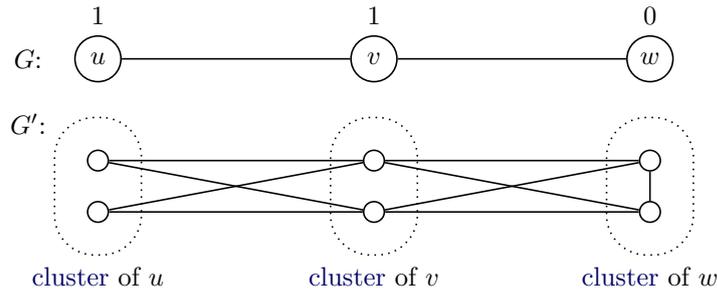
By considering reductions that do not necessarily preserve the topology of the input graph, we allow reductions to more natural graph properties. We now illustrate this using EULERIAN, the property of graphs that contain an *Eulerian cycle*, i.e., a cycle that uses each edge exactly once. The complement of this property will be denoted by NON-EULERIAN.

▶ **Proposition 15.** EULERIAN *is* **LP***-complete.*

Proof. By a famous theorem due to Euler (see, e.g., [6, Thm. 1.8.1]) a connected graph is Eulerian if and only if all its nodes have even degree. This characterization makes it straightforward to decide EULERIAN with a local-polynomial machine.

To show that the property is also **LP**-hard, we now describe a local-polynomial reduction to it from the property ALL-SELECTED, which is itself **LP**-complete by Remark 14. Given an arbitrary graph G, we construct a graph G' whose nodes all have even degree precisely if all nodes of G have label 1. Let us assume without loss of generality that G has at least two nodes. (Single-node graphs can easily be treated as a special case.) For each node $u \in V^G$, the new graph G' has two copies u_0 and u_1, and for each edge $\{u, v\} \in E^G$, it contains the four edges $\{u_i, v_j\}_{i,j \in \{0,1\}}$. In addition, for each node $u \in V^G$ whose label is not 1, the new graph also contains the edge $\{u_0, u_1\}$. An example is shown in Figure 7. Notice that $G \in \text{ALL-SELECTED}$ if and only if $G' \in \text{EULERIAN}$, that G' is always connected (as required by our definition of graphs), and that the nodes of G can compute G' in constant round time and polynomial step time.   ◀

While it is easy to determine if a given graph is Eulerian, it is much harder to determine if it is Hamiltonian. Indeed, a characterization of HAMILTONIAN similar to Euler's characterization of EULERIAN is neither known nor expected to exist (see, e.g., [6, Ch. 10]). The next two propositions show that this is reflected in the complexity of HAMILTONIAN in our model of computation: the property is both **LP**-hard and **coLP**-hard, which implies that it lies neither in **NLP** nor in **coNLP** (see Corollary 26).

Figure 7 Example illustrating the reduction from ALL-SELECTED to EULERIAN used in the proof of Proposition 15. The graph G has all node labels equal to 1 if and only if the graph G' has an Eulerian cycle. In this particular case, if node w of G had label 1, then its cluster in G' would lack the "vertical" edge, and thus all nodes of G' would have even degree, making G' Eulerian.

■ **Figure 8** *(repeated from Figure 2)* Example illustrating the reduction from ALL-SELECTED to HAMILTONIAN used in the proof of Proposition 16. The graph G has all node labels equal to 1 if and only if the graph G' has a Hamiltonian cycle. The thick edges in G form a spanning tree, which is replicated by the thick edges in G'. In this particular case, if node u_2 of G had label 1, then its cluster in G' would lack the "central" node, and thus the thick edges in G' would form a Hamiltonian cycle.

▶ **Proposition 16.** HAMILTONIAN *is* **LP**-*hard.*

Proof. By Remark 14, it suffices to provide a local-polynomial reduction from ALL-SELECTED to HAMILTONIAN. Given an arbitrary graph G, we construct a graph G' that has a Hamiltonian cycle if and only if all nodes of G have label 1. The main idea is that a Hamiltonian cycle in G' will represent a depth-first traversal of a spanning tree of G, using a method known as the Euler tour technique. For this purpose, each edge of G is represented by two edges in G', so that it can be traversed twice by a Hamiltonian cycle in G'. More precisely, each node $u \in G$ of degree d with neighbors $v_1, \ldots, v_d$ is represented in G' by a cycle of length $\max\{3, 2d\}$ of the form $u_{\to v_1}, u_{\leftarrow v_1}, \ldots, u_{\to v_d}, u_{\leftarrow v_d}, u_{\to v_1}$. For each neighbor v_i of u, this cycle contains two adjacent nodes $u_{\to v_i}$ and $u_{\leftarrow v_i}$, which can be thought of as the "ports" that allow us to "go to" and "come from" v_i, respectively. (To ensure that there are enough nodes to form a cycle, we add three dummy nodes if $d = 0$, and one dummy node if $d = 1$.) If u's label differs from 1, then G' additionally contains a node u_{bad} that is connected to exactly one node of the cycle representing u. Moreover, for each edge $\{u, v\}$ of G, the graph G'

contains the two edges $\{u_{\to v}, v_{\leftarrow u}\}$ and $\{u_{\leftarrow v}, v_{\to u}\}$. An example is shown in Figure 8.

Now, if all nodes of G are labeled with 1, then any spanning tree of G yields a Hamiltonian cycle of G'. This cycle includes all edges of G' of the form $\{u_{\leftarrow v}, u_{\to w}\}$ with $v \neq w$, and additionally, for each edge $\{u, v\}$ of G, either the edges $\{u_{\to v}, v_{\leftarrow u}\}$ and $\{u_{\leftarrow v}, v_{\to u}\}$ if $\{u, v\}$ belongs to the spanning tree, and otherwise the edges $\{u_{\to v}, u_{\leftarrow v}\}$ and $\{v_{\to u}, v_{\leftarrow u}\}$ (see Figure 8). However, if at least one node u of G has a label different from 1, then G' is not Hamiltonian, because the node u_{bad}, which has degree 1, cannot be part of any cycle. Hence, $G \in$ ALL-SELECTED if and only if $G' \in$ HAMILTONIAN.
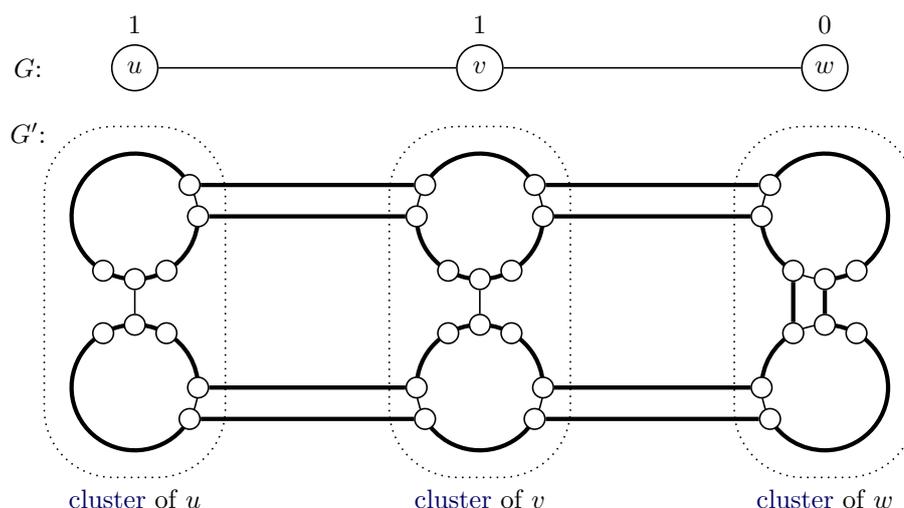
Notice that G' is guaranteed to be connected (as required by our definition of graphs), and that the nodes of G can compute G' in constant round time and polynomial step time. As an aside, note that we could reduce the number of nodes of G' by a factor of two by contracting edges of the form $\{u_{\leftarrow v}, u_{\to w}\}$ with $v \neq w$, but this would make the graphical representation somewhat less intuitive.   ◀

▶ **Proposition 17.** HAMILTONIAN *is* **coLP**-*hard.*

Proof. By Remark 14 and duality, NOT-ALL-SELECTED is **coLP**-complete, so it suffices to provide a local-polynomial reduction from NOT-ALL-SELECTED to HAMILTONIAN. In essence, given any graph G, we use the construction from the proof of Proposition 16 twice to create two subgraphs G'_{top} and G'_{bot}, and then connect them in such a way that the resulting graph G' has a Hamiltonian cycle if and only if G has at least one unselected node (i.e., a node whose label is not 1). By construction, G'_{top} and G'_{bot} each admit a Hamiltonian cycle, and the presence of an unselected node in G will ensure the presence of two edges in G' by which the two cycles can be connected to form a Hamiltonian cycle of G'. An example is provided in Figure 9.

More formally, each node $u \in G$ of degree d is represented in G' by two cycles of length $(2d + 3)$, which we will call the "top" and "bottom" cycles. As in the proof of Proposition 16, for each neighbor v of u, each of the two cycles contains two adjacent nodes, which can be thought of as the "ports" that allow us to "go to" and "come from" the corresponding cycle of v. That is, there are two adjacent nodes $u_{\to v}^{top}$ and $u_{\leftarrow v}^{top}$ in the "top" cycle of u, which are connected to the corresponding nodes $v_{\to u}^{top}$ and $v_{\leftarrow u}^{top}$ of the "top" cycle of v by means of the two edges $\{u_{\to v}^{top}, v_{\leftarrow u}^{top}\}$ and $\{u_{\leftarrow v}^{top}, v_{\to u}^{top}\}$. The "bottom" cycles are connected by analogous nodes and edges. Moreover, to ensure that G' is connected (as required by our definition of graphs), each cycle contains a sequence of three additional nodes, named $u_{\downarrow 1}^{top}, u_{\downarrow 2}^{top}, u_{\downarrow 3}^{top}$ in the "top" cycle, and $u_{\uparrow 1}^{bot}, u_{\uparrow 2}^{bot}, u_{\uparrow 3}^{bot}$ in the "bottom" cycle. For each node u of the original graph G, the graph G' contains at least the edge $\{u_{\downarrow 2}^{top}, u_{\uparrow 2}^{bot}\}$. In addition, if u has a label other than 1, then G' also contains the edge $\{u_{\downarrow 1}^{top}, u_{\uparrow 1}^{bot}\}$.

Clearly, the new graph G' can be computed in constant round time and polynomial step time by the nodes of G. To show the correctness of the construction, let G'_{top} and G'_{bot} be the induced subgraphs of G' that contain all "top" cycles and all "bottom" cycles, respectively. By the same argument as in the proof of Proposition 16, G'_{top} and G'_{bot} each have a Hamiltonian cycle, say H_{top} and H_{bot}. Now, if some node u of G has a label different from 1, then we can connect H_{top} and H_{bot} to form a Hamiltonian cycle of G'. This can be achieved by adding the edges $\{u_{\downarrow 1}^{top}, u_{\uparrow 1}^{bot}\}$ and $\{u_{\downarrow 2}^{top}, u_{\uparrow 2}^{bot}\}$, and removing the edges $\{u_{\downarrow 1}^{top}, u_{\downarrow 2}^{top}\}$ and $\{u_{\uparrow 1}^{bot}, u_{\uparrow 2}^{bot}\}$ (see Figure 9). However, if all nodes of G are labeled with 1, then G' does not have a Hamiltonian cycle. To see why, observe that in this case, all nodes of the form $u_{\downarrow 1}^{top}, u_{\downarrow 3}^{top}$, $u_{\uparrow 1}^{bot}$, or $u_{\uparrow 3}^{bot}$ have degree 2. This implies that all edges incident to these nodes must belong to any hypothetical Hamiltonian cycle H of G', and hence that none of the edges of the form $\{u_{\downarrow 2}^{top}, u_{\uparrow 2}^{bot}\}$ can be part of H. Since there are no other edges connecting G'_{top} and G'_{bot}, the cycle H cannot exist.   ◀

Figure 9 Example illustrating the reduction from NOT-ALL-SELECTED to HAMILTONIAN used in the proof of Proposition 17. The graph G has at least one node with a label different from 1 if and only if the graph G' has a Hamiltonian cycle. In this particular case, since node w of G has label 0, there is indeed a Hamiltonian cycle in G', represented by the thick edges. This Hamiltonian cycle can be thought of as the result of connecting the two "horizontally stretched" cycles (the "top" one and the "bottom" one), using the two "vertical" edges of w's cluster. If w had label 1, then one of these edges would be missing, making it impossible to connect the two cycles.

We now climb up one level in the local-polynomial hierarchy and investigate the notion of **NLP**-completeness. Our treatment of centralized computing as a special case of distributed computing is particularly helpful here, as it allows us to build directly on classical results from complexity theory. We begin by recalling the Cook–Levin theorem, which concerns the problem of determining whether a given Boolean formula is satisfiable.

▶ **Theorem 18** (Cook and Levin [5, 25])**.** SAT *is* **NP**-*complete.*

While this result was discovered a few years before Fagin's theorem, it can also be obtained as a simple corollary of the latter (see, e.g., [19, Thm. 3.2.6]). In the following, we will show that this observation extends to the distributed setting by using Theorem 11 to obtain a generalized version of the Cook–Levin theorem. But first, we need to generalize the Boolean satisfiability problem to graphs.

Boolean graph satisfiability. A *Boolean graph* is a graph G whose nodes are labeled with (encodings of) Boolean formulas. We call G *satisfiable* if there exists a function *val* that assigns to each node $u \in G$ a valuation of the Boolean variables occurring in u's formula $\lambda^G(u)$ such that

- $val(u)$ satisfies $\lambda^G(u)$, and
- $val(u)$ is consistent with the valuations of u's neighbors, i.e., $val(u)(P) = val(v)(P)$ for every neighbor v of u and every Boolean variable P that occurs in both $\lambda^G(u)$ and $\lambda^G(v)$.

We denote the set of all satisfiable Boolean graphs by SAT-GRAPH. The standard *Boolean satisfiability* problem SAT is simply the restriction of SAT-GRAPH to single-node graphs, i.e., SAT = SAT-GRAPH ∩ NODE.

Now we are ready to generalize the Cook–Levin theorem from **NP** to **NLP**. The original result (Theorem 18) can be recovered by restricting the following statement to single-node graphs, as noted in Remark 13.

▶ **Theorem 19.** SAT-GRAPH *is* **NLP**-*complete. This holds even if we impose that all local-polynomial reductions must be topology-preserving.*

Proof. Obviously, SAT-GRAPH lies in **NLP**, since each node can check in one communication round and polynomial step time whether a given valuation is both locally satisfying and consistent with the valuations of its neighbors.

It remains to show that SAT-GRAPH is **NLP**-hard, i.e., that $L \leq^{\mathrm{LP}}_{\mathrm{GRAPH}}$ SAT-GRAPH for every $L \in$ **NLP**, and to note that the involved local-polynomial reductions are topology-preserving. By Theorem 11, we know that L can be defined by a Σ_1^{LFO}-formula $\exists R_1 \ldots \exists R_n \, \forall x \, \varphi(x)$. For each graph $G \in$ GRAPH, we now construct a Boolean graph G' such that $G \in L$ if and only if $G' \in$ SAT-GRAPH. The nodes and edges of G' are the same as those of G, and the labeling function $\lambda^{G'}$ assigns to each node $u \in G$ a Boolean formula φ_u^G which states that, for a given interpretation of $R_1, \ldots, R_n$, the BF-formula φ is satisfied at the element of \$G representing u and at all the elements representing u's labeling bits. To represent the relations $R_1, \ldots, R_n$, we introduce Boolean variables of the form $P_{R(a_1,\ldots,a_k)}$, with the intended meaning that the tuple of elements $(a_1, \ldots, a_k) \in (D^{\$G})^k$ lies in the relation R.

