

Uniform-Circuit and Logarithmic-Space Approximations of Refined Combinatorial Optimization Problems*

TOMOYUKI YAMAKAMI[†]

Abstract: A significant progress has been made in the past three decades over the study of combinatorial NP optimization problems and their associated optimization and approximate classes, such as NPO, PO, APX (or APXP), and PTAS. Unfortunately, a collection of problems that are simply placed inside the P-solvable optimization class PO never have been studiously analyzed regarding their exact computational complexity. To improve this situation, the existing framework based on polynomial-time computability needs to be expanded and further refined for an insightful analysis of various approximation algorithms targeting optimization problems within PO. In particular, we deal with those problems characterized in terms of logarithmic-space computations and uniform-circuit computations. We are focused on nondeterministic logarithmic-space (NL) optimization problems or NPO problems. Our study covers a wide range of optimization and approximation classes, dubbed as, NLO, LO, APXL, and LSAS as well as new classes NC^1O , $APXNC^1$, NC^1AS , and AC^0O , which are founded on uniform families of Boolean circuits. Although many NL decision problems can be naturally converted into NL optimization (NLO) problems, few NLO problems have been studied vigorously. We thus provide a number of new NLO problems falling into those low-complexity classes. With the help of NC^1 or AC^0 approximation-preserving reductions, we also identify the most difficult problems (known as complete problems) inside those classes. Finally, we demonstrate a number of collapses and separations among those refined optimization and approximation classes with or without unproven complexity-theoretical assumptions.

Keywords: optimization problem, approximation-preserving reduction, approximation algorithm, NC^1 circuit, AC^0 circuit, logarithmic space, complete problem

1 Refined Combinatorial Optimization Problems

1.1 NL Optimization Problems

Many combinatorial problems can be understood as sets of constraints (or requirements), which specify certain relations between *admissible instances* and *feasible solutions*. Of such problems, a *combinatorial optimization problem*, in particular, asks to find an “optimal” solution that satisfies certain constraints specified by each given admissible instance, where the optimality usually takes a form of either “maximization” or “minimization” according to a predetermined ordering over all feasible solutions. When finding such optimal solutions is costly, we often resort to look for solutions that are close enough to the desired optimal solutions. A significant progress had been made in a field of fundamental research on these combinatorial optimization problems during 1990s and its trend has continued promoting our understandings of the approximability of the problems. In particular, *NP optimization problems* (or *NPO problems*, in short) have been a centerfold of our interests because of their direct connection to NP (nondeterministic polynomial time) decision problems.

NPO problems are naturally derived from NP decision problems. As a typical NP problem, let us consider the *CNF Boolean formula satisfiability problem* (SAT) of determining whether a satisfying assignment exists for a given Boolean formula in conjunctive normal form. It is easy to convert SAT to its corresponding optimization problem, MAXIMUM WEIGHTED SATISFIABILITY, of finding a satisfying assignment having the maximal weight. This problem is an NPO problem. As is customary, the notation NPO also denotes the collection of such optimization problems. Of those NPO problems, those that can be solved exactly in polynomial time form a “tractable” optimization class PO, whereas an approximation class APX (which is hereafter denoted by APXP to emphasize its feature of “polynomial time” in comparison with “logarithmic space” and “circuits”) consists of NPO problems whose optimal solutions are *relatively approximated* within constant factors in polynomial time. Another optimization problem, MAXIMUM CUT, of finding a partition of a given graph into two disjoint sets that maximize the number of crossing edges falls into this approximation class APXP.

*A preliminary report appeared in the Proceedings of the 7th International Conference on Combinatorial Optimization and Applications (COCOA 2013), Chengdu, China, December 12–14, 2013, Lecture Notes in Computer Science, Springer-Verlag, vol.8287, pp.318–329, 2013.

[†]Present Affiliation: Graduate School of Engineering, University of Fukui, 3-9-1 Bunkyo, Fukui 910-8507, Japan

Up to now, a large number of NPO problems have been nicely classified into those classes of optimization problems (see, e.g., [6, Compendium]). Among those optimization and approximation classes, PO is the smallest class and has been proven to contain a number of intriguing optimization problems, including a minimization problem, MIN WEIGHT-ST-CUT, of finding a minimal s - t cut of a given directed graph. In a study on NPO problems, the use of *approximation-preserving reductions* helps us identify the most difficult optimization problems in a given class of optimization problems and many natural problems have been classified as the computationally hardest problems for NPO, PO, or APXP. Those problems are known as “complete” problems. MAXIMUM WEIGHTED SATISFIABILITY and MAXIMUM CUT are respectively proven to be complete for NLO and APXP.

The above classification of optimization problems is all described from a single viewpoint of “polynomial-time” computability and approximability and, as a result, a systematic discussion on optimization problems inside PO has been vastly neglected although PO contains numerous intriguing problems of various complexities. For instance, the *minimum path weight problem* (MIN PATH-WEIGHT) is to find in a given directed graph G a path $\mathcal{S} = (v_1, v_2, \dots, v_k)$ with $k \geq 2$ from given vertex v_1 to another vertex v_k having its (biased) path weight having binary representation of the form $\text{bin}(w(v_1))\text{bin}(w(v_2)) \cdots \text{bin}(w(v_k))$, where $\text{bin}(a)$ denotes the binary representation of a nonnegative integer a . This minimization problem MIN PATH-WEIGHT belongs to PO. Another example is the *maximum Boolean formula value problem* (MAX BFVP) of finding a maximal subset of a given set of Boolean formulas that are satisfied by a given truth assignment. This simple problem also resides inside PO; however, it apparently looks much easier to solve than MIN PATH-WEIGHT. This circumstantial evidence leads us to ponder that there might exist a finer and richer structure inside PO. Consequently, we may raise a natural question of whether it is possible to find such a finer structure within PO.

To achieve this goal, we first seek to develop a *new, finer framework*—a low-complexity world of optimization problems—and reexamine the computational complexity of such optimization problems within this new framework. For this purpose, we need to reshape the existing framework of expressing optimization complexity classes by clarifying the scope and complexity of verification processes used for solutions using objective (or measure) functions. While MIN WEIGHT-ST-CUT is known to be one of the most difficult problems in PO under P-reductions (even under NC^1 -reductions, shown in Proposition 3.2), the computational complexity of MIN PATH-WEIGHT seems to be significantly lower than MIN WEIGHT-ST-CUT residing in PO. To study the fine structures inside PO, we wish to shift our interest from a paradigm of polynomial-time optimization to much lower-complexity optimization, notably logarithmic-space or uniform-circuit optimization.

In the past decades, *logarithmic-space* (or *log-space*) computation has exhibited intriguing features, which are often different from those of polynomial-time computation. A notable result is the closure property of NL (nondeterministic logarithmic space) under complementation [22, 38].

Álvarez and Jenner [4, 5] first studied optimization problems from a viewpoint of log-space computability and discussed a class OptL of functions that compute optimal solutions using only a logarithmic amount of memory storage. In contrast, along the line of a study on NP optimization problems, Tantau [39] investigated *nondeterministic logarithmic-space (NL) optimization problems* or *NLO problems*. Intuitively, an NLO problem Q is asked to find its optimal solutions among all possible feasible solutions of size polynomial in input size n , provided that, (i) we can check, using only $O(\log n)$ memory space, whether any given solution candidate y is indeed a solution of the problem Q and, if so, (ii) we can calculate the objective value of y using $O(\log n)$ memory space. We simply write NLO for the collection of all NLO problems. It turns out that significant differences actually exist between two optimization classes NPO and NLO. One of the crucial differences is caused by the way that an underlying Turing machine produces its output strings on its output tape. When a log-space machine writes such a string, the machine must produce it *obliviously* because the output string is usually longer than the machine’s memory size. In short, log-space computation cannot remember polynomially-many symbols. As a result, unlike NPO problems, such machines do not seem to implement a typical approximation-preserving reduction between minimization problems and maximization problems inside NLO (see Section 4). When we discuss NLO problems, we need to heed the size of objective functions. An optimization problem is *polynomially bounded* if its objective (or measure) function outputs only polynomially-large integers.

Throughout this paper, we shall target those intriguing NLO problems. As unfolded in later sections, NLO problems occupy a substantial portion of PO and they include numerous important and natural problems. The aforementioned problems MIN PATH-WEIGHT and MAX BFVP are typical examples of the NLO problems. As other examples, the class NLO contains a restricted knapsack problem, called MAX 2BCU-KNAPSACK, and a restricted algebraic problem, called MAX AGEN (see Section 3.3 for their definitions). When we refer to PO, APXP, and PTAS in the existing framework based on NPO problems, we need to clarify their underlying framework; therefore, we intend to use new notations PO_{NPO} (instead of

PO), PTAS_{NPO} (instead of PTAS), and APXP_{NPO} (instead of APXP), when we discuss exact solvability and approximability of “NPO problems.”

1.2 Optimization Problems Inside NLO

By shifting the paradigm of optimization problems, we wish to look into a world of NLO problems and to unearth rich and complex structures underlying in this world. Of all NLO problems, those that can be *L-solvable* (i.e., solvable exactly by multi-tape deterministic Turing machines using logarithmic space) form an optimization class LO_{NLO} . If we restrict input graphs of MIN PATH-WEIGHT onto undirected forests, then the resulted problem, called $\text{MIN FOREST-PATH-WEIGHT}$, belongs to LO_{NLO} . Using uniform families of NC^1 -circuits and AC^0 -circuits in place of log-space Turing machines used in the existing notion of AP reduction, respectively, we can introduce two extra optimization classes $\text{NC}^1\text{O}_{\text{NLO}}$ and $\text{AC}^0\text{O}_{\text{NLO}}$, where NC^1 refers to $O(\log n)$ -depth polynomial-size circuits of bounded fan-in AND and OR gates and AC^0 indicates constant-depth polynomial-size circuits of unbounded fan-in AND and OR gates.

In analogy with APXP_{NPO} , another refined approximation class APXL_{NLO} is introduced using log-space approximation algorithms for NLO problems. Between APXL_{NLO} and LO_{NLO} exists a special class of optimization problems that have log-space approximation schemes. We call this class LSAS_{NLO} , similar to PTAS_{NPO} . In a similar way, we define $\text{APXNC}_{\text{NLO}}^k$, $\text{APXAC}_{\text{NLO}}^k$, $\text{NC}^k\text{AS}_{\text{NLO}}$, and $\text{AC}^k\text{AS}_{\text{NLO}}$ for each index $k \in \{0, 1\}$.

To compare the complexity of NLO problems, we consider *approximation-preserving (AP) reduction*, *exact (EX) reduction*, and *strong AP (sAP) reduction* using logarithmic space or by NC^1 -circuits (or even AC^0 -circuits). Using those weak reductions, we shall present in Section 3–4 a number of concrete optimization problems that are complete for the aforementioned refined classes of optimization problems. As discussed in Section 3.1, those weak reductions are necessary for low-complexity optimization problems, because strong reductions tend to obscure the essential characteristics of “complete” problems. Because of their fundamental nature, approximation classes are quite sensitive to the use of weak reductions. To use such reductions, we need to guarantee the existence of certain approximation bounds that must be easy to estimate.

Unlike NPO problems, a special attention is required for “complete” problems among NLO problems. Because of its logarithmic space-constraint, at this moment, it is unknown that complete problems actually exist in NLO. What we do know is the existence of complete problems for the class MaxNL of all maximization NLO problems (or the class MinNL of all minimization NLO problems) as shown in Section 3. More specifically, we manage to demonstrate that MIN PATH-WEIGHT is indeed complete for MinNL . A similar situation is observed also for APXL_{NLO} . In contrast, the class LO_{NLO} of L-solvable NLO problems possesses complete problems. When we limit our attention to polynomially-bounded NLO problems, each of NLO , APXL_{NLO} , LO_{NLO} , $\text{APXNC}_{\text{NLO}}^1$, and $\text{NC}^1\text{O}_{\text{NLO}}$ actually owns complete problems (Section 4).

Among the aforementioned refined classes, we shall also prove relationships concerning collapses and separations in Section 6. If we limit our optimization problems onto NLO, then PO_{NLO} , PTAS_{NLO} , APXP_{NLO} , and $\text{AC}^1\text{O}_{\text{NLO}}$ all coincide with NLO (Lemmas 2.2 and 6.2(1)). For polynomially-bounded NLO problems, in contrast, we can characterize them in terms of LO problems if their underlying log-space Turing machines are further allowed to access NL oracles. Following [39], $\text{L} \neq \text{NL}$ if and only if the polynomially-bounded subclasses of NLO , APXL_{NLO} , LSAS_{NLO} , and LO_{NLO} are all distinct (Theorem 6.5(1)). Similarly, we can show that $\text{NC}^1 \neq \text{L}$ if and only if the polynomially-bounded subclasses of LO_{NLO} , $\text{APXNC}_{\text{NLO}}^1$, $\text{NC}^1\text{AS}_{\text{NLO}}$, and $\text{NC}^1\text{O}_{\text{NLO}}$ are all different (Theorem 6.5(2)). For much lower-complexity optimization problems, we can separate $\text{AC}^0\text{O}_{\text{NLO}}$, $\text{AC}^0\text{AS}_{\text{NLO}}$, $\text{APXAC}_{\text{NLO}}^0$, and $\text{NC}^1\text{O}_{\text{NLO}}$ one from another (Theorem 6.6). Those separations directly follow from the well-known separation $\text{AC}^0 \neq \text{NC}^1$ [1, 15].

To help the reader overview intrinsic relationships among the aforementioned optimization (complexity) classes, we include Figure 1, which illustrates class containments and class separations obtained in Section 6. The last section provides a short list of open problems.

In the subsequent section, we shall provide a set of basic terminology on the approximation complexity of optimization problems.

2 Optimization and Approximation Preliminaries

We aim at *refining* an existing framework for studying combinatorial optimization problems of, in particular, low computational complexity. Throughout this paper, the notation \mathbb{N} denotes the set of all *natural numbers* (i.e., nonnegative integers) and \mathbb{N}^+ indicates $\mathbb{N} - \{0\}$. Moreover, \mathbb{Q} (resp., \mathbb{R}) indicates the set of all *rational numbers* (resp., *real numbers*). Two special notations $\mathbb{Q}^{>1}$ and $\mathbb{Q}^{\geq 1}$ respectively express the sets




Figure 1: A hierarchy of optimization and approximation classes. Single lines indicate simple class inclusions, whereas double lines indicate proper inclusions.

$\{q \in \mathbb{Q} \mid q > 1\}$ and $\{q \in \mathbb{Q} \mid q \geq 1\}$. Given two numbers $m, n \in \mathbb{N}$ with $m \leq n$, an *integer interval* $[m, n]_{\mathbb{Z}}$ is the set $\{m, m+1, m+2, \dots, n\}$. In the case of $[1, n]_{\mathbb{Z}}$ for $n \geq 1$, we abbreviate it as $[n]$. A (*multi-variate*) *polynomial* is always assumed to have nonnegative integer coefficients. We also assume that all logarithms are to *base 2*.

For any set A , $\mathcal{P}(A)$ denotes the *power set* of A , i.e., the set of all subsets of A . Given two sequences $\mathcal{S} = (s_1, s_2, \dots, s_k)$ and $\mathcal{S}' = (t_1, t_2, \dots, t_m)$, the notation $\mathcal{S} * \mathcal{S}'$ denotes a concatenated sequence $(s_1, s_2, \dots, s_k, t_1, t_2, \dots, t_m)$.

An *alphabet* is a finite nonempty set of “symbols” and a *string* (or a *word*) over alphabet Σ is a finite series of symbols taken from Σ . In particular, the *empty string* is denoted λ . Let $|x|$ denote the *length* of string x . The set Σ^* is composed of all strings over Σ and Σ^+ denotes $\Sigma^* - \{\lambda\}$. A *language* over Σ is a subset of Σ^* . Given two languages A and B , their *disjoint union* $A \oplus B$ is the set $\{0x \mid x \in A\} \cup \{1x \mid x \in B\}$. Given each number $n \in \mathbb{N}^+$ and $\Sigma = \{0, 1\}$, $\text{bin}(n)$ represents a string w in $1\Sigma^*$ ($= \{1x \mid x \in \Sigma^*\}$) that represents n in binary. Additionally, we set $\text{bin}(0) = \lambda$. For example, we obtain $\text{bin}(1) = 1$, $\text{bin}(2) = 10$, $\text{bin}(5) = 101$, and $\text{bin}(7) = 111$. Note that $|\text{bin}(n)| = \lceil \log_2(n+1) \rceil$ for every number $n \in \mathbb{N}^+$. By the contrary, for any string w in $\{\lambda\} \cup 1\Sigma^*$, $\text{rep}(w)$ denotes a positive integer satisfying $\text{bin}(n) = w$. For a number $n \in \mathbb{N}^+$, $\text{bin}(n)^{(-)}$ denotes the binary string obtained from $\text{bin}(n)$ by removing its first bit “1.” We also set $\text{bin}(0)^{(-)} = \lambda$. A function $f : \Sigma_1^* \rightarrow \Sigma_2^*$ (resp., $f : \Sigma_1^* \rightarrow \mathbb{R}^{\geq 0}$) for two alphabets Σ_1 and Σ_2 is *polynomially bounded* if there exists a polynomial p satisfying $|f(x)| \leq p(|x|)$ (resp., $f(x) \leq p(|x|)$) for all inputs $x \in \Sigma_1^*$.

A *directed graph* is a pair (V, E) for which V is a finite set of vertices and E is a binary relation on V and each element (a, b) in E is an edge. An *undirected graph* is similarly defined but E is required to be a symmetric relation, and we treat (a, b) and (b, a) in E equivalently. Given any graph $G = (V, E)$ with vertex set V and edge set E , a *path* of G is a sequence (v_1, v_2, \dots, v_k) of vertices in V satisfying that (v_i, v_{i+1}) is an edge for every index $i \in [k-1]$. Such a path is called *simple* exactly when there are no repeated vertices in it; that is, a simple path has no loop. The *length* of a path is the number of edges in it. A *tree* is an undirected connected graph with no cycle whereas a *forest* is an acyclic undirected graph. A *weighted graph* is a graph $G = (V, E)$ for which each edge (or vertex) has an associated weight given by a weight function $w : E \rightarrow \mathbb{R}$ (or $w : V \rightarrow \mathbb{R}$).

A *Boolean formula* is made up of (Boolean) variables and three logical connectives: \wedge (AND), \vee (OR), and \neg (NOT) in *infix notation*. Given a (Boolean) truth assignment σ , which maps $\{T, F\}$ to variables, a Boolean formula ϕ is said to be *satisfied* by σ if ϕ is evaluated to be true after σ assigns truth values to the variables in the formula.

Representation of Graphs, Matrices, Circuits, and Boolean Formulas: When we consider weak computations, it is often critical to choose what types of representation of input instances, such as graphs, matrices, circuits, and Boolean formulas. For example, as noted in [24], if we describe trees and forests using *bracketed expressions* as part of inputs, then the connectivity problem between two designated nodes in a given forest becomes solvable even on NC^1 -circuits. With respect to logarithmic-space computation, however, the representation of graphs via incidence matrices, adjacency matrices, or sets of ordered pairs are

all “equivalent” [27]. Unless otherwise specified, we assume that every graph is expressed by a *listing of its edge relation*, such as $\{(v_1, v_3), (v_2, v_5), (v_5, v_4)\}$; namely, all ordered pairs of vertices that define edges of a given graph. For isolated vertices, we list them as the names of those vertices, such as $\{(v_3), (v_6)\}$, instead of $\{(v_3, v_3), (v_6, v_6)\}$, which indicate self-loops in directed graphs. Boolean circuits are viewed as directed acyclic graphs. Boolean formulas are expressed in infix notation.[‡]

2.1 Basic Models of Computation

As a mechanical model of computation, we shall use the following basic form of (*multi-tape*) *deterministic Turing machine*. For the formal definition of Turing machine, refer to, e.g., [14, 21]. Our machine is equipped with a read-only input tape, multiple work tapes, and possibly an output tape. An input x of length n is given on the input tape, surrounded by two endmarkers: \clubsuit (left endmarker) and $\$$ (right endmarker) and all input tape cells are consecutively indexed by integers between 0 and $n + 1$, where \clubsuit is at cell 0 and $\$$ at cell $n + 1$. At any moment, a tape head working on the input tape either stays still on the same tape cell or moves to the left or the right. The *running time* (or *runtime*) of a Turing machine is the total number of steps (or moves) taken by the machine starting with the input, whereas its (*tape*) *space* is the maximum number of distinct tape cells visited by a tape head during the machine’s computation.

The behaviors of tapes and their tape heads are quite important in this paper; thus, we wish to pay our special attention to the following terminology. A tape is said to be *read-once* if it is a read-only tape and its tape head does not scan the same cell more than once; namely, it either stays at the same cell without reading any information (known as a λ -*move* or an ε -*move*) or moves instantly to the right cell. In contrast, a *write-only* tape indicates that, whenever its tape head writes a nonempty symbol in a tape cell, the head should move immediately to its right cell. In this paper, “output tapes” are always assumed to be write-only tapes. Turing machines with write-only output tapes are considered to compute (*multi-valued partial*) *functions*, by viewing strings left on the output tapes (when the machines halt) as “outputs.”

To describe low-complexity classes, we also use a notion of “random access” input tapes. In this mode, a machine is further equipped with an index tape and tries to write on this index tape a string of the form $\text{bin}(k)$ for a certain number $k \in [0, n + 1]_{\mathbb{Z}}$. Whenever the machine enters a specific inner state (an input-query state), the input-tape head jumps in a single step to the cell indexed by k and reads a symbol written in this particular cell. If k is not in the range $[0, n + 1]_{\mathbb{Z}}$, then M simply reads a blank symbol as an “out of range” symbol. Whenever we need to clarify a use of this special model, we refer to it as *random-access Turing machines*.

To express “nondeterminism” in our framework, we introduce a special tape called a *read-once auxiliary input tape* and equip Turing machines with such auxiliary tapes. An *auxiliary Turing machine* is the above-mentioned deterministic Turing machine equipped with an extra read-once auxiliary input tape on which a sequence of (nonblank) symbols (called an *auxiliary input*) is provided as an extra input (other than an ordinary input given on the input tape). In the rest of this paper, we shall understand that “auxiliary tapes” means read-only auxiliary input tapes unless otherwise stated. Such an auxiliary input given on the auxiliary input tape is surrounded by the two endmarkers. This machine can therefore read off two symbols (except for work-tape symbols) at once, one of which is from the input tape and the other from the auxiliary tape at each step in order to make a deterministic move. As our convention, when a tape head on the auxiliary tape reaches $\$$, the head must remain at this endmarker in the rest of a computation.

More formally, a $(k + 2)$ -tape auxiliary Turing machine M is a tuple $(Q, \Sigma, \{\clubsuit, \$\}, \Gamma, \Theta, \Phi, q_0, q_{acc}, q_{rej})$, where Q is a finite set of inner states, Σ is an input alphabet, Γ is a work alphabet, Θ is an auxiliary input alphabet, Φ is an output alphabet, q_0 is the initial state in Q , q_{acc} (resp., q_{rej}) is an accepting (resp., a rejecting) state in Q , and δ is a transition function from $(Q - \{q_{acc}, q_{rej}\}) \times (\Sigma \cup \{\clubsuit, \$, \lambda\}) \times \Gamma^k \times (\Theta \cup \{\$\})$ to $Q \times \Gamma^k \times (\Phi \cup \{\lambda\}) \times D \times D_1 \times \cdots \times D_k$, where D and each D_i ($i \in [k]$) are sets of head directions, $\{-1, 0, +1\}$, of an input tape and the i th work tape. Notice that, since tape heads on an auxiliary tape and an output tape move in one direction, we do not need to include their head directions.