Formally, we set

$$\varphi_u^G = \overbrace{\tau_{\{x \mapsto u\}}(\varphi)}^{\text{node } u} \wedge \overbrace{\bigwedge_{a \in (D^{\$G} \setminus V^G):\, a \rightleftharpoons^{\$G} u} \tau_{\{x \mapsto a\}}(\varphi)}^{u\text{'s labeling bits}},$$

where the translation function $\tau_\sigma(\psi)$ is defined inductively as follows for every BF-formula ψ and every variable assignment $\sigma \colon \mathrm{free}_{\mathrm{FO}}(\psi) \to D^{\$G}$:

- Atomic formulas that do not involve a second-order variable are replaced by their truth value in \$G, i.e.,

$$\tau_\sigma(\odot_1 y) = \begin{cases} \top & \text{if } \sigma(y) \in \odot_1^{\$G}, \\ \bot & \text{otherwise,} \end{cases} \qquad \tau_\sigma(y \rightharpoonup_i z) = \begin{cases} \top & \text{if } \sigma(y) \rightharpoonup_i^{\$G} \sigma(z), \\ \bot & \text{otherwise,} \end{cases}$$

$$\tau_\sigma(y \doteq z) = \begin{cases} \top & \text{if } \sigma(y) = \sigma(z), \\ \bot & \text{otherwise,} \end{cases}$$

  for $i \in \{1, 2\}$.
- Atomic formulas that involve a second-order variable are replaced by the corresponding Boolean variable, i.e., $\tau_\sigma(R(y_1, \ldots, y_k)) = P_{R(\sigma(y_1), \ldots, \sigma(y_k))}$.
- Boolean connectives are preserved: $\tau_\sigma(\neg\psi) = \neg\tau_\sigma(\psi)$, and $\tau_\sigma(\psi_1 \vee \psi_2) = \tau_\sigma(\psi_1) \vee \tau_\sigma(\psi_2)$.
- First-order quantification is expressed through a case distinction over all possible variable assignments, i.e.,

$$\tau_\sigma(\exists z \rightleftharpoons y \; \psi) = \bigvee_{a \in \$G:\, a \rightleftharpoons^{\$G} \sigma(y)} \tau_{\sigma[z \mapsto a]}(\psi).$$

Clearly, every interpretation of a k-ary relation variable R on \$G can be translated to a valuation of the Boolean variables $(P_{R(a_1,\ldots,a_k)})_{a_1,\ldots,a_k \in \$G}$, and vice versa. Hence, the structure \$G satisfies $\exists R_1 \ldots \exists R_n \, \forall x \, \varphi(x)$ if and only if there exists a (global) valuation of all Boolean variables occurring in G' that simultaneously satisfies all the formulas of G'.

The only issue is that a distributed Turing machine would require globally unique identifiers to compute G' from G, since the Boolean variables in G' allow to distinguish between all elements of \$G. So instead of G', we compute an equisatisfiable Boolean graph G''

for which locally unique identifiers suffice. Let r be the maximum nesting depth of bounded first-order quantifiers in φ (intuitively, the distance up to which φ can "see"). For a given $(r+1)$-locally unique identifier assignment id of G, we define G'' as the graph that one obtains from G' by rewriting all the Boolean formulas in such a way that each node u and each of its labeling bits are referred to using the identifier $id(u)$. Note that G'' can be computed from G in round time $(r+1)$ and polynomial step time. (The size of φ is constant with respect to G.) Furthermore, G'' is satisfiable if and only if G' is so, because elements that share the same identifier are never referred to by the same formula or by two formulas belonging to adjacent nodes. ◄
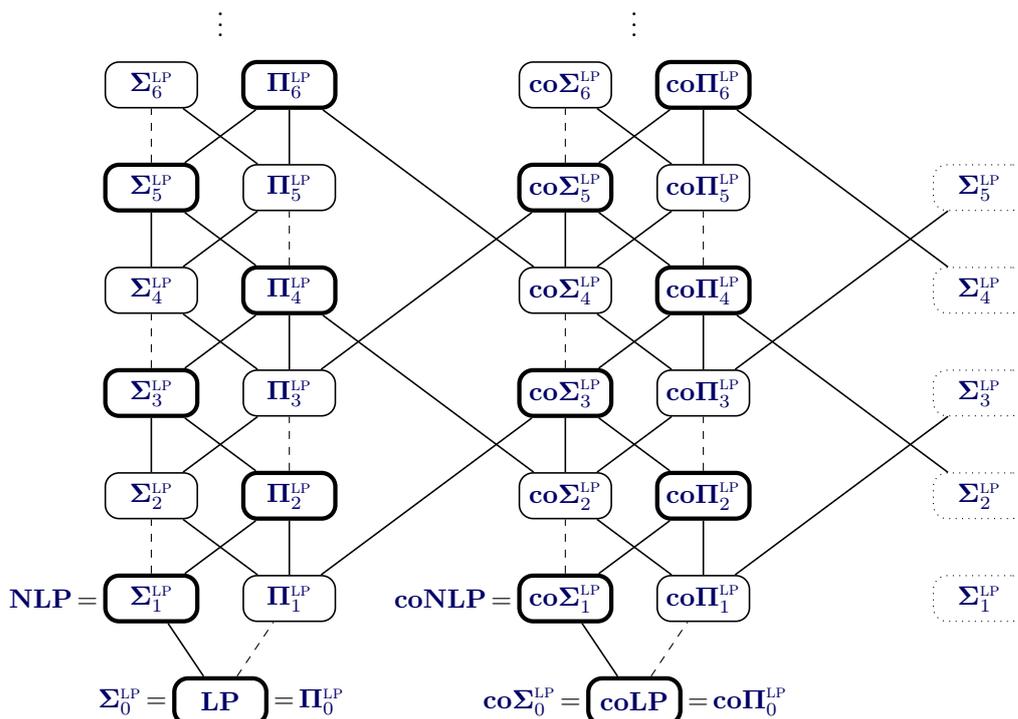
While SAT-GRAPH itself may not be a particularly natural graph property, we can use it as a basis for establishing the **NLP**-completeness of more natural properties, in the same way that Karp [22] used SAT to prove many other problems **NP**-complete. Again, this is made possible by the flexibility of local-polynomial reductions, which do not necessarily have to be topology-preserving. We now apply this approach to 3-colorability.

► **Theorem 20.** 3-COLORABLE *is* **NLP***-complete.*

Proof sketch. By Example 3 and Theorem 11, the property 3-COLORABLE lies in **NLP**. To show that it is also **NLP**-hard, we first reduce SAT-GRAPH to 3-SAT-GRAPH, the set of all satisfiable Boolean graphs in which every node is labeled with a 3-CNF formula, i.e., a Boolean formula in conjunctive normal form consisting of clauses with at most three literals. This is straightforward because the standard reduction from SAT to 3-SAT (i.e., to 3-SAT-GRAPH ∩ NODE) can be trivially generalized: Given a Boolean graph G and an identifier assignment id of G, we create an equisatisfiable graph G' by replacing the formula φ at each node $u \in G$ by an equisatisfiable 3-CNF formula φ' whose size is proportional to the size of φ. (This can be done using the Tseytin transformation.) The new formula φ' may contain additional Boolean variables that do not occur in φ, but every satisfying valuation of φ can be extended to a satisfying valuation of φ', and conversely, every satisfying valuation of φ' can be restricted to a satisfying valuation of φ. To ensure that the graphs G and G' are indeed equisatisfiable, we make the new variables' names depend on the identifier $id(u)$. Thus, the valuations of adjacent nodes in G' need to be consistent only for the original variables, not for the newly introduced ones.

Next, we reduce 3-SAT-GRAPH to 3-COLORABLE, again by generalizing to arbitrary graphs the corresponding reduction on single-node graphs. The standard reduction from 3-SAT to the string-encoded version of 3-COLORABLE converts any given 3-CNF-formula φ into a graph H_φ, which we will refer to as a formula gadget. This graph contains two special nodes called *false* and *ground*, two nodes P and $\bar{P}$ for each variable P occurring in φ (to represent the corresponding positive and negative literals), and a small gadget for each clause of φ. The edges are chosen in such a way that H_φ is 3-colorable if and only if φ is satisfiable. Moreover, if a 3-coloring exists, then we may assume without loss of generality that *false* and *ground* are colored 0 and 2, respectively, and each node representing a literal is colored by that literal's truth value (0 for false, and 1 for true). For details, see, e.g., [17, Prp. 2.27].

We now generalize this construction to arbitrary graphs. An example is provided in Figure 10. Given a Boolean graph G, we construct a graph G' that is 3-colorable if and only if G is satisfiable. For each node $u \in G$ labeled with a Boolean formula φ, the cluster representing u in G' contains a copy of the formula gadget H_φ. We denote this copy by H^u and mark its nodes with a superscript u. To enforce that the formula gadgets of adjacent clusters are colored consistently, we connect some of their nodes by means of an additional

gadget. More precisely, for $\{u, v\} \in E^G$, if we require that two nodes $w^u \in H^u$ and $w^v \in H^v$ have the same color, then we connect them using the following connector gadget:

Note that any valid 3-coloring of the connector gadget has to assign the same color to w^u and w^v. Using this, we connect $false^u$ to $false^v$, $ground^u$ to $ground^v$, and P^u to P^v for any Boolean variable P that occurs in the formulas of both u and v (see Figure 10). Thereby we ensure that each color has the same meaning in both formula gadgets and that variables shared by u and v are assigned the same truth value. Besides, this also ensures that the graph G' is connected (as required by our definition of graphs), even if some adjacent nodes of G do not share any variables.

Clearly, G' can be computed from G by a distributed Turing machine in two rounds and polynomial step time. (In the first round, the nodes send their label and identifier to their neighbors; in the second round, they perform only local computations.) The cluster map from G' to G can be chosen such that each cluster contains half of the nodes of each connector gadget attached to it, as shown in Figure 10. ◀

■ **Figure 10** *(repeated from Figure 3)* Example illustrating the reduction from 3-SAT-GRAPH to 3-COLORABLE used in the proof of Theorem 20. The Boolean graph G is satisfiable if and only if the graph G' is 3-colorable. The labels in the depiction of G' serve explanatory purposes only and are not part of the graph.

9 Infiniteness of the local-polynomial hierarchy

In this section, we show that the local-polynomial hierarchy does not collapse. Intuitively, this means that the more alternations we allow between Eve and Adam, the more graph properties we can express. More precisely, we prove that all inclusions represented by solid lines in Figure 11 are strict, and that classes represented on the same level are pairwise

■ **Figure 11** *(extended from Figure 1)* The local-polynomial hierarchy (left) and its complement hierarchy (right). The dotted classes on the far right are the same as those on the far left, repeated for the sake of readability. Only the lowest seven levels are shown, but the pattern extends infinitely. Each line (whether solid or dashed) indicates an inclusion of the lower class in the higher class. (This holds by definition for classes in the same hierarchy, and by Proposition 39 and duality for classes in separate hierarchies.) All inclusions represented by solid lines are proved to be strict, and classes located on the same level (regardless of which hierarchy) are proved to be pairwise distinct, even when restricted to graphs of bounded structural degree (by Proposition 21, Theorem 33, Corollaries 36, 38 and 40, and duality). The inclusions represented by dashed lines are in fact equalities on graphs of bounded structural degree (by Proposition 35), but this statement is unlikely to generalize to arbitrary graphs, where it holds if and only if $\mathbf{P} = \mathbf{coNP}$ (by Remark 37). This means that from a distributed computing perspective, the classes depicted with thick borders are the most meaningful.

distinct. For the remaining inclusions, represented by dashed lines, we expect the proof of strictness to be difficult, since they are equalities if and only if $\mathbf{P} = \mathbf{coNP}$.

The picture becomes simpler when restricted to graphs of bounded maximum degree and label length. On such graphs, the inclusions represented by dashed lines in Figure 11 are in fact equalities, so the local-polynomial hierarchy boils down to the classes depicted with thick borders, yielding a strict linear order of the following form:

$$\mathbf{\Pi}_0^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)} \subsetneq \mathbf{\Sigma}_1^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)} \subsetneq \mathbf{\Pi}_2^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)} \subsetneq \mathbf{\Sigma}_3^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)} \subsetneq \cdots$$

Here, $\mathrm{GRAPH}(\Delta)$ denotes the set of graphs of Δ-bounded structural degree, for some $\Delta \in \mathbb{N}$. Formally, the *structural degree* of a node u in a graph G is the number of elements that are connected to u in the graph's structural representation $\$G$, i.e., $\mathrm{card}(\{a \in \$G \mid a \rightleftharpoons^{\$G} u\})$. In other words, u's structural degree is the sum of its degree and its label length. We say that G is of Δ-*bounded structural degree* if the structural degree of every node $u \in G$ is at most Δ.

The remainder of this section is organized as follows. In Section 9.1, we separate **LP** from **NLP** and **coLP** by identifying simple graph properties that lie in one class but not the other. Then, in Section 9.2, we separate all higher levels of the local-polynomial hierarchy that end with an existential quantifier. Our proof uses an analogous result about monadic second-order logic on pictures, as well as an automaton model characterizing the existential fragment of that logic. Finally, the picture is completed in Section 9.3, where we establish the remaining separations and inclusions shown in Figure 11 and identify graph properties that lie outside the hierarchy.
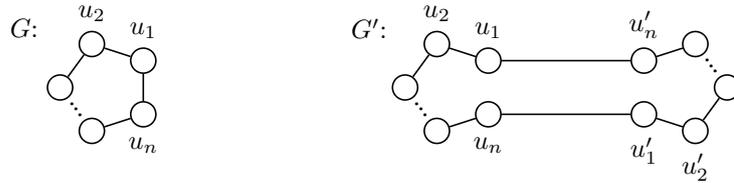
9.1 Warming up at ground level

The connection to logic established in Section 7 will be useful for separating the higher levels of the local-polynomial hierarchy and its complement hierarchy. But since this connection does not hold for the lowest level, we must use a different approach to separate **LP** from **NLP** and **coLP**. In this subsection, we do so using elementary arguments based on symmetry breaking and the pigeonhole principle. This also gives us the opportunity to gain a better intuition for the lower levels of the local-polynomial hierarchy by revisiting the graph properties presented in Section 5.2 and taking advantage of the hardness and completeness results established in Section 8.

▶ **Proposition 21.** *Some graph properties can be verified but not decided by a local-polynomial machine, even when restricted to graphs of bounded structural degree. More precisely,* $\mathbf{LP}|_{\text{GRAPH}(\Delta)} \subsetneq \mathbf{NLP}|_{\text{GRAPH}(\Delta)}$ *for all* $\Delta \geq 2$*, and a fortiori* $\mathbf{LP} \subsetneq \mathbf{NLP}$*.*

Proof. The graph property 2-COLORABLE clearly lies in **NLP** (simply adapt the formula from Example 3 and apply Theorem 11). In the following, we show that it does not lie in **LP**, even when restricted to graphs of 2-bounded structural degree.

Suppose, for the sake of contradiction, that 2-COLORABLE ∩ GRAPH(2) ∈ $\mathbf{LP}|_{\text{GRAPH}(2)}$. By Lemma 8, this means that there is a restrictive **LP**-decider M for 2-COLORABLE on GRAPH(2). Assume that M operates under r_{id}-locally unique identifiers, and consider an unlabeled cycle graph G with node set $\{u_1, \ldots, u_n\}$ such that n is odd and greater than $2r_{\text{id}}$. Since G is of odd length, it is not 2-colorable. We now construct a new cycle graph G' with node set $\{u_1, \ldots, u_n, u'_1, \ldots, u'_n\}$ by "gluing together" two copies of G as follows:

Since G' is of even length, it is 2-colorable. Given any r_{id}-locally unique identifier assignment id of G, let $id' \colon V^{G'} \to \{0,1\}^*$ be such that $id'(u_i) = id'(u'_i) = id(u_i)$ for $i \in [1 : n]$. As $n \geq 2r_{\text{id}}$, the function id' is an r_{id}-locally unique identifier assignment of G'. Moreover, the verdict of u_i and u'_i in $M(G', id')$ is the same as the verdict of u_i in $M(G, id)$, because the tape contents of these nodes are the same in every communication round and computation step. Hence, M accepts G if and only if it accepts G'. This contradicts our assumption that M is a restrictive **LP**-decider for 2-COLORABLE on GRAPH(2), since both G and G' are of 2-bounded structural degree, but only G' is 2-colorable. ◀

▶ **Corollary 22.** 3-COLORABLE *does not lie in* **LP**.