We say that an auxiliary Turing machine *uses log space* if there exist two constants $a, b > 0$ for which, on every input x and every auxiliary input y , M uses the total of at most $a \log |x| + b$ cells of all work tapes (where an auxiliary input tape is not a work tape). Such a machine is succinctly called a *log-space auxiliary Turing machine*. Similarly, we define the notion of *polynomial-time auxiliary Turing machine*.

We assume that the reader is familiar with the foundation of computational complexity theory, in particular, the definitions and properties of those fundamental classes. See, e.g., [14] for their fundamental properties. The complexity class P (deterministic polynomial time) is composed of all decision problems (or

[‡]Boolean formulas in infix notation are defined inductively as follows: (i) 0 and 1 are Boolean formulas and (ii) if α and β are Boolean formulas, then $(\neg\alpha)$, $(\alpha \vee \beta)$, and $(\alpha \wedge \beta)$ are Boolean formulas.

languages) solved by deterministic Turing machines in polynomial time, whereas L (deterministic logarithmic space) contains decision problems solved by log-space deterministic Turing machines. The notation FP (resp., FL) refers to a functional version of P (resp., L), provided that all functions in FL output only strings of size polynomial in the lengths of inputs. Thus, $FL \subseteq FP$ holds.

For later convenience, we denote by auxP (resp., auxL) the collection of all sets $A \subseteq \Sigma^* \times \Sigma^*$ over alphabet Σ for which there exist a polynomial p and a polynomial-time (resp., log-space) auxiliary Turing machine M such that, for every x and y , (i) $(x, y) \in A$ implies $|y| \leq p(|x|)$ and (ii) whenever $|y| \leq p(|x|)$, M accepts (x, y) iff $(x, y) \in A$, where y is given on M 's auxiliary tape. Their functional versions with polynomially-bounded outputs (i.e., the size of output strings is bounded from above by a suitable polynomial in the input size) are denoted by auxFP (resp., auxFL). These classes auxP and auxL are respectively associated with nondeterministic classes NP and NL in the following fashion. Given a set $A \subseteq \Sigma^* \times \Sigma^*$ and any polynomial p , let $A_p = \{(x, y) \in A \mid |y| \leq p(|x|)\}$ and $A_p^\exists = \{x \in \Sigma^* \mid \exists y [(x, y) \in A_p]\}$. When p is clear from the context, we tend to drop subscript “ p ” and write A^\exists instead of A_p^\exists . The nondeterministic class NP (resp., NL) is composed of all languages of the form A_p^\exists for all $A \subseteq \Sigma^* \times \Sigma^*$ and all polynomials p satisfying $A_p \in \text{auxP}$ (resp., auxL). In other words, $A_p \in \text{auxP}$ (resp., auxL) if and only if $A_p^\exists \in \text{NP}$ (resp., NL).

In addition, the notation DLOGTIME is used to express the collection of all languages recognized by random-access Turing machines in $O(\log n)$ time. A function $f : \Sigma_1^* \rightarrow \Sigma_2^*$ is *DLOGTIME-computable* if the output size of f is polynomially bounded and the language $A_f = \{(x, i, b) \mid \text{the } i\text{th bit of } f(x) \text{ equals } b\}$ belongs to DLOGTIME.

In the subsequent sections, we shall concentrate mostly on functions and languages (which can be viewed as Boolean functions) whose domains are limited to certain subsets I of Σ^* (for alphabets Σ), and thus any given input to those functions and languages are always assumed, as a “promise,” to be taken from those domains I . Our functions and languages are therefore promise problems. To simplify our discussion in the later sections, however, we tend to treat those promise problems F as if they have no promise and we explicitly write, e.g., $F \in \text{FL}$ and $F \in \text{auxL}$ unless there is no confusion.

To describe circuit-based complexity classes, we use a standard notion of *Boolean circuits* (or just *circuits*), which is a labeled acyclic directed graph whose nodes of indegree 0 are called inputs and the other nodes are called gates. In our setting, a circuit is made up only of two basic gates *AND* and *OR* with inputs, which are labeled by *literals* (that is, either Boolean variables or their negations). A *fan-in* of a gate is the number of incoming edges. A fan-in is said to be *bounded* (resp., *unbounded*) if it is smaller than or equal to 2 (resp., it has no upper bound). The *size* of a circuit is the number of its nodes and the *depth* is the number of the longest path from an input to an output. A *family of circuits* is a set $\{C_n \mid n \in \mathbb{N}\}$, where each C_i is a Boolean circuit with n distinct variables.

There have been a number of uniformity notions proposed in the past literature, e.g., [7, 13, 37]. The different choice of uniformity endows circuit families with (possibly) different computational power. To explain such uniformity, we define the *direct connection language* of a circuit family $\{C_n\}_{n \in \mathbb{N}}$ as a set of all tuples $\langle t, a, b, y \rangle$, where a and b are numbers of nodes in C_n , b is a child of a , t is the type (e.g., literals, *AND*, *OR*, *NOT*, etc.) of a , and y is any string of length n . The *standard encoding* of C_n is a string, each symbol of which is of the form (a, t, b_L, b_R) , where a, b_L, b_R are gate numbers, b_L (resp., b_R) is the left (resp., right) child of a , and t is the type of a .

A family $\{C_n\}_{n \in \mathbb{N}}$ of Boolean circuits is *log-space uniform* (or *L-uniform*) if there exists a log-space deterministic Turing machine computing a function that maps 1^n to the standard encoding of C_n . We say that a family $\{C_n\}_{n \in \mathbb{N}}$ of Boolean circuits is *DLOGTIME-uniform*[§] if the directed connection language of $\{C_n\}_{n \in \mathbb{N}}$ can be recognized by a log-time random-access Turing machines. Other uniformity notions include *U_{E*}-uniformity* and *P-uniformity* [37]. For each $k \in \mathbb{N}$, NC^k (resp., AC^k) denotes the class of decision problems (or languages) solvable by DLOGTIME-uniform families of bounded (resp., unbounded) fan-in Boolean circuits of polynomial size and $O(\log^k n)$ depth. To refer to NC^1 of different uniformity, when clarification is necessary, we tend to describe it as “L-uniform NC^1 ” or “P-uniform NC^1 .” To describe their functional versions, we intentionally use the notations FAC^k and FNC^k , respectively. It is known that ALOGTIME (alternating logarithmic time) coincides with DLOGTIME-uniform NC^1 [9], which also equals NC^1 -uniform NC^1 [7].

Another characterization of NC^1 is given in [7] as follows. The *formula language* of a Boolean formula family $\{F_n\}_{n \in \mathbb{N}}$ is composed of all tuples $\langle c, i, y \rangle$ such that $|y| = n$ and the i th character of the n th formula F_n is c . A language A is in NC^1 iff there exists a family $\{F_n\}_{n \in \mathbb{N}}$ of Boolean formulas with depth $O(\log n)$ such that (i) for every x , $F_{|x|}(x)$ is true exactly when $x \in A$ and (ii) there exists a log-time deterministic Turing machine recognizes the formal language of $\{F_n\}_{n \in \mathbb{N}}$.

[§]As shown in [7, Theorem 9.1], this definition is equivalent to the one used in [9, 7] using formula languages.

Known inclusion relationships among the aforementioned complexity classes are shown as: $\text{NC}^0 \subsetneq \text{AC}^0 \subsetneq \text{TC}^0 \subseteq \text{NC}^1 \subseteq \text{L} \subseteq \text{NL} = \text{co-NL} \subseteq \text{AC}^1 \subseteq \text{NC}^2 \subseteq \text{P} \subseteq \text{NP}$. For more details, refer to, e.g., [14].

It is important to note that, on an output tape of a machine, a natural number is represented in binary, where the least significant bit is always placed at the right end of the output bits. In the rest of paper, a generic but informal term of “algorithm” will be often used to refer to either a deterministic Turing machine or a uniform family of circuits.

Lemma 2.1 *For n -bit numbers $x, y, x_i \in \mathbb{N}$ with $i \in [n]$, the operations $x + y$ and $\max\{0, x - y\}$ are in AC^0 , and $\lfloor x/y \rfloor$, $\sum_{i=1}^n x_i$, and $\prod_{i=1}^n x_i$ are in TC^0 [19, 20].*

2.2 Refined Optimization Problems

An *optimization problem* is simply a search problem, in which we are asked to look for a best possible feasible solution of the problem for each given admissible input. In the past literature, NP optimization problems have been a centerfold of the intensive study and low-complexity optimization problems have been mostly neglected except for [39]. To deal with those problems, we intend to refine the existing framework of NP optimization problems in terms of log-space and uniform-circuit computations.

In what follows, we shall formally introduce 14 different classes of refined combinatorial optimization problems, including 4 well-known classes NPO , PONPO , APXP_{NPO} , and PTAS_{NPO} , in order to justify the correctness of our definitions.

NPO and NLO. As a starting point of our study, we formally introduce *NP optimization problems* or *NPO problems* in the style of [6]. Since our purpose is to investigate low-complexity optimization problems, it is better for us to formulate a notion of NPO problems using auxiliary Turing machines instead of non-deterministic Turing machines. An NPO problem P is formally a quadruple $(I, \text{SOL}, m, \text{goal})$ whose entries satisfy the following properties.

- I is a finite set of *admissible instances*. There must be a deterministic Turing machine that recognizes I in polynomial time; that is, I belongs to P .
- SOL is a function mapping I to a collection of certain finite sets, where $\text{SOL}(x)$ is a set of *feasible solutions* of input instance x . There must be a polynomial q such that (i) for every $x \in I$ and every $y \in \text{SOL}(x)$, it holds that $|y| \leq q(|x|)$ and (ii) the set $I \circ \text{SOL} = \{(x, y) \mid x \in I, y \in \text{SOL}(x)\}$ is in auxP ; namely, $I \circ \text{SOL}$ is recognized in time polynomial in $|x|$ by a certain auxiliary Turing machine starting with x on an input tape and y on an auxiliary tape. By the definition of SOL , the set $\{x \in I \mid \text{SOL}(x) \neq \emptyset\}$ matches $(I \circ \text{SOL})_q^{\exists}$, and thus it belongs to NP .
- goal is either MAX or MIN . When $\text{goal} = \text{MAX}$, P is called a *maximization problem*; when $\text{goal} = \text{MIN}$, it is a *minimization problem*.
- m is a *measure function* (or an *objective function*) from $I \circ \text{SOL}$ to \mathbb{N}^+ whose value $m(x, y)$ is computed in time polynomial in $|x|$ by a certain auxiliary Turing machine starting with x written on an input tape and y on an auxiliary tape. Technically speaking, m is a promise problem; however, by abusing notations, we often express m as a member of auxFP (i.e., $m \in \text{auxFP}$). For any instance $x \in I$, $m^*(x)$ denotes the optimal value $\text{goal}\{m(x, y) \mid y \in \text{SOL}(x)\}$. Moreover, $\text{SOL}^*(x)$ expresses the set $\{y \in \text{SOL}(x) \mid m(x, y) = m^*(x)\}$ of optimal solutions of x .

Notice that, in polynomial time, an auxiliary Turing machine can copy any string y given on an auxiliary tape into its work tape and then manipulate it freely. This makes the read-once requirement of an auxiliary tape redundant. Therefore, the above definition logically matches the existing notion of NPO problems in, e.g., [6]. Let the notation NPO also express the class of all NPO problems.

A measure function m is called *polynomially bounded* if there exists a polynomial p such that $m(x, y) \leq p(|x|, |y|)$ holds for all pairs $(x, y) \in I \circ \text{SOL}$. An optimization problem is also said to be *polynomially bounded* if its measure function is polynomially bounded. For convenience, a succinct notation PBO indicates the collection of all optimization problems that are polynomially bounded.

To analyze the behaviors of low-complexity optimization problems, Tantau [39] formulated a notion of *NL optimization problems* (or *NLO problems*, in short), which are obtained simply by replacing the term “polynomial time” in the above definition of NPO problems with “logarithmic space.” For those NLO problems, the use of auxiliary Turing machine is essential and it may not be replaced by any Turing machine having no read-once auxiliary input tapes.

Here, we draw our attention to the read-once requirement posed on an auxiliary input tape. This requirement is quite severe for Turing machines. To see this fact, let us consider the following maximization

problem MAX WEIGHT-2SAT. In the *maximum weighted 2-satisfiability problem* (MAX WEIGHT-2SAT), we seek a truth assignment σ satisfying a given 2CNF formula on a set X of variables and a variable weight function $w : X \rightarrow \mathbb{N}^+$ such that the sum $\sum_{x \in X} \sigma(x)w(x) + 1$ must be maximized. Although its associated decision problem 2SAT, in which we are asked to decide whether a given 2CNF formula is satisfiable, is NL-complete (from a result of [27]), it is not clear whether MAX WEIGHT-2SAT belongs to NLO.

To express the class of all NLO problems, we use the notation of NLO. It follows that $\text{NLO} \subseteq \text{NPO}$. Moreover, MinNL (resp., MaxNL) denotes the class of all minimization (resp., maximization) problems in NLO; thus, NLO equals the union $\text{MinNL} \cup \text{MaxNL}$.

PO, LO, NCⁱO, and ACⁱO. We say that an NPO problem $P = (I, SOL, m, goal)$ is *P-solvable* if there exists a polynomial-time deterministic Turing machine M such that, for every instance $x \in I$, if $SOL(x) \neq \emptyset$, then M returns an optimal solution y in $SOL(x)$ and, otherwise, M returns “no solution” (or a designated symbol \perp). Moreover, the values $m(x, M(x)) (= m^*(x))$ must be computed in polynomial time from inputs x . As a result, the set $\{x \in I \mid SOL(x) \neq \emptyset\}$ must be in P. Given a class \mathcal{D} of optimization problems, the notation $\text{PO}_{\mathcal{D}}$ expresses the class of all optimization problems in \mathcal{D} that are P-solvable. Similarly, we can define the notations of $\text{LO}_{\mathcal{D}}$, $\text{NC}^i\text{O}_{\mathcal{D}}$, and $\text{AC}^i\text{O}_{\mathcal{D}}$ by replacing the term “P-solvable” with “L-solvable,” “NCⁱ-solvable,” and “ACⁱ-solvable,” respectively, for each index $i \in \mathbb{N}$. Conventionally, PO_{NPO} is written as PO and LO_{NLO} is noted briefly as LO in [39]. Notice that $\text{AC}^0\text{O}_{\mathcal{D}} \subseteq \text{NC}^1\text{O}_{\mathcal{D}} \subseteq \text{LO}_{\mathcal{D}} \subseteq \text{PO}_{\mathcal{D}}$ for any reasonable class \mathcal{D} .

It is important to note that, as in the case of LO_{NPO} , for example, when a problem P is L-solvable, its log-space algorithm, say, M that solves P does not need to check whether an input x given to M is actually admissible instance (i.e., $x \in I$), because such a task may be in general impossible for log-space machines. Hence, P is technically a promise problem and we normally allow M to behave arbitrarily on inputs outside of I or $I \circ SOL$.

APXP, APXL, APXNCⁱ, and APXACⁱ. Next, we shall define approximation classes using a notion of γ -approximation. Given an optimization problem $P = (I, SOL, m, goal)$, the *performance ratio* of solution y with respect to instance x is defined as

$$R(x, y) = \max \left\{ \left| \frac{m(x, y)}{m^*(x)} \right|, \left| \frac{m^*(x)}{m(x, y)} \right| \right\},$$

provided that neither $m(x, y)$ nor $m^*(x)$ is zero. Notice that $R(x, y) = 1$ iff $y \in SOL^*(x)$. Let $\gamma > 1$ be a constant indicating an upper bound of the performance ratio. With this constant γ , we say that P is *polynomial-time γ -approximable* if there exists a polynomial-time deterministic Turing machine M such that, for any instance $x \in I$, if $SOL(x) \neq \emptyset$, then $M(x) \in SOL(x)$ and $R(x, M(x)) \leq \gamma$; otherwise, $M(x)$ outputs “no solution” (or a symbol \perp); in addition, the values[¶] $m(x, M(x))$ must be computed in polynomial time from inputs x . Such a machine is referred to as a *γ -approximate algorithm*. The γ -approximability clearly implies that the set $\{x \in I \mid SOL(x) \neq \emptyset\}$ belongs to P. The notation $\text{APXP}_{\mathcal{D}}$ denotes a class consisting of problems P in class \mathcal{D} of optimization problems such that, for a certain fixed constant $\gamma > 1$, P is polynomial-time γ -approximable. Notice that APXP_{NPO} is conventionally expressed as APX (see, e.g., [6]).

Likewise, we define three extra notions of “log-space γ -approximation” [39], “NCⁱ γ -approximation,” and “ACⁱ γ -approximation” by replacing “polynomial-time Turing machine” in the above definition with “logarithmic-space (auxiliary) Turing machine,” “uniform family of NCⁱ-circuits,” and “uniform family of ACⁱ-circuits,” respectively, for every index $i \in \mathbb{N}$. We then introduce the notations of $\text{APXL}_{\mathcal{D}}$, $\text{APXNC}^i_{\mathcal{D}}$, and $\text{APXAC}^i_{\mathcal{D}}$ using “log-space γ -approximation,” “NCⁱ γ -approximation,” and “ACⁱ γ -approximation,” respectively. It follows that $\text{APXAC}^0_{\mathcal{D}} \subseteq \text{APXNC}^1_{\mathcal{D}} \subseteq \text{APXL}_{\mathcal{D}} \subseteq \text{APXP}_{\mathcal{D}}$ for any reasonable optimization/approximation class \mathcal{D} .

PTAS, LSAS, NCⁱAS, and ACⁱAS. A deterministic Turing machine M is called a *polynomial-time approximation scheme* (or a PTAS) if, for any “fixed constant” $r \in \mathbb{Q}^{>1}$, there exists a polynomial $p_r(n)$ such that, for every admissible instance $x \in I$, if $SOL(x) \neq \emptyset$, then M takes (x, r) as its input and outputs an r -approximate solution of x in time at most $p_r(|x|)$; otherwise, $M(x)$ outputs “no solution” (or a symbol \perp). Examples of such polynomial $p_r(n)$ are $\lceil \frac{r}{r-1} \rceil n^3$ and $n^{\lceil 1/(r-1) \rceil}$. Any approximation scheme is also a γ -approximate algorithm for any chosen constant $\gamma > 1$. The approximation class PTAS_{NPO} denotes a collection of all NPO problems that admit PTAS’s. In a similar manner, we can define a notion

[¶]The polynomial-time computability of the value $m(x, M(x))$ is trivial; however, the computability requirement for this value is quite important for the log-space computability and the NC¹ computability.

of *logarithmic-space approximation scheme* (or LSAS) and the associated approximation class LSAS_{NLO} by replacing “polynomial time” and “polynomial” with “logarithmic space” and “logarithmic function,” respectively.

The definitions of $\text{NC}^i\text{AS}_{\text{NLO}}$ and $\text{AC}^i\text{AS}_{\text{NLO}}$ are given essentially in the same way with a slight technical complication on uniformity condition. $\text{NC}^i\text{AS}_{\text{NLO}}$ (resp., $\text{AC}^i\text{AS}_{\text{NLO}}$) can be introduced using circuits of size $p_r(n)$ and depth $\ell_r(n)$ with bounded (resp., unbounded) fan-in gates, where $p_r(n)$ is a polynomial and $\ell_r(n)$ is a logarithmic function as long as r is treated as a fixed constant. Here, the uniformity requires $\text{DTIME}(\ell'_r(n))$ for another logarithmic function ℓ'_r with r being treated as a constant.

We have so far given 14 classes of optimization problems, which we shall discuss in details in the subsequent sections. Given an arbitrary nonempty class \mathcal{D} of optimization problems, it holds that $\text{NC}^1\text{O}_{\mathcal{D}} \subseteq \text{LO}_{\mathcal{D}} \subseteq \text{PO}_{\mathcal{D}}$ and $\text{APXNC}^1_{\mathcal{D}} \subseteq \text{APXL}_{\mathcal{D}} \subseteq \text{APXP}_{\mathcal{D}}$. It also follows that $\text{NC}^1\text{O}_{\mathcal{D}} \subseteq \text{NC}^1\text{AS}_{\mathcal{D}} \subseteq \text{APXNC}^1_{\mathcal{D}}$, $\text{LO}_{\mathcal{D}} \subseteq \text{LSAS}_{\mathcal{D}} \subseteq \text{APXL}_{\mathcal{D}}$, and $\text{PO}_{\mathcal{D}} \subseteq \text{PTAS}_{\mathcal{D}} \subseteq \text{APXP}_{\mathcal{D}}$. When $\mathcal{D} = \text{NLO}$, in particular, three classes APXP_{NLO} , PTAS_{NLO} , and PO_{NLO} coincide with NLO . Since the proof of this fact is short, we include it here.

Lemma 2.2 $\text{APXP}_{\text{NLO}} = \text{PTAS}_{\text{NLO}} = \text{PO}_{\text{NLO}} = \text{NLO}$.

Proof. Note that $\text{PO}_{\text{NLO}} \subseteq \text{PTAS}_{\text{NLO}} \subseteq \text{APXP}_{\text{NLO}}$. First, we claim that $\text{APXP}_{\text{NLO}} \subseteq \text{NLO}$. By the definition of APXP_{NLO} , all problems in APXP_{NLO} must be NLO problems, and hence they are in NLO .

Next, we show that $\text{NLO} \subseteq \text{PO}_{\text{NLO}}$. Let $P = (I, \text{SOL}, m, \text{goal})$ be any problem in NLO . Here, we consider only the case of $\text{goal} = \text{MAX}$ because the case of MIN is analogous. We want to show that P belongs to PO_{NLO} . Let x be any instance in I . Consider the following algorithm on input x . Here, we define $D = \{(x, y) \in I \circ \text{SOL} \mid \exists z \in \text{SOL}(x) [z \geq y \wedge m(x, z) \geq m(x, y)]\}$, where the notation \geq used for strings x and y is the lexicographic ordering. Note that $D \in \text{NL} \subseteq \text{P}$. Now, we can use a binary search technique using D to find a maximal solution $y \in \text{SOL}^*(x)$ in polynomial time. Therefore, we conclude that $\text{NLO} \subseteq \text{PO}_{\text{NLO}} \subseteq \text{PTAS}_{\text{NLO}} \subseteq \text{APXP}_{\text{NLO}} \subseteq \text{NLO}$. This implies the lemma. \square

Taking a slightly different approach toward a study on NPO problems, Krentel [32] introduced a class OptP of optimization functions. Let MaxP (resp., MinP) denote the class of all functions from Σ_1^* to Σ_2^* , each of which satisfies the following property: there exists a polynomial-time nondeterministic Turing machine M such that, for every input $x \in \Sigma_1^*$, $f(x)$ denotes the maximal (resp., minimal) string (in the lexicographic order) generated by M on x [31], where Σ_1 and Σ_2 are alphabets. The class OptP is simply defined as $\text{MaxP} \cup \text{MinP}$. We further define OptL in a similar way but using log-space nondeterministic Turing machines. Notice that Álvarez and Jenner [4] originally defined OptL as the set of *only* maximization problems and that we need to pay a special attention to their results whenever we apply them in our setting.

2.3 Approximation-Preserving Reductions

To compare the computational complexity of two optimization problems, we wish to use three types of reductions between those two problems. We follow well-studied reductions, known as *approximation-preserving (AP) reductions* and *exact (EX) reductions*. Given two optimization problems $P = (I_1, \text{SOL}_1, m_1, \text{goal})$ and $Q = (I_2, \text{SOL}_2, m_2, \text{goal})$, P is *polynomial-time AP-reducible* (or more conveniently, *APP-reducible*) to Q , denoted $P \leq_{\text{AP}}^{\text{P}} Q$, if there are two functions f and g and a constant $c \geq 1$ such that the following *APP-condition* is satisfied:

- for any instance $x \in I_1$ and any $r \in \mathbb{Q}^{>1}$, it holds that $f(x, r) \in I_2$,
- for any $x \in I_1$ and any $r \in \mathbb{Q}^{>1}$, if $\text{SOL}_1(x) \neq \emptyset$ then $\text{SOL}_2(f(x, r)) \neq \emptyset$,
- for any $x \in I_1$, any $r \in \mathbb{Q}^{>1}$, and any $y \in \text{SOL}_2(f(x, r))$, it holds that $g(x, y, r) \in \text{SOL}_1(x)$,
- $f(x, r)$ is computed by a deterministic Turing machine and $g(x, y, r)$ is computed by an auxiliary Turing machine, both of which run in time polynomial in $(|x|, |y|)$ for any $(x, y) \in I \circ \text{SOL}$ and any number $r \in \mathbb{Q}^{>1}$, and
- for any $x \in I_1$, any $r \in \mathbb{Q}^{>1}$, and any $y \in \text{SOL}_2(f(x, r))$, $R_2(f(x, r), y) \leq r$ implies $R_1(x, g(x, y, r)) \leq 1 + c(r - 1)$, where R_1 and R_2 respectively express the performance ratios for P_1 and P_2 .

Notice that the above APP-condition makes us concentrate only on instances of $\{x \in I \mid \text{SOL}(x) \neq \emptyset\}$ and that, for other instances x , we might possibly set the value $f(x, r)$ arbitrarily (as long as $x \in I_1$ iff $f(x, r) \in I_2$). When this APP-condition holds, we also say that P *APP-reduces* to Q . The triplet (f, g, c) is called a *polynomial-time AP-reduction* (or an *APP-reduction*) from P to Q . For more details, refer to, e.g., [6].

To discuss optimization problems within PO_{NPO} , we further need to introduce another type of reduction (f, g) , in which g “exactly” transforms in polynomial time an optimal solution for Q to another optimal solution for P so that “ $Q \in \text{PO}_{\text{NPO}}$ ” directly implies “ $P \in \text{PO}_{\text{NPO}}$.” We write $P \leq_{\text{EX}}^{\text{P}} Q$ when the following *EX-condition* holds:

- for any instance $x \in I_1$, it holds that $f(x) \in I_2$,
- for any $x \in I_1$, if $\text{SOL}_1(x) \neq \emptyset$ then $\text{SOL}_2(f(x)) \neq \emptyset$,
- for any $x \in I_1$ and any $y \in \text{SOL}_2(f(x))$, it holds that $g(x, y) \in \text{SOL}_1(x)$,
- $f(x)$ is computed by deterministic Turing machine and $g(x, y)$ is computed by an auxiliary Turing machine, both of which run in time polynomial in $(|x|, |y|)$, and
- for any $x \in I_1$ and any $y \in \text{SOL}_2(f(x))$, $R_2(f(x), y) = 1$ implies $R_1(x, g(x, y)) = 1$, where R_1 and R_2 respectively express the performance ratios for P_1 and P_2 .

The above pair (f, g) is called a *polynomial-time EX-reduction* (or an *EXP-reduction*) from P to Q .

It is quite useful to introduce a notion that combines both $\leq_{\text{AP}}^{\text{P}}$ and $\leq_{\text{EX}}^{\text{P}}$. Let us define the notion of *polynomial-time strong AP-reduction* (strong APP-reduction or sAPP-reduction), denoted $\leq_{\text{sAP}}^{\text{P}}$, obtained from $\leq_{\text{AP}}^{\text{P}}$ by allowing r (used in the above definition of APP-reduction) to be chosen from $\mathbb{Q}^{\geq 1}$ (instead of $\mathbb{Q}^{> 1}$).

Next, we weaken the behaviors of polynomial-time (strong) APP-reductions by modifying the “polynomial-time” requirement imposed on the aforementioned definition of (strong) APP-condition. When we replace “polynomial-time” by “logarithmic-space,” “uniform family of NC^1 -circuits,” and “uniform family of AC^0 -circuits,” we respectively obtain the corresponding notions of (*strong*) *APL-reduction* ($\leq_{\text{AP}}^{\text{L}}$, $\leq_{\text{sAP}}^{\text{L}}$), (*strong*) *APNC¹-reduction* ($\leq_{\text{AP}}^{\text{NC}^1}$, $\leq_{\text{sAP}}^{\text{NC}^1}$), and (*strong*) *APAC⁰-reduction* ($\leq_{\text{AP}}^{\text{AC}^0}$, $\leq_{\text{sAP}}^{\text{AC}^0}$). Notice that the notion of error-preserving reduction (or E-reduction), which was used in [39], essentially matches sAPL-reduction. Likewise, we define *EXL-reduction* ($\leq_{\text{EX}}^{\text{L}}$), *EXNC¹-reduction* ($\leq_{\text{EX}}^{\text{NC}^1}$), and *EXAC⁰-reduction* ($\leq_{\text{EX}}^{\text{AC}^0}$) from EXP-reduction.

The following two lemmas are immediate from the definition of sAP-reductions and we omit their proofs.

Lemma 2.3 *For any two reduction type $e_1, e_2 \in \{\text{P}, \text{L}, \text{NC}^1, \text{AC}^0\}$, if $e_1 \subseteq e_2$ (seen as complexity classes), then $P_1 \leq_{\text{sAP}}^{e_1} P_2$ implies $P_1 \leq_{\text{sAP}}^{e_2} P_2$. The same statement holds for \leq_{AP}^e and \leq_{EX}^e .*

Lemma 2.4 *For any reduction type $e \in \{\text{P}, \text{L}, \text{NC}^1, \text{AC}^0\}$, $P_1 \leq_{\text{sAP}}^e P_2$ implies both $P_1 \leq_{\text{AP}}^e P_2$ and $P_1 \leq_{\text{EX}}^e P_2$.*

In the next lemma, we shall present a useful property, called a *downward closure property*, for $\leq_{\text{AP}}^{\text{L}}$ - and $\leq_{\text{EX}}^{\text{L}}$ -reductions. A similar property holds also for $\leq_{\text{AP}}^{\text{NC}^1}$ - and $\leq_{\text{EX}}^{\text{NC}^1}$ -reductions.

Lemma 2.5 [*downward closure property*] *Let P and Q be any two optimization problems in NLO.*

1. *Let $\mathcal{D} \in \{\text{NLO}, \text{APXL}, \text{LSAS}, \text{APXNC}^1, \text{NC}^1\text{AS}\}$. If $P \leq_{\text{AP}}^{\text{NC}^1} Q$ and $Q \in \mathcal{D}_{\text{NLO}}$, then $P \in \mathcal{D}_{\text{NLO}}$, where NLO_{NLO} is understood as NLO.*
2. *Let $\mathcal{D} \in \{\text{LO}, \text{NC}^1\text{O}\}$. If $P \leq_{\text{EX}}^{\text{NC}^1} Q$ and $Q \in \mathcal{D}_{\text{NLO}}$, then $P \in \mathcal{D}_{\text{NLO}}$.*

An immediate consequence of Lemma 2.5 is the following corollary. In comparison, by setting $\mathcal{D} \in \{\text{NLO}, \text{APXL}, \text{LSAS}, \text{LO}\}$, for any $P, Q \in \text{NLO}$, if $P \leq_{\text{sAP}}^{\text{L}} Q$ and $Q \in \mathcal{D}_{\text{NLO}}$, then $P \in \mathcal{D}_{\text{NLO}}$ [39].

Corollary 2.6 *Let $\mathcal{D} \in \{\text{NLO}, \text{APXL}, \text{LSAS}, \text{LO}\}$. For any $P, Q \in \text{NLO}$, if $P \leq_{\text{sAP}}^{\text{NC}^1} Q$ and $Q \in \mathcal{D}_{\text{NLO}}$, then $P \in \mathcal{D}_{\text{NLO}}$, where NLO_{NLO} is NLO.*

Here, we shall briefly give the proof of Lemma 2.5.

Proof of Lemma 2.5. Take any two optimization problems $P = (I_1, \text{SOL}_1, m_1, \text{goal}_1)$ and $Q = (I_2, \text{SOL}_2, m_2, \text{goal}_2)$ in NLO. In what follows, we shall prove only the case of $\mathcal{D} = \text{APXNC}_{\text{NLO}}^1$, because the other cases can be similarly treated.

(1) Assume that $P \leq_{\text{sAP}}^{\text{NC}^1} Q$ via an APNC¹-reduction (f, g, c) and that Q is in $\text{APXNC}_{\text{NLO}}^1$. Given any constant $r' > 1$, let $C_{r'}$ be an NC¹ r' -approximate algorithm solving Q . To show that $P \in \text{APXNC}_{\text{NLO}}^1$, it suffices to construct, for each constant $r > 1$, an appropriate NC¹ circuit, say, D_r that finds r -approximate solutions for P .

Given a constant $r > 1$, let us define $r' = 1 + (r - 1)/c > 1$ and consider $C_{r'}$. Since $C_{r'}$ is an NC¹ r' -approximate algorithm, it follows that the performance ratio R_2 for $C_{r'}$ satisfies $R_2(z, C_{r'}(z)) \leq r'$ for any $z \in I_2$. Next, we define the desired algorithm N_r as follows: on input $x \in I_1$, compute simultaneously

$z = f(x, r)$ and $y = C_{r'}(z)$ and then output $g(x, y, r')$. Since $R_2(z, C_{r'}(z)) \leq r'$, it follows by the definition of $\leq_{\text{AP}}^{\text{NC}^1}$ that $R_1(x, N_r(x)) = R_1(x, g(x, C_{r'}(z), r')) \leq 1 + c(r' - 1) = r$. Hence, N_r is an r -approximate algorithm for P .

We still need to show that N_r can be realized by an NC^1 -circuit. For this purpose, we prepare an NC^1 circuit C_f that, on input $(x, e) \in \Sigma^* \times 1\{0, 1\}^*$, outputs the $\text{rep}(e)$ -th bit of $f(x, r)$. Notice that $|f(x, r)|$ is polynomially bounded. Moreover, let C_g denote an NC^1 circuit computing g . We construct an NC^1 -circuit M' that, on input $(x, e) \in \Sigma^* \times 1\{0, 1\}^*$, computes the $\text{rep}(e)$ -th bit of $C_{r'}(f(x, r'))$. During this procedure, whenever $C_{r'}$ tries to access the j th bit of $f(x, r')$, we run C_f on $(x, \text{bin}(j))$. The desired algorithm N_r is executed as follows. We first run C_g using the first and third input tapes for x and r' and leaving the second tape blank. Whenever C_g tries to access the i th bit y_i of $y = C_{r'}(f(x, r'))$, we run M' on $(x, \text{bin}(i))$. It is not difficult to show that this procedure can be implemented on an appropriate NC^1 -circuit.

(2) Assume that $P \leq_{\text{EX}}^{\text{NC}^1} Q$ via an $\leq_{\text{EX}}^{\text{NC}^1}$ -reduction (f, g) with $Q \in \text{NC}^1\text{O}_{\text{NLO}}$. Since $Q \in \text{NC}^1\text{O}_{\text{NLO}}$, there exists an NC^1 circuit M for which $M(x) \in \text{SOL}_2(x)$ and $R_2(x, M(x)) = 1$ for any $x \in I_2$. To show that P is in $\text{NC}^1\text{O}_{\text{NLO}}$, let us consider the following algorithm N . On input x , compute $y = M(f(x))$ and output $w = g(x, y)$. For a similar reason to (1), N can be implemented by a certain NC^1 circuit. Since $R_1(x, N(x)) = R_1(x, g(x, y))$, by the definition of an $\leq_{\text{EX}}^{\text{NC}^1}$ -reduction, we obtain $R_2(f(x), M(f(x))) = 1$. Therefore, N exactly solves P . \square

Our AP-, EX-, and sAP-reductions can help us identify the most difficult problems in a given optimization/approximation class. Such problems are generally called ‘‘complete problems,’’ which have played a crucial role in understanding the structural features of optimization and approximation classes.

Formally, let \leq be any reduction discussed in this section, and let \mathcal{D} be any class of optimization problems. An optimization problem P is called \leq -hard for \mathcal{D} if, for every problem Q in \mathcal{D} , $Q \leq P$ holds. Moreover, P is said to be \leq -complete for \mathcal{D} if P is in \mathcal{D} and it is \leq -hard for \mathcal{D} . This completeness will be a central subject in Sections 3–4.

3 General Complete Problems

Complete problems represent a certain structure of a given optimization or approximation class and they provide useful insights into specific features of the class. To develop a coherent theory of NLO problems, it is essential to study such complete problems. In the subsequent subsections, we shall present numerous complete problems for various optimization and approximation classes.

3.1 Why APNC¹- and EXNC¹-Reductions?

To discuss complete problems for refined optimization and approximation classes under certain reductions, it is crucial to choose reasonable types of reductions. By Lemma 2.3, for example, any $\leq_{\text{AP}}^{\text{NC}^1}$ -complete problem for an optimization/approximation class \mathcal{D} is also $\leq_{\text{AP}}^{\text{L}}$ -complete, but the converse may not be true in general. In what follows, we briefly argue that the $\leq_{\text{AP}}^{\text{L}}$ - and $\leq_{\text{EX}}^{\text{L}}$ -reductions are so powerful that all problems in APXL_{NLO} and LO_{NLO} respectively become reducible to similar problems residing even in $\text{APXAC}_{\text{NLO}}^0 \cap \text{PBO}$ and $\text{AC}^0\text{O}_{\text{NLO}} \cap \text{PBO}$.

Proposition 3.1 1. $\text{APXL}_{\text{NLO}} = \{P \in \text{NLO} \mid \exists Q \in \text{APXAC}_{\text{NLO}}^0 \cap \text{PBO} [P \leq_{\text{sAP}}^{\text{L}} Q]\}$.
 2. $\text{LO}_{\text{NLO}} = \{P \in \text{NLO} \mid \exists Q \in \text{AC}^0\text{O}_{\text{NLO}} \cap \text{PBO} [P \leq_{\text{EX}}^{\text{L}} Q]\}$.

Proof. (1) This claim is split into two opposite containments.

(\supseteq) Let $P \in \text{NLO}$ and $Q \in \text{APXAC}_{\text{NLO}}^0$, and assume that $P \leq_{\text{sAP}}^{\text{L}} Q$. Notice that Q also belongs to APXL_{NLO} . Lemma 2.5(1) therefore implies that $P \in \text{APXL}_{\text{NLO}}$.

(\subseteq) Since $\text{APXL}_{\text{NLO}} = \text{APXL}_{\text{MaxNL}} \cup \text{APXL}_{\text{MinNL}}$, we first consider the case of $\text{APXL}_{\text{MaxNL}}$. Take any maximization problem $P = (I_1, \text{SOL}_2, m_1, \text{MAX})$ in $\text{APXL}_{\text{MaxNL}}$. We want to define a new maximization problem $Q = (I_2, \text{SOL}_2, m_2, \text{MAX})$ and show that $Q \in \text{APXAC}_{\text{NLO}}^0$ and $P \leq_{\text{sAP}}^{\text{L}} Q$.

Since $P \in \text{APXL}_{\text{NLO}}$, there exists a constant $e \in \mathbb{Q}^{>1}$ and a log-space deterministic Turing machine M that produces e -approximate solutions of P ; namely, the performance ratio R_1 of M 's outcome for P satisfies $R_1(x, M(x)) \leq e$ for every $x \in (I_1 \circ \text{SOL}_1)^\exists$. First, we set I_2 to be composed of all instances of the form $(x, M(x))$ for $x \in I_1$. Notice that, whenever $\text{SOL}_1(x) = \emptyset$, M outputs the designated symbol \perp . Since M uses only log space, $I_2 \in \text{L}$ follows. Next, we define $\text{SOL}_2(x, y) = \text{SOL}_1(x)$ and $m_2((x, y), z) = m_1(x, z)$ for any $x \in I_1$ and $y, z \in \text{SOL}_1(x)$. By those definitions, Q is a problem in NLO .

Let us consider an AC^0 -circuit that outputs y on admissible instance (x, y) in $(I_2 \circ SOL_2)^\exists$. For any $(x, y) \in (I_2 \circ SOL_2)^\exists$, it follows that $R_2((x, y), C(x, y)) = \frac{m_2^*(x, y)}{m_1((x, y), C(x, y))} = \frac{m_1^*(x)}{m_1(x, y)} = R_1(x, M(x)) \leq e$, where R_2 means the performance ratio for Q . Hence, $C(x, y)$ is an e -approximate solution of Q . Thus, Q belongs to $APXAC_{NLO}^0$.

Next, we want to show that $P \leq_{sAP}^L Q$ via $(f, g, 1)$. Take any number $r \in \mathbb{Q}^{\geq 1}$ and define $f(x, r) = (x, M(x))$ and $g((x, y), z, r) = z$ for $x \in I_1$ and $y, z \in SOL_1(x)$. It follows that $R_2(f(x, r), z) = \frac{m_1^*(x)}{m_1(x, z)} = \frac{m_1^*(x)}{m_1(x, g((x, y), z, r))} = R_1(x, g((x, y), z, r))$. Since f is in FL and g is in FAC^0 , P indeed $sAPL$ -reduces to Q . Because P is arbitrary, we conclude that every maximization problem in $APXL_{NLO}$ is \leq_{sAP}^L -reducible to Q .

In a similar fashion, we can show that every minimization problem P' in $APXL_{MinNL}$ can be reduced to a certain minimization problem Q' in $APXAC_{NLO}^0$.

(2) This claim can be proven in a similar way to (1).

(\supseteq) Take two optimization problems $P \in NLO$ and $Q \in AC^0_{ONLO}$. Assume that $P \leq_{EX}^L Q$ via (f, g) . Lemma 2.5(2) then ensures that P belongs to LO_{NLO} .

(\subseteq) We begin with the case of LO_{MaxNL} . Let $P = (I_1, SOL_1, m_1, MAX)$ be an arbitrary problem in LO_{MaxNL} and take a deterministic Turing machine M that solves P using log space. We intend to construct another problem $Q = (I_2, SOL_2, m_2, MAX)$ so that P is \leq_{EX}^L -reducible to Q . For this desired problem Q , we set $I_2 = I_1 \circ SOL_1$ and $SOL_2(x, y) = \{y\}$ for any $(x, y) \in I_2$. The measure function m_2 is defined as $m_2((x, y), z) = 2$ if $y = z$, and 1 otherwise. Obviously, m_2 is polynomially bounded and is in FAC^0 . For a particular input $(x, M(x))$, since $m_2((x, M(x)), M(x)) = m_2^*(x, M(x))$, we obtain $M(x) \in SOL_2^*(x, M(x))$. Thus, it follows that $Q \in AC^0_{ONLO}$. Finally, we define a reduction (f, g) as $f(x) = (x, M(x))$ and $g(x, z) = M(x)$ for any $z \in SOL_2(f(x))$. Note that $R_2(f(x), z) = 1$ implies $z = M(x)$, and thus the performance ratio $R_1(x, g(x, z))$ for P satisfies $R_1(x, g(x, M(x))) = R_1(x, M(x)) = 1$. Therefore, $(f, g) \leq_{EX}^L$ -reduces P to Q . \square

Proposition 3.1 suggests that the notions of \leq_{sAP}^L - and \leq_{EX}^L -completeness do not capture the essential difficulty of the optimization complexity class $APXL_{NLO}$ and LO_{NLO} . Therefore, in what follows, we intend to use weaker types of reductions. In particular, we limit our interest within $\leq_{sAP}^{NC^1}$ -reductions and $\leq_{EX}^{NC^1}$ -reductions.

As a quick example of $\leq_{EX}^{NC^1}$ -complete problems, let us consider the *minimum weighted s - t cut problem* (MIN WEIGHT-ST-CUT), which is to find an s - t cut of a given weighted directed graph so that the (*weighted*) *capacity* of the cut (i.e., the total weight of edges from S_0 to S_1) is minimized, where an s - t cut for two distinct vertices $s, t \in V$ is a partition (S_0, S_1) of the vertices for which $s \in S_0$ and $t \in S_1$. We represent this cut (S_0, S_1) by an assignment σ from V to $\{0, 1\}$ satisfying the following condition: for every $v \in V$ and every $i \in \{0, 1\}$, $\sigma(v) = i$ iff $v \in S_i$.

MINIMUM WEIGHTED S - T CUT PROBLEM (MIN WEIGHT-ST-CUT):

- INSTANCE: a directed graph $G = (V, E)$, two distinguished vertices $s, t \in V$, where s is a *source* and t is a *sink* (or a *target*), and an edge weight function $c : E \rightarrow \mathbb{N}^+$.
- SOLUTION: an s - t cut (S_0, S_1) , specified by an assignment $\sigma : V \rightarrow \{0, 1\}$ as described above.
- MEASURE: the (*weighted*) *capacity* of the s - t cut (i.e., $\sum_{(v, w) \in E \wedge v \in S_0 \wedge w \in S_1} c(v, w)$).

Note that the capacity of any s - t cut is at most $\max_{e \in E} \{c(e)\} |E|$. It is possible to prove that MIN WEIGHT-ST-CUT is $\leq_{EX}^{NC^1}$ -complete for PO_{NPO} .

Proposition 3.2 MIN WEIGHT-ST-CUT is $\leq_{EX}^{NC^1}$ -complete for PO_{NPO} .

The proof of Proposition 3.2 can be obtained by an appropriate modification of the P-completeness proof of Goldschlager et al. [17] for the “decision version” of the *maximum s - t flow problem*. The proof of Proposition 3.2 is placed in Appendix for readability. The proposition will be used in Section 6.

3.2 Complete Problems Concerning Path Weight

We have seen in Section 3.1 the importance of $\leq_{sAP}^{NC^1}$ - and $\leq_{EX}^{NC^1}$ -reductions for discussing the computational complexity of our refined optimization problems. In this and the next subsections under those special reductions, we shall present a few complete problems for various optimization and approximation classes.

There are two categories of NLO problems to distinguish in our course of studying the complexity of NLO problems. The first category contains NLO problems $(I, SOL, m, goal)$ for which the set $(I \circ SOL)^\exists$ ($= \{x \in I \mid SOL(x) \neq \emptyset\}$) belongs to NL but may not fall into L unless $L = NL$. The second category, in

contrast, requires the set $(I \circ SOL)^\exists$ to be in L. Many of the optimization problems of the first category are unlikely to fall into $APXL_{NLO}$ or L_{NLO} .

First, we shall look into an optimization analogue of the well-known *directed s-t connectivity problem* (also known as the *graph accessibility problem* and the *graph reachability problem* in the past literature), denoted by DSTCON, in which, for any directed graph $G = (V, E)$ and two vertices $s, t \in V$, we are asked to determine whether there is a path from s to t in G . Earlier, Jones [26] showed that DSTCON is NL-complete under \leq_m^L (log-space many-one) reductions. These reductions can be replaced by appropriate $\leq_m^{NC^1}$ -reductions, and thus DSTCON becomes $\leq_m^{NC^1}$ -complete for NL. Let us consider a series of problems associated with minimum path weights of graphs. First, recall the minimum path weight problem (MIN PATH-WEIGHT) introduced in Section 1.

MINIMUM PATH WEIGHT PROBLEM (MIN PATH-WEIGHT):

- INSTANCE: a directed graph $G = (V, E)$, two distinguished vertices $s, t \in V$, and a (vertex) weight function $w : V \rightarrow \mathbb{N}$.
- SOLUTION: a path $\mathcal{S} = (v_1, v_2, \dots, v_k)$ from s to t (i.e., $s = v_1$ and $t = v_k$).
- MEASURE: “biased” path weight $w(\mathcal{S}) = \max\{1, \text{rep}(\text{bin}(w(v_1))\text{bin}(w(v_2)) \cdots \text{bin}(w(v_k)))\}$.