Proof. By Theorem 20, we know that 3-COLORABLE is **NLP**-hard, and by Proposition 21, we know that **NLP** $\not\subseteq$ **LP**. This implies that 3-COLORABLE cannot lie in **LP**, since otherwise we could show that **NLP** $\subseteq$ **LP**. To do so, it would suffice to transform a hypothetical **LP**-decider M for 3-COLORABLE into an **LP**-decider M' for an arbitrary property $L \in$ **NLP**. The nodes running M' would first apply a local-polynomial reduction from L to 3-COLORABLE. Then, each node would simulate M on its cluster and accept precisely if all nodes of the cluster do so. (See, e.g., Figure 10 on page 46 for an illustration of two clusters.) ◀

▶ **Proposition 23.** *The classes* **coLP** *and* **NLP** *are incomparable, even when restricted to graphs of bounded structural degree. More precisely,* **coLP**$|_{\text{GRAPH}(\Delta)} \not\subseteq$ **NLP**$|_{\text{GRAPH}(\Delta)}$ *for all* $\Delta \geq 3$*, and a fortiori* **coLP** $\not\subseteq$ **NLP**.

Proof. Since NOT-ALL-SELECTED lies in **coLP** (by Remark 14 it is even **coLP**-complete), it suffices to show that NOT-ALL-SELECTED $\cap$ GRAPH(3) $\notin$ **NLP**$|_{\text{GRAPH}(3)}$. Indeed, this statement implies that **coLP**$|_{\text{GRAPH}(3)} \not\subseteq$ **NLP**$|_{\text{GRAPH}(3)}$, and by duality that **LP**$|_{\text{GRAPH}(3)} \not\subseteq$ **coNLP**$|_{\text{GRAPH}(3)}$, and hence also that **NLP**$|_{\text{GRAPH}(3)} \not\subseteq$ **coLP**$|_{\text{GRAPH}(3)}$.

Suppose then, for the sake of contradiction, that NOT-ALL-SELECTED $\cap$ GRAPH(3) $\in$ **NLP**$|_{\text{GRAPH}(3)}$. By Lemma 8, this means that there exists a restrictive **NLP**-verifier M for NOT-ALL-SELECTED on GRAPH(3). Assume that M operates under r_{id}-locally unique identifiers and (r_1, p)-bounded certificates and runs in round time r_2, and let $r = \max\{2r_{\text{id}}, r_1, r_2\}$. In the following, we restrict our attention to cycle graphs whose nodes are all labeled with a single bit and whose length is a multiple of $(r + 1)$. Notice that on such graphs, we can construct an r_{id}-locally unique identifier assignment by cyclically assigning each node an identifier corresponding to a number in $[0:r]$, encoded as a binary string. If we do so, then the length of an (r_1, p)-bounded certificate is at most $m = p\big((2r+1)\cdot(\lceil \log_2 r \rceil + 3)\big)$. Thus, for our choice of graphs and identifier assignments, there are no more than $n = (r+1)\cdot(2^{m+2} - 2)^{2r+1}$ possible ways to assign labels, identifiers, and certificates to the r_2-neighborhood of any node.

First, let G be a cycle graph whose length is greater than n and a multiple of $(r + 1)$ such that exactly one node u has label 0 (the unselected node) and all others have label 1. Let id be a cyclic r_{id}-locally unique identifier assignment of G as described above. Since $G \in$ NOT-ALL-SELECTED $\cap$ GRAPH(3), there exists an (r_1, p)-bounded certificate assignment κ of (G, id) such that $M(G, id, \kappa) \equiv$ ACCEPT. By the pigeonhole principle, there must be two distinct nodes v and v' in G whose r_2-neighborhoods are indistinguishable because the labels, identifiers, and certificates therein are all the same. Now consider the cycle graph G' obtained from G by taking the path between v and v' that does not contain the 0-labeled node u, and identifying v with v'. Let id' and κ' be the restrictions of id and κ to G'. Notice that id' is r_{id}-locally unique, κ' is (r_1, p)-bounded, and $M(G', id', \kappa') \equiv$ ACCEPT because every node of G' reaches the same verdict in $M(G', id', \kappa')$ as in $M(G, id, \kappa)$. We conclude that $G' \in$ NOT-ALL-SELECTED $\cap$ GRAPH(3), which contradicts the fact that all nodes of G' are labeled with 1. ◀

▶ **Corollary 24.** *The class of graph properties decidable by a local-polynomial machine is not closed under complementation, even when restricted to graphs of bounded structural degree. More precisely,* **LP**$|_{\text{GRAPH}(\Delta)} \neq$ **coLP**$|_{\text{GRAPH}(\Delta)}$ *for all* $\Delta \geq 3$*, and a fortiori* **LP** $\neq$ **coLP**.

Proof. This follows immediately from Proposition 23. ◀

▶ **Corollary 25.** NON-3-COLORABLE *does not lie in* **NLP**.

Proof. By Theorem 20, we know that 3-COLORABLE is **NLP**-hard, which by duality means that NON-3-COLORABLE is **coNLP**-hard. This implies that NON-3-COLORABLE cannot lie in **NLP**, since otherwise we could show that **coNLP** ⊆ **NLP**, contradicting Proposition 23. The argument is analogous to the proof of Corollary 22, with the additional observation that a node can simulate an **NLP**-verifier on its cluster by interpreting its own certificate as an encoding of the certificates of all nodes of the cluster. ◀

▶ **Corollary 26.** *None of the following graph properties lies in* **NLP**: NON-EULERIAN, HAMILTONIAN, NON-HAMILTONIAN.

Proof. By the duals of Propositions 15 and 16, and by Proposition 17, all these graph properties are **coLP**-hard. Using the same argument as in the proofs of Corollaries 22 and 25, we conclude from Proposition 23 that none of them lies in **NLP**. ◀

9.2   Climbing up the hierarchy

We now come to the more technical part of our separation proof, which uses the connection to logic provided by Theorem 12 to leverage two results about monadic second-order logic on pictures (stated in Theorems 27 and 29).

Monadic second-order logic is the fragment of second-order logic that can only quantify over sets instead of arbitrary relations. That is, for second-order quantifications of the form $\exists R\,\varphi$, the second-order variable R must necessarily be of arity 1. We analogously define *local monadic second-order logic* as the corresponding fragment of local second-order logic.
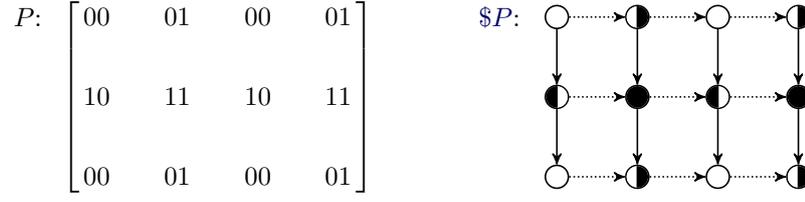
Accordingly, the *monadic* versions of the second-order and local second-order hierarchies are obtained by restricting second-order quantification in each class to unary relations. To denote the monadic classes, we add the letter "m" as a prefix to the corresponding non-monadic classes. This gives us the classes of formulas $\mathrm{m}\Sigma_\ell^{\mathrm{FO}}$, $\mathrm{m}\Pi_\ell^{\mathrm{FO}}$, $\mathrm{m}\Sigma_\ell^{\mathrm{LFO}}$, $\mathrm{m}\Pi_\ell^{\mathrm{LFO}}$, for $\ell \in \mathbb{N}$, and the classes of structure properties $\mathbf{m\Sigma}_\ell^{\mathrm{FO}}$, $\mathbf{m\Pi}_\ell^{\mathrm{FO}}$, $\mathbf{m\Sigma}_\ell^{\mathrm{LFO}}$, $\mathbf{m\Pi}_\ell^{\mathrm{LFO}}$ that can be defined by formulas of the respective classes. We will refer to $\mathrm{m}\Sigma_1^{\mathrm{FO}}$ and $\mathrm{m}\Sigma_1^{\mathrm{LFO}}$ as the the *existential fragments* of monadic second-order logic and local monadic second-order logic.

9.2.1   A digression on pictures

In the context of this paper, pictures are matrices of fixed-length binary strings. More precisely, for any $t \in \mathbb{N}$ and $m, n \in \mathbb{N}_{>0}$, a t-bit *picture* P of *size* (m, n) is an $(m \times n)$-matrix whose entries are bit strings of length t. We refer to the positions $(i, j) \in \langle m] \times \langle n]$ as P's *pixels*. Alternatively, a picture can also be viewed as a function $P \colon \langle m] \times \langle n] \to \{0, 1\}^t$. All the usual terminology of matrices applies; for instance, the pixel $(1, 1)$ is referred to as the top-left corner. The class of all t-bit pictures (of arbitrary size) is denoted by PIC(t), and any subset $L \subseteq$ PIC(t) is called a *picture property*.

As with graphs, we can evaluate logical formulas on pictures by identifying each t-bit picture P with a structure. Formally, the *structural representation* of P is the structure $\$P = (D^{\$P}, \odot_1^{\$P}, \ldots, \odot_t^{\$P}, \rightharpoonup_1^{\$P}, \rightharpoonup_2^{\$P})$ of signature $(t, 2)$ with domain $D^{\$P} = \langle m] \times \langle n]$, unary relations $\odot_1^{\$P}, \ldots, \odot_t^{\$P} \subseteq D^{\$P}$ such that $(i, j) \in \odot_k^{\$P}$ precisely when the k-th bit of $P(i, j)$ is 1, and "vertical" and "horizontal" successor relations $\rightharpoonup_1^{\$P}, \rightharpoonup_2^{\$P} \subseteq (D^{\$P} \times D^{\$P})$ such that $(i, j) \rightharpoonup_1^{\$P} (i+1, j)$ and $(i, j) \rightharpoonup_2^{\$P} (i, j+1)$ for all suitable i, j. An example is provided in Figure 12.

Monadic second-order logic on pictures has been fairly well-understood since the early 2000s. In particular, Matz, Schweikardt, and Thomas [29, Thm. 1] have shown that the

$$P: \begin{bmatrix} 00 & 01 & 00 & 01 \\ \\ 10 & 11 & 10 & 11 \\ \\ 00 & 01 & 00 & 01 \end{bmatrix} \qquad \$P:$$

■ **Figure 12** *(repeated from Figure 5)* A 2-bit picture P of size $(3,4)$ and its structural representation $\$P$. The sets $\odot_1^{\$P}$ and $\odot_2^{\$P}$ are represented by elements whose left and right halves, respectively, are colored black, and the relations $\rightharpoonup_1^{\$P}$ and $\rightharpoonup_2^{\$P}$ are represented by solid and dotted arrows, respectively.

monadic second-order hierarchy is infinite on several kinds of structures, including pictures. Here, we state only the part of their result we need, in a stronger form obtained by Matz [28, Thm. 2.26].

▶ **Theorem 27** (Matz, Schweikardt, Thomas [28, 29])**.** *The monadic second-order hierarchy on pictures is infinite. More precisely,* $\mathbf{m\Sigma}_\ell^{\mathrm{FO}}|_{\mathrm{PIC}(t)}$ *and* $\mathbf{m\Pi}_\ell^{\mathrm{FO}}|_{\mathrm{PIC}(t)}$ *are incomparable, which implies that* $\mathbf{m\Sigma}_\ell^{\mathrm{FO}}|_{\mathrm{PIC}(t)} \subsetneqq \mathbf{m\Pi}_{\ell+1}^{\mathrm{FO}}|_{\mathrm{PIC}(t)}$ *and* $\mathbf{m\Pi}_\ell^{\mathrm{FO}}|_{\mathrm{PIC}(t)} \subsetneqq \mathbf{m\Sigma}_{\ell+1}^{\mathrm{FO}}|_{\mathrm{PIC}(t)}$, *for all* $\ell \in \mathbb{N}_{>0}$ *and* $t \in \mathbb{N}$.

The ultimate goal of Section 9.2 is to transfer part of Theorem 27 from monadic second-order logic on pictures to local second-order logic on graphs (and thus to the local-polynomial hierarchy). This will culminate in Theorem 33 on page 59. As a first milestone towards this goal, we establish a partial levelwise equivalence between the two logics when restricted to pictures. We do this in two steps, showing roughly that the expressive power of local second-order logic on pictures remains unaffected if we first weaken second-order quantification and then strengthen first-order quantification. The outcome is presented in Theorem 31.

We begin by reducing unrestricted second-order quantification to quantification over sets, exploiting the fact that local second-order logic on pictures allows us to represent arbitrary relations as collections of sets.

▶ **Proposition 28.** *When restricted to pictures, each level of the local second-order hierarchy is equivalent to the corresponding level of the local monadic second-order hierarchy. More precisely,* $\mathbf{\Sigma}_\ell^{\mathrm{LFO}}|_{\mathrm{PIC}(t)} = \mathbf{m\Sigma}_\ell^{\mathrm{LFO}}|_{\mathrm{PIC}(t)}$ *and* $\mathbf{\Pi}_\ell^{\mathrm{LFO}}|_{\mathrm{PIC}(t)} = \mathbf{m\Pi}_\ell^{\mathrm{LFO}}|_{\mathrm{PIC}(t)}$ *for all* $\ell, t \in \mathbb{N}$.

Proof. We only have to show inclusions from left to right, since $\mathrm{m\Sigma}_\ell^{\mathrm{LFO}}$ is a syntactic fragment of $\Sigma_\ell^{\mathrm{LFO}}$, and $\mathrm{m\Pi}_\ell^{\mathrm{LFO}}$ is a syntactic fragment of Π_ℓ^{LFO}. Consider any formula $\varphi \in \Sigma_\ell^{\mathrm{LFO}} \cup \Pi_\ell^{\mathrm{LFO}}$, and let $r \in \mathbb{N}$ be the maximum nesting depth of bounded first-order quantifiers in φ. Intuitively, r is the Manhattan distance up to which each pixel can "see" when evaluating φ on (the structural representation of) a picture. Given a picture P of size (m, n), a Manhattan distance $d \in \mathbb{N}$, and a pixel (i, j) of P, we denote by $N_d^P(i, j)$ the pixel's von Neumann neighborhood of range d, i.e., $N_d^P(i, j) = \{(i', j') \in \langle m] \times \langle n] \mid \mathrm{abs}(i' - i) + \mathrm{abs}(j' - j) \le d\}$. We now show how to use multiple sets of pixels to encode a single relation of arbitrary arity, taking advantage of the fact that the number of pixels in $N_{2r}^P(i, j)$ is bounded and that each pixel in this set can be addressed by relative coordinates with respect to (i, j). More precisely, in the picture's structural representation $\$P$, each element a can address the elements of $N_{2r}^P(a)$ using relative coordinate pairs from the finite set $C = \{(\Delta_i, \Delta_j) \in \mathbb{Z}^2 \mid \mathrm{abs}(i) + \mathrm{abs}(j) \le 2r\}$. (The first coordinate gives the "vertical" distance and the second coordinate the "horizontal" distance with respect to a.) Based on that, if φ contains a second-order variable R of arity $k \ge 2$, we proceed as follows to represent R by the collection of unary variables

$(X_{R(*,c_2,\ldots,c_k)})_{c_2,\ldots,c_k \in C}$: for any elements $a_1, \ldots, a_k \in \$P$ such that $a_2, \ldots, a_k$ belong to $N_{2r}^P(a_1)$ and are located at positions $c_2, \ldots, c_k \in C$ with respect to a_1, we stipulate that a_1 lies in the set $X_{R(*,c_2,\ldots,c_k)}$ if and only if the tuple $(a_1, \ldots, a_k)$ lies in the relation R. Notice that we do not encode the entire relation R, but only its restriction to tuples whose elements lie within a Manhattan distance of at most $2r$ from the first element a_1. This is sufficient for our purposes, since φ cannot refer to any other tuples. (Intuitively, when evaluating φ, each element a can only "see" tuples whose elements are all in $N_r^P(a)$, so the elements cannot be further than $2r$ from each other, and in particular from the first element.)