In the above definition, we generally do not demand that s is a *source* (i.e., a node of indegree 0) and t is a *sink* (i.e., a node of outdegree 0) although such a restriction does not change the completeness of the problem.

Here, we need to remark that the choice of our measure function for MIN PATH-WEIGHT is quite artificial. As a quick example, if $\mathcal{S} = (v_1, v_2, v_3, v_4)$ with $w(v_1) = 3$, $w(v_2) = 0$, $w(v_3) = 2$, and $w(v_4) = 4$, then $w(\mathcal{S}) = \text{rep}(1110100)$ since $\text{bin}(0) = \lambda$ (the empty string). It is important to note that we use the *biased path weight* instead of a *standard path weight* defined as $\sum_{i \in [k]} w(v_i)$. This comes from the fact that, because log-space computation cannot store super-logarithmically many bits, it cannot sum up all super-logarithmically large weights of vertices. However, if we set all vertices of a given input graph have weights of exactly 1, then MIN PATH-WEIGHT is essentially identical to a problem of finding the “shortest” s - t path in the graph.

In comparison, we also define a polynomially-bounded form of MIN PATH-WEIGHT simply by demanding that $1 \leq w(v) \leq |V|$ for all $v \in V$ and by changing $w(\mathcal{S})$ to the *total path weight* $w'(\mathcal{S}) = \sum_{i=1}^k w(v_i)$. Notice that $w'(\mathcal{S}) \leq k|V|$. For our later reference in Section 4, we call this modified problem the *minimum bounded path weight problem* (MIN BPATH-WEIGHT) to emphasize the polynomially-boundedness of the problem.

Hereafter, we shall prove that MIN PATH-WEIGHT is $\leq_{sAP}^{AC^0}$ -complete for MinNL.

Theorem 3.3 MIN PATH-WEIGHT is $\leq_{sAP}^{AC^0}$ -complete for MinNL.

For the $\leq_{sAP}^{AC^0}$ -hardness part of Theorem 3.3, we want to introduce a useful notion of *configuration graph* of a log-space auxiliary Turing machine M on a given input x together with any possible auxiliary input y , which describes an entire computation tree of M working on x and y . This is a weighted directed graph, which will be used in later proofs, establishing the hardness of target optimization problems; however, in those proofs, we may need to modify the original configuration graph given below. Since each vertex of a configuration graph is labeled by a “partial configuration,” we first define such partial configurations of M on x . To simplify the following description, we consider the case where M has only one work tape.

Recall from Section 2.1 an auxiliary Turing machine $M = (Q, \Sigma, \{\phi, \$\}, \Gamma, \Theta, \Phi, q_0, q_{acc}, q_{rej})$ with its transition function δ mapping $(Q - \{q_{acc}, q_{rej}\}) \times (\Sigma \cup \{\phi, \$, \lambda\}) \times \Gamma \times (\Theta \cup \{\$\})$ to $Q \times \Gamma \times (\Phi \cup \{\lambda\}) \times D \times D_1$. The current tape situation is encoded into uhw , which indicates that the tape content is uv and the tape head is scanning the leftmost symbol of w , where h is a special symbol representing the tape heard. A *partial configuration* of M on input x is a tuple $v = \langle q, w, u, \tau, \xi, k \rangle$, which intuitively indicates a snapshot of M 's computation at time k ($k \in \mathbb{N}$) when q is an inner state, w is an encoding of M 's input tape, u is an encoding of M 's work tape, τ is a scanning auxiliary input symbol, and ξ is an output symbol or λ to write.

We connect each partial configuration $v = \langle q, w, u, \tau, \xi, k \rangle$ to others $\langle p, w', u', \tau', \xi', k+1 \rangle$ for all $\tau' \in \Gamma \cup \{\$\}$ by applying a transition “ $\delta(q, \sigma'_1, \sigma'_2, \tau) = (p, \sigma'_2, \xi', d_1, d_2)$,” where w' (resp., u') is an encoding of the input (resp., work) tape obtained from w (resp., u) by this transition. When a machine makes a λ -move on the output tape, we use the same symbol “ λ ” in place of ξ and ξ' . Here, we encode such partial configurations into binary strings of the same length by padding extra garbage bits (if necessary).

The *weight* of this vertex v is defined as ξ (expressed in binary). We can view a *computation path* y of M on x together with a series of nondeterministic choices of M , as a sequence of partial configurations. For two vertices u and v , (u, v) is a direct edge if, seen as partial configurations, v is obtained from u by a single

application of δ and a choice of auxiliary input symbol. Since each vertex is represented by $O(n)$ symbols, the total number of vertices is at most a polynomial in n . We denote by G_x^M the obtained configuration graph of M on x since M halts in polynomial time. Note that the size of G_x^M is bounded from above by a polynomial in the size of input instance x of M .

It is important to note that, from a given encoding of a computation path y , we can easily extract an associated auxiliary input, because each partial configuration in y contains a piece of information on the auxiliary input and M 's head on the auxiliary tape moves in only one direction.

Proof of Theorem 3.3. For notational convenience, in the following argument, MIN PATH-WEIGHT is expressed as $(I_0, SOL_0, m_0, \text{MIN})$. Firstly, we want to claim that $I_0 \in \text{L}$. This follows from the facts that $\text{DSTCON} \in \text{NL}$ and that $(I_0 \circ SOL_0)^\exists$ (more accurately, $(I_0 \circ SOL_0)_q^\exists$ for a suitable polynomial q) is essentially “equivalent” to DSTCON , except for the presence of a weight function w . Next, we claim that MIN PATH-WEIGHT belongs to MinNL . This claim comes from the following facts. On input $x = (G, s, t, w)$, let $\mathcal{S} = (v_1, v_2, \dots, v_k)$ denote an arbitrary path from s to t in G . Since $m_0(x, \mathcal{S})$ equals $w(\mathcal{S})$ by definition, the value $m_0(x, \mathcal{S})$ can be computed by an appropriate auxiliary Turing machine that writes down $\text{bin}(w(v_i))$ sequentially on a write-only output tape using $O(\log n)$ space-bounded work tapes. Similarly, given $x = (G, s, t, w)$ and an arbitrary sequence \mathcal{S} of vertices, we can decide whether $\mathcal{S} \in SOL_0(x)$ by checking whether \mathcal{S} is a path from s to t using a certain log-space auxiliary Turing machine.

Secondly, we shall claim that MIN PATH-WEIGHT is $\leq_{\text{sAP}}^{\text{AC}^0}$ -hard for MinNL ; namely, every minimization problem in NLO is $\leq_{\text{sAP}}^{\text{AC}^0}$ -reducible to MIN PATH-WEIGHT. To prove this claim, let $P = (I, SOL, m, \text{MIN})$ be any minimization problem in NLO . Note that $I \in \text{L}$, $I \circ SOL \in \text{auxL}$, and $m \in \text{auxFL}$. Since $m \in \text{auxFL}$, we take an appropriate log-space auxiliary Turing machine M (with three tapes) computing m , where any solution candidate to P is provided on an auxiliary read-once tape. Notice that there is a unique initial partial configuration. To ensure that M has a *unique* accepting partial configuration, it suffices to force M to clear out all tapes just before entering a unique accepting state.

Let us define an $\leq_{\text{sAP}}^{\text{AC}^0}$ -reduction $(f, g, 1)$ from P to MIN PATH-WEIGHT as follows. Let $r \geq 1$ and define $f(x, r)$ to be a configuration graph G_x^M of M on input x . If $x \in I$, then $f(x, r) \in I_0$. Let s denote the initial partial configuration of M on x and let t be the unique accepting partial configuration of M_2 on x . As a solution to MIN PATH-WEIGHT, let y be any path in the graph $f(x, r)$ starting with s . Each vertex in y contains the information on content τ_i of the tape cell at which the auxiliary-tape head scans at time i . Hence, from y , we can recover the content of the auxiliary tape as follows. Given $y = (y_0, y_1, \dots, y_m)$ with $y_i = (q_i, w_i, u_i, \tau_i, \xi_i, k_i)$, we retrieve τ_i for all indices $i \in [0, m]_{\mathbb{Z}}$ and output $\tau_0 \tau_1 \tau_2 \dots \tau_m$. This procedure requires only an AC^0 circuit. Let $g(x, y, r)$ denote the entire content of the auxiliary tape that is reconstructed from y as described above. Clearly, g is in FAC^0 and, for any $y \in SOL(x)$, we obtain $g(x, y, r) \in SOL_0(f(x, r))$. It is not difficult to show that $m(f(x, r), y) = m_0(x, g(x, y, r))$. Hence, $R_2(f(x, r), y)$ equals $R_1(x, g(x, y, r))$.

To complete the proof, we still need to verify that f belongs to FAC^0 . For this, consider the following procedure. Recall that a graph is represented by a list of edges (i.e., vertex pairs). Starting with any input $x = (G, s, t, w)$ and $r \geq 1$, generate all pairs (u, v) of partial configurations and mark (u, v) whenever it is an edge of G_x^M . This procedure needs to wire only a finite number of bits between u and v . Hence, f can be computed by an AC^0 circuit.

Therefore, MIN PATH-WEIGHT is $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete for MinNL . □

In contrast to MIN PATH-WEIGHT, it is possible to define a maximization problem, MAX PATH-WEIGHT, simply by taking the maximally-weighted s - t path for the minimally-weighted one in the definition of MIN PATH-WEIGHT. A similar argument in the proof of Theorem 3.3 establishes the $\leq_{\text{sAP}}^{\text{NC}^1}$ -completeness of MAX PATH-WEIGHT for MaxNL .

Corollary 3.4 MAX PATH-WEIGHT is $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete for MaxNL .

Is MIN PATH-WEIGHT also $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete for MaxNL and thus for NLO ($= \text{MaxNL} \cup \text{MinNL}$)? Unlike NPO problems, the log-space limitation of work tapes of Turing machines complicates the circumstances around NLO problems. At present, we do not know that MIN PATH-WEIGHT is $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete for NLO . This issue will be discussed later in Section 4. Under a certain assumption on auxFL , nevertheless, it is possible to achieve the $\leq_{\text{sAP}}^{\text{NC}^1}$ -completeness of MIN PATH-WEIGHT for NLO .

We say that auxFL is *closed under division* if, for any two functions $f, g \in \text{auxFL}$ outputting natural numbers in binary, the function h defined by $h(x, y) = \lceil f(x, y)/g(x, y) \rceil$ for all inputs x and all auxiliary inputs y is in auxFL , provided that $g(x, y) > 0$ for all inputs (x, y) .

Proposition 3.5 Assume that auxFL is closed under division. MIN PATH-WEIGHT is $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete for NLO.

We have already proven that MIN PATH-WEIGHT is $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete for MinNL in Theorem 3.3. In Lemma 3.6, we shall demonstrate that every problem P_1 in MaxNL is sAPAC^0 -reducible to an appropriately chosen problem P_2 in MinNL if auxFL is closed under division. Since $\text{NLO} = \text{MaxNL} \cup \text{MinNL}$, this implies that MIN PATH-WEIGHT is also $\leq_{\text{sAP}}^{\text{NC}^1}$ -hard for MaxNL, completing the proof of Proposition 3.5.

We shall prove the remaining lemma, Lemma 3.6.

Lemma 3.6 Assume that auxFL is closed under division. Every problem P_1 in MaxNL is $\leq_{\text{sAP}}^{\text{AC}^0}$ -reducible to an appropriate problem P_2 in MinNL.

Proof. Let $P_1 = (I_1, \text{SOL}_1, m_1, \text{MAX})$ be any optimization problem in MaxNL. Take an appropriate polynomial p satisfying $2^{p(|x|)} \geq m_1^*(x)$ for every instance $x \in I_1$. For brevity, we set $b(x) = 2^{p(|x|)}$ for all $x \in I_1$. We shall construct the desired minimization problem $P_2 = (I_2, \text{SOL}_2, m_2, \text{MIN})$ in MinNL. Let $I_2 = I_1$ and $\text{SOL}_2 = \text{SOL}_1$. Moreover, for every pair $(x, y) \in I_2 \circ \text{SOL}_2$, define $m_2(x, y) = \lceil \frac{b(x)^2}{m_1(x, y)} \rceil$. It is important to note that, by our definition of measure function, m_1 always returns *positive values*. From this definition, it follows that $\frac{b(x)^2}{m_2(x, y)} \leq m_1(x, y) \leq \frac{b(x)^2}{m_2(x, y)+1}$ for any $(x, y) \in I_2 \circ \text{SOL}_2$.

Clearly, $I_2 \in \text{L}$ and $I_2 \circ \text{SOL}_2 \in \text{auxFL}$. From our assumption on the closure property of auxFL under division, m_2 falls into auxFL. Therefore, P_2 belongs to NLO.

Let us define an sAPAC^0 -reduction (f, g, c) from P_1 to P_2 as follows. Let $f(x, r) = x$ and $g(x, y, r) = y$ for $r \in \mathbb{Q}^{\geq 1}$, $x \in I_1$, and $y \in \text{SOL}_2(f(x, r))$. Obviously, $f, g \in \text{FAC}^0$ follows. If $R_2(f(x, r), y) \leq r$ for $r \geq 1$; namely, $m_2^*(x)/r \leq m_2(x, y) \leq m_2^*(x)$, then the performance ratio $R_1(x, g(x, y, r))$ for P_1 is upper-bounded as

$$R_1(x, g(x, y, r)) = \frac{m_1^*(x)}{m_1(x, y)} \leq \frac{b(x)^2}{m_2^*(x) + 1} \div \frac{b(x)^2}{m_2(x, y)} = \frac{m_2(x, y)}{m_2^*(x)}.$$

The last term is further upper-bounded by $\frac{rm_2^*(x)}{m_2^*(x)+1} \leq 1 + c(r-1)$, where $c = 1$, since $m_2(x, y) \leq rm_2^*(x)$. Overall, we obtain $R_1(x, g(x, y, r)) \leq 1 + c(r-1)$. We then conclude that (f, g, c) is indeed an $\leq_{\text{sAP}}^{\text{AC}^0}$ -reduction from P_1 to P_2 . \square

Henceforth, we shall discuss several variants of MIN PATH-WEIGHT. A simple variant is an *undirected-graph version* of MIN PATH-WEIGHT, denoted by MIN UPATH-WEIGHT. It is possible to demonstrate that MIN UPATH-WEIGHT is log-space $n^{O(1)}$ -approximable because, by the result of Reingold [36], using only log space, we not only determine the existence of a certain feasible solution for MIN UPATH-WEIGHT but also find at least one feasible solution if any. The special case where the weights of all vertices are exactly 1 is the problem of finding the *shortest s-t path*. This problem was discussed in [39]; nonetheless, it is unknown that MIN UPATH-WEIGHT belongs to APXL_{NLO} .

As another variant of MIN PATH-WEIGHT, we consider *forests*. Cook and McKenzie [12] showed that the *s-t connectivity problem* for forests is complete for L under L-uniform NC^1 many-one reductions. Similarly, when all admissible input graphs of MIN UPATH-WEIGHT are restricted to be forests, we call the corresponding problem MIN FOREST-PATH-WEIGHT. As shown in the following proposition, MIN FOREST-PATH-WEIGHT turns out to be one of the most difficult problems in LO_{NLO} . It is of importance that, unlike NLO, the class LO_{NLO} does possess complete problems.

Proposition 3.7 MIN FOREST-PATH-WEIGHT is $\leq_{\text{sAP}}^{\text{AC}^0}$ -complete for LO_{NLO} .

To simplify the proof of Proposition 3.7, we first give a useful lemma that helps us pay central attention to optimization problems of particular form. Here, we say that an optimization problem $P = (I, \text{SOL}, m, \text{goal})$ admits *unique solutions* if $|\text{SOL}(x)| \leq 1$ holds for all $x \in I$.

Lemma 3.8 For any maximization problem $Q \in \text{LO}_{\text{NLO}}$ (resp., $\text{LO}_{\text{NLO}} \cap \text{PBO}$), there are another maximization problem $P = (I, \text{SOL}, m, \text{MAX})$ in LO_{NLO} (resp., $\text{LO}_{\text{NLO}} \cap \text{PBO}$) and a log-space deterministic Turing machine M_P such that, for any $x \in I$, (i) $Q \leq_{\text{sAP}}^{\text{AC}^0} P$, (ii) $m(x, z) = m(x, M_P(x))$ for all $z \in \text{SOL}(x)$, and (iii) P admits unique solutions. The same statement holds for minimization problems.

Proof. Let $Q = (I_1, \text{SOL}_1, m_1, \text{MAX})$ be any maximization problem in LO_{NLO} . Let M_Q be a log-space deterministic Turing machine producing optimal solutions of Q . We then define $P = (I_2, \text{SOL}_2, m_2, \text{MAX})$ as follows. First, we set $I_2 = I_1$ and $\text{SOL}_2(x) = \{M_Q(x)\}$ if $M_Q(x) \neq \perp$, and $\text{SOL}_2(x) = \emptyset$ otherwise. From this definition follows $|\text{SOL}_2(x)| \leq 1$ for all $x \in I_2$. Moreover, we define m_2 by setting

$m_2(x, z) = m_1(x, M_Q(x))$ for any $(x, z)x \in I_2 \circ SOL_2$. Here, we set M_P to be the same as M_Q . Obviously, $m_2(x, M_P(x)) = m_2^*(x)$ holds if $SOL_1(x) \neq \emptyset$ since $M_Q(x) \in SOL_1^*(x)$.

For the desired $\leq_{sAP}^{AC^0}$ -reduction (f, g, c) , we define $c = 1$, $f(x, r) = x$, and $g(x, y, r) = y$. Clearly, $f, g \in FAC^0$. Consider the performance ratio R_1 and R_2 for Q and P , respectively, and assume that $R_2(f(x, r), y) \leq r$ for any $y \in SOL_2(f(x, r))$ and $r \in \mathbb{Q}^{\geq 1}$. This assumption yields $y = M_P(x)$. Hence, $R_1(x, g(x, y, r)) = R_1(x, y) = R_1(x, M_Q(x)) = 1 \leq r$. Therefore, (f, g, c) reduces Q to P . \square

Let us begin the proof of Proposition 3.7.

Proof of Proposition 3.7. MIN FOREST-PATH-WEIGHT is assumed to have the form $(I_0, SOL_0, m_0, \text{MIN})$. The membership relation MIN FOREST-PATH-WEIGHT $\in LO_{NLO}$ essentially comes from a simple fact that, by the forest property of a given graph G , two nodes s and t are connected in G if and only if a *unique* path exists between them. We can search such a unique path by starting from s and following recursively adjacent edges to next nodes until either no more edges remain unsearched or t is found. At the same time, we progressively write down the weight, in binary, of each node along this found path. The recursive part of this procedure works as follows. Let u be the currently visiting node. We then pick each neighbor, say, v and check if there is a path between v and t in a graph obtained from G by deleting the edge (u, v) . This procedure needs no more than log space.

Let $P = (I, SOL, m, \text{MIN})$ be any minimization problem in LO_{NLO} . We assume that P satisfies Conditions (ii)–(iii) of Lemma 3.8. Our goal is to show that P is $\leq_{sAP}^{AC^0}$ -reducible to MIN FOREST-PATH-WEIGHT via a suitably constructed $\leq_{sAP}^{AC^0}$ -reduction (f, g, c) . Choose a log-space deterministic Turing machine M_1 that produces optimal solutions of P . For convenience, we set $b(x)$ to be $m(x, M_1(x))$ for any instance $x \in (I \circ SOL)^{\exists}$. Since $m \in \text{auxFL}$, there is a log-space auxiliary Turing machine M_m computing m . By combining M_1 and M_m properly, we can design another log-space deterministic Turing machine, say, M_2 that computes b with no auxiliary tape.

To make all final partial configurations unique, we want to force M_2 to erase all symbols on all tapes just before entering a halting state. To avoid the same partial configurations to be reached along a single computation, we additionally equip an *internal clock* to M_2 .

Let us consider partial configurations of M_2 . Note that M_2 is deterministic and the internal clock marks all partial configurations of M_2 on x with different time stamps. Note also that each symbol of the string $M_1(x)$ appears as a symbol read from the auxiliary input tape encoded into certain partial configurations of M_1 . Hence, if we have a valid series y of partial configurations associated with an accepting computation path of M_2 on x , then we can recover the string $M_1(x)$ correctly. Notationally, $\eta(y)$ denotes this unique string obtained from a valid series y of partial configurations.

Take a configuration graph $G_x^{M_2} = (V, E)$ from M_2 . Note that there is at most one correct computation path of M_2 . We set s to be the initial partial configuration of M_2 on x and set t be a unique accepting partial configuration of M_2 on x . The resulted graph forms an *acyclic undirected graph*, namely a forest, because, otherwise, there are two accepting computation paths on the same input x . Given any partial configuration $v \in V$, we define $w(v)$ to be one bit written down newly on the output tape in this partial configuration v . For the desired reduction, we define $c = 1$, $f(x, r) = \langle G, s, t, w \rangle$, and $g(x, y, r) = \eta(y)$ for any $y \in SOL_0(f(x, r))$. We obtain $f, g \in \text{FNC}^1$. It follows that $m_0(f(x, r), y) = m(x, \eta(y))$ for any $x \in I$ and $y \in SOL_0(f(x, r))$. In particular, $g(x, y, r)$ is a minimal solution of x if and only if y is a minimal solution of $f(x, r)$. Thus, $(f, g, c) \leq_{sAP}^{AC^0}$ -reduces P to MIN FOREST-PATH-WEIGHT.

Next, we consider any maximization problem P in LO_{NLO} . Since P also satisfies Condition (ii)–(iii) of Lemma 3.8, the above argument also works for this P and thus establishes the $\leq_{sAP}^{AC^0}$ -reducibility of P to MIN FOREST-PATH-WEIGHT. \square

As other variants of MIN PATH-WEIGHT, Nickelsen and Tantau [34] studied *series-parallel graphs* and *tournaments*.

3.3 Complete Problems Concerning Finite Automata

We shall leave graph problems behind and look into problems associated with finite automata. Álvarez and Jenner [4] and later Tantau [39] discussed an intimate relationship between accepting computations of nondeterministic finite automata and log-space search procedures for optimal solutions. Those problems are also closely related to *maximal word problems (or functions)* for fixed underlying machines. Allender, Bruschi, and Pighizzini [2], for instance, discussed the maximal word problems of various types of auxiliary pushdown automata. Within our framework of NLO problems, Tantau [39] presented a maximization problem finding

the maximal input strings accepted by nondeterministic finite automata and demonstrated that this problem is $\leq_{\text{sAP}}^{\text{L}}$ -complete for MaxNL. Here, we shall show that a restricted version of this problem is $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete for $\text{APXL}_{\text{MaxNL}}$.

A *one-way one-head nondeterministic finite automaton with λ -moves* (or a λ -Infa, in short) M is a tuple $(Q, \{0, 1\}, \{\$, \#\}, \delta, q_0, F)$ working with the input alphabet $\{0, 1\}$ and a transition function $\delta : (Q - F) \times \{0, 1, \lambda\} \rightarrow \mathcal{P}(Q)$, where $q_0 \in Q$, and $F \subseteq Q$. Initially, an input $x \in \{0, 1\}^*$ is written on an input tape, surrounded by two endmarkers $\$$ (left) and $\#$ (right). If M makes a λ -move simply by applying $p \in \delta(q, \lambda)$, then M 's read-only tape head stays still; otherwise, the tape head moves to the next right cell. A *configuration* of M is a pair (q, σ) of current inner state q and scanning symbol σ . An *accepting computation path* $p_{M,x}$ of M on input x is a series of configurations starting with an initial configuration $(q_0, \$)$ and ending with a final configuration $(q_f, \#)$ with $q_f \in F$ and, for any consecutive two elements (q_i, σ_i) and (q_{i+1}, σ_{i+1}) in $p_{M,x}$, q_{i+1} is obtained in a single step from (q_i, σ_i) by applying a transition of the form $q_{i+1} \in \delta(q_i, \sigma_i)$ with $\sigma_i \in \{0, 1, \lambda\}$, where $\sigma_1 \sigma_2 \cdots \sigma_k$ is a partition of the input string $\$x\#$. If M enters a certain final state in F along a certain accepting computation path, then M is said to *accept* x ; otherwise, M *rejects* x . Associated with such λ -Infa's, we consider the following optimization problem. For succinctness, we hereafter express a transition " $p \in \delta(q, \sigma)$ " as a triplet (q, σ, p) .

- MAXIMUM FIXED-LENGTH λ -NONDETERMINISTIC FINITE AUTOMATA PROBLEM (MAX FL- λ -NFA):
- INSTANCE: a λ -Infa $M = (Q, \{0, 1\}, \{\$, \#\}, \delta, q_0, F)$ and a string 0^n for a length parameter n , provided that $0^n \in L(M)$.
 - SOLUTION: an accepting computation path of M of length at most $|Q|$ on a certain input y of length exactly n .
 - MEASURE: an integer $\text{rep}(1y)$.

In the above definition, if we remove the requirement " $0^n \in L(M)$ " and we allow y to have any length up to n , then we obtain MAX λ -NFA, which is $\leq_{\text{sAP}}^{\text{L}}$ -complete for MaxNL [39].

Proposition 3.9 MAX FL- λ -NFA is $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete for $\text{APXL}_{\text{MaxNL}}$.

Before proving this proposition, we show a useful supporting lemma. The lemma helps us concentrate only on optimization problems in APXL_{NLO} that have a certain simple structure.

Lemma 3.10 For any maximization problem Q in APXL_{NLO} (resp., $\text{APXL}_{\text{NLO}} \cap \text{PBO}$), there exist another maximization problem $P = (I, \text{SOL}, m, \text{MAX})$ in APXL_{NLO} (resp., $\text{APXL}_{\text{NLO}} \cap \text{PBO}$) and a log-space deterministic Turing machine M_P such that, for all $x \in I$, (i) $Q \leq_{\text{sAP}}^{\text{AC}^0} P$, (ii) $\max_{y \in \text{SOL}(x)} \{m(x, y)\} \leq 2 \min_{y \in \text{SOL}(x)} \{m(x, y)\}$ holds for all $x \in I$, (iii) there exists a function $b \in \text{FL}$ such that $m(x, z) \geq 2^{\lfloor \log b(x) \rfloor}$ for all $z \in \text{SOL}(x)$, and (iv) for any $x \in I$ with $\text{SOL}(x) \neq \emptyset$, $m(x, M_P(x)) = 2^{\lfloor \log b(x) \rfloor}$. A similar statement holds for minimization problems; however, we need to replace (iii) by (iii') $m(x, z) \leq 2^{\lfloor \log b(x) \rfloor}$.

Proof. Let $Q = (I_1, \text{SOL}_1, m_1, \text{MAX})$ be any maximization problem in APXL_{NLO} . In what follows, we shall modify Q to obtain the desired problem $P = (I_2, \text{SOL}_2, m_2, \text{MAX})$.

Take a polynomial p such that, for any $(x, y) \in I_1 \circ \text{SOL}_1$, $|y| \leq p(|x|)$ holds. Since $Q \in \text{APXL}_{\text{NLO}}$, take a log-space deterministic Turing machine M_P producing β -approximate solutions of P for a certain constant $\beta > 1$. We obtain $m_1(x, M_Q(x)) \leq m_1^*(x) \leq \beta m_1(x, M_Q(x))$ for all $x \in I_1$ with $\text{SOL}_1(x) \neq \emptyset$. For such an x , we further set $b_0(x) = 2^{\lfloor \log m_1(x, M_Q(x)) \rfloor + 1}$. Note that $b_0 \in \text{FL}$ since $m_1 \in \text{auxFL}$. Since $b_0(x) \leq b(x) \leq 2b_0(x)$, it follows that $b_0(x) \leq m_1^*(x) \leq 2\beta b_0(x)$. Let us consider a configuration graph $G_x^{M_Q}$ of M_Q on x . Note that we can compute a string $M_Q(x)$ from $(x, G_x^{M_Q})$ using only log space. For later use, we set $\alpha = 2^{\lfloor \log 2\beta \rfloor + 1} + 1$, which implies $b_0(x) \leq m_1^*(x) \leq \alpha b_0(x)$.

Here, we define the desired problem $P = (I_2, \text{SOL}_2, m_2, \text{MAX})$. For convenience, set $\Delta_x = (\alpha - 2)b_0(x)$. Let $I_2 = \{x \# G_x^{M_Q} \mid x \in I_1\}$. Given $\tilde{x} = x \# G_x^{M_Q} \in I_2$, $\text{SOL}_2(\tilde{x})$ contains the following strings: (i) $y \# M_Q(x)$ for all $y \in \text{SOL}_1(x)$ satisfying $m_1(x, y) \geq b_0(x)$ and (ii) $By \# M_Q(x)$ for every $y \in \text{SOL}_1(x)$ satisfying $m_1(x, y) < b_0(x)$, where B is a special symbol. Obviously, $I_2 \circ \text{SOL}_2$ is a member of auxL .

For any $\tilde{y} \in \text{SOL}_2(\tilde{x})$, if $\tilde{y} = y \# M_Q(x)$, then we set $m_2(\tilde{x}, \tilde{y}) = m_1(x, y) + \Delta_x$; if $\tilde{y} = By \# M_Q(x)$, then we set $m_2(\tilde{x}, \tilde{y}) = m_1(x, y) \cdot 2^t + \Delta_x$, where $t = |\text{bin}(b_0(x))| - |\text{bin}(m_1(x, y))|$. It follows that $m_2(\tilde{x}, \tilde{y}) \geq b_0(x) + \Delta_x$ for all $(\tilde{x}, \tilde{y}) \in I_2 \circ \text{SOL}_2$. Note that $m_2(\tilde{x}, M_Q(x) \# M_Q(x))$ equals $b_0(x) + \Delta_x$, which is $2^{\lfloor \log 2\beta \rfloor + \lfloor \log b_0(x) \rfloor + 2}$. Choose a function b so that $\lfloor \log b(x) \rfloor = \lfloor \log 2\beta \rfloor + \lfloor \log b_0(x) \rfloor + 2$ for all $x \in (I_1 \circ \text{SOL}_1)^\exists$. Clearly, $m_2 \in \text{auxFL}$ holds.

The desired $M_P(\tilde{x})$ outputs $M_Q(x) \# M_Q(x)$ for any $\tilde{x} \in I_2$ if $\text{SOL}_2(\tilde{x}) \neq \emptyset$, and it outputs \perp otherwise. Let $\tilde{x} \in (I_2 \circ \text{SOL}_2)^\exists$. It follows that $m_2^*(\tilde{x}) = \max_{\tilde{y} \in \text{SOL}_2(\tilde{x})} \{m_2(\tilde{x}, \tilde{y})\} = \max_{y \in \text{SOL}_1(x)} \{m_1(x, y) + \Delta_x\} =$

$m_1^*(x) + \Delta_x$ since $m_1^*(x) \geq b_0(x)$. Moreover, we obtain $m_1(x, g(x, \tilde{y}, r)) \geq b_0(x)$. It follows that

$$\frac{\max_{y \in SOL_2(\tilde{x})} \{m_2(\tilde{x}, y)\}}{\min_{y \in SOL_2(\tilde{x})} \{m_2(\tilde{x}, y)\}} = \frac{m_1^*(x) + \Delta_x}{b_0(x) + \Delta_x} \leq \frac{\alpha b_0(x) + (\alpha - 2)b_0(x)}{b_0(x) + (\alpha - 2)b_0(x)} = 2. \quad (1)$$

We remark that, by the construction of P from Q , if Q is polynomially bounded, then so is P . We wish to define an $\leq_{sAP}^{NC^1}$ -reduction (f, g, c) from Q to P . First, we set $c = \alpha - 1$. We then define $f(x, r) = x \downarrow G_x^{M_Q}$ ($= \tilde{x}$), and $g(x, \tilde{y}, r) = y$ if \tilde{y} is of the form $y \downarrow z$, $g(x, \tilde{y}, r) = z$ if \tilde{y} is of the form $By \downarrow z$, and $g(x, \tilde{y}, r) = x$ otherwise. If $x \in I_1$, then $f(x, r) \in I_2$ since $f(x, r) = x \downarrow G_x^{M_Q}$. If $y \in SOL_1(x)$, then $y \downarrow M_P(x) \in SOL_2(f(x, r))$ because $M_Q(x)$ can be constructed from $(x, G_x^{M_Q})$ using log space. Next, we assume that the performance ratio R_2 for P satisfies $R_2(f(x, r), \tilde{y}) \leq r$ for $x \in I_1$ and $\tilde{y} \in SOL_2(f(x, r))$. Note that $m_1(x, g(x, \tilde{y}, r)) = m_1(x, y)$ if $\tilde{y} = y \downarrow M_Q(x)$. If $\tilde{y} = By \downarrow M_Q(x)$, then $m_1(x, g(x, \tilde{y}, r)) = m_1(x, M_Q(x))$. We obtain $m_1(x, g(x, \tilde{y}, r)) \geq m_1(x, M_Q(x))$. Write $u = g(x, \tilde{y}, r)$ for simplicity. As for the performance ratio R_1 for Q , it follows that

$$R_1(x, g(x, \tilde{y}, r)) - 1 = \frac{m_1^*(x) - m_1(x, u)}{m_1(x, u)} \leq (\alpha - 1) \cdot \frac{m_2^*(\tilde{x}) - m_2(\tilde{x}, \tilde{y})}{(\alpha - 1)m_1(x, u)}. \quad (2)$$

The last term is further calculated as

$$c \cdot \frac{m_2^*(\tilde{x}) - m_2(\tilde{x}, \tilde{y})}{m_1(x, u) + (\alpha - 2)m_1(x, u)} \leq c \cdot \frac{m_2^*(\tilde{x}) - m_2(\tilde{x}, \tilde{y})}{m_1(x, u) + \Delta_x} = c \cdot \frac{m_2^*(\tilde{x}) - m_2(\tilde{x}, \tilde{y})}{m_2(\tilde{x}, \tilde{y})}. \quad (3)$$

The last term equals $c(R_2(f(x, r), \tilde{y}) - 1)$. Since $R_2(f(x, r), \tilde{y}) \leq r$, it follows that $R_1(x, g(x, \tilde{y}, r)) \leq 1 + c(r - 1)$. Therefore, (f, g, c) reduces Q to P . \square

Let us begin the proof of Proposition 3.9.

Proof of Proposition 3.9. First, we shall argue that MAX FL- λ -NFA belongs to APXL_{MaxNL}. For simplicity, we set MAX FL- λ -NFA as $(I_0, SOL_0, m_0, \text{MIN})$. Let $x = (M, 0^n)$ be any instance in I_0 with $M = (Q, \{0, 1\}, \{\downarrow, \$\}, \delta, q_0, F)$, where M is demanded to accept 0^n . Since we feed only inputs y of length n , the value $rep(1y)$ varies from 2^n to $2^{n+1} - 1$. It thus follows that, for any $u \in SOL_0(x)$, $m_0^*(x)/2 \leq m_0(x, u) \leq m_0^*(x)$. Consider the following algorithm N : take x as input and simulate M on input 0^n (also by checking the size of 0^n). This algorithm requires only log space. We then obtain $m_0^*(x)/2 \leq m_0(x, N(x)) \leq m_0^*(x)$ for any $x \in I_0$. These bounds imply that MAX FL- λ -NFA is a member of APXL_{MaxNL}.

Next, we shall show the $\leq_{sAP}^{NC^1}$ -hardness of MAX FL- λ -NFA. Let $P = (I, SOL, m, \text{MAX})$ be any problem in APXL_{MaxNL}. Without loss of generality, we assume that P satisfies Conditions (ii)–(iv) of Lemma 3.10, and thus P admits a 2-approximate algorithm. Our goal is to show that P is $\leq_{sAP}^{NC^1}$ -reducible to MAX FL- λ -NFA via a suitable reduction (f, g, c) . Let M_m be a log-space auxiliary Turing machine computing m and let M_P denote a log-space 2-approximate algorithm for P . There is a function $b \in \text{FL}$ such that $b(x) \leq m^*(x) \leq 2b(x)$ and $m(x, y) \geq b(x)$ for all $x \in (I \circ SOL)^\exists$, where $b(x)$ is of the form $2^{\lceil \log C(x) \rceil}$ for appropriate function $C \in \text{FL}$. Moreover, we assume that, for each $x \in (I \circ SOL)^\exists$, there is a solution $y \in SOL(x)$ such that $m(x, y) = b(x)$.

As in the proof of Proposition 3.3, we consider partial configurations of M_m . We define Q to be the set of all possible partial configurations of M_m . Fix $x \in (I \circ SOL)^\exists$ arbitrarily and let n be the size of binary string $bin(b(x))$.

We want to define a λ -1nfa N , which ‘‘mimics’’ a computation of M_m . N ’s inner states are partial configurations of M_m . An input to N is $bin(m(x, y))^{(-)}$ for a certain auxiliary input $y \in SOL(x)$. A move of N is described as follows. Given a string u and a number $k \in \mathbb{N}^+$, u_k denotes the k th symbol of u . On such an input, N nondeterministically guesses a string $y \in \{0, 1\}^n$. By reading an input symbol $\tau' \in \{0, 1\}$ from $bin(m(x, y))$ one by one from left to right, N changes an inner state $v = (q, x, j, \xi, u, k, \tau)$ to another inner state $v' = (p, x, j + d_1, \xi', w, k + d_2, \tau')$ in a single step by applying M_m ’s transition ‘‘ $(p, \tau', w_k, d_1, d_2) \in \delta(q, x_j, \xi, u_k)$,’’ where ξ' is the next bit of ξ in y . More precisely, we define N ’s transition as $\langle p, x, j + d_1, \xi', w, k + d_2, \tau' \rangle \in \delta_N(\langle q, x, j, \xi, u, k, \tau \rangle, \tau')$ iff $(p, \tau', w_k, d_1, d_2) \in \delta(q, x_j, \xi, u_k)$, where w is obtained from u by changing u_k to w_k .

We modify N so that it simultaneously checks whether its input is of the form 0^n . If so, N enters a designated accepting state. Hence, $0^n \in L(N)$.

Here, let us define the desired reduction (f, g, c) . First, we set $f(x, r) = \langle N \rangle$. Given an accepting computation path e in $SOL_0(f(x, r))$, we define $g(x, e, r)$ to be an auxiliary input $y_{(e)}$ fed into M_m , which

can be obtained from e . It follows that $y_{(e)} \in SOL(x)$ and that $m(x, y_{(e)}) = m_0(f(x, r), e)$. Let $r \in \mathbb{Q}^{\geq 1}$. The performance ratio R_P for P satisfies that $R_P(x, g(x, e, r)) = \frac{m^*(x)}{m(x, g(x, e, r))} = \frac{m_0^*(f(x, r))}{m_0(f(x, r), e)} = R_0(f(x, r), e)$.

Therefore, the reduction (f, g, c) ensures $P \leq_{\text{sAP}}^{\text{NC}^1} \text{MAX FL-}\lambda\text{-NFA}$. \square

As a simple variant of MAX λ -NFA, we shall consider *one-way one-head deterministic finite automata with λ -moves* (or λ -1dfa's, in short). A λ -1dfa M is a tuple $(Q, \{0, 1\}, \{\uparrow, \$\}, \delta, q_0, F)$ working with input alphabet $\{0, 1\}$ and $\delta : (Q_\lambda \times \{\lambda\}) \cup (Q_+ \times \{0, 1, \uparrow\}) \rightarrow Q$, where $Q = Q_\lambda \cup Q_+$, $Q_\lambda \cap Q_+ = \emptyset$, $q_0 \in Q$, and $F \subseteq Q$. This M must satisfy the following condition: if M is in state $q \in Q_\lambda$, then M 's read-only tape head stays still; otherwise, the tape head moves to the next right cell. Here, each transition " $\delta(q, \sigma) = p$ " is succinctly expressed as (q, σ, p) .

MAXIMUM INPUT-RESTRICTED λ -DETERMINISTIC FINITE AUTOMATA PROBLEM (MAX IR- λ -DFA):

- INSTANCE: a λ -1dfa $M = (Q, \{0, 1\}, \{\uparrow, \$\}, \delta, q_0, F)$ and a list $Y = (y_1, y_2, \dots, y_k)$ of strings over $\{0, 1\}$, where δ is given as a list of (partial) transitions of the form (q, σ, p) .
- SOLUTION: an accepting computation path of M of length at most $|Q|$ on a certain input $y \in Y$, which is surrounded by \uparrow and $\$$.
- MEASURE: an integer $rep(1y)$.

Proposition 3.11 MAX IR- λ -DFA is $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete for LO_{NLO} .

If we take $(M, 1^n)$ as an instance and demand y to have length at most n , then we obtain another problem, MAX λ -DFA. It is not clear that MAX λ -DFA is $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete for either LO_{NLO} or MaxNL .

Proof of Proposition 3.11. For convenience, let MAX IR- λ -DFA have the form $(I_0, SOL_0, m_0, \text{MAX})$. Let us claim that MAX IR- λ -DFA is in NLO. Note that it is easy to check using log space whether a given instance $M = (Q, \{0, 1\}, \{\uparrow, \$\}, \delta, q_0, F)$ is indeed a λ -1dfa; thus, $I_0 \in \text{L}$ follows. To see $I_0 \circ SOL_0 \in \text{auxL}$, on input (M, Y) together with a sequence $p = (p_1, p_2, \dots, p_m)$ of configurations of M as an auxiliary input, we can check using only log space whether p is indeed an accepting computation path of M (by checking that $p_1 = (q_0, \uparrow)$, $p_m = (q_f, \$)$, (p_i, σ, p_{i+1}) is a transition for each $i \in [m]$ with $m \leq |Q|$ for a certain $\sigma \in \{0, 1, \lambda\}$, and a series of such symbols matches one of y 's in Y). As for $m_0 \in \text{auxFL}$, it is possible to retrieve an input y from p and output $1y$ using log space if y is in Y .

Next, we shall show that MAX IR- λ -DFA belongs to LO_{NLO} . Recursively, we pick each y in Y in the lexicographic order and simulate a given λ -1dfa M on this input y to check if M accepts y within $|Q|$ steps. This process determines the maximal accepted input y in Y . Finally, we generate an accepting computation path of M on this y . This whole procedure requires log space. Hence, MAX IR- λ -DFA can be solved using only log space.

Hereafter, we shall show that MAX IR- λ -DFA is $\leq_{\text{sAP}}^{\text{NC}^1}$ -hard for LO_{NLO} . Let us consider any maximization problem $P = (I, SOL, m, \text{MAX})$ in LO_{NLO} . We assume that P satisfies Conditions (ii)–(iii) of Lemma 3.8. Since optimal solutions and their objective values are both computed using log space, for the purpose of defining g , we can build a log-space deterministic Turing machine M that, on input x , records each symbol y_i of a solution $y = y_1 y_2 \dots y_k \in SOL^*(x)$ in one cell of one work tape (after erasing the previous symbol y_{i-1} if any) and produces $bin(m(x, y))^{(-)}$ on an output tape (by removing the first bit "1" from $bin(m(x, y))$).

We shall construct a pair (f, g) of functions that $\leq_{\text{sAP}}^{\text{NC}^1}$ -reduces P to MAX IR- λ -DFA. Let x be any instance in I . We construct a λ -1dfa $N_x = (Q, \{0, 1\}, \{\uparrow, \$\}, \delta, q_0, F)$ as follows. We view each configuration of M (including the content of its output symbol and an input symbol) as an inner state of N_x . Let Q be a set of all such configurations. Note that $|Q| \geq |x|$. We roughly treat an output tape of M as an input tape of N_x . More precisely, when M writes a symbol $\sigma \in \{0, 1\}$ on its output tape, N_x reads σ on the input tape. When M does not write any non-blank output symbol, N_x makes its associated λ -move. Finally, we set $f(x, r) = \langle N_x \rangle$. For N_x , let $p = (p_0, p_1, \dots, p_m)$ be an accepting computation path of N_x of length $\leq |Q|$. We also define $g(x, p, r)$ to be an input y_p to N_x that is recovered from p as stated above. Note that $m_0(f(x, r), p) = rep(bin(m(x, y_p))^{(-)}) = m(x, y_p)$. Concerning the performance ratio R_P and R_0 for P and MAX IP- λ -DFA, respectively, it follows that $R_P(x, g(x, p, r)) = \frac{m^*(x)}{m(x, y_p)} = \frac{m_0^*(f(x, r))}{m_0(f(x, r), p)} = R_0(f(x, r), p)$. Therefore, $(f, g, 1)$ is an $\leq_{\text{sAP}}^{\text{NC}^1}$ -reduction from P to MAX IR- λ -DFA, as requested. The case where P is a minimization problem can be similarly treated. \square

4 Polynomially-Bounded Complete Problems

Let us recall from Section 2.2 that an optimization problem is said to be *polynomially bounded* exactly when its measure function is polynomially bounded. Recall also the notation PBO, which expresses the set of all polynomially-bounded optimization problems. For many low-complexity optimization/approximation classes below PO_{NPO} , polynomially-bounded optimization problems play a quite special role. With respect to log-space computation, it appears more natural to deal with polynomially-bounded optimization problems than polynomially-unbounded ones because, through Section 3, we have been unable to present any complete problem in NLO and APXL_{NLO} but, as we shall see shortly, we can exhibit complete problems in $\text{NLO} \cap \text{PBO}$ and $\text{APXL}_{\text{NLO}} \cap \text{PBO}$.

In the subsequent subsections, we shall present polynomially-bounded optimization problems, which turn out to be complete for various NL optimization and approximation classes.

4.1 Maximization Versus Minimization

Assume that we wish to show the completeness of a certain optimization problem P for a target optimization/approximation class \mathcal{D} . Since \mathcal{D} may be composed of maximization problems as well as minimization problems, it is necessary to construct desirable reductions to P from all maximization problems in \mathcal{D} and also from all minimization problems in \mathcal{D} . Regarding NPO problems, it is well-known that every minimization problem Q in APXP_{NPO} has its maximization counterpart Q' in APXP_{NPO} whose complexity is at least as hard as Q (see, e.g., [6, Theorem 8.7] for the proof).

A similar statement holds for polynomially-bounded problems in APXL_{NLO} . This is because a log-space auxiliary Turing machine that computes a polynomially-bounded measure function can freely manipulate the outcome of the function using its space-bounded work tapes before writing it down onto an output tape.

- Lemma 4.1**
1. For any minimization (resp., maximization) problem P in $\text{NLO} \cap \text{PBO}$, there exists a maximization (resp., minimization) problem Q in $\text{NLO} \cap \text{PBO}$ such that P is $\leq_{\text{sAP}}^{\text{AC}^0}$ -reducible to Q .
 2. Let $\mathcal{D} \in \{\text{APXL}, \text{APXNC}^1\}$. For every minimization (resp., maximization) problem P in $\mathcal{D}_{\text{NLO}} \cap \text{PBO}$, there exists a maximization (resp., minimization) problem Q in $\mathcal{D}_{\text{NLO}} \cap \text{PBO}$ such that P is $\leq_{\text{sAP}}^{\text{AC}^0}$ -reducible to Q .
 3. Let $\mathcal{D} \in \{\text{LO}, \text{NC}^1\text{O}\}$. For any minimization (resp., maximization) problem P in $\mathcal{D}_{\text{NLO}} \cap \text{PBO}$, there exists a maximization (resp., minimization) problem Q in $\mathcal{D}_{\text{NLO}} \cap \text{PBO}$ such that P is $\leq_{\text{EX}}^{\text{AC}^0}$ -reducible to Q .