Formally, given r, we define a translation τ_r from local second-order logic to local monadic second-order logic, proceeding by structural induction:

- Atomic formulas that do not involve a relation variable of higher arity are kept unchanged, i.e., $\tau_r(\psi) = \psi$ if ψ is of the form $\odot_i x$, $x \rightharpoonup_i y$, $x \doteq y$, or $R(x)$, where R is of arity 1.
- Atomic formulas involving a relation variable R of arity $k \geq 2$ are rewritten in terms of the corresponding unary variables:

$$\tau_r\big(R(x_1,\ldots,x_k)\big) = \bigvee_{c_2,\ldots,c_k \in C}\Big(X_{R(*,c_2,\ldots,c_k)}(x_1) \wedge Pos_{c_2}(x_1,x_2) \wedge \ldots \wedge Pos_{c_k}(x_1,x_k)\Big),$$

where for $(\Delta_i, \Delta_j) \in C$, the formula $Pos_{(\Delta_i,\Delta_j)}(x,y)$ states that y is located at position (Δ_i, Δ_j) with respect to x. It can be defined inductively as follows:

$$Pos_{(\Delta_i,\Delta_j)}(x,y) = \begin{cases} \exists z \rightleftharpoons x\big(z \rightharpoonup_1 x \wedge Pos_{(\Delta_i+1,\Delta_j)}(z,y)\big) & \text{if } \Delta_i < 0, \\ \exists z \rightleftharpoons x\big(x \rightharpoonup_1 z \wedge Pos_{(\Delta_i-1,\Delta_j)}(z,y)\big) & \text{if } \Delta_i > 0, \\ \exists z \rightleftharpoons x\big(z \rightharpoonup_2 x \wedge Pos_{(\Delta_i,\Delta_j+1)}(z,y)\big) & \text{if } \Delta_i = 0 \text{ and } \Delta_j < 0, \\ \exists z \rightleftharpoons x\big(x \rightharpoonup_2 z \wedge Pos_{(\Delta_i,\Delta_j-1)}(z,y)\big) & \text{if } \Delta_i = 0 \text{ and } \Delta_j > 0, \\ x \doteq y & \text{if } \Delta_i = \Delta_j = 0. \end{cases}$$

- Boolean connectives, first-order quantifiers, and second-order quantifiers over unary relations are preserved: $\tau_r(\neg\psi) = \neg\tau_r(\psi)$, $\tau_r(\psi_1 \vee \psi_2) = \tau_r(\psi_1) \vee \tau_r(\psi_2)$, $\tau_r(\exists x\, \psi) = \exists x\, \tau_r(\psi)$, $\tau_r(\exists x \rightleftharpoons y\, \psi) = \exists x \rightleftharpoons y\, \tau_r(\psi)$, and $\tau_r(\exists R\, \psi) = \exists R\, \tau_r(\psi)$ if R is of arity 1.
- Each second-order quantifier binding a relation variable R of arity $k \geq 2$ is replaced by a collection of second-order quantifiers binding the corresponding unary variables: $\tau_r(\exists R\, \psi) = \exists \big(X_{R(*,c_2,\ldots,c_k)}\big)_{c_2,\ldots,c_k \in C}\, \tau_r(\psi)$.

Applying this translation to the initial formula φ, we obtain a formula $\tau_r(\varphi)$ that is equivalent to φ on pictures. Notice that τ_r preserves the alternation level of second-order quantifiers and that each subformula $Pos_{(\Delta_i,\Delta_j)}(x,y)$ lies in BF. Hence, if $\varphi \in \Sigma_\ell^{\text{LFO}}$, then $\tau_r(\varphi) \in \text{m}\Sigma_\ell^{\text{LFO}}$, and if $\varphi \in \Pi_\ell^{\text{LFO}}$, then $\tau_r(\varphi) \in \text{m}\Pi_\ell^{\text{LFO}}$.   ◀

Next, we want to strengthen first-order quantification. To do this, we take advantage of an automaton model for pictures introduced by Giammarresi and Restivo [15], which is closely related to monadic second-order logic. A "machine" in this model, called a t-bit *tiling system*, is defined as a tuple $T = (Q, \Theta)$, where t is a nonnegative integer, Q is a finite set of *states*, and $\Theta \subseteq \big((\{0,1\}^t \times Q) \cup \{\#\}\big)^4$ is a set of (2×2)-*tiles*. Each tile in Θ consists of entries that are either a t-bit string accompanied by a state in Q, or the special boundary symbol $\#$ (assumed not to be contained in $\{0,1\}^t \times Q$).

A t-bit tiling system $T = (Q, \Theta)$ operates similarly to a nondeterministic finite automaton generalized to two dimensions: given a picture P, it first nondeterministically assigns a state of Q to each pixel of P, and then checks that this assignment of states respects the

"transitions" that are allowed by Θ. More precisely, a t-bit picture P of size (m, n) is *accepted* by T if there exists an assignment $\langle m] \times \langle n] \to Q$ such that each (2×2)-subblock of P matches some tile of Θ, assuming that the entire picture is surrounded by a frame consisting of $\#$'s (to detect the borders), and that a pixel matches $(s, q) \in \{0, 1\}^t \times Q$ precisely if its value is s and its assigned state is q. The picture property *recognized* by T consists of those t-bit pictures that are accepted by T. We write **TS** for the class of picture properties that are recognized by some tiling system.

Exploiting a locality property of first-order logic, Giammarresi, Restivo, Seibert, and Thomas [16, Thm. 3.1] have shown that tiling systems capture precisely the nondeterministic level of the monadic second-order hierarchy on pictures:

▶ **Theorem 29** (Giammarresi, Restivo, Seibert, Thomas [16]). *Tiling systems are equivalent to the existential fragment of monadic second-order logic on pictures. That is,* $\mathbf{TS}|_{\mathrm{PIC}(t)} = \mathbf{m\Sigma}_1^{\mathrm{FO}}|_{\mathrm{PIC}(t)}$ *for all* $t \in \mathbb{N}$.

This result gives us the key to move from bounded to arbitrary first-order quantification. The following corollary is based on the observation that tiling systems can be easily described in local monadic second-order logic.

▶ **Corollary 30.** *When restricted to pictures, the existential fragment of local monadic second-order logic is equivalent to the existential fragment of monadic second-order logic. That is,* $\mathbf{m\Sigma}_1^{\mathrm{LFO}}|_{\mathrm{PIC}(t)} = \mathbf{m\Sigma}_1^{\mathrm{FO}}|_{\mathrm{PIC}(t)}$ *for all* $t \in \mathbb{N}$.

Proof. Since $\mathrm{m\Sigma}_1^{\mathrm{LFO}}$ can be seen as a syntactic fragment of $\mathrm{m\Sigma}_1^{\mathrm{FO}}$, it suffices to show that $\mathbf{TS}|_{\mathrm{PIC}(t)} \subseteq \mathbf{m\Sigma}_1^{\mathrm{LFO}}|_{\mathrm{PIC}(t)}$ and then apply Theorem 29. We thus require a translation τ from tiling systems to $\mathrm{m\Sigma}_1^{\mathrm{LFO}}$-sentences on pictures such that a picture P is accepted by a tiling system T if and only if its structural representation $\$P$ satisfies $\tau(T)$.

By inspecting the proof of Theorem 29 in [16, Thm. 3.1], it is easy to see that the $\mathrm{m\Sigma}_1^{\mathrm{FO}}$-sentence provided there can be rewritten as an equivalent $\mathrm{m\Sigma}_1^{\mathrm{LFO}}$-sentence, essentially by replacing unbounded first-order quantifiers with their bounded counterparts. We therefore only give a high-level description of the construction. For $T = (Q, \Theta)$, the formula $\tau(T)$ is of the form

$$\exists (X_q)_{q \in Q} \, \forall x \big(OneState(x) \wedge LegalTiling(x) \big),$$

where each X_q is a unary relation variable intended to represent the set of pixels in state q, $OneState(x)$ is a BF-formula stating that exactly one state has been assigned to pixel x, and $LegalTiling(x)$ is another BF-formula stating that each of the (2×2)-subblocks containing pixel x corresponds to some tile of Θ. Since the boundary markers $\#$ surrounding the picture are not represented by any elements in the structure $\$P$, the formula $LegalTiling(x)$ performs a case distinction on whether x lies in one of the four corners, along one of the four borders, or somewhere else inside the picture. This can be written as a conjunction of the form

$$LegalTiling(x) = TopLeft(x) \wedge TopRight(x) \wedge BottomLeft(x) \wedge BottomRight(x) \wedge$$
$$Top(x) \wedge Bottom(x) \wedge Left(x) \wedge Right(x) \wedge Inside(x),$$

where, for example, the conjunct $TopLeft(x)$ states that if x lies in the top-left corner (i.e., if it has neither a "vertical" nor a "horizontal" predecessor), then there must be some tile $\left(\begin{smallmatrix} \# & \# \\ \# & s,q \end{smallmatrix} \right) \in \Theta$ such that x has value s and lies in state q. The other conjuncts are similar. ◀

By combining Proposition 28 (which weakens second-order quantification) and Corollary 30 (which strengthens first-order quantification), we can now derive a partial levelwise equivalence between local second-order logic and monadic second-order logic on pictures.

▶ **Theorem 31.** *When restricted to pictures, every level of the local second-order hierarchy that ends with a block of existential quantifiers is equivalent to the corresponding level of the monadic second-order hierarchy. That is, $\mathbf{\Sigma}_\ell^{\mathrm{LFO}}|_{\mathrm{PIC}(t)} = \mathbf{m\Sigma}_\ell^{\mathrm{FO}}|_{\mathrm{PIC}(t)}$ if ℓ is odd, and $\mathbf{\Pi}_\ell^{\mathrm{LFO}}|_{\mathrm{PIC}(t)} = \mathbf{m\Pi}_\ell^{\mathrm{FO}}|_{\mathrm{PIC}(t)}$ if ℓ is even, for all $\ell \in \mathbb{N}_{>0}$ and $t \in \mathbb{N}$.*

Proof. We proceed by induction on ℓ. For $\ell = 1$, it suffices to apply Proposition 28 and then Corollary 30, i.e.,

$$\mathbf{\Sigma}_1^{\mathrm{LFO}}|_{\mathrm{PIC}(t)} \overset{\mathrm{Prp.28}}{=} \mathbf{m\Sigma}_1^{\mathrm{LFO}}|_{\mathrm{PIC}(t)} \overset{\mathrm{Cor.30}}{=} \mathbf{m\Sigma}_1^{\mathrm{FO}}|_{\mathrm{PIC}(t)}.$$

For $\ell \geq 2$, let us assume that ℓ is even, the other case being completely analogous. We have

$$\mathbf{\Pi}_\ell^{\mathrm{LFO}}|_{\mathrm{PIC}(t)} \overset{\mathrm{Prp.28}}{=} \mathbf{m\Pi}_\ell^{\mathrm{LFO}}|_{\mathrm{PIC}(t)} = \mathbf{m\Pi}_\ell^{\mathrm{FO}}|_{\mathrm{PIC}(t)},$$

by first applying Proposition 28 and then using the fact that $\mathrm{m\Pi}_\ell^{\mathrm{LFO}}$ and $\mathrm{m\Pi}_\ell^{\mathrm{FO}}$ are defined in terms of $\mathrm{m\Sigma}_{\ell-1}^{\mathrm{LFO}}$ and $\mathrm{m\Sigma}_{\ell-1}^{\mathrm{FO}}$, for which the induction hypothesis already provides an equivalence. More precisely, if $\forall X_1 \ldots \forall X_n(\varphi)$ is an $\mathrm{m\Pi}_\ell^{\mathrm{LFO}}$-sentence evaluated on t-bit pictures, where φ starts with a block of existential quantifiers over sets, then φ can be evaluated as an $\mathrm{m\Sigma}_{\ell-1}^{\mathrm{LFO}}$-sentence on $(t + n)$-bit pictures. The idea is simply to interpret each atomic formula $X_i(x)$ as $\odot_{t+i}\, x$, for $i \in [1:n]$. The analogous observation holds for $\mathrm{m\Pi}_\ell^{\mathrm{FO}}$-sentences on t-bit pictures, whose subformulas can be interpreted as $\mathrm{m\Sigma}_{\ell-1}^{\mathrm{FO}}$-sentences on $(t + n)$-bit pictures. Since we already know that $\mathbf{m\Sigma}_{\ell-1}^{\mathrm{LFO}}|_{\mathrm{PIC}(t+n)} = \mathbf{m\Sigma}_{\ell-1}^{\mathrm{FO}}|_{\mathrm{PIC}(t+n)}$, this implies that $\mathbf{m\Pi}_\ell^{\mathrm{LFO}}|_{\mathrm{PIC}(t)} = \mathbf{m\Pi}_\ell^{\mathrm{FO}}|_{\mathrm{PIC}(t)}$. ◀

9.2.2   From pictures to graphs

With the partial levelwise equivalence obtained in Theorem 31, we can already transfer part of the separation result of Matz, Schweikardt, and Thomas (Theorem 27) from monadic second-order logic to local second-order logic, while remaining in the realm of pictures. To further transfer the result from pictures to graphs, we now show how to encode 0-bit pictures as graphs of 4-bounded structural degree, and how to translate formulas from one type of structure to the other.

The *graph encoding* of the 0-bit picture P of size (m, n) is a graph G_P that represents each pixel of P by five nodes: one main node (pxl), and four auxiliary nodes (in_1, in_2, out_1, out_2), which can be thought of as the incoming and outgoing "ports" of the pixel. Each main node is connected to its four ports, and the ports of any adjacent pixels are connected in such a way as to represent the relations $\rightarrow_1^{\$P}$ and $\rightarrow_2^{\$P}$. Formally, G_P is defined by the set of nodes
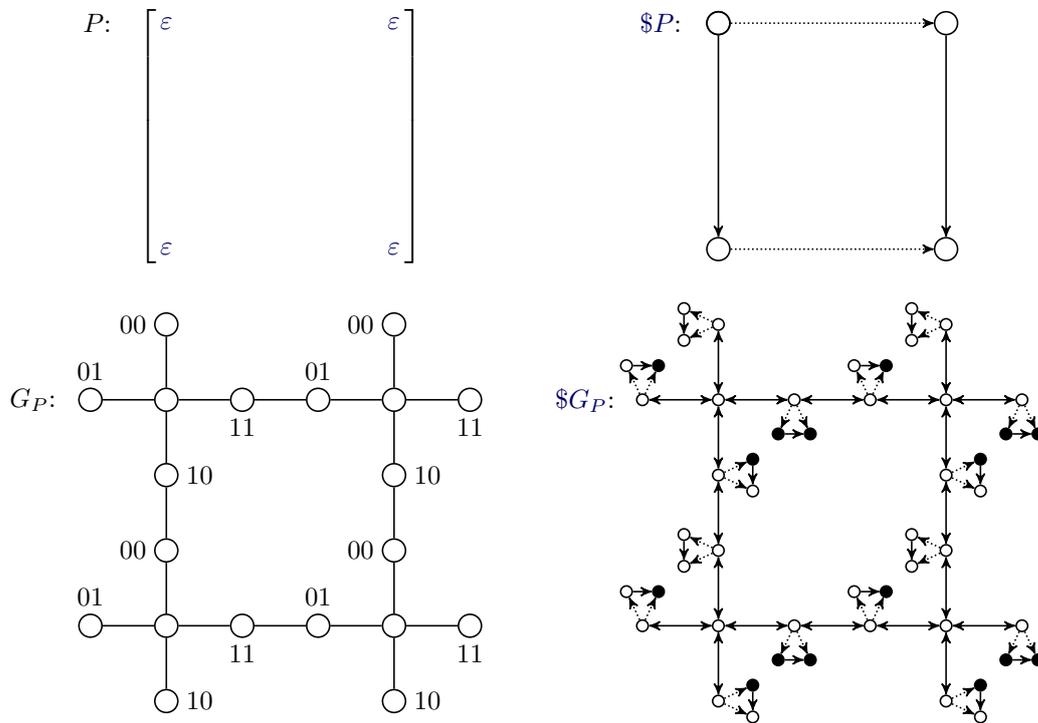
$$V^{G_P} = \langle m] \times \langle n] \times \{pxl, in_1, in_2, out_1, out_2\},$$

the set of edges

$$\begin{aligned} E^{G_P} = \;& \big\{\, \{(i,j,pxl),(i,j,\alpha)\} \;\big|\; i \in \langle m], j \in \langle n], \alpha \in \{in_1, in_2, out_1, out_2\} \,\big\} \\ & \cup \big\{\, \{(i,j,out_1),(i+1,j,in_1)\} \;\big|\; i \in \langle m), j \in \langle n] \,\big\} \\ & \cup \big\{\, \{(i,j,out_2),(i,j+1,in_2)\} \;\big|\; i \in \langle m], j \in \langle n) \,\big\}, \end{aligned}$$

and the labeling function

$$\begin{aligned} \lambda^{G_P}: \quad & (i,j,pxl) \mapsto \varepsilon, & (i,j,in_1) \mapsto 00, & \quad (i,j,out_1) \mapsto 10, \\ & & (i,j,in_2) \mapsto 01, & \quad (i,j,out_2) \mapsto 11, \end{aligned}$$

Figure 13 The 0-bit picture P of size $(2, 2)$, its graph encoding G_P, and their respective structural representations $\$P$ and $\$G_P$. We follow the same graphical conventions as in Figures 4 and 12.