Proof. (1) This proof is similar in essence to that of Lemma 3.6. Given an arbitrary problem $P_1 = (I_1, \text{SOL}_1, m_1, \text{MIN})$ in $\text{NLO} \cap \text{PBO}$, we aim at constructing a maximization problem $P_2 = (I_2, \text{SOL}_2, m_2, \text{MAX})$ in $\text{NLO} \cap \text{PBO}$ to which P_1 is $\leq_{\text{sAP}}^{\text{AC}^0}$ -reducible. Since $P_1 \in \text{PBO}$, there is a polynomial p satisfying $m_1(x, y) \leq p(|x|)$ for every $(x, y) \in I_1 \circ \text{SOL}_1$. Let $I_2 = I_1$, $\text{SOL}_2 = \text{SOL}_1$, and $m_2(x, y) = \lceil \frac{p(|x|)^2}{m_1(x, y)} \rceil$ for every $(x, y) \in I_2 \circ \text{SOL}_2$. Since m_1 is polynomially-bounded, we can generate the entire value $m_1(x, y)$ on one of log-space work tapes and manipulate it freely as if a normal input. (If m_2 is not polynomially bounded, then there is no guarantee that a log-space auxiliary Turing machines can compute m_2 .) Since “division” can be implemented on TC^0 circuit [19] (and thus, by a log-space machine), $m_2(x, y)$ can be generated on a log-space work tape. Here, we define $f(x, r) = x$ and $g(x, y, r) = y$ so that f and g belong to FAC^0 . For the performance ratio R_2 for P_2 , assume that $R_2(f(x, r), y) \leq r$ for any $r \geq 1$; that is, $m_2^*(x) \leq m_2(x, y) \leq r m_2^*(x)$. The ratio $R_1(x, g(x, y, r)) = \frac{m_1(x, y)}{m_1^*(x)}$ is thus at most $1 + c(r - 1)$, where $c = 1$. Hence, (f, g, c) reduces P_1 to P_2 .

(2) We shall show only the case of $\mathcal{D} = \text{APXNC}^1$. Let $P_1 = (I_1, \text{SOL}_1, m_1, \text{MIN})$ be any polynomially-bounded minimization problem in $\text{APXNC}_{\text{NLO}}^1$. We want to construct a polynomially-bounded maximization problem $P_2 = (I_2, \text{SOL}_2, m_2, \text{MAX})$, which is in $\text{APXNC}_{\text{NLO}}^1$ and is $\leq_{\text{sAP}}^{\text{AC}^0}$ -reduced from P_1 . Take a constant $\gamma > 1$ and an NC^1 γ -approximate algorithm C for P_1 . For convenience, we set $b(x) = m_1(x, C(x))$ for each instance $x \in (I_1 \circ \text{SOL}_1)^\exists$. By the definition of $\text{APXNC}_{\text{NLO}}^1$, b must belong to FNC^1 . Note that $m_1^*(x) \leq b(x) \leq \gamma m_1^*(x)$ for all $x \in (I_1 \circ \text{SOL}_1)^\exists$.

We choose a constant Δ so that $\Delta > \gamma$ holds. Note that $\Delta - 1 > 0$ since $\gamma \geq 1$. Next, we define $I_2 = I_1$, $\text{SOL}_2(x) = \{y \in \text{SOL}_1(x) \mid m_1(x, y) \leq b(x)\}$, and $m_2(x, y) = \Delta b(x) - \gamma m_1(x, y)$ for any $(x, y) \in I_2 \circ \text{SOL}_2$. Notice that m_2 is computed by an appropriate log-space auxiliary Turing machine, and thus $I_2 \circ \text{SOL}_2$ is also computed by a certain log-space auxiliary Turing machine. This implies that P_2 is an NLO problem.

We shall claim that $P_2 \in \text{APXNC}_{\text{NLO}}^1$. Note that $m_2^*(x) = \Delta b(x) - \gamma m_1^*(x)$ for every $x \in (I_2 \circ \text{SOL}_2)^\exists$. Moreover, $m_2(x, C(x)) = (\Delta - \gamma)b(x)$ holds. Since $m_1^*(x) \leq b(x)$, we obtain $(\Delta - \gamma)b(x) \leq m_2^*(x)$, which

implies $m_2(x, C(x)) \leq m_2^*(x)$. Since $b \in \text{FNC}^1$, the value $m_2(x, C(x))$ can be computed from x by a certain NC^1 -circuit. Since $b(x) \leq \gamma m_1^*(x)$, it follows that $m_2^*(x) \leq (\Delta - 1)b(x)$; thus, we obtain $m_2^*(x) \leq \frac{\Delta-1}{\Delta-\gamma} m_2(x, C(x))$. In summary, it holds that $\frac{\Delta-\gamma}{\Delta-1} m_2^*(x) \leq m_2(x, C(x)) \leq m_2^*(x)$. Hence, P_2 belongs to $\text{APXNC}_{\text{NLO}}^1$.

As for the desired $\leq_{\text{sAP}}^{\text{AC}^0}$ -reduction, we define $f(x, r) = x$ and $g(x, y, r) = y$, where r is any number in $\mathbb{Q}^{\geq 1}$. It is obvious that f and g are in FAC^0 . For any fixed solution $y \in \text{SOL}_2(x)$, we assume that $R_2(f(x, r), y) \leq r$; that is, $m_2(x, y) \leq m_2^*(x) \leq r m_2(x, y)$. We want to claim that $m_1(x, y) \leq [1 + (\Delta - 1)(r - 1)] m_1^*(x)$. First, we note that

$$m_1(x, y) = (1/\gamma)[\Delta b(x) - m_2(x, y)] \leq (1/r\gamma)[(r-1)\Delta b(x) + \gamma m_1^*(x)]$$

because $m_2^*(x) \leq r m_2(x, y)$ and $m_2^*(x) = \Delta b(x) - \gamma m_1^*(x)$. Since $b(x) \leq \gamma m_1^*(x)$, we obtain $m_1(x, y) \leq (1/r)[(r-1)\Delta + 1] m_1^*(x)$. Hence, it follows that $m_1(x, y) \leq [1 + \theta(r-1)] m_1^*(x)$, where $\theta = (\Delta - 1)/r$, which is at most $\Delta - 1$ since $r \geq 1$. This shows that P_1 is $\leq_{\text{sAP}}^{\text{AC}^0}$ -reducible to P_2 .

(3) Consider the case of $\mathcal{D} = \text{LO}$. The construction of the desired reduction is similar to (2). Assuming that P_1 is in $\text{LO}_{\text{NLO}} \cap \text{PBO}$, we want to show that P_2 is also in $\text{LO}_{\text{NLO}} \cap \text{PBO}$. This is obtained simply by mapping minimal solutions for P_1 to maximal solutions for P_2 . \square

4.2 Completeness of Graph Problems

As our starting point, we recall from Section 3.2 a bounded variant of MIN PATH-WEIGHT , called MIN BPATH-WEIGHT , which uses the *total path weight* (i.e., $w(\mathcal{S}) = \sum_{i=1}^k w(v_i)$ for $\mathcal{S} = (v_1, v_2, \dots, v_k)$ with $k \leq |V|$) as a measure function with an extra condition that $1 \leq w(v) \leq |V|$ for all $v \in V$. Earlier, Tantau [39, Theorem 5.1] discussed the case when all vertices of a given graph have weight exactly 1. To verify that MIN BPATH-WEIGHT is $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete for $\text{MinNL} \cap \text{PBO}$, we need to modify the proof of Proposition 3.3 in the following way. For this purpose, we first modify a given measure function m so that its log-space auxiliary Turing machine M_m produces $m(x, y)$ on one of its work tapes and then copies each bit (including “0”) from the lower bit to the higher bit at each step. We further modify M_m so that its internal clock helps it halt in exactly $p(n)$ steps, for a suitable polynomial p . We then define $w(v)$ to be $b_1 2^{e-1} + 1$ if v contains a string $b_1 b_2 \dots b_e$ written on an output tape of M_m in a target partial final configuration. Otherwise, define $w(v) = 1$. It follows that $\sum_{i=1}^{q(n)} w(v_i) = q(n) + m(x, y)$. By reducing minimization problems to maximization problems by Lemma 4.1(1), we can prove that every minimization problem in $\text{NLO} \cap \text{PBO}$ is also $\leq_{\text{sAP}}^{\text{NC}^1}$ -reducible to MIN BPATH-WEIGHT . Therefore, we obtain the following completeness result.

Lemma 4.2 MIN BPATH-WEIGHT is $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete for $\text{NLO} \cap \text{PBO}$.

In Section 3.2, we have mostly dealt with optimization problems that are associated with the path weights of graphs. Another natural type of optimization problems is a problem of searching a path of a directed graph starting at a given source toward an appropriately chosen vertex whose weight is well-defined and must be maximal. Tantau [39, Theorem 3.2] earlier demonstrated that this maximization problem is $\leq_{\text{sAP}}^{\text{L}}$ -complete for $\text{NLO} \cap \text{PBO}$. Moreover, its slightly modified version was shown to be $\leq_{\text{sAP}}^{\text{L}}$ -complete for $\text{APXL}_{\text{NLO}} \cap \text{PBO}$ [39, Theorem 3.7]. In what follows, we shall discuss a similar optimization problem using “undirected” graphs with “total” weight functions, particularly, in the case where vertex weights are all bounded.

Let us define formally this problem as follows.

MAXIMUM UNDIRECTED BOUNDED VERTEX WEIGHT PROBLEM (MAX UB-VERTEX):

- **INSTANCE:** an undirected graph $G = (V, E)$, a source $s \in V$, and a (vertex) weight function $w : V \rightarrow \mathbb{N}^+$ satisfying $w(v) \leq |V|$ for every $v \in V$.
- **SOLUTION:** a path of G starting at s and ending at a certain vertex t in V .
- **MEASURE:** the weight $w(t)$ of t .

A *directed-graph version* of MAX UB-VERTEX , called MAX B-VERTEX , has a log-space $n^{O(1)}$ -approximate algorithm [39] but is not known to fall into APXL_{NLO} .

Proposition 4.3 MAX UB-VERTEX is $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete for $\text{LO}_{\text{NLO}} \cap \text{PBO}$.

Even if the vertex t in the above definition of MAX UB-VERTEX is restricted to a vertex of degree exactly 1 as the following proof shows, the obtained problem is still $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete for $\text{LO}_{\text{NLO}} \cap \text{PBO}$.

Proof. We begin with setting MAX UB-VERTEX as $(I_0, SOL_0, m_0, \text{MAX})$. Let $G = (V, E)$ be any graph given as an instance in I_0 . notice that $SOL_0(x) \neq \emptyset$ for all $x \in I_1$. Since the weight of every vertex in G is at most the input size, MAX UB-VERTEX is polynomially bounded. Thus, it is not difficult to show that optimal solutions for MAX UB-VERTEX can be found using log space by making a series of nonadaptive queries to oracle $A = \{\langle G, s, w, 1^k \rangle \mid \exists t \in V [w(t) \geq k \wedge s \text{ and } t \text{ are connected}]\}$ by incrementing k from 1 to $|V|$. To see that A is in L, we sequentially pick a different $t \in V$, check if s and t are connected using Reingold's log-space algorithm for DSTCON, and finally check if $w(t) \geq k$. Therefore, MAX UB-VERTEX is L-solvable; that is, MAX UB-VERTEX belongs to LONLO .

Concerning the $\leq_{\text{EX}}^{\text{NC}^1}$ -hardness of MAX UB-VERTEX, let us consider a polynomially-bounded maximization problem $P = (I, SOL, m, \text{MAX})$ in LONLO . The case of minimization problems follows from Lemma 4.1(3). By Lemma 3.8, it is possible to assume that P satisfies Conditions (ii)–(iii) of the lemma. Take a log-space auxiliary Turing machine M_m that computes m . Moreover, we take a log-space deterministic Turing machine M_P that produces maximal solutions for P . Define $b(x) = m(x, M_P(x))$ for each instance $x \in (I \circ SOL)^{\exists}$. Take a polynomial p such that $b(x) \leq p(|x|)$ for all $x \in (I \circ SOL)^{\exists}$. For technicality, we demand that $p(n) > 1$ for all n . Notice that b is computed using log space by the following simple machine M' equipped with a special work tape, called a *solution tape*, in which we use only one tape cell. Starting with input x , M' simulates M_P on x . Whenever M_P writes a symbol σ on its output tape, M' writes σ on the solution tape, simulates M_m on x while an auxiliary-tape head is scanning σ , and erases σ from the solution tape. This deletion of the symbol σ is necessary because the output $M_P(x)$ may be super-logarithmically long. Finally, M' outputs $b(x)$.

Here, we construct a configuration graph G in a way similar to the proof of Proposition 3.7 using M' ; however, we use a quite different weight function. Since all weights are polynomially bounded, we can embed an entire content of an output tape of M' into a partial configuration. To be more precise, we can force M' to use one of its work tapes to compute $b(x)$ and, in the end, copy the content of this tape into its write-only output tape.

For a vertex v representing a certain partial halting configuration, its weight $w(v)$ is set to be the value written on the output tape in this partial configuration unless the value is not zero. For any other vertex associated with partial non-halting configurations, we simply set $w(v) = 1$.

We further define $f(x, r)$ to be the above-mentioned configuration graph G together with the weight function w . Given a computation path y of M' on x , since each partial configuration in y contains the information on an output symbol produced by M_P , it is possible to recover from y an entire output string $M_P(x)$. We thus set $g(x, y, r)$ to be $M_P(x)$ reconstructed from y . It is easy to check that $f, g \in \text{FNC}^1$. Note that $m_0(f(x, r), y) = b(x)$ if y is in $SOL_0(f(x, r))$. This implies that the performance ratio R of $g(x, y, r)$ always satisfies $R(x, g(x, y, r)) = 1$ for any $y \in SOL_0(f(x, r))$. As a result, $(f, g, 1)$ reduces P to MAX UB-VERTEX, as requested. \square

To obtain an $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete problem for $\text{APXNC}_{\text{NLO}}^1 \cap \text{PBO}$, we place a restriction on the behavior of a (vertex) weight function of MAX UB-VERTEX. The *maximum undirected 2 vertex weight problem* (MAX U2-VERTEX) is a variant of MAX UB-VERTEX with an additional requirement: for any instance $\langle G, s, w \rangle$ with $G = (V, E)$, it holds that $\max_{v \in V} \{w(v)\} \leq 2 \min_{v \in V} \{w(v)\}$.

Proposition 4.4 MAX U2-VERTEX is $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete for $\text{APXNC}_{\text{NLO}}^1 \cap \text{PBO}$.

Proof. To simplify the following proof, we write MAX U2-VERTEX as $(I_0, SOL_0, m_0, \text{MAX})$. Similar to MAX UB-VERTEX, MAX U2-VERTEX can be shown to be in $\text{NLO} \cap \text{PBO}$. In particular, to verify the extra condition that (*) $\max_{v \in V} \{w(v)\} \leq 2 \min_{v \in V} \{w(v)\}$ for an instance (G, s, w) , we pick each pair (v_1, v_2) of G 's vertices and check that either $w(v_1) \leq 2w(v_2)$ or $w(v_2) \leq 2w(v_1)$ holds. This procedure requires only log space.

The following algorithm C confirms that MAX U2-VERTEX belongs to $\text{APXNC}_{\text{NLO}}^1$. On input $x = (G, s, w) \in I_0$, output a length-1 path of G from vertex s to the vertex that appears first in an instance (where G is given in binary as a list of edge relations). This process can be implemented by NC^1 -circuits. It thus follows from Condition (*) that $m_0(x, C(x)) \leq m_0^*(x) \leq 2m_0(x, C(x))$. Therefore, C is a 2-approximate algorithm for MAX U2-VERTEX. This shows that MAX U2-VERTEX falls into $\text{APXNC}_{\text{NLO}}^1$.

Next, we want to show that MAX U2-VERTEX is $\leq_{\text{sAP}}^{\text{NC}^1}$ -hard for $\text{APXNC}_{\text{NLO}}^1 \cap \text{PBO}$. Let $P = (I, SOL, m, \text{MAX})$ be any maximization problem in $\text{APXNC}_{\text{NLO}}^1 \cap \text{PBO}$. The case of minimization is handled by Lemma 4.1(2). Let M_m be an auxiliary Turing machine computing m using log space. Since m is polynomially bounded, M_m first produces $m(x, y)$ on one of its work tapes and then copies it onto an output tape just before halting. Let C_P be an NC^1 -circuit producing α -approximate solutions of P for a certain

constant $\alpha > 1$. For convenience, let $b(x) = m(x, C_P(x))$ for each $x \in (I \circ SOL)^\exists$ and $b(x) = \perp$ for all the others x . It follows that $b \in \text{FNC}^1$ and that $b(x) \leq m^*(x) \leq \alpha b(x)$ for all $x \in (I \circ SOL)^\exists$. Because of the definition of partial configurations, from any accepting computation path of $M_m(x, y)$, we can easily recover an auxiliary input y .

Let us prove $P \leq_{\text{sAP}}^{\text{NC}^1} \text{MAX U2-VERTEX}$ via a certain $\leq_{\text{sAP}}^{\text{AC}^0}$ -reduction (f, g, c) . Let $r \geq 1$ and $x \in (I \circ SOL)^\exists$. Firstly, we define $f(x, r)$ to be (G, s, w) , where G is a configuration graph of M_m using normal input x , s is the initial partial configuration of M_m , and w is defined later. Given a computation path $u = (v_1, v_2, \dots, v_k)$ of M_m using x , let y_u denote a unique auxiliary input used for M_m constructed from u . Note that y_u is in $SOL(x)$. Here, we further define $g(x, u, r) = y_u$ and set $c = \alpha - 1$. In addition, we set $\Delta_x = (\alpha - 2)b(x)$. For an accepting configuration v , let $w(v) = w' + \Delta_x$, where w' is the number written on M_m 's output tape. For other configurations v , let $w(v) = b(x) + \Delta_x$. Note that $b(x) + \Delta_x \leq w(v) \leq m^*(x) + \Delta_x$. Similarly to Eqn.(1) in the proof of Lemma 3.10, it holds that $\max_{v \in V} \{w(v)\} / \min_{v \in V} \{w(v)\} \leq 2$. Since $b \in \text{FNC}^1$, $f(x, r)$ is computed by an NC^1 -circuit. For any $u = (v_1, v_2, \dots, v_k) \in SOL_0(f(x, r))$, it follows that $m_0(f(x, r), u) = w(v_k) = m(x, y_u) + \Delta_x = m(x, g(x, u, r)) + \Delta_x$. Moreover, we obtain $m_0^*(f(x, r)) = m^*(x) + \Delta_x$. Given a number $r \in \mathbb{Q}^{\geq 1}$, consider the performance ratio R of $g(x, u, r)$ with respect to x . Since $m(x, g(x, u, r)) = b(x)$ for any $u \in SOL(x)$, calculations similar to Eqns.(2)–(3) lead to $R(x, g(x, u, r)) - 1 \leq c(R_0(f(x, r), u) - 1)$. We therefore conclude that (f, g, c) reduces P to MAX U2-VERTEX . \square

4.3 Completeness of Algebraic and Combinatorial Problems

Apart from graph problems, we shall study algebraic and combinatorial problems. We begin with an algebraic problem, which turns out to be complete for $\text{NLO} \cap \text{PBO}$. As MAX UB-VERTEX and MAX U2-VERTEX have orderly structures induced by “edge relations,” the algebraic problem that we shall consider below has a similar structure induced by “operations.”

Given a finite set X , we consider a binary operation $\circ : X \times X \rightarrow X$. A binary operation \circ is *associative* if $(x \circ y) \circ z = x \circ (y \circ z)$ holds for all $x, y, z \in X$. For a subset S of X , we say that a set $G(S)$ is *generated by \circ from S* if $G(S)$ is the smallest set that contains S and is closed under \circ . The decision problem, called AGEN , of determining whether t is in $G(S)$ for a given instance (X, S, \circ, t) with $t \in X$ is \leq_m^{L} -complete for NL [27]. Let us consider its optimization counterpart, which we call MIN AGEN .

MINIMUM ASSOCIATIVE GENERATION PROBLEM (MIN AGEN):

- **INSTANCE:** a finite set X , an associative binary operation $\circ : X \times X \rightarrow X$, a set $S \subseteq X$, an element $t \in X$, and a weight function $w : X \rightarrow \mathbb{N}^+$ satisfying $w(x) \leq |X|$ for all $x \in X$.
- **SOLUTION:** a sequence (x_1, x_2, \dots, x_m) of elements in X with $1 \leq m \leq |X|$ so that the element $x_1 \circ x_2 \circ \dots \circ x_m$ in $G(S)$ equals t .
- **MEASURE:** the value $\sum_{i=1}^m w(x_i)$.

Note that s belongs to $G(S)$ iff there is a sequence (x_1, x_2, \dots, x_m) of elements in S with $1 \leq m \leq |X|$ for which $s = x_1 \circ x_2 \circ \dots \circ x_m$ holds [27]. In what follows, we demonstrate the $\leq_{\text{sAP}}^{\text{NC}^1}$ -completeness of MIN AGEN for $\text{NLO} \cap \text{PBO}$.

Proposition 4.5 *MIN AGEN is $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete for $\text{NLO} \cap \text{PBO}$.*

Proof. For convenience, we set $\text{MIN AGEN} = (I_0, SOL_0, m_0, \text{MIN})$. First, we show that MIN AGEN is in $\text{NLO} \cap \text{PBO}$. Let (X, S, \circ, w, t) be any instance to MIN AGEN . We can build an auxiliary Turing machine M as follows. Let $u = (x_1, x_2, \dots, x_m)$ be a sequence of X and is given to an auxiliary tape of M . To see that $I_0 \in \text{L}$, it suffices to check, using log space, whether (i) $S \subseteq X$, (ii) \circ is associative, and (iii) $w(x) \leq |X|$ for all $x \in X$. It is also easy to see that $I_0 \circ SOL_0 \in \text{auxFL}$. To obtain $m_0((X, S, \circ, w, t), u)$, we need to compute the value $x_1 \circ x_2 \circ \dots \circ x_m$ and then output the value $\sum_{i=1}^m w(x_i)$. Since m_0 is polynomially bounded, this value can be obtained using log space. This m_0 therefore belongs to auxFL . This show that MIN AGEN is in NLO .

Next, we want to show that MIN AGEN is $\leq_{\text{sAP}}^{\text{NC}^1}$ -hard for $\text{NLO} \cap \text{PBO}$. For this purpose, we shall show that $\text{MIN BPATH-WEIGHT} \leq_{\text{sAP}}^{\text{NC}^1} \text{MIN AGEN}$ since, by Lemma 4.2, MIN BPATH-WEIGHT is $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete for $\text{NLO} \cap \text{PBO}$. Let $\text{MIN BPATH-WEIGHT} = (I_1, SOL_1, m_1, \text{MIN})$. Let $h = (G, s, t, w)$ be any instance in I_1 with $G = (V, E)$. Note that $w(v) \leq |V|$ for all $v \in V$. For a sequence $\mathcal{S} = (v_1, v_2, \dots, v_k)$ with $k \leq |V|$, since $m_1(h, \mathcal{S}) = \sum_{i=1}^k w(v_i)$, we obtain $m_1(h, \mathcal{S}) \leq |V|^2$.

We want to define an $\leq_{\text{sAP}}^{\text{NC}^1}$ -reduction (f, g, c) from MIN BPATH-WEIGHT to MIN AGEN as follows. Our construction essentially follows from the proof of [27, Theorem 5]. Let $c = 1$. Letting $f(h, r) = (X, S, \circ, \tilde{w}, t)$, we define X, S, \circ, \tilde{w} as follows. Let $X = V \cup (V \times V) \cup \{\natural\}$ and $S = E \cup \{s\}$. A binary operation \circ admits the following rules: for all $u \in X$ and $x, y, z \in V$, (i) $u \circ \natural = \natural \circ u = \natural$, (ii) $x \circ y = \natural$ for all $x, y \in V$, (iii) $(x, y) \circ z = \natural$, $x \circ (x, z) = z$, and $x \circ (y, z) = \natural$ if $y \neq z$, and (iv) $(x, y) \circ (y, v) = (x, v)$ and $(x, y) \circ (u, v) = \natural$ if $y \neq u$. The desired weight function \tilde{w} is defined as $\tilde{w}(\natural) = |X|$, $\tilde{w}(x, y) = w(x) + w(y)$, and $\tilde{w}(x) = w(x)$ for all $x, y \in V$. Note that $\tilde{w}(x) \leq \max\{|X|, 2|V|\} \leq |X|$ for all $x \in X$ and $\tilde{w}(x_1) + \sum_{i=1}^k \tilde{w}(x_i, x_{i+1}) + \tilde{w}(x_k) = 2 \sum_{i=1}^k w(x_i)$ for $x_1, x_2, \dots, x_k \in V$.

Note that, given an s - t path (x_1, x_2, \dots, x_k) with $s = x_1$ and $t = x_k$, since $s \in S$, it is possible for us to prove recursively the membership $x_i \in G(S)$ for every $i \in [2, k]_{\mathbb{Z}}$; hence, $t \in G(S)$ follows. For any sequence $u = (x_1, (x_1, x_2), (x_2, x_3), \dots, (x_{k-1}, x_k))$ in $\text{SOL}_0(f(h, r))$ with $s = x_1$ and $t = x_k$, we define $g(h, u, r) = (x_1, x_2, \dots, x_k)$. It follows that $m_0(f(h, r), u) = \tilde{w}(x_1) + \sum_{i=1}^k \tilde{w}(x_i, x_{i+1}) = 2 \sum_{i=1}^k w(x_i) - w(t)$ and $m_1(h, g(h, u, r)) = \sum_{i=1}^k w(x_i)$. From those equalities, we obtain $2m_1(h, g(h, u, r)) = m_0(f(h, r), u) + w(t)$.

Here, we intend to verify that (f, g, c) correctly reduces MIN BPATH-WEIGHT to MIN AGEN. Take any $r \in \mathbb{Q}^{\geq 1}$ and any $u \in \text{SOL}_0(f(h, r))$. Assume that the performance ratio R_0 for MIN AGEN satisfies $R_0(f(h, r), u) \leq r$. It then follows that $R_1(h, g(h, u, r)) = \frac{m_1(h, g(h, u, r))}{m_1^*(h)} = \frac{m_0(f(h, r), u) + w(t)}{m_0^*(f(h, r)) + w(t)} \leq \frac{m_0(f(h, r), u)}{m_0^*(f(h, r))} \leq r$. Therefore, (f, g, c) is a correct $\leq_{\text{sAP}}^{\text{NC}^1}$ -reduction. \square

Jenner [23] studied a few variants of the well-known *knapsack problem*. In particular, she introduced a decision problem, called CUK, of determining whether, given unary string pairs $(0^w, 0^p), (0^{w_0}, 0^{p_0}), (0^{w_1}, 0^{p_1}), \dots, (0^{w_n}, 0^{p_n})$, where $w, w_i \in \mathbb{N}^+$ and $p, p_i \in \mathbb{N}$ for $i \in [n]$, there is a $\{0, 1\}$ -sequence $(z_0, z_1, z_2, \dots, z_n)$ satisfying $w2^p = \sum_{i=0}^n z_i \cdot w_i 2^{p_i}$. She showed that this problem is (L-uniform) $\leq_m^{\text{NC}^1}$ -complete for NL. Here, we turn this decision problem into an optimization problem, which will be proven to be $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete for $\text{APXL}_{\text{NLO}} \cap \text{PBO}$.

MAXIMIZATION 2-BOUNDED CLOSE-TO-UNARY KNAPSACK PROBLEM (MAX 2BCU-KNAPSACK):

- INSTANCE: a pair $(0^w, 0^p)$, a sequence $(0^{w_0}, 0^{p_0}), (0^{w_1}, 0^{p_1}), \dots, (0^{w_n}, 0^{p_n})$ of pairs, and a weight sequence (c_0, c_1, \dots, c_n) of positive integers, where $w, w_i \in \mathbb{N}^+$, $p, p_i \in \mathbb{N}$, $1 \leq c_i \leq nwp$ for all $i \in [0, n]_{\mathbb{Z}}$, and $\max_{0 \leq i \leq n} \{c_i\} \leq 2 \min_{0 \leq i \leq n} \{c_i\}$, provided that $w_0 = w$ and $p_0 = p$. Here, the notation 0^0 expresses the empty string λ .
- SOLUTION: a sequence $z = (z_0, z_1, \dots, z_n)$ of Boolean values satisfying $w2^p = \sum_{i=0}^n z_i \cdot w_i 2^{p_i}$.
- MEASURE: $\max_{0 \leq i \leq n} \{c_i z_i\}$.

A trivial solution $z = (1, 0, \dots, 0)$, which indicates the choice of $(0^{w_0}, 0^{p_0})$, is needed to ensure that MAX 2BCU-KNAPSACK is indeed in APXL_{NLO} . Regarding each value $w_i 2^{p_i}$, the following simple inequalities hold: $|\text{bin}(w_i 2^{p_i})| \leq \log |0^{w_i}| + |0^{p_i}| + 1 = \log w_i + p_i + 1$.

In what follows, we show the completeness of MAX 2BCU-KNAPSACK for $\text{APXL}_{\text{NLO}} \cap \text{PBO}$.

Lemma 4.6 MAX 2BCU-KNAPSACK is $\leq_{\text{sAP}}^{\text{NC}^1}$ -complete for $\text{APXL}_{\text{NLO}} \cap \text{PBO}$.

Proof. Let MAX 2BCU-KNAPSACK = $(I_0, \text{SOL}_0, m_0, \text{MAX})$. First, we argue that MAX 2BCU-KNAPSACK is in $\text{APXL}_{\text{NLO}} \cap \text{PBO}$. It is obvious that m_0 is polynomially bounded. Earlier, Jenner [23] demonstrated that CUK \in NL. A similar argument shows that $I_0 \circ \text{SOL}_0 \in \text{auxL}$ and $m_0 \in \text{auxFL}$; therefore, MAX 2BCU-KNAPSACK falls into $\text{NLO} \cap \text{PBO}$. Let $x = (K, C)$ be any instance to MAX 2BCU-KNAPSACK, where $C = (c_0, c_1, \dots, c_n)$ and K is composed of $(0^w, 0^p), (0^{w_0}, 0^{p_0}), (0^{w_1}, 0^{p_1}), \dots, (0^{w_n}, 0^{p_n})$. Choose $z_0 = (1, 0, \dots, 0) \in \{0, 1\}^{n+1}$. It follows that $z_0 \in \text{SOL}_0(x)$ and that $R_0(x, z_0) = \frac{m_0^*(x)}{m_0(x, z_0)} \leq 2$ since $\min_{0 \leq i \leq n} \{c_i\} \leq m_0^*(x) \leq 2 \min_{0 \leq i \leq n} \{c_i\}$. This implies that MAX 2BCU-KNAPSACK is in APXL_{NLO} .

For the $\leq_{\text{sAP}}^{\text{NC}^1}$ -hardness of MAX 2BCU-KNAPSACK, let us consider an arbitrary maximization problem $P = (I, \text{SOL}, m, \text{MAX})$ in $\text{APXL}_{\text{NLO}} \cap \text{PBO}$ that satisfies Conditions (ii)–(iv) of Lemma 3.10. To construct an $\leq_{\text{sAP}}^{\text{NC}^1}$ -reduction $(f, g, 2)$ from P to MAX 2BCU-KNAPSACK, we first consider a log-space deterministic Turing machine M_P that produces 2-approximate solutions of P . In addition, we take a function $b \in \text{FL}$ for which $b(x) \leq m^*(x) \leq 2b(x)$ and $m(x, y) \geq b(x)$ for all $(x, y) \in I \circ \text{SOL}$ by Lemma 3.10. Let M_m be a log-space auxiliary Turing machine computing m .

For the purpose of defining an appropriate reduction, we assume that M_m on input (x, y) writes each symbol of y on a designated cell of a particular work tape so that,

The definition of partial configurations makes it possible to retrieve the entire string y from any halting computation path e of M_m on input (x, y) . We write $h(e)$ for the auxiliary input y fed into M_m . Since m

is polynomially bounded, we force M_m to calculate $m(x, y)$ on one of its work tapes and copies it onto an output tape just before halting so that, for any $y \in SOL(x)$, $m(x, y)$ equals the sum of the numbers written on the output tape of configuration along a computation path of M_m on input (x, y) .

Let us recall the notion of configuration graph from the proof of Proposition 3.3 and consider a configuration graph $G_x^{M_m}$ of M_m on input $x \in I$. In what follows, we consider only the case where $SOL(x) \neq \emptyset$. For convenience, we set $G_x^{M_m} = (V, E)$ with $V = \{v_1, v_2, \dots, v_m\}$ for a certain integer $m \geq 1$, where v_1 expresses the initial partial configuration of M_m . Note that the first move of M_m is fixed and does not depend on the choice of inputs. We further make M_m terminate in exactly $q(|x|)$ steps for a suitable polynomial q , independent of the choice of y satisfying $|y| \leq p(|x|)$.

Fix $x \in (I \circ SOL)^\exists$ and write q for $q(|x|)$. Take an arbitrary number $r \in \mathbb{Q}^{\geq 1}$. Following [23, Theorem 1], let $w = 2^t$ and $p = 2(q+1)t$, where $t = 2\lceil \log q \rceil$. If $(v_i, v_j) \in E$, then let $w_{ij} = (2^t - j)2^{2t} - (2^t - i)$ and $p_{ij,k} = 2kt$ for $0 \leq k \leq q-1$. Let $w_0 = w$, $p_0 = p$, $w_1 = 2^t$, $p_1 = 0$, $w_q = q2^{2t}$, and $p_q = 2(q-1)t$. Let K be composed of the following pairs: $(0^w, 0^p)$, $(0^{w_0}, 0^{p_0})$, $(0^{w_1}, 0^{p_1})$, $(0^{w_q}, 0^{p_q})$, and $(0^{w_{ij}}, 0^{p_{ij}})$ for all $(v_i, v_j) \in E$. A series $C = (c_0, c_1, c_q, c_{ij})_{(v_i, v_j) \in E}$ is defined as follows. For any partial configuration pair (v_i, v_j) , if v_j is an accepting partial configuration, then let c_{ij} be the number written on an output tape of M_m ; otherwise, let $c_{ij} = b(x)$. Note that $\min_{(v_i, v_j) \in E} \{c_0, c_1, c_q, c_{ij}\} \geq b(x)$. For any $x \in I$ and $r \in \mathbb{Q}^{\geq 1}$, we define $f(x, r)$ to be the pair (K, C) given above.

Given any accepting computation path $e = (v_1, v_2, \dots, v_q)$ of M_m on input x ($|x| = n$), we define a series $z = (z_0, z_1, z_q, z_{ij})_{(v_i, v_j) \in E}$ as $z_0 = 0$, $z_1 = z_q = 1$, and $z_{ij} = 1$ if $(v_i, v_j) \in e$, and $z_{ij} = 0$ otherwise. To emphasize e , we write $z_{(e)}$ for this series z . As was shown in [23], we obtain $z \in SOL_0(f(x, r))$ with $z \neq (1, 0, \dots, 0)$ iff there exists an accepting computation path e satisfying $z = z_{(e)}$, namely, $w2^p = w_12^{p_1} + w_q2^{p_q} + \sum_{(v_i, v_j) \in E} \sum_{0 \leq k \leq q-1} w_{ij}2^{p_{ij,k}}$ and thus $m_0(f(x, r), z_{(e)}) = \max_{(v_i, v_j) \in E} \{c_0, c_1, c_q, c_{ij}\} = m(x, h(e))$ since $m(x, y) \geq b(x)$ for all $y \in SOL(x)$. From this fact, we define $g(x, z_{(e)}, r) = h(e)$ for every solution $z_{(e)} \in SOL_0(f(x, r))$. Since $z \in SOL_0(f(x, r))$, there exists a suitable accepting computation path e in $G_x^{M_m}$ satisfying $z = z_{(e)}$. Thus, we obtain

$$R(x, g(x, z_{(e)}, r)) = \frac{m^*(x)}{m(x, g(x, z_{(e)}, r))} = \frac{m_0^*(f(x, r))}{m_0(f(x, r), z_{(e)})} = R_0(f(x, r), z_{(e)}).$$

Therefore, $(f, g, 2)$ reduces P to MAX 2BCU-KNAPSACK. \square

Next, we shall present a complete problem in $NC^1_{ONLO} \cap PBO$. Our problem is a simple extension of the Boolean formula value problem. The *Boolean formula value problem* (BFVP) is to determine whether a given Boolean formula ϕ is satisfied by a given truth assignment σ . This problem can be compared with CVP (circuit value problem), which is known to be P-complete [33]. Buss [9] showed the membership of BFVP to ALOGTIME and the $\leq_m^{DLOGTIME}$ -hardness of BFVP for ALOGTIME is relatively easy to verify. Therefore, since ALOGTIME equals NC^1 , BFVP is $\leq_m^{DLOGTIME}$ -complete for NC^1 .

Let us consider its optimization counterpart.

MAXIMUM BOOLEAN FORMULA VALUE PROBLEM (MAX BFVP):

- INSTANCE: a set $\Phi = \{\phi_1, \phi_2, \dots, \phi_n\}$ of Boolean formulas and a Boolean assignment σ for all variables in the formulas in Φ with $n \geq 1$.
- SOLUTION: a *nonempty* subset $S \subseteq \Phi$ of formulas satisfied by σ .
- MEASURE: the cardinality $|S|$ of S .

Recall that the formula language of a Boolean formula family $\{\phi_n\}_{n \in \mathbb{N}}$ is composed of all tuples $\langle c, i, y \rangle$ such that $|y| = n$ and the i th character of the n th formula ϕ_n is c . Following [7], a language L is in NC^1 iff there exists a family of Boolean formulas representing L whose formula language is in DLOGTIME.

Lemma 4.7 MAX BFVP is $\leq_{sAP}^{AC^0}$ -complete for $NC^1_{ONLO} \cap PBO$.

Proof. For our convenience, express MAX BFVP as $(I_0, SOL_0, m_0, \text{MAX})$. Given a set $\Phi = \{\phi_1, \phi_2, \dots, \phi_n\}$ of Boolean formulas and its truth assignment σ , since the number of all Boolean formulas ϕ_i satisfied by σ is upper-bounded by the input size, MAX BFVP is polynomially bounded. Next, we argue that MAX BFVP is in NC^1_{ONLO} . It is obvious that MAX BFVP belongs to NLO. On input (Φ, σ) , we check in parallel whether σ satisfies ϕ_i by making nonadaptive queries to oracle BFVP (i.e., $(\phi_i, \sigma) \in \text{BFVP}$), and finally we count the number of satisfied formulas ϕ_i . The last counting process requires another NC^1 -circuit. Since $\text{BFVP} \in NC^1$ [9], we can implement the whole procedure using NC^1 -circuits. Thus, MAX BFVP belongs to NC^1_{ONLO} .

To see the $\leq_{\text{sAP}}^{\text{AC}^0}$ -hardness of MAX BFVP for $\text{NC}^1_{\text{ONLO}} \cap \text{PBO}$, let $P = (I, \text{SOL}, m, \text{MAX})$ be any maximization problem in $\text{NC}^1_{\text{ONLO}} \cap \text{PBO}$. The minimization case can be similarly treated. Let us take a uniform family $\{C_n\}_{n \in \mathbb{N}}$ of NC^1 -circuits computing maximal solutions of P , that is, $C_{|x|}(x) \in \text{SOL}^*(x)$ for any $x \in (I \circ \text{SOL})^\exists$. Define $\text{SOL}'(x) = \{C_{|x|}(x) \mid C_{|x|}(x) \neq \perp\}$ for each $x \in I$. Note that $|\text{SOL}'(x)| \leq 1$ holds for all $x \in I$. For simplicity, we assume that $I \subseteq \{0, 1\}^*$ and take any binary string $x = x_1x_2 \cdots x_n$ in I . Let the output size of C_n be k . For each index $i \in [k]$, we define a new AC^0 -circuit $D_k^{(i)}(y_1, \dots, y_k) = y_i$. It follows that $D_k^{(i)}(C_n(z_1, \dots, z_n))$ is an NC^1 -circuit, where z_1, z_2, \dots, z_n are Boolean variables.

Owing to [7, Theorem 9.1], we can replace a uniform family of NC^1 -circuits by a family of Boolean formulas whose formula language is in DLOGTIME . For each $i \in [k]$, let ϕ_i be a Boolean formula expressing the circuit $D_k^{(i)}(C_n(z_1, \dots, z_n))$. Note that $D_k^{(i)}(C_n(x_1 \cdots x_n)) = u_i$ if $C_n(x_1x_2 \cdots x_n) = u_1u_2 \cdots u_k$. Let σ_x be the assignment that assigns value x_i to variable z_i .

Finally, we define $f(x, r) = \langle \langle \phi_1, \dots, \phi_m \rangle, \sigma_x \rangle$ and $g(x, y, r) = u_1u_2 \cdots u_k$ for each $y \in \text{SOL}_0(f(x, r))$, where y is of the form $(\phi_{i_1}, \phi_{i_2}, \dots, \phi_{i_t})$, and $u_i = 1$ if $\phi_i \in y$ and $u_i = 0$ otherwise. It follows that $|\text{SOL}_0(f(x, r))| \leq 1$ for all $x \in I_0$. If $y = (\phi_{i_1}, \phi_{i_2}, \dots, \phi_{i_t}) \in \text{SOL}_0(f(x, r))$ with $f(x, r) = \langle \langle \phi_1, \dots, \phi_m \rangle, \sigma_x \rangle$, then σ_x satisfies $\{\phi_{i_1}, \phi_{i_2}, \dots, \phi_{i_t}\}$. Hence, we obtain $g(x, y, r) \in \text{SOL}^*(x)$. Hence, $\frac{m^*(x)}{m(x, g(x, y, r))} = \frac{m^*(x)}{m(x, u_1u_2 \cdots u_k)} = 1$. It thus follows that P is $\leq_{\text{sAP}}^{\text{AC}^0}$ -reducible to MAX BFVP. \square

5 Weak Approximation Schemes

Approximability of optimization problems is a crucial concept in a course of our study on the computational complexity of those problems. The classes APXL_{NLO} and $\text{APXNC}^1_{\text{NLO}}$ both admit approximation algorithms; however, the effectiveness of approximation is often far from the desirable one used in practice. Approximation schemes, on the contrary, provide a much finer notion of approximability because the performance ratio of used approximation algorithms can be made arbitrarily smaller; in other words, we can find solutions that are arbitrarily close to optimal solutions.

We shall discuss optimization problems that admit such approximation schemes. Concerning LSAS_{NLO} , Nickelsen and Tantau [34] and also Tantau [39] proposed NLO problems that admit an LSAS_{NLO} ; in particular, a polynomially-bounded maximization problem, called MAX-HPP, was shown in [39, Theorem 5.7] as a member of LSAS_{NLO} . Here, we rephrase this problem in terms of *complete graphs* in comparison with MIN BPATH-WEIGHT .

MAXIMUM COMPLETE-GRAPH PATH WEIGHT PROBLEM (MAX CPATH-WEIGHT):

- INSTANCE: a directed complete graph $G = (V, E)$ with self-loops, a source $s \in V$, and an edge weight function $w : V \times V \rightarrow \mathbb{N}^+$ with $w(v_1, v_2) \leq |V|$ for any $v_1, v_2 \in V$.
- SOLUTION: a path $\mathcal{S} = (v_1, v_2, \dots, v_k)$ of length $k \leq |V|$ in G starting at s (i.e., $s = v_1$).
- MEASURE: total path weight $w(\mathcal{S}) = \sum_{i=1}^{k-1} w(v_i, v_{i+1})$.

Proposition 5.1 1. [39] MAX CPATH-WEIGHT is in $\text{LSAS}_{\text{NLO}} \cap \text{PBO}$.

2. MAX CPATH-WEIGHT is $\leq_{\text{EX}}^{\text{NC}^1}$ -hard for $\text{LO}_{\text{NLO}} \cap \text{PBO}$.

Proof. (2) Let MAX CPATH-WEIGHT = $(I_0, \text{SOL}_0, m_0, \text{MAX})$. Given any maximization problem $P = (I, \text{SOL}, m, \text{MAX})$ in $\text{LO}_{\text{NLO}} \cap \text{PBO}$, we want to define an $\leq_{\text{EX}}^{\text{NC}^1}$ -reduction (f, g) from P to MAX CPATH-WEIGHT in the following fashion. Let M_P be a log-space deterministic Turing machine computing optimal solutions of P ; in other words, $m^*(x) = m(x, M_P(x))$ for all $x \in (I \circ \text{SOL})^\exists$ with $\text{SOL}(x) \neq \emptyset$. Let x be any instance in $(I \circ \text{SOL})^\exists$ and consider a configuration graph with $G_x^{M_P} = (V_x, E_x)$ of M_P on x . Let s denote an initial partial configuration of M_P on x .

Define $V = V_x$ and $E = V \times V$ and set $G = (V, E)$. Since M is deterministic, there exists a unique computation path (p_1, p_2, \dots, p_k) starting at $s = p_1$ and ending at an accepting configuration $t = p_k$. As for an edge weight function w , we define $w(v, v') = |V|$ if either (i) v is an accepting configuration and $v = v'$ or (ii) v' is obtained from v in a single step by M_P ; otherwise, let $w(v, v') = 1$. The desired f is obtained by setting $f(x) = (G, s, w)$. This f can be computed by an appropriate AC^0 -circuit. It is not difficult to see that any optimal solution of MAX CPATH-WEIGHT is $u_x = (p_1, p_2, \dots, p_k, p_k, \dots, p_k)$ of length exactly $|V|$. Since (p_1, p_2, \dots, p_k) is a computation path, we can retrieve from it an output string produced by M_P on x . More generally, given a path u starting at s and ending at a certain accepting configuration, we reconstruct from u an output string y_u of M_P on x . Using these strings, we set $g(x, u) = y_u$. This function g can be implemented by a certain AC^0 -circuit. Note that $m_0^*(f(x)) = (|V| - 1)|V|$.

If $u \in SOL_0^*(f(x))$, then, since $m_0^*(f(x)) = (|V| - 1)|V|$, u coincides with u_x . Hence, y_u equals $M_P(x)$. Thus, (f, g) reduces P to MAX CPATH-WEIGHT. \square

Next, we shall look into another approximation class $NC^1AS_{NLO} \cap PBO$. The next example is slightly artificial but it can be proven to fall into NC^1AS_{NLO} . Recall from Section 4 MAX U2-VERTEX, which is a member of $APXNC_{NLO}^1 \cap PBO$. We modify this problem by changing its requirement of $\max_{v \in V} \{w(v)\} \leq 2 \min_{v \in V} \{w(v)\}$ for each instance (G, s, w) to $\max_{v \in V} \{w(v)\} \leq (1 + 1/k_G) \min_{v \in V} \{w(v)\}$ with $k_G = \log \log |V|$. We call the resulted problem MAX UAPP-VERTEX (maximum undirected approximable vertex weight problem).