Figure 14 The gadget Gad representing a single pixel whose value is the empty string. Gad occurs four times in the structure $\$G_P$ shown in Figure 13. We follow the same graphical conventions as in Figure 4. The names of the elements serve explanatory purposes only and are not part of the gadget.

where $i \in \langle m \rangle$, $j \in \langle n \rangle$, and ε denotes the empty string. An example is provided in Figure 13. Notice that G_P is always of 4-bounded structural degree.

If we look at the graph encoding's structural representation $\$ G_P$ (also illustrated in Figure 13), we see that each pixel is represented by the gadget shown in Figure 14. Formally, the *gadget Gad* representing any pixel (i, j) of a 0-bit picture P corresponds to the structural representation of the subgraph of G_P that is induced by $\{i\} \times \{j\} \times \{pxl, in_1, in_2, out_1, out_2\}$. It is convenient to identify the domain of *Gad* with the set

$$D^{Gad} = \{pxl, in_1, in_2, out_1, out_2\} \cup \left(\{in_1, in_2, out_1, out_2\} \times \langle 2 \rangle\right),$$

where the elements in $\{pxl, in_1, in_2, out_1, out_2\}$ represent nodes, and the remaining elements represent labeling bits. Considering Cartesian products to be associative, this allows us to identify the domain of the entire structure $\$ G_P$ with the set of elements $D^{\$ G_P} = D^{\$ P} \times D^{Gad}$, where the first component specifies the pixel, and the second component specifies the gadget element.

The following lemma states that the expressive power of the local second-order hierarchy is basically the same whether we consider pictures or graph encodings of pictures. This is because we can translate formulas from one type of structure to the other without changing the alternation level of second-order quantifiers. Consequently, any separation result established for pictures can be transferred to graphs.

▶ **Lemma 32.** *Let $\ell \in \mathbb{N}$.*
1. *For every $\Sigma_\ell^{\mathrm{LFO}}$-sentence φ evaluated on 0-bit pictures, there is a $\Sigma_\ell^{\mathrm{LFO}}$-sentence φ' evaluated on graphs such that $\$ P \models \varphi$ if and only if $\$ G_P \models \varphi'$ for all $P \in \mathrm{PIC}(0)$.*
2. *Conversely, for every $\Sigma_\ell^{\mathrm{LFO}}$-sentence φ evaluated on graphs, there is a $\Sigma_\ell^{\mathrm{LFO}}$-sentence φ' evaluated on 0-bit pictures such that $\$ G_P \models \varphi$ if and only if $\$ P \models \varphi'$ for all $P \in \mathrm{PIC}(0)$.*
The analogous statements hold for Π_ℓ^{LFO}-sentences.

Proof. The forward direction (Statement 1) is straightforward. First-order quantification on $\$ G_P$ is relativized to elements of $D^{\$ P} \times \{pxl\}$, which correspond to the central element of the gadget, and atomic formulas that refer to the relations $\rightharpoonup_1^{\$ P}$ and $\rightharpoonup_2^{\$ P}$ of $\$ P$ are rewritten in terms of the representation of these relations in $\$ G_P$. There is no need to explicitly relativize second-order quantification, since relations can only be evaluated for elements represented by first-order variables anyway.

We start by defining some helper formulas. The formula

$$IsPixel(x) = \neg \exists y \rightleftharpoons x \left(x \rightharpoonup_2 y \lor y \rightharpoonup_2 x\right)$$

states that x is a "pixel center", corresponding to element pxl of the gadget. For $i \in \{1, 2\}$, the formulas

$$HasBitFalse_i(x) = \exists y \rightleftharpoons x \left(AtBitPos_i(x, y) \land \neg \odot_1 y\right) \quad \text{and}$$

$$HasBitTrue_i(x) = \exists y \rightleftharpoons x \left(AtBitPos_i(x, y) \land \odot_1 y\right)$$

state that the i-th labeling bit of node x has value 0 and 1, respectively, where the subformulas

$$AtBitPos_i(x, y) = \begin{cases} x \rightharpoonup_2 y \ \land \ \exists z \rightleftharpoons x \left(y \rightharpoonup_1 z\right) & \text{if } i = 1, \\ x \rightharpoonup_2 y \ \land \ \exists z \rightleftharpoons x \left(z \rightharpoonup_1 y\right) & \text{if } i = 2 \end{cases}$$

identify y as the i-th labeling bit of x. Building on that,

$$IsIn_1(x) = HasBitFalse_1(x) \land HasBitFalse_2(x)$$

states that x is a "vertical input port", corresponding to element in_1 of the gadget. Similarly, we define $IsIn_2(x)$, $IsOut_1(x)$, and $IsOut_2(x)$, which identify the "horizontal input port", "vertical output port", and "horizontal output port", respectively.

We now show by structural induction that there is a translation τ from formulas on 0-bit pictures to formulas on graphs such that for every formula φ, every picture $P \in \text{PIC}(0)$, and every variable assignment σ of free(φ) on $\$P$, we have $\$P, \sigma \models \varphi$ if and only if $\$G_P, \sigma' \models \tau(\varphi)$, where σ' is the variable assignment of free(φ) on $\$G_P$ such that $\sigma'(x) = \big(\sigma(x), pxl\big)$ for all $x \in \text{free}_{\text{FO}}(\varphi)$, and

$$\sigma'(R) = \big\{ \big((a_1, pxl), \ldots, (a_k, pxl)\big) \mid (a_1, \ldots, a_k) \in \sigma(R) \big\}$$

for all $R \in \text{free}_{\text{SO}}(\varphi)$ of arity $k \in \mathbb{N}_{>0}$.

- To express that one pixel of P is the "vertical" or "horizontal" successor of another, we state that the appropriate ports of the corresponding nodes of G_P are connected, i.e.,

$$\tau(x \rightharpoonup_i y) \ = \ \exists z_1, z_2 \overset{\leq 2}{\rightleftharpoons} x \big(IsOut_i(z_1) \wedge IsIn_i(z_2) \wedge x \rightharpoonup_1 z_1 \rightharpoonup_1 z_2 \rightharpoonup_1 y \big).$$

- Atomic formulas for equality and other relations are kept unchanged: $\tau(x \doteq y) = (x \doteq y)$, and $\tau\big(R(x_1, \ldots, x_k)\big) = R(x_1, \ldots, x_k)$ for any second-order variable R of arity $k \in \mathbb{N}_{>0}$.
- Boolean connectives are also preserved: $\tau(\neg\varphi) = \neg\tau(\varphi)$, and $\tau(\varphi \vee \psi) = \tau(\varphi) \vee \tau(\psi)$.
- First-order quantifiers are relativized to "pixel centers". More precisely, the unbounded universal quantifier at the outermost scope of an LFO-formula is translated by $\tau(\forall x\, \varphi) = \forall x \big(IsPixel(x) \rightarrow \tau(\varphi) \big)$. For bounded first-order quantifiers, the idea is the same, but we have to take into account that two adjoining "pixel centers" lie at a distance of 3 from each other, thus $\tau(\exists y \rightleftharpoons x\, \varphi) \ = \ \exists y \overset{\leq 3}{\rightleftharpoons} x \big(y \not\equiv x \wedge IsPixel(y) \wedge \tau(\varphi) \big)$.
- Second-order quantifiers are not affected by the translation, i.e., $\tau(\exists R\, \varphi) = \exists R \tau(\varphi)$ for any second-order variable R of arity $k \in \mathbb{N}_{>0}$. This works because $\$P, \sigma \models \exists R\, \varphi$ is equivalent to the existence of $A \subseteq (D^{\$P})^k$ such that $\$P, \sigma[R \mapsto A] \models \varphi$, which by induction is equivalent to the existence of $A' \subseteq (D^{\$P} \times \{pxl\})^k$ such that $\$G_P, \sigma'[R \mapsto A'] \models \tau(\varphi)$, where σ' is defined as above. This in turn is equivalent to the existence of

$$A' \subseteq \big(D^{\$P} \times \{pxl\}\big)^k \qquad \text{and} \qquad A'' \subseteq \big(D^{\$G_P}\big)^k \setminus \big(D^{\$P} \times \{pxl\}\big)^k$$

such that $\$G_P, \sigma'[R \mapsto A' \cup A''] \models \tau(\varphi)$, where A'' is irrelevant because we ensure that all first-order variables refer to "pixel centers". Finally, since $A' \cup A''$ can be any k-ary relation on $D^{\$G_P}$, the last condition is equivalent to $\$G_P, \sigma' \models \exists R\, \tau(\varphi)$.

Note that our translation preserves the alternation level of second-order quantifiers and the fact that there is exactly one unbounded universal quantifier nested directly below the second-order quantifiers. That is, if φ is a Σ_ℓ^{LFO}- or Π_ℓ^{LFO}-formula, then so is $\tau(\varphi)$. Hence, Statement 1 corresponds to the special case of the induction hypothesis where φ does not have any free variables.

The backward direction (Statement 2) is a bit more tedious because $\$G_P$ has card(Gad) times as many elements as $\$P$. To simulate the additional elements when translating a formula φ from graphs to pictures, we introduce a "virtual variable assignment" $f \colon \text{free}_{\text{FO}}(\varphi) \rightarrow D^{Gad}$ that tells us for each free first-order variable of φ to which element of the gadget it corresponds. In combination with the actual variable assignment σ' on $\$P$, which tells us the corresponding pixel, this allows us to reference every element of $\$G_P$. Our translation is thus

parameterized by f. We handle first-order quantification by combining actual quantification with a case distinction over all possible values of f, and second-order quantification by representing each k-ary relation variable R of φ by a collection of variables $R_{(a_1,\ldots,a_k)}$, one for each k-tuple $(a_1,\ldots,a_k)$ of gadget elements.

Formally, we show by structural induction that there is a parameterized translation τ_f from formulas on graphs to formulas on 0-bit pictures such that for every formula φ, every picture $P \in \text{PIC}(0)$, and every variable assignment σ of free(φ) on $\$G_P$, we have $\$G_P, \sigma \models \varphi$ if and only if $\$P, \sigma' \models \tau_f(\varphi)$, where σ' is the variable assignment of free(φ) on $\$P$ and f is the "virtual variable assignment" free$_{\text{FO}}(\varphi) \to D^{Gad}$ such that $\sigma(x) = \big(\sigma'(x), f(x)\big)$ for all $x \in \text{free}_{\text{FO}}(\varphi)$, and

$$\sigma(R) = \bigcup_{a_1,\ldots,a_k \in Gad} \big\{ \big((b_1, a_1), \ldots, (b_k, a_k)\big) \mid (b_1, \ldots, b_k) \in \sigma'(R_{(a_1,\ldots,a_k)}) \big\}$$

for all $R \in \text{free}_{\text{SO}}(\varphi)$ of arity $k \in \mathbb{N}_{>0}$.

- To express that an element of $\$G_P$ lies in the set $\odot_1^{\$G_P}$, we state that the corresponding pixel of $\$P$ is mapped by f to a gadget element that lies in $\odot_1^{Gad}$. Hence,

$$\tau_f(\odot_1 x) = \begin{cases} \top & \text{if } f(x) \in \odot_1^{Gad}, \\ \bot & \text{otherwise.} \end{cases}$$

- To express that two elements of $\$G_P$ are connected by the relation $\rightarrowtail_i^{\$G_P}$, we need to distinguish the case where they belong to the same gadget from the case where they are in two adjacent gadgets. In the first case, the corresponding pixels of $\$P$ must coincide and be mapped by f to two gadget elements that are connected by $\rightarrowtail_i^{Gad}$. In the second case, the connection must necessarily be a "$\rightarrowtail_1$"-link from an "input port" to an "output port", or vice versa, and the corresponding pixels of $\$P$ must be connected accordingly by the "vertical" or "horizontal" successor relation:

$$\tau_f\big(x \rightarrowtail_i y\big) = \begin{cases} x \doteq y & \text{if } f(x) \rightarrowtail_i^{Gad} f(y), \\ x \rightarrowtail_j y & \text{if } i = 1,\ f(x) = out_j,\ \text{and } f(y) = in_j,\ \text{where } j \in \{1, 2\}, \\ y \rightarrowtail_j x & \text{if } i = 1,\ f(x) = in_j,\ \text{and } f(y) = out_j,\ \text{where } j \in \{1, 2\}, \\ \bot & \text{otherwise.} \end{cases}$$

- In order for two elements of $\$G_P$ to be equal, they must correspond to the same pixel of $\$P$ and the same element of the gadget, i.e.,

$$\tau_f\big(x \doteq y\big) = \begin{cases} x \doteq y & \text{if } f(x) = f(y), \\ \bot & \text{otherwise.} \end{cases}$$

- To express that k elements of $\$G_P$ are R-related, for some second-order variable R of arity k, we state that the corresponding pixels of $\$P$ are related by the appropriate copy of R, which is determined by the gadget elements that f assigns to each pixel. That is, $\tau_f\big(R(x_1, \ldots, x_k)\big) = R_{(f(x_1), \ldots, f(x_k))}(x_1, \ldots, x_k)$.
- Boolean connectives are preserved: $\tau_f(\neg\psi) = \neg\tau_f(\psi)$, and $\tau_f(\psi_1 \vee \psi_2) = \tau_f(\psi_1) \vee \tau_f(\psi_2)$.
- First-order quantification over $\$G_P$ is expressed through a combination of first-order quantification over $\$P$ and a case distinction over the gadget element to which f maps the quantified variable. For the (unique) unbounded first-order quantifier of LFO, this simply means

$$\tau_f\big(\forall x\, \varphi\big) = \forall x \bigwedge_{a \in Gad} \tau_{f[x \mapsto a]}(\varphi).$$

For bounded first-order quantifiers, the case distinction is a bit more involved because we must take into account the topology of $\$G_P$: each element of $\$G_P$ is connected to its neighbors within the same gadget, and additionally, an "input" or "output port" is also connected to its counterpart in the appropriate adjacent gadget. This leads to

$$\tau_f\big(\exists y \rightleftharpoons x\; \varphi\big) = \exists y \overset{\leq 1}{\rightleftharpoons} x\big(IntraGadget_\varphi(x,y) \vee InterGadget_\varphi(x,y)\big),$$

where

$$IntraGadget_\varphi(x,y) = (y \doteq x) \wedge \bigvee_{a \rightleftharpoons^{Gad} f(x)} \tau_{f[y\mapsto a]}(\varphi),$$

and

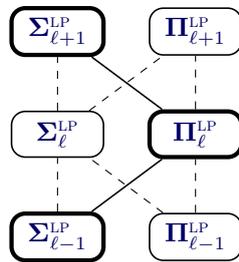$$InterGadget_\varphi(x,y) = \begin{cases} y \rightharpoonup_i x \wedge \tau_{f[y\mapsto out_i]}(\varphi) & \text{if } f(x) = in_i, \text{ for } i \in \{1,2\}, \\ x \rightharpoonup_i y \wedge \tau_{f[y\mapsto in_i]}(\varphi) & \text{if } f(x) = out_i, \text{ for } i \in \{1,2\}, \\ \bot & \text{otherwise.} \end{cases}$$

- Each second-order quantification over $\$G_P$ is expressed through multiple second-order quantifications over $\$P$. More precisely, each k-ary relation A on $\$G_P$ is represented as the union of $\mathrm{card}(Gad)^k$ pairwise disjoint relations such that each relation $A_{(a_1,\dots,a_k)}$ contains precisely those k-tuples of A whose components correspond to the gadget elements $a_1,\dots,a_k$ (in that order). Hence,

$$\tau_f\big(\exists R\,\varphi\big) = \exists (R_{(a_1,\dots,a_k)})_{a_1,\dots,a_k\in Gad}\big(\tau_f(\varphi)\big).$$

Again, our translation preserves the alternation level of second-order quantifiers and the fact that there is exactly one unbounded universal quantifier nested directly below the second-order quantifiers. Hence, Statement 2 corresponds to the special case of the induction hypothesis where φ does not have any free variables (which means, in particular, that the "virtual variable assignment" f is empty). ◀

We now have everything at hand to transfer part of Theorem 27 from monadic second-order logic on pictures to local second-order logic on graphs, and thus to the local-polynomial hierarchy. The result is stated in the following theorem and illustrated in Figure 15.

Figure 15 A partial separation result for the local-polynomial hierarchy obtained in Theorem 33 for every even integer $\ell \geq 2$. Each line indicates an inclusion of the lower class in the higher class. The inclusions represented by solid lines are proved to be strict, even when restricted to graphs of bounded structural degree. This forms the basis for the fuller separation result shown in Figure 11.

▶ **Theorem 33.** *The local-polynomial hierarchy is infinite, even when restricted to graphs of bounded structural degree. More precisely, $\mathbf{\Sigma}^{\mathrm{LP}}_{\ell-1}|_{\mathrm{GRAPH}(\varDelta)} \subsetneq \mathbf{\Pi}^{\mathrm{LP}}_\ell|_{\mathrm{GRAPH}(\varDelta)} \subsetneq \mathbf{\Sigma}^{\mathrm{LP}}_{\ell+1}|_{\mathrm{GRAPH}(\varDelta)}$, and a fortiori $\mathbf{\Sigma}^{\mathrm{LP}}_{\ell-1} \subsetneq \mathbf{\Pi}^{\mathrm{LP}}_\ell \subsetneq \mathbf{\Sigma}^{\mathrm{LP}}_{\ell+1}$, for every even integer $\ell \geq 2$ and every integer $\varDelta \geq 4$.*

Proof. By Theorem 12, the statement can be equivalently formulated in terms of the local second-order hierarchy on graphs: $\mathbf{\Sigma}^{\text{LFO}}_{\ell-1}|_{\text{GRAPH}(4)} \subsetneqq \mathbf{\Pi}^{\text{LFO}}_{\ell}|_{\text{GRAPH}(4)} \subsetneqq \mathbf{\Sigma}^{\text{LFO}}_{\ell+1}|_{\text{GRAPH}(4)}$ for every even integer $\ell \geq 2$. To prove it, we start with the analogous separation result for the monadic second-order hierarchy on 0-bit pictures, which holds by Theorem 27:

$$\mathbf{m\Sigma}^{\text{FO}}_{\ell-1}|_{\text{PIC}(0)} \subsetneqq \mathbf{m\Pi}^{\text{FO}}_{\ell}|_{\text{PIC}(0)} \subsetneqq \mathbf{m\Sigma}^{\text{FO}}_{\ell+1}|_{\text{PIC}(0)}$$

By Theorem 31, this can be rewritten in terms of local second-order logic:

$$\mathbf{\Sigma}^{\text{LFO}}_{\ell-1}|_{\text{PIC}(0)} \subsetneqq \mathbf{\Pi}^{\text{LFO}}_{\ell}|_{\text{PIC}(0)} \subsetneqq \mathbf{\Sigma}^{\text{LFO}}_{\ell+1}|_{\text{PIC}(0)} \tag{$*$}$$

We now transfer this result from 0-bit pictures to graphs of 4-bounded structural degree. The first inequality of $(*)$ tells us that there exists a picture property $L \in \mathbf{\Pi}^{\text{LFO}}_{\ell}|_{\text{PIC}(0)}$ that does not lie in $\mathbf{\Sigma}^{\text{LFO}}_{\ell-1}|_{\text{PIC}(0)}$. Applying the forward translation provided by Lemma 32.1 and the fact that graph encodings of 0-bit pictures are of 4-bounded structural degree, we infer the existence of a graph property $L' \in \mathbf{\Pi}^{\text{LFO}}_{\ell}|_{\text{GRAPH}(4)}$ such that $P \in L$ if and only if $G_P \in L'$, for every 0-bit picture P. Similarly, the backward translation provided by Lemma 32.2 lets us deduce that $L' \notin \mathbf{\Sigma}^{\text{LFO}}_{\ell-1}|_{\text{GRAPH}(4)}$, because otherwise we would have $L \in \mathbf{\Sigma}^{\text{LFO}}_{\ell-1}|_{\text{PIC}(0)}$. Hence, $\mathbf{\Sigma}^{\text{LFO}}_{\ell-1}|_{\text{GRAPH}(4)} \subsetneqq \mathbf{\Pi}^{\text{LFO}}_{\ell}|_{\text{GRAPH}(4)}$. Analogously, we can conclude from the second inequality of $(*)$ that $\mathbf{\Pi}^{\text{LFO}}_{\ell}|_{\text{GRAPH}(4)} \subsetneqq \mathbf{\Sigma}^{\text{LFO}}_{\ell+1}|_{\text{GRAPH}(4)}$. ◀

9.3   Completing the picture

In this subsection, we establish all the remaining separations and inclusions shown in Figure 11 on page 47, and then conclude by identifying graph properties that lie outside the local-polynomial hierarchy.

Our first goal is to prove that the inclusions represented by dashed lines in Figure 11 are equalities when restricted to graphs of bounded structural degree. To do this, we first show that on such graphs, we can refine the notion of restrictive arbiters introduced in Section 6 to require that identifier assignments are necessarily small.

Let ℓ be a nonnegative integer, r_{id} and r be positive integers, p be a polynomial function, K be an **LP**-property, and $M_1, \ldots, M_\ell$ be certificate restrictors for (r, p)-bounded certificates under r_{id}-locally unique identifiers. A *small-restrictive $\mathbf{\Sigma}^{\text{LP}}_{\ell}$-arbiter* for a graph property L on K under r_{id}-locally unique identifiers and (r, p)-bounded certificates restricted by $M_1, \ldots, M_\ell$ is a local-polynomial machine M that satisfies the same equivalence as a restrictive $\mathbf{\Sigma}^{\text{LP}}_{\ell}$-arbiter (see on page 28) for every graph $G \in K$ and every *small* r_{id}-locally unique identifier assignment *id* of G. That is, the equivalence does not have to hold for arbitrary-sized identifiers. We analogously define *small-restrictive $\mathbf{\Pi}^{\text{LP}}_{\ell}$-arbiters*.

The following lemma is a refinement of Lemma 8 for small-restrictive arbiters on graphs of bounded structural degree.

▶ **Lemma 34.** *Let $\ell, \Delta \in \mathbb{N}$ and $L \subseteq \text{GRAPH}$. The graph property $L \cap \text{GRAPH}(\Delta)$ belongs to $\mathbf{\Sigma}^{\text{LP}}_{\ell}|_{\text{GRAPH}(\Delta)}$ if and only if L has a small-restrictive $\mathbf{\Sigma}^{\text{LP}}_{\ell}$-arbiter on $\text{GRAPH}(\Delta)$. The analogous statement holds for $\mathbf{\Pi}^{\text{LP}}_{\ell}|_{\text{GRAPH}(\Delta)}$.*

Proof. We prove only the first statement, since the proof for $\mathbf{\Pi}^{\text{LP}}_{\ell}|_{\text{GRAPH}(\Delta)}$ is completely analogous. By definition, if $L \cap \text{GRAPH}(\Delta)$ belongs to $\mathbf{\Sigma}^{\text{LP}}_{\ell}|_{\text{GRAPH}(\Delta)}$, then there exists a permissive $\mathbf{\Sigma}^{\text{LP}}_{\ell}$-arbiter M for a graph property L' such that $L' \cap \text{GRAPH}(\Delta) = L \cap \text{GRAPH}(\Delta)$, and thus M is also a small-restrictive $\mathbf{\Sigma}^{\text{LP}}_{\ell}$-arbiter for L on $\text{GRAPH}(\Delta)$ under unrestricted certificates.

For the converse, since GRAPH(Δ) is an **LP**-property, it suffices by Lemma 8 to convert a small-restrictive arbiter into a restrictive arbiter that operates under arbitrary-sized identifiers. Let M be a small-restrictive $\Sigma_\ell^{\mathrm{LP}}$-arbiter for L on GRAPH(Δ) under r_{id}-locally unique identifiers and (r, p)-bounded certificates restricted by $M_1, \ldots, M_\ell$. We need to distinguish two cases:

1. If $\ell = 0$, then M is in fact a small-restrictive **LP**-decider for L on GRAPH(Δ). This means that every node reaches its verdict simply by examining a portion of the input graph, i.e., without having to consider any certificates that could potentially depend on a particular choice of identifiers. Suppose that M runs in round time t. We construct M' such that when it is executed on a graph G, the nodes first communicate for $(t + 2r_{\mathrm{id}})$ rounds to reconstruct their $(t + 2r_{\mathrm{id}})$-neighborhoods. Then, each node u simulates M locally on its t-neighborhood under every possible small r_{id}-locally unique identifier assignment of G restricted to $N_t^G(u)$, and finally accepts if it did so in every simulation. The condition of being small r_{id}-locally unique can be respected because u knows the $2r_{\mathrm{id}}$-neighborhood of every node in $N_t^G(u)$. Note that by running the simulation under every possible identifier assignment, we avoid the problem of the nodes having to agree on a particular one. This can be done in constant step time because the restriction to graphs of Δ-bounded structural degree entails a constant upper bound on the number of (partial) identifier assignments each node has to consider. The machine M' obtained this way is a restrictive **LP**-decider for L on GRAPH(Δ) under arbitrary-sized r'_{id}-locally unique identifiers, where $r'_{\mathrm{id}} = t + 2r_{\mathrm{id}}$.

2. If $\ell > 0$, then M can be simulated by a machine M' that uses the first certificate assignment to encode small identifiers. More precisely, let the given small-restrictive arbiter M be such that for every graph $G \in$ GRAPH(Δ) and every small r_{id}-locally unique identifier assignment id of G,

$$G \in L \iff \exists \kappa_1 \, \forall \kappa_2 \ldots \rotatebox[origin=c]{180}{Q}\kappa_\ell : M(G, \mathit{id}, \kappa_1 \cdot \kappa_2 \cdot \ldots \cdot \kappa_\ell) \equiv \text{ACCEPT},$$

where all quantifiers range over (r, p)-bounded certificate assignments of (G, id) with the additional restriction that $M_i(G, \mathit{id}, \kappa_1 \cdot \ldots \cdot \kappa_i) \equiv$ ACCEPT for all $i \in [1 : \ell]$. We now construct M' with appropriately chosen constants $r'_{\mathrm{id}}, r' \in \mathbb{N}_{>0}$, polynomial p', and certificate restrictors $M'_1, \ldots, M'_\ell$ such that for every graph $G \in$ GRAPH(Δ) and every arbitrary-sized r'_{id}-locally unique identifier assignment id' of G,

$$G \in L \iff \exists \kappa'_1 \, \forall \kappa_2 \ldots \rotatebox[origin=c]{180}{Q}\kappa_\ell : M'(G, \mathit{id}', \kappa'_1 \cdot \kappa_2 \cdot \ldots \cdot \kappa_\ell) \equiv \text{ACCEPT},$$

where all quantifiers range over (r', p')-bounded certificate assignments of (G, id') with the additional restriction that $M'_i(G, \mathit{id}', \kappa'_1 \cdot \kappa_2 \cdot \ldots \cdot \kappa_i) \equiv$ ACCEPT for all $i \in [1 : \ell]$. The certificate restrictors are chosen such that κ'_1 encodes both a small r_{id}-locally unique identifier assignment id of G and an (r, p)-bounded certificate assignment κ_1 of (G, id) satisfying the restrictions imposed by M_1, and the remaining certificate assignments $\kappa_2, \ldots, \kappa_\ell$ are (r, p)-bounded with respect to (G, id) and satisfy the restrictions imposed by $M_2, \ldots, M_\ell$. The machine M' itself then simply simulates M on G under id and $\kappa_1 \cdot \kappa_2 \cdot \ldots \cdot \kappa_\ell$.

Note that it is easy to construct $M'_1, \ldots, M'_\ell$ such that they satisfy local repairability (and thus the definition of a certificate restrictor). In particular, the restriction of id being small r_{id}-locally unique is compatible with local repairability because if a node has an invalid identifier (i.e., too large or not r_{id}-locally unique), then by the proof of Remark 1, we can assign it a valid identifier without affecting the validity of the other nodes' identifiers. Also note that we need the restriction to graphs of Δ-bounded structural degree only for the case $\ell = 0$, not for the case $\ell > 0$. ◄

Using the notion of small-restrictive arbiters, we now show that on graphs of bounded structural degree, it is useless to let Adam choose the last certificate assignment. Intuitively, this is because on such graphs, the step running time of the nodes can be arbitrarily large with respect to their local input and the messages they receive, so they can use brute force to perform universal quantification over the last certificate. This is analogous to the observation made by Feuilloley, Fraigniaud, and Hirvonen [10] for their alternation hierarchy (which does not impose any restrictions on the processing power of the nodes).

▶ **Proposition 35.** *When restricted to graphs of bounded structural degree, every level of the local-polynomial hierarchy that ends with a universal quantifier is equivalent to the level directly below it that lacks that final quantifier. Formally, $\Sigma_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)} = \Sigma_{\ell+1}^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)}$ if ℓ is odd, and $\Pi_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)} = \Pi_{\ell+1}^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)}$ if ℓ is even, for all $\ell, \Delta \in \mathbb{N}$.*

Proof. We prove only the second equality, the proof of the first being completely analogous. Let L be a graph property in $\Pi_{\ell+1}^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)}$, and let M be a restrictive $\Pi_{\ell+1}^{\mathrm{LP}}$-arbiter for L on $\mathrm{GRAPH}(\Delta)$ under r_{id}-locally unique identifiers and unrestricted (r, p)-bounded certificates. By definition, for every graph $G \in \mathrm{GRAPH}(\Delta)$ and every r_{id}-locally unique identifier assignment id of G, we have

$$G \in L \iff \forall \kappa_1 \exists \kappa_2 \dots \forall \kappa_{\ell+1} : M(G, id, \kappa_1 \cdot \kappa_2 \cdot \dots \cdot \kappa_{\ell+1}) \equiv \mathrm{ACCEPT},$$

where all quantifiers range over (r, p)-bounded certificate assignments of (G, id). Now, to show that $L \in \Pi_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)}$, it suffices by Lemma 34 to provide a small-restrictive Π_ℓ^{LP}-arbiter M' for L on $\mathrm{GRAPH}(\Delta)$. More precisely, we construct M' with an appropriately chosen constant r'_{id} such that for every graph $G \in \mathrm{GRAPH}(\Delta)$ and every small r'_{id}-locally unique identifier assignment id of G, we have

$$G \in L \iff \forall \kappa_1 \exists \kappa_2 \dots \exists \kappa_\ell : M'(G, id, \kappa_1 \cdot \kappa_2 \cdot \dots \cdot \kappa_\ell) \equiv \mathrm{ACCEPT},$$

where all quantifiers range over (r, p)-bounded certificate assignments of (G, id).

Given the certificates assigned to them by $\kappa_1, \dots, \kappa_\ell$, the nodes running M' simulate M for every possible choice of $\kappa_{\ell+1}$. To do so, they first communicate for $r' = r_M + r$ rounds, where r_M is a (constant) bound on the round running time of M. Setting $r'_{\mathrm{id}} = r'$, this allows each node u of G to reconstruct its r'-neighborhood $N_{r'}^G(u)$ and the identifiers and certificates of all nodes therein. Then, u simulates M locally for every (r, p)-bounded certificate assignment $\kappa_{\ell+1}$ of (G, id) restricted to $N_{r_M}^G(u)$. Note that u can respect the condition of (r, p)-boundedness because it knows the r-neighborhood of every node in $N_{r_M}^G(u)$. Finally, u accepts precisely if it has accepted in every simulation.

Intuitively, the reason why this approach works is that there is a universal quantification on both $\kappa_{\ell+1}$ and on the nodes in the acceptance criterion of distributed Turing machines. Therefore, we can reverse the order of quantification (by letting the nodes perform the quantification over $\kappa_{\ell+1}$) without changing the semantics.

The step running time of M' at u is clearly exponential in

$$\sum_{v \in N_{r_M}^G(u)} p\Big(\sum_{w \in N_r^G(v)} 1 + \mathrm{len}(\lambda^G(w)) + \mathrm{len}(id(w)) \Big).$$

However, since we require G to be of Δ-bounded structural degree and id to be small r'-locally unique, this value is bounded by a constant that depends only on Δ, r', and p. Hence, M' runs in constant and thus polynomial step time.   ◀

The equalities established in Proposition 35 not only simplify the local-polynomial hierarchy on graphs of bounded structural degree, but also imply, in combination with the previous separation results, that the remaining inclusions in the hierarchy are strict. This is particularly relevant on arbitrary graphs.

▶ **Corollary 36.** *The separation results for the local-polynomial hierarchy stated in Proposition 21 and Theorem 33 can be extended as follows:* $\mathbf{\Pi}_{\ell-1}^{\mathrm{LP}} \subsetneq \ngeq \mathbf{\Pi}_{\ell}^{\mathrm{LP}}$, *and* $\mathbf{\Sigma}_{\ell}^{\mathrm{LP}} \subsetneq \ngeq \mathbf{\Sigma}_{\ell+1}^{\mathrm{LP}}$, *and* $\mathbf{\Pi}_{\ell-1}^{\mathrm{LP}} \subsetneq \ngeq \mathbf{\Sigma}_{\ell}^{\mathrm{LP}} \subsetneq \ngeq \mathbf{\Pi}_{\ell+1}^{\mathrm{LP}}$, *for every even integer $\ell \geq 2$.*

Proof. By Proposition 21 and Theorem 33 (for the inequalities), and by Proposition 35 (for the equalities), we have

$$
\mathbf{\Pi}_{\ell-2}^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)} \quad \subsetneq \ngeq \quad \mathbf{\Sigma}_{\ell-1}^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)} \quad \subsetneq \ngeq \quad \mathbf{\Pi}_{\ell}^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)} \quad \subsetneq \ngeq \quad \mathbf{\Sigma}_{\ell+1}^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)}
$$
$$
\| \qquad\qquad\qquad \| \qquad\qquad\qquad \|
$$
$$
\mathbf{\Pi}_{\ell-1}^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)} \qquad\qquad \mathbf{\Sigma}_{\ell}^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)} \qquad\qquad \mathbf{\Pi}_{\ell+1}^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)},
$$

which by transitivity yields $\mathbf{\Pi}_{\ell-1}^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)} \subsetneq \ngeq \mathbf{\Pi}_{\ell}^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)}$, and $\mathbf{\Sigma}_{\ell}^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)} \subsetneq \ngeq \mathbf{\Sigma}_{\ell+1}^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)}$, and $\mathbf{\Pi}_{\ell-1}^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)} \subsetneq \ngeq \mathbf{\Sigma}_{\ell}^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)} \subsetneq \ngeq \mathbf{\Pi}_{\ell+1}^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)}$. This implies the desired inequalities on arbitrary graphs. ◀

Since the proof of Proposition 35 relies on exhaustive search, it is unlikely to generalize to arbitrary graphs. This can be restated as follows:

▶ Remark 37. The statement of Proposition 35 generalizes to arbitrary graphs if and only if $\mathbf{P} = \mathbf{coNP}$.

Proof. If $\mathbf{P} = \mathbf{coNP}$, then the proof of Proposition 35 does not require the restriction to graphs of bounded structural degree. Indeed, the local simulations described there can then be performed by a polynomial-time algorithm that is equivalent to testing all possible certificate assignments in parallel.

Conversely, if the statement of Proposition 35 holds for arbitrary graphs, then in particular we have $\mathbf{\Pi}_0^{\mathrm{LP}} = \mathbf{\Pi}_1^{\mathrm{LP}}$, which entails $\mathbf{\Pi}_0^{\mathrm{LP}}|_{\mathrm{NODE}} = \mathbf{\Pi}_1^{\mathrm{LP}}|_{\mathrm{NODE}}$, and thus $\mathbf{P} = \mathbf{coNP}$. ◁

Next, we focus on the relationship between the local-polynomial hierarchy and its complement hierarchy. Using our previous results, it is easy to show that the two hierarchies are completely distinct.

▶ **Corollary 38.** *Classes on the same level of the local-polynomial hierarchy are neither complement classes of each other, nor are they closed under complementation, even when restricted to graphs of bounded structural degree. More precisely,* $\mathbf{\Sigma}_{\ell}^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)} \nsubseteq \mathbf{co\Pi}_{\ell}^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)}$ *if ℓ is odd, and* $\mathbf{\Pi}_{\ell}^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)} \nsubseteq \mathbf{co\Sigma}_{\ell}^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)}$ *if ℓ is even, for all $\ell \in \mathbb{N}$ and $\Delta \geq 4$. Moreover,* $\mathbf{\Sigma}_{\ell}^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)} \neq \mathbf{co\Sigma}_{\ell}^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)}$ *and* $\mathbf{\Pi}_{\ell}^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)} \neq \mathbf{co\Pi}_{\ell}^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)}$, *for all $\ell \in \mathbb{N}$ and $\Delta \geq 4$.*

Proof. The statement for $\ell = 0$ reduces to $\mathbf{LP}|_{\mathrm{GRAPH}(4)} \neq \mathbf{coLP}|_{\mathrm{GRAPH}(4)}$, which holds by Corollary 24.

For $\ell \geq 1$, we first show that none of the classes is closed under complementation, building on the analogous result for the monadic second-order hierarchy on 0-bit pictures. By Theorem 27, we know that $\mathbf{m\Sigma}_{\ell}^{\mathrm{FO}}|_{\mathrm{PIC}(0)}$ and $\mathbf{m\Pi}_{\ell}^{\mathrm{FO}}|_{\mathrm{PIC}(0)}$ are incomparable for all $\ell \in \mathbb{N}_{>0}$. Since these two classes are complement classes of each other (FO being closed under negation), this means that neither class is closed under complementation. By Theorem 31, this implies that $\mathbf{\Sigma}_{\ell}^{\mathrm{LFO}}|_{\mathrm{PIC}(0)}$ is not closed under complementation if ℓ is odd, and that $\mathbf{\Pi}_{\ell}^{\mathrm{LFO}}|_{\mathrm{PIC}(0)}$ is not closed under complementation if ℓ is even. We can transfer this result from 0-bit pictures to

graphs of 4-bounded structural degree by using Lemma 32 in the same way as in the proof of Theorem 33. This in turn allows us to conclude by Theorem 12 that $\mathbf{\Sigma}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)}$ is not closed under complementation if ℓ is odd, and that $\mathbf{\Pi}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)}$ is not closed under complementation if ℓ is even. The analogous statement for the remaining cases follows by Proposition 35, which tells us that $\mathbf{\Sigma}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)} = \mathbf{\Sigma}_{\ell-1}^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)}$ if ℓ is even, and $\mathbf{\Pi}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)} = \mathbf{\Pi}_{\ell-1}^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)}$ if ℓ is odd. Hence, we have $\mathbf{\Sigma}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)} \neq \mathbf{co\Sigma}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)}$ and $\mathbf{\Pi}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)} \neq \mathbf{co\Pi}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)}$ for all $\ell \in \mathbb{N}$.

It remains to show for $\ell \geq 1$ that the classes on level ℓ are not complement classes of each other. If ℓ is odd, suppose for the sake of contradiction that $\mathbf{\Sigma}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)} \subseteq \mathbf{co\Pi}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)}$. In combination with Proposition 35, this allows us to write the chain of inclusions

$$\mathbf{\Sigma}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)} \;\subseteq\; \mathbf{co\Pi}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)} \;\overset{\mathrm{Prp.35}}{=}\; \mathbf{co\Pi}_{\ell-1}^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)} \;\subseteq\; \mathbf{co\Sigma}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)},$$

which contradicts the already established fact that $\mathbf{\Sigma}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)}$ is not closed under complementation. Analogously, we can show that $\mathbf{\Pi}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)} \nsubseteq \mathbf{co\Sigma}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)}$ if ℓ is even. ◄

Although the local-polynomial hierarchy is distinct from its complement hierarchy, there are inclusions between the two. This can be shown by generalizing Examples 4 and 5 from Section 5.2. The strictness of these inclusions is immediate by Corollary 38.

▶ **Proposition 39.** *In the local-polynomial hierarchy, complementation can be achieved at the cost of two or three additional quantifier alternations. More precisely, $\mathbf{co\Sigma}_\ell^{\mathrm{LP}} \subseteq \mathbf{\Pi}_{\ell+2}^{\mathrm{LP}}$ if ℓ is even, and $\mathbf{co\Pi}_\ell^{\mathrm{LP}} \subseteq \mathbf{\Sigma}_{\ell+2}^{\mathrm{LP}}$ if ℓ is odd, for all $\ell \in \mathbb{N}_{>0}$.*

Proof. Let ℓ be even and L be a graph property in $\mathbf{\Sigma}_\ell^{\mathrm{LP}}$. By Theorem 12, L can be defined by a $\Sigma_\ell^{\mathrm{LFO}}$-formula of the form $\exists \bar{R}_1 \forall \bar{R}_2 \ldots \forall \bar{R}_\ell \, \forall x \, \psi(x)$. Based on that, the complement $\bar{L}$ can be defined by the formula $\forall \bar{R}_1 \exists \bar{R}_2 \ldots \exists \bar{R}_\ell \, ExistsBadNode$, where *ExistsBadNode* is a Σ_3^{LFO}-formula with free variables in $\bar{R}_1, \ldots, \bar{R}_\ell$ that is equivalent to $\exists x \, \neg\psi(x)$. The definition of *ExistsBadNode* is the same as in Example 5, except that now we use $\neg\psi(x)$ instead of $\neg WellColored(x)$ to instantiate the formula schema $PointsTo[\vartheta](x)$. Hence, again by Theorem 12, $\bar{L} \in \mathbf{\Pi}_{\ell+2}^{\mathrm{LP}}$.

The proof for ℓ odd and $L \in \mathbf{\Pi}_\ell^{\mathrm{LP}}$ is completely analogous. ◄

▶ **Corollary 40.** *The inclusions stated in Proposition 39 are strict, even when restricted to graphs of bounded structural degree.*

Proof. By Corollary 38, we know that $\mathbf{\Pi}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)} \nsubseteq \mathbf{co\Sigma}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)}$ for all even $\ell \in \mathbb{N}$ and $\Delta \geq 4$. This implies that $\mathbf{\Pi}_{\ell+2}^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)} \nsubseteq \mathbf{co\Sigma}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)}$. Analogously, we obtain $\mathbf{\Sigma}_{\ell+2}^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)} \nsubseteq \mathbf{co\Pi}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)}$ for all odd $\ell \in \mathbb{N}$ and $\Delta \geq 4$. ◄

We end this section by showing that no level of the local-polynomial hierarchy is capable of expressing graph properties that intuitively require counting the total number of nodes. To show this, the connection to logic again proves valuable, as it gives us access to classical results from automata theory.

In the following proposition, we use two examples of counting properties: SQUARE denotes the property of graphs whose cardinality is a perfect square, and PRIME denotes the property of graphs whose cardinality is a prime number.

▶ **Proposition 41.** *There are graph properties, such as SQUARE and PRIME, that lie outside the local-polynomial hierarchy, even when restricted to graphs of bounded structural degree. More precisely, $L \cap \mathrm{GRAPH}(\Delta) \notin \mathbf{\Sigma}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(\Delta)}$ for all $\ell \in \mathbb{N}$ and $\Delta \geq 4$, where L can be SQUARE or PRIME.*

Proof sketch. We focus on the statement for PRIME, as the proof for SQUARE is the same. Assume, for the sake of contradiction, that PRIME ∩ GRAPH(4) ∈ $\mathbf{\Sigma}_\ell^{\mathrm{LP}}|_{\mathrm{GRAPH}(4)}$ for some $\ell \in \mathbb{N}$, which we may suppose to be odd without loss of generality. By Theorem 12, this means there is a $\Sigma_\ell^{\mathrm{LFO}}$-sentence φ that defines the graphs of prime cardinality on GRAPH(4). We can now translate φ into a $\Sigma_\ell^{\mathrm{LFO}}$-sentence φ' such that the graph encoding G_P of a 0-bit picture P satisfies φ' if and only if the number of pixels of P is prime. This can be done similarly to the proof of Lemma 32.1, essentially by relativizing the first-order quantifiers of φ to "pixel centers" and rewriting the relation $\rightharpoonup_1$ in terms of the two "ports" through which adjacent "pixel centers" are connected (see Figure 13). Intuitively, evaluating φ' on G_P corresponds to evaluating φ on the graph obtained from G_P by removing all "ports" and connecting "pixel centers" directly instead. Then, by Lemma 32.2, we can translate φ' into a $\Sigma_\ell^{\mathrm{LFO}}$-sentence φ'' that, when evaluated on PIC(0), defines the property PRIME-GRID of 0-bit pictures whose number of pixels is prime. Hence, PRIME-GRID ∈ $\mathbf{\Sigma}_\ell^{\mathrm{LFO}}|_{\mathrm{PIC}(0)}$, which by Theorem 31 is equivalent to PRIME-GRID ∈ $\mathbf{m\Sigma}_\ell^{\mathrm{FO}}|_{\mathrm{PIC}(0)}$. Since the 0-bit picture of size $(1, n)$ can be identified with the unary word of length n, this means that there is a sentence of monadic second-order logic defining the language of unary words of prime length. By the Büchi-Elgot-Trakhtenbrot theorem (see, e.g., [38, Thm. 3.1]), we conclude that the same language is recognized by a finite-state automaton. This, however, is a contradiction because it can be shown using the pumping lemma for regular languages that no such automaton exists (see, e.g., [21, § 4.1]). ◀

10 Discussion

We have extended the polynomial hierarchy to the LOCAL model of distributed computing. Some major results of complexity theory generalize well to this setting, including Fagin's theorem and the Cook–Levin theorem. Moreover, we could go beyond what is known in the centralized setting by showing that the local-polynomial hierarchy is infinite. Descriptive complexity was very helpful in this regard, as it allowed us to build directly on sophisticated results from logic and automata theory, in particular the infiniteness of the monadic second-order hierarchy on pictures.

It seems highly unlikely that this paper will provide new insights into major open problems in complexity theory, such as **P** versus **NP**. This is because our separation results rely on the distributed nature of the LOCAL model. They do not hold in cases where distributedness is irrelevant, such as on single-node graphs, or when certificates can be replaced by local computation (see Proposition 35). However, our findings may provide a new perspective on the concept of locality in distributed computing.

Measuring locality. Within the LOCAL model, round-time complexity is certainly the most natural and widely studied measure of locality. It tells us the radius up to which each node must see in order to solve a given problem. But, as pointed out by Feuilloley [8], if we require the radius to be constant, and compensate for this by introducing nondeterminism, then certificate size becomes a natural measure of locality. Intuitively, certificate size tells us how much global information about the graph each node must receive from the prover (Eve) in order to verify a given property. Now, if we go one step further and additionally restrict the size of certificates to depend only on a constant-radius neighborhood of the nodes, and compensate for this in turn by introducing quantifier alternation, then the level of alternation arguably becomes our new measure of locality. Its meaning is more abstract, as it represents the number of moves in a two-player game, but the longer the game, the more

global information the two players can prove or disprove.

Since the local-polynomial hierarchy is infinite, it provides, at least in principle, a fine-grained measure of locality based on alternation. What remains to be seen is how meaningful its different levels are, given that the properties used to separate them involve graph encodings of pictures, which make little sense from a distributed computing perspective. In this paper, we have seen some initial clues. At the bottom of the hierarchy, the class $\mathbf{\Pi}_0^{\mathrm{LP}} = \mathbf{LP}$ represents, by definition, purely local properties. A canonical example of such a property is Eulerianness, which is $\mathbf{LP}$-complete (by Proposition 15). One level higher, $\mathbf{\Sigma}_1^{\mathrm{LP}} = \mathbf{NLP}$ represents properties that are almost, but not quite, local. A canonical example of such a property is 3-colorability, which is $\mathbf{NLP}$-complete (by Theorem 20). Since this property requires only constant-size certificates to be verified, it can also be considered quasi-local when using certificate size as the measure of locality. In contrast, the complements of Eulerianness and 3-colorability are more global, as neither of them lies in $\mathbf{\Sigma}_1^{\mathrm{LP}}$ (by Corollaries 25 and 26). We could only place non-Eulerianness in $\mathbf{\Sigma}_3^{\mathrm{LP}}$ and non-3-colorability in $\mathbf{\Pi}_4^{\mathrm{LP}}$ (by Proposition 39), leaving open whether there are matching lower bounds. Perhaps even higher in the hierarchy, we have Hamiltonicity, which can be viewed as a combination of a quasi-local condition (having a 2-regular spanning subgraph) and a more global condition (the subgraph must be connected). Again, we leave open the exact complexity of this property, but we have seen that it is at most $\mathbf{\Sigma}_5^{\mathrm{LP}}$ (by Example 6), and strictly greater than $\mathbf{\Sigma}_1^{\mathrm{LP}}$ (by Corollary 26). In addition, we have identified some graph properties that lie outside the hierarchy, a fact that indicates their inherently global nature. We have shown that this is the case for cardinality-dependent properties, such as the number of nodes being a perfect square or a prime (by Proposition 41), and we conjecture that the same holds for the property of having a nontrivial automorphism.

To gain a better intuition for the higher levels of the hierarchy, the notions of hardness and completeness under local-polynomial reductions could be helpful. In the centralized setting, Meyer and Stockmeyer [30] generalized the Cook–Levin theorem to classes of quantified Boolean formulas, thus providing complete problems for all levels of the polynomial hierarchy (see, e.g., [2, §5.2.2]). Although these problems are rather artificial in themselves, they have been used to prove the completeness of more natural problems, especially on the second and third levels of the polynomial hierarchy (see [35]). A similar strategy could be pursued in the distributed setting. It should be straightforward to further generalize our distributed version of the Cook–Levin theorem (Theorem 19) to cover the entire local-polynomial hierarchy, and based on that, we may find more natural complete properties for higher levels of the hierarchy. Given the $\mathbf{NLP}$-completeness of 3-colorability, a promising candidate would be the generalization of 3-colorability to a family of multi-round games, as introduced by Ajtai, Fagin, and Stockmeyer [1, §11]. While it is to be expected that many graph properties of interest are not complete for any level of the local-polynomial hierarchy, we may still be able to derive lower bounds for them by proving their hardness for certain levels of the hierarchy. For example, although Hamiltonicity is probably not complete for any level of the local-polynomial hierarchy,[4] we were still able to show that it does not lie in $\mathbf{\Sigma}_1^{\mathrm{LP}}$ by

---

[4] This is because from the work of Ajtai, Fagin, and Stockmeyer [1, §11] we can conclude that each level of the local-polynomial hierarchy contains a graph property whose string-encoded version is complete for the corresponding level of the classical polynomial hierarchy. For example, it is easy to see that $\mathbf{\Pi}_2^{\mathrm{LP}}$ contains the property 2-round-3-colorable, which holds for a given graph G if every 3-color assignment to the leaves of G can be extended to a valid 3-coloring of G. Thus, if hamiltonian were $\mathbf{\Pi}_2^{\mathrm{LP}}$-hard, then by simulating a distributed Turing machine with a centralized one, we could get a polynomial-time reduction from $enc(\text{2-round-3-colorable})$ to $enc(\text{hamiltonian})$, where $enc\colon \textsc{graph} \to \textsc{node}$ is some encoding of graphs as strings. But this would imply the collapse of the polynomial hierarchy to $\mathbf{NP}$, since $enc(\text{2-round-3-colorable})$ is $\mathbf{\Pi}_2^{\mathrm{P}}$-complete, and $enc(\text{hamiltonian})$ lies in $\mathbf{NP}$.

establishing its **coLP**-hardness. More generally, just as distributedness made it easier to separate the different levels of the hierarchy, it can also make it easier to prove unconditional lower bounds for individual graph properties.

Beyond polynomial bounds.   As the primary goal of this paper was to explore the connections between standard complexity theory and local distributed decision, a natural starting point was to impose polynomial bounds on the processing time and certificate sizes of the nodes. This allowed us to build on classical results with the help of descriptive complexity and to take the view that major open questions in standard complexity theory concern a particularly difficult special case of network computing. However, it could be argued that polynomial bounds are not the most canonical choice if one wishes to use quantifier alternation purely as a measure of locality. In that case, the main concern is not to limit the individual processing power of the nodes, but rather to restrict the certificates in a way that preserves the local nature of the arbitrating algorithm. As explained in Section 1.3, the three alternation hierarchies based on **LD** do not meet this requirement, since they allow certificate sizes to depend on the entire input graph.

It turns out that we can generalize the local-polynomial hierarchy without compromising its potential as a measure of locality, simply by replacing polynomial bounds with arbitrary bounds. This leads us to define the class **LB** (for *local-bounded time*), which consists of all graph properties that can be decided by a distributed Turing machine operating under locally unique identifiers and running in constant round time and arbitrary step time (i.e., step time bounded by some arbitrary computable function). Based on this, we obtain the *local-bounded hierarchy* $\{\mathbf{\Sigma}_\ell^{\mathrm{LB}}, \mathbf{\Pi}_\ell^{\mathrm{LB}}\}_{\ell \in \mathbb{N}}$, which is defined analogously to the local-polynomial hierarchy, but where the certificate assignments are (r, f)-bounded for some arbitrary computable function $f : \mathbb{N} \to \mathbb{N}$.

Most of our results carry over directly to this generalized setting. This includes all of our separation results, in particular the infiniteness of the hierarchy, since all these separations already hold on graphs of bounded structural degree, where the local-bounded and the local-polynomial hierarchies are equivalent. On arbitrary graphs, the local-bounded hierarchy even exhibits a "cleaner" structure, in the sense that it forms a strict linear order, while the local-polynomial hierarchy presumably does so only on graphs of bounded structural degree (see Proposition 35 and Remark 37). Moreover, if we generalize local-polynomial reductions to reductions computable in constant round time and arbitrary step time, then all our hardness and completeness results can be extended to the corresponding classes of the local-bounded hierarchy. This even holds for our distributed version of the Cook–Levin theorem (Theorem 19), although it would have to be proved directly instead of relying on descriptive complexity.[5] Indeed, descriptive complexity is the only aspect of this paper for which there does not seem to be a direct generalization to the local-bounded hierarchy. This is quite striking, given that Fagin's theorem (in its generalized form) has been our guide and a helpful tool throughout the paper. We used it first as an indicator of the robustness of the local-polynomial hierarchy, then as a shortcut to easily derive the first completeness result

---

[5] A similar observation can be made in the centralized setting: the standard proof of the Cook–Levin theorem already shows how to construct a Boolean formula that encodes the possible space-time diagrams of any given nondeterministic Turing machine whose running time is bounded by some known computable function. While not particularly useful for classical complexity theory, this implies, for instance, that Boolean satisfiability is **NEXPTIME**-complete under exponential-time reductions. (Note, however, that exponential-time reductions are not closed under composition, and that they allow us to reduce any problem in **2EXPTIME** to a problem in **EXPTIME**.)

for **NLP**, and finally as a bridge to the realm of logic and automata theory, where we proved most of our separation results. In a way, we lose Fagin's theorem by further generalizing the hierarchy, but the insights gained from it remain fully applicable.

─── **References** ───

1   Miklós Ajtai, Ronald Fagin, and Larry J. Stockmeyer. The closure of monadic NP. *J. Comput. Syst. Sci.*, 60(3):660–716, 2000. `doi:10.1006/jcss.1999.1691`.

2   Sanjeev Arora and Boaz Barak. *Computational Complexity - A Modern Approach.* Cambridge University Press, 2009. URL: `http://www.cambridge.org/catalogue/catalogue.asp?isbn=9780521424264`.

3   Alkida Balliu, Gianlorenzo D'Angelo, Pierre Fraigniaud, and Dennis Olivetti. What can be verified locally? *J. Comput. Syst. Sci.*, 97:106–120, 2018. `doi:10.1016/j.jcss.2018.05.004`.

4   Benedikt Bollig, Patricia Bouyer, and Fabian Reiter. Identifiers in registers - describing network algorithms with logic. In Mikolaj Bojanczyk and Alex Simpson, editors, *Foundations of Software Science and Computation Structures - 22nd International Conference, FOSSACS 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings*, volume 11425 of *Lecture Notes in Computer Science*, pages 115–132. Springer, 2019. `doi:10.1007/978-3-030-17127-8\_7`.

5   Stephen A. Cook. The complexity of theorem-proving procedures. In Michael A. Harrison, Ranan B. Banerji, and Jeffrey D. Ullman, editors, *Proceedings of the 3rd Annual ACM Symposium on Theory of Computing, May 3-5, 1971, Shaker Heights, Ohio, USA*, pages 151–158. ACM, 1971. `doi:10.1145/800157.805047`.

6   Reinhard Diestel. *Graph Theory, 5th Edition*, volume 173 of *Graduate texts in mathematics.* Springer, 2017. URL: `http://diestel-graph-theory.com`, `doi:10.1007/978-3-662-53622-3`.

7   Ronald Fagin. Generalized first-order spectra and polynomial-time recognizable sets. In Richard M. Karp, editor, *Complexity of Computation*, volume 7 of *SIAM-AMS Proceedings*, pages 43–73, 1974.

8   Laurent Feuilloley. Introduction to local certification. *Discret. Math. Theor. Comput. Sci.*, 23(3), 2021. `doi:10.46298/dmtcs.6280`.

9   Laurent Feuilloley and Pierre Fraigniaud. Survey of distributed decision. *Bulletin of the EATCS*, 119, 2016. URL: `http://eatcs.org/beatcs/index.php/beatcs/article/view/411`.

10 Laurent Feuilloley, Pierre Fraigniaud, and Juho Hirvonen. A hierarchy of local decision. *Theor. Comput. Sci.*, 856:51–67, 2021. `doi:10.1016/j.tcs.2020.12.017`.

11 Laurent Feuilloley and Juho Hirvonen. Local verification of global proofs. In Ulrich Schmid and Josef Widder, editors, *32nd International Symposium on Distributed Computing, DISC 2018, New Orleans, LA, USA, October 15-19, 2018*, volume 121 of *LIPIcs*, pages 25:1–25:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. `doi:10.4230/LIPIcs.DISC.2018.25`.

12 Pierre Fraigniaud. Distributed computational complexities: are you volvo-addicted or nascar-obsessed? In Andréa W. Richa and Rachid Guerraoui, editors, *Proceedings of the 29th Annual ACM Symposium on Principles of Distributed Computing, PODC 2010, Zurich, Switzerland, July 25-28, 2010*, pages 171–172. ACM, 2010. `doi:10.1145/1835698.1835700`.

13 Pierre Fraigniaud, Magnús M. Halldórsson, and Amos Korman. On the impact of identifiers on local decision. In Roberto Baldoni, Paola Flocchini, and Binoy Ravindran, editors, *Principles of Distributed Systems, 16th International Conference, OPODIS 2012, Rome, Italy, December 18-20, 2012. Proceedings*, volume 7702 of *Lecture Notes in Computer Science*, pages 224–238. Springer, 2012. `doi:10.1007/978-3-642-35476-2\_16`.

14 Pierre Fraigniaud, Amos Korman, and David Peleg. Towards a complexity theory for local distributed computing. *J. ACM*, 60(5):35:1–35:26, 2013. `doi:10.1145/2499228`.

15 Dora Giammarresi and Antonio Restivo. Recognizable picture languages. *IJPRAI*, 6(2&3):241–256, 1992. `doi:10.1142/S021800149200014X`.

16 Dora Giammarresi, Antonio Restivo, Sebastian Seibert, and Wolfgang Thomas. Monadic second-order logic over rectangular pictures and recognizability by tiling systems. *Inf. Comput.*, 125(1):32–45, 1996. `doi:10.1006/inco.1996.0018`.

17 Oded Goldreich. *Computational Complexity - A Conceptual Perspective*. Cambridge University Press, 2008. `doi:10.1017/CBO9780511804106`.

18 Mika Göös and Jukka Suomela. Locally checkable proofs in distributed computing. *Theory Comput.*, 12(1):1–33, 2016. `doi:10.4086/toc.2016.v012a019`.

19 Erich Grädel, Phokion G. Kolaitis, Leonid Libkin, Maarten Marx, Joel Spencer, Moshe Y. Vardi, Yde Venema, and Scott Weinstein. *Finite Model Theory and Its Applications*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2007. `doi:10.1007/3-540-68804-8`.

20 Lauri Hella, Matti Järvisalo, Antti Kuusisto, Juhana Laurinharju, Tuomo Lempiäinen, Kerkko Luosto, Jukka Suomela, and Jonni Virtema. Weak models of distributed computing, with connections to modal logic. *Distributed Comput.*, 28(1):31–53, 2015. `doi:10.1007/s00446-013-0202-3`.

21 John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. *Introduction to automata theory, languages, and computation, 3rd Edition*. Pearson international edition. Addison-Wesley, 2007.

22 Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller and James W. Thatcher, editors, *Proceedings of a symposium on the Complexity of Computer Computations, held March 20-22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, USA*, The IBM Research Symposia Series, pages 85–103. Plenum Press, New York, 1972. `doi:10.1007/978-1-4684-2001-2\_9`.

23 Amos Korman, Shay Kutten, and David Peleg. Proof labeling schemes. *Distributed Comput.*, 22(4):215–233, 2010. `doi:10.1007/s00446-010-0095-3`.

24 Antti Kuusisto. Modal logic and distributed message passing automata. In Simona Ronchi Della Rocca, editor, *Computer Science Logic 2013 (CSL 2013), CSL 2013, September 2-5, 2013, Torino, Italy*, volume 23 of *LIPIcs*, pages 452–468. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2013. `doi:10.4230/LIPIcs.CSL.2013.452`.

25 Leonid A. Levin. Universal sequential search problems (in Russian). *Problemy Peredachi Informatsii*, 9(3):115–116, 1973. URL: `http://mi.mathnet.ru/ppi914`.

26 Leonid Libkin. *Elements of Finite Model Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2004. URL: `http://www.cs.toronto.edu/%7Elibkin/fmt`, `doi:10.1007/978-3-662-07003-1`.

27 Nathan Linial. Locality in distributed graph algorithms. *SIAM J. Comput.*, 21(1):193–201, 1992. `doi:10.1137/0221015`.

28 Oliver Matz. Dot-depth, monadic quantifier alternation, and first-order closure over grids and pictures. *Theor. Comput. Sci.*, 270(1-2):1–70, 2002. `doi:10.1016/S0304-3975(01)00277-8`.

29 Oliver Matz, Nicole Schweikardt, and Wolfgang Thomas. The monadic quantifier alternation hierarchy over grids and graphs. *Inf. Comput.*, 179(2):356–383, 2002. `doi:10.1006/inco.2002.2955`.

30 Albert R. Meyer and Larry J. Stockmeyer. The equivalence problem for regular expressions with squaring requires exponential space. In *13th Annual Symposium on Switching and Automata Theory, College Park, Maryland, USA, October 25-27, 1972*, pages 125–129. IEEE Computer Society, 1972. `doi:10.1109/SWAT.1972.29`.

31 Moni Naor and Larry J. Stockmeyer. What can be computed locally? *SIAM J. Comput.*, 24(6):1259–1277, 1995. `doi:10.1137/S0097539793254571`.

32 Christos H. Papadimitriou. *Computational complexity*. Addison-Wesley, 1994.

33 David Peleg. *Distributed Computing: A Locality-Sensitive Approach*, volume 5 of *SIAM Monographs on Discrete Mathematics and Applications*. Society for Industrial and Applied Mathematics (SIAM), 2000. `doi:10.1137/1.9780898719772`.

34 Fabian Reiter. Asynchronous distributed automata: A characterization of the modal mu-fragment. In Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl, editors, *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July*

10-14, 2017, Warsaw, Poland, volume 80 of *LIPIcs*, pages 100:1–100:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. `doi:10.4230/LIPIcs.ICALP.2017.100`.

35   Marcus Schaefer and Christopher Umans. Completeness in the polynomial time hierarchy - A compendium. *SIGACT News, Complexity Theory Column*, 33(3/4):32–49/22–36, 2002. URL: `https://ovid.cs.depaul.edu/documents/phcom.pdf`, `doi:10.1145/582475.582484`.

36   Larry J. Stockmeyer. The polynomial-time hierarchy. *Theor. Comput. Sci.*, 3(1):1–22, 1976. `doi:10.1016/0304-3975(76)90061-X`.

37   Jukka Suomela. Landscape of locality (invited talk). In Susanne Albers, editor, *17th Scandinavian Symposium and Workshops on Algorithm Theory, SWAT 2020, June 22-24, 2020, Tórshavn, Faroe Islands*, volume 162 of *LIPIcs*, pages 2:1–2:1. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. URL: `https://jukkasuomela.fi/landscape-of-locality`, `doi:10.4230/LIPIcs.SWAT.2020.2`.

38   Wolfgang Thomas. Languages, automata, and logic. In Grzegorz Rozenberg and Arto Salomaa, editors, *Handbook of Formal Languages, Volume 3: Beyond Words*, pages 389–455. Springer, 1997. `doi:10.1007/978-3-642-59126-6\_7`.

39   Eden Aldema Tshuva and Rotem Oshman. Brief announcement: On polynomial-time local decision. In Alessia Milani and Philipp Woelfel, editors, *PODC '22: ACM Symposium on Principles of Distributed Computing, Salerno, Italy, July 25 - 29, 2022*, pages 48–50. ACM, 2022. `doi:10.1145/3519270.3538463`.