Proposition 5.2 1. MAX UAPP-VERTEX is in $NC^1AS_{NLO} \cap PBO$.

2. MAX UAPP-VERTEX is $\leq_{SAP}^{AC^0}$ -hard for $LO_{NLO} \cap PBO$.

Proof. (1) We express MAX UAPP-VERTEX as $(I_0, SOL_0, m_0, \text{MAX})$ and we plan to define its approximation scheme, which can be implemented by certain AC^0 -circuits. Let $r \in \mathbb{Q}^{>1}$ and set $k = \frac{1}{r-1}$. Take any instance (x, r) in I_0 given to MAX UAPP-VERTEX with $G = (V, E)$. We consider two cases separately, depending on the value of k .

(i) Assume that $k \leq \log \log |V|$. It follows from the requirement for I_0 that $\max_{v \in V} \{w(v)\} \leq (1 + 1/k) \min_{v \in V} \{w(v)\}$. In this case, our desired approximation algorithm first searches the nearest vertex, say, v_0 connected from s and directly outputs it. This vertex v_0 gives $R(x, v_0) = \frac{m_0^*(x)}{w(v_0)} \leq \frac{\max_{v \in V} \{w(v)\}}{\min_{v \in V} \{w(v)\}} \leq 1 + 1/k = r$. This procedure provides an r -approximate solution and can be easily implemented by an appropriate AC^0 -circuit.

(ii) Assume that $k > \log \log |V|$; in other words, $|V| < 2^{2^k}$. Consider the following brute force algorithm: pick a vertex y one by one, check if s and y are connected, and choose the one y that has the largest value $w(y)$. This procedure requires space $O(|V| \log |V|)$, which equals $O(2^{2^k} \log n)$.

(2) Similarly to the proof of Proposition 4.3, let $P = (I, SOL, m, \text{MAX})$ denote any maximization problem in $LO_{NLO} \cap PBO$. Write M_m for a log-space auxiliary Turing machine computing m and write M_P for a log-space deterministic Turing machine finding maximal solutions of P . Let $b(x) = m(x, M_P(x))$ for all $x \in (I \circ SOL)^\exists$. Since b is in FL, choose a log-space deterministic Turing machine M' that computes b together with generating each symbol of $M_P(x)$ for the purpose of the later reconstruction of $M_P(x)$ from any halting computation patj of M' on x . For the problem P , we assume Conditions (ii)–(iii) of Lemma 3.8. Note that $|SOL(x)| \leq 1$ holds for all instances $x \in I$.

Let us define $f(x, r) = (G, s, w)$, where $G = (V, E)$ is a configuration graph of M' on x and s is the initial configuration. Here, we assume that M' is $c \log n$ space-bounded for a certain constant $c > 0$; moreover, by modifying M' slightly, we also assume that M' has one work tape with tape alphabet of cardinality at least 2. Because of this space bound, there are at least $2^{c \log n} = n^c$ vertices in G . Thus, $\log \log |V| \geq \log \log n$ follows. Let $w(v)$ be $\lceil \log \log n \rceil + 1$ if v is a halting partial configuration of M' on x ; otherwise, let $w(v) = \lceil \log \log n \rceil$. It follows that $\lceil \log \log n \rceil \leq m(x, y) \leq \lceil \log \log n \rceil + 1$ for all path y in G . Hence, $\max_{y \in SOL(f(x, r))} \{m(x, y)\} \leq \lceil \log \log n \rceil + 1 \leq (1 + 1/\log \log n) \lceil \log \log n \rceil \leq \min_{y \in SOL(f(x, r))} \{m(x, y)\}$. Next, we define $g(x, y, r)$ to be the string $M_P(x)$ recovered from y . This is possible by the assumption on the behavior of M' . Finally, we set $c' = 1$. It is not difficult to show that (f, g, c') reduces P to MAX UAPP-VERTEX. \square

6 Relations among Refined Optimization/Approximation Classes

Let us turn our attention to relationships among classes of refined optimization problems introduced in Section 2. Optimization problems are, by definition, associated directly with their *underlying decision problems*. It is therefore natural to ask what relationships exist between classes of optimization problems and classes of decision problems. We intend to investigate relationships between optimization problems and decision problems. First, let us recall from Lemma 2.2 the collapse of classes PO_{NLO} , $PTAS_{NLO}$, and $APXP_{NLO}$ down to NLO. The class NLO may further collapse to much lower-complexity classes if AC^1 collapses to L or NC^1 . Here, we prove the following assertions concerning the class NLO.

Proposition 6.1 1. [39] If $L = AC^1$, then $LO_{NLO} = NLO$.

2. If $NC^1 = AC^1$, then $NC^1O_{NLO} = NLO$.

The opposite direction of this proposition is not known to hold. The proposition comes directly from a more general result stated below. In particular, Lemma 6.2(1) extends Lemma 2.2.

- Lemma 6.2**
1. $AC^1O_{NLO} = NLO$.
 2. If $L = AC^1$, then $LO_{NLO} = AC^1O_{NLO}$.
 3. If $NC^1 = AC^1$, then $NC^1O_{NLO} = AC^1O_{NLO}$.

Proof. (1) Obviously, AC^1O_{NLO} is included in NLO . For the opposite inclusion, consider $MAX \lambda$ -NFA, given in Section 3.3, which is $\leq_{sAP}^{NC^1}$ -complete for $MaxNL$ [39, Theorem 3.1]. It is sufficient to show that (i) $MAX \lambda$ -NFA is in AC^1O_{NLO} and (ii) AC^1O_{NLO} is closed under \leq_{sAP}^L (and thus $\leq_{sAP}^{NC^1}$).

(i) Álvarez and Jenner [5] also Allender et al. [2] showed the containment $OptL \subseteq FAC^1$. Take any λ -Infa M as an instance to $MAX \lambda$ -NFA. We want to find a maximal solution using AC^1 circuits. For this purpose, using AC^1 circuits, we convert M into an equivalent regular grammar G in *Chomsky Normal Form*. As done in the proof of [2, Corollary 4.6], we can find the lexicographically first string over $\{0, 1\}$ of length at most n produced by G . This process can be implemented on AC^1 -circuits. Therefore, $MAX \lambda$ -NFA can be solved by AC^1 -circuits.

(ii) Assume that $P \leq_{sAP}^L Q$ and $Q \in AC^1O_{NLO}$. Take an \leq_{sAP}^L -reduction (f, g, c) from P to Q . Let $P = (I_1, SOL_1, m_1, MAX)$ and $Q = (I_2, SOL_2, m_2, MAX)$. Let C_Q be an AC^1 circuit computing optimal solutions of Q . Consider the following algorithm D . On input $x \in I_1$, compute $f(x, 1)$ and output $C_Q(f(x, 1))$. Clearly, $D(x) \in SOL_1(x)$. Note that FAC^1 is closed under functional composition. Since $FL \subseteq FAC^1$, D is also in FAC^1 . Hence, P is in AC^1O_{NLO} .

(2) Assume that $L = AC^1$. This implies that $FL = FAC^1$. Take $P = (I, SOL, m, goal)$ in AC^1O_{NLO} . There is an AC^1 circuit C_P computing optimal solutions of P and another AC^1 -circuit C_m that computes $m(x, C_P(x))$ for all $x \in (I \circ SOL)^\exists$. Treating circuits as functions they compute, we obtain $C_P, C_m \in FAC^1$. Since $FL = FAC^1$, C_P and C_m fall into FL . This implies that P belongs to LO_{NLO} . Thus, $LO_{NLO} = AC^1O_{NLO}$ follows.

(3) Here, we assume that $NC^1 = AC^1$. Note that this assumption implies $FNC^1 = FAC^1$. Let us consider an arbitrary optimization problem $P = (I, SOL, m, goal)$ in AC^1O . Take AC^1 -circuits that compute optimal solutions of P as well as their values. By a way similar to (2), since $FNC^1 = FAC^1$, P 's optimal solutions can be computed using appropriate NC^1 -circuits. As a result, P belongs to NC^1O_{NLO} . \square

As revealed in Section 4, the polynomial-boundedness property is crucial for optimization problems in discussing their computational complexity. If we fix our focal point on polynomially-bounded optimization problems, we can manage to give another characterization of $NLO \cap PBO$ in terms of adaptive relativization.

To explain a notion of relativization, we empower each log-space Turing machine M with *query mechanism*, in which M makes a series of “queries” to a given “oracle” A (which is simply a language) and A returns its answers to M in a single step. The machine M has an extra *query tape* on which M writes a query string z by moving its tape head from the left to the right. After entering a designated inner state, called a *query state*, the string z is transmitted to the oracle A . The oracle erases all symbols on the query tape and it writes 1 if $z \in A$ and writes 0 otherwise. The machine’s inner state is changed to an answer state, from which M resumes its computation. We make the length of every query string polynomially bounded. This machine M is generally known as an *oracle Turing machine*. We write LO_{NLO}^A to denote the class of all NLO problems whose optimal solutions and their values are computed by log-space oracle Turing machines using oracles A . The union of LO_{NLO}^A for all sets $A \in NL$ is denoted by LO_{NLO}^{NL} .

Lemma 6.3 $LO_{NLO}^{NL} \cap PBO = NLO \cap PBO$.

Proof. (\subseteq) This inclusion is trivial because $LO_{NLO}^{NL} \subseteq NLO$ holds by the definition of LO_{NLO}^{NL} .

(\supseteq) Let $P = (I, SOL, m, goal)$ be any NLO problem. It suffices to show the existence of log-space oracle Turing machines that compute optimal solutions and their values relative to oracles in NL . Here, we consider only the case of $goal = MAX$. We define two sets A and B . Take a polynomial p for which $y \in SOL(x)$ implies both $|y| \leq p(|x|)$ and $m(x, y) \leq p(|x|)$. Let A be composed of all strings (x, k) such that there exists a string $z \in SOL(x)$ with $k \leq p(|x|)$ satisfying $m(x, z) = k$. By making a series of queries $(x, 1), (x, 2), \dots, (x, p(|x|))$ one by one to oracle A , we can find the maximal number k_0 satisfying $(x, k_0) \in B$. Next, we define B as the set of all strings (x, k, u, b) with $b \in \{0, 1\}$ and $k, |ub| \leq p(|x|)$ such that there is a string $z \in SOL(x)$ for which $m(x, z) = k$, ub is an initial segment of z . Given the maximal number k_0 , by making the i th query $(x, k_0, u_i, 0)$ and $(x, k_0, u_i, 1)$ for all $i \in [0, p(|x|)]_{\mathbb{Z}}$ sequentially, we reconstruct z_0 satisfying $m(x, z_0) = k_0$. As the desired oracle, we take $A \oplus B$, which clearly belongs to NL . Therefore, we find a maximal solution

z_0 using only log space. This yields $P \in \text{LO}_{\text{NLO}}^{A \oplus B} \subseteq \text{LO}_{\text{NLO}}^{\text{NL}}$. \square

Here, we present three close connections between classes of decision problems and classes of optimization problems.

Proposition 6.4 1. [39] $L = \text{NL}$ iff $\text{LO}_{\text{NLO}} \cap \text{PBO} = \text{NLO} \cap \text{PBO}$.
 2. $\text{NC}^1 = L$ iff $\text{NC}^1_{\text{ONLO}} \cap \text{PBO} = \text{LO}_{\text{NLO}} \cap \text{PBO}$ iff $\text{NC}^1_{\text{ONLO}} = \text{LO}_{\text{NLO}}$.
 3. $L = P$ iff $\text{LO}_{\text{NPO}} \cap \text{PBO} = \text{PO}_{\text{NPO}} \cap \text{PBO}$ iff $\text{LO}_{\text{NPO}} = \text{PO}_{\text{NPO}}$.

For Proposition 6.4(1)–(2), we use DSTCON and USTCON on unweighted directed graphs. Earlier, Jones [26] showed that DSTCON is NL-complete under log-space reduction. His reduction can be improved to NC^1 -reduction, and thus DSTCON is $\leq_m^{\text{NC}^1}$ -complete for NL. Reingold [36] proved that USTCON belongs to L. From its $\leq_m^{\text{NC}^1}$ -hardness for L, we conclude that USTCON is $\leq_m^{\text{NC}^1}$ -complete for L. For Proposition 6.4(3), we use the *circuit value problem (CVP)* of determining whether a given Boolean circuit outputs 1 on a certain input. Ladner [33] pointed out that this is \leq_m^L -complete for P; more strongly, it is $\leq_m^{\text{NC}^1}$ -complete for P.

Proof of Proposition 6.4. (1) Here, we give a proof different from the one given in [39]. Our proof relies on Lemma 6.3. If $L = \text{NL}$, then $\text{LO}_{\text{NLO}}^{\text{NL}}$ equals LO_{NLO}^L , which coincides with LO_{NLO} . Since $\text{LO}_{\text{NLO}}^{\text{NL}} \cap \text{PBO} = \text{NLO} \cap \text{PBO}$ by Lemma 6.3, we obtain $\text{LO}_{\text{NLO}} \cap \text{PBO} = \text{NLO} \cap \text{PBO}$. The desired consequence follows because $\text{LO}_{\text{NLO}} = \text{LO}_{\text{NLO}}^L$.

Conversely, assume that $\text{LO}_{\text{NLO}} \cap \text{PBO} = \text{NLO} \cap \text{PBO}$. Lemma 6.3 implies that $\text{LO}_{\text{NLO}}^{\text{NL}} \cap \text{PBO} = \text{LO}_{\text{NLO}} \cap \text{PBO}$. Consider DSTCON, which is $\leq_m^{\text{AC}^0}$ -complete for NL. Define $P = (I, \text{SOL}, m, \text{MAX})$ as follows. Let I consist of all strings (G, s, t) for which $G = (V, E)$, $s, t \in V$, and there is an s - t path in G . Let $x = (G, s, t)$. Define $\text{SOL}(x)$ to be a collection of all s - t path in G . Let $m(x, y) = 2$ if $y \in \text{SOL}(x)$ and $m(x, y) = 1$ otherwise. Note that P is in $\text{NLO} \cap \text{PBO}$. By our assumption, P belongs to LO_{NLO} . Hence, take a log-space Turing machine M that computes optimal solutions of P . It follows that $M(x) \neq \perp$ iff $x \in \text{DSTCON}$. Hence, DSTCON can be recognized using log space, yielding $L = \text{NL}$.

(2) For this claim, we shall show separately that (i) if $\text{NC}^1_{\text{ONLO}} \cap \text{PBO} = \text{LO}_{\text{NLO}} \cap \text{PBO}$, then $\text{NC}^1 = L$ and (ii) if $\text{NC}^1 = L$, then $\text{NC}^1_{\text{ONLO}} = \text{LO}_{\text{NLO}}$. Since $\text{NC}^1_{\text{ONLO}} = \text{LO}_{\text{NLO}}$ implies $\text{NC}^1_{\text{ONLO}} \cap \text{PBO} = \text{LO}_{\text{NLO}} \cap \text{PBO}$, the claim follows immediately.

(i) First, we assume that $\text{NC}^1_{\text{ONLO}} \cap \text{PBO} = \text{LO}_{\text{NLO}} \cap \text{PBO}$. Since USTCON is $\leq_m^{\text{NC}^1}$ -complete for L, it suffices for us to verify that USTCON falls into NC^1 using our assumption. We intend to define a polynomially-bounded maximization problem $P = (I, \text{SOL}, m, \text{MAX})$ in the following manner. An instance w to P is a tuple (G, s, t) for an undirected graph $G = (V, E)$, including an edge (s, t) , and two vertices $s, t \in V$. A feasible solution of w is a path $y = (y_1, y_2, \dots, y_k)$ in G starting at $s = y_1$. If y is a path from s to t with $k \geq 2$, we set its objective value $m(w, y)$ to be 2; otherwise, we set $m(w, y) = 1$. Obviously, m is polynomially bounded and belongs to auxFL. Note that $w \in \text{USTCON}$ iff there exists an s - t path y for which $m(w, y) = 2$. Next, we claim that P is in LO_{NLO} . Consider a log-space deterministic Turing machine M that computes 2-approximate solutions by simply outputting the special edge (s, t) . Hence, P falls into $\text{LO}_{\text{NLO}} \cap \text{PBO}$.

Since $P \in \text{PBO}$ and $\text{LO}_{\text{NLO}} \subseteq \text{NC}^1_{\text{ONLO}}$, our assumption implies that P is in $\text{NC}^1_{\text{ONLO}}$. There exists a uniform family $\{C_n\}_{n \in \mathbb{N}}$ of NC^1 -circuits finding maximal solutions of P (if any). Using those C_n 's, we construct another uniform NC^1 -circuit family $\{D_n\}_{n \in \mathbb{N}}$ such that $D_n(w) = 1$ iff $C_n(w)$ is an s - t path by checking that every consecutive elements in y are connected by a single edge and the first and the last elements are s and t , respectively. Clearly, $\{D_n\}_{n \in \mathbb{N}^+}$ recognizes USTCON. This implies that $\text{USTCON} \in \text{NC}^1$, as requested.

(ii) Assuming $\text{NC}^1 = L$, let us consider any problem $P = (I, \text{SOL}, m, \text{goal})$ in LO_{NLO} . To simplify the following argument, we assume that all feasible solutions of P are expressed *in binary*. Consider the case of $\text{goal} = \text{MIN}$. Let p be a polynomial that upper-bounds the size of any feasible solution of P . Let M be a log-space deterministic Turing machine computing minimal solutions of P . To build a solver for P , we first define an oracle A as a collection of $\langle x, 1^i, b \rangle$ for which there exists a binary string y with $|y| \leq p(|x|)$ such that $M(x)$ outputs y and the i th bit of y is $b \in \{0, 1\}$. Since A belongs to L, our assumption implies that A falls into NC^1 . Note that, if $M(x) = y$ with $y = y_1 y_2 y_3 \dots y_e$, then y coincides with $A(\langle x, 1, y_1 \rangle) A(\langle x, 1^2, y_2 \rangle) A(\langle x, 1^3, y_3 \rangle) \dots A(\langle x, 1^e, y_e \rangle)$, where A is treated as its characteristic function.

To recover the i th bit y_i of $M(x)$ for all $i \in [e]$, we compute both $A(\langle x, 1^i, 0 \rangle)$ and $A(\langle x, 1^i, 1 \rangle)$ and decide that y_i takes a value $b \in \{0, 1\}$ exactly when $A(\langle x, 1^i, b \rangle) = 1$ and $A(\langle x, 1^i, \bar{b} \rangle) = 0$. This procedure can be done in parallel with accesses to the oracle A using an NC^1 circuit. Since $A \in \text{NC}^1$, the procedure

can be implemented using uniform NC^1 -circuits without any oracle Since $\text{NC}^1 = \text{L}$ implies $\text{FNC}^1 = \text{FL}$, the measure function m is in auxFNC^1 . From this, we conclude that $m(x, M(x))$ is computed from x by a suitable NC^1 -circuit. Thus, P belongs to $\text{NC}^1_{\text{ONLO}}$.

(3) Notice that, in this proof, we consider NPO problems instead of NLO problems. Similarly to (2), we wish to prove that (i) $\text{LO}_{\text{NPO}} \cap \text{PBO} = \text{PO}_{\text{NPO}} \cap \text{PBO}$ implies $\text{L} = \text{P}$ and (ii) $\text{L} = \text{P}$ implies $\text{LO}_{\text{NPO}} = \text{PO}_{\text{NPO}}$.

(i) Assume that $\text{LO}_{\text{NPO}} \cap \text{PBO} = \text{PO}_{\text{NPO}} \cap \text{PBO}$. Consider CVP, which is $\leq_m^{\text{NC}^1}$ -complete for P. Founded on CVP, we intend to build a maximization problem P of the form $(I, \text{SOL}, m, \text{MAX})$ lying in $\text{PO}_{\text{NPO}} \cap \text{PBO}$. Define I as the set of all Boolean circuits and, for any circuit $C \in I$, define $\text{SOL}(x)$ to be a set of all variable assignments σ setting all input variables of C to either 0 or 1 for which σ forces C to output 1. Since $\text{CVP} \in \text{P}$, P belongs to $\text{PO}_{\text{NPO}} \cap \text{PBO}$. By our assumption, P falls into $\text{LO}_{\text{NPO}} \cap \text{PBO}$. This yields a log-space deterministic Turing machine solving P . By running M and checking M 's outputs, we can solve CVP using only log space. Hence, CVP is in L, implying that $\text{L} = \text{P}$.

(ii) Assuming $\text{L} = \text{P}$, let us consider MIN WEIGHT-ST-CUT , defined in Section 3.1, which is $\leq_{\text{EX}}^{\text{NC}^1}$ -complete for PO_{NPO} by Proposition 3.2. Let M be any polynomial-time deterministic Turing machine computing minimal solutions of MIN WEIGHT-ST-CUT . Let p be a polynomial satisfying that, for all instances x to MIN WEIGHT-ST-CUT , $p(|x|)$ bounds M 's output size. We define a problem A to be composed of all strings of the form $\langle x, 1^i, b \rangle$ for which $M(x)$ outputs a certain string y and y 's i th bit is b . It is not difficult to show that A is in P. Our assumption of $\text{L} = \text{P}$ then implies that A indeed belongs to L.

The following log-space algorithm can compute optimal solutions of MIN WEIGHT-ST-CUT : on input x to MIN WEIGHT-ST-CUT , reconstruct $M(x)$ bit by bit by incrementing i by one from 1 to $p(|x|)$, check whether $\langle x, 1^i, 0 \rangle \in A$ or $\langle x, 1^i, 1 \rangle \in A$ or neither, and determine the i th bit of $M(x)$ to be $b \in \{0, 1\}$ if $\langle x, 1^i, b \rangle \in A$ and $\langle x, 1^i, \bar{b} \rangle \notin A$. Since $A \in \text{L}$, this procedure requires only log space. Hence, MIN WEIGHT-ST-CUT belongs to LO_{NPO} . \square

Here, we further refine Proposition 6.4.

Theorem 6.5 1. $\text{L} \neq \text{NL}$ iff $\text{LO}_{\text{NLO}} \cap \text{PBO} \neq \text{LSAS}_{\text{NLO}} \cap \text{PBO} \neq \text{APXL}_{\text{NLO}} \cap \text{PBO} \neq \text{NLO} \cap \text{PBO}$.
 2. $\text{NC}^1 \neq \text{L}$ iff the following classes are mutually distinct: $\text{NC}^1_{\text{ONLO}}$, $\text{NC}^1_{\text{AS}_{\text{NLO}}}$, $\text{APXNC}^1_{\text{NLO}}$, and LO_{NLO} .
 3. $\text{L} \neq \text{P}$ iff the following classes are mutually distinct: LO_{NPO} , LSAS_{NPO} , APXL_{NPO} , and PO_{NPO} .

Given any string $y \in \{0, 1\}^+$, we set $\text{del}(y)$ to be the string obtained from y by deleting all 0s from the leftmost bit of y until the first 1. For example, $\text{del}(0010) = 10$, $\text{del}(01001) = 1001$, and $\text{del}(00) = \lambda$. We then define $\text{rep}_+(y)$ to be *one plus* the natural number represented in binary as $\text{del}(y)$; that is, $\text{rep}_+(y) = 1 + \text{rep}(\text{del}(y))$.

Proof of Theorem 6.5. (1) Tantau [39, Theorem 4.1] showed that, under the assumption of $\text{L} \neq \text{NL}$, the following four classes are mutually distinct: $\text{LO}_{\text{NLO}} \cap \text{PBO}$, $\text{LSAS}_{\text{NLO}} \cap \text{PBO}$, $\text{APXL}_{\text{NLO}} \cap \text{PBO}$, and $\text{NLO} \cap \text{PBO}$. To see the converse, assume that those classes are different. If $\text{L} = \text{NL}$, then Proposition 6.4(1) implies $\text{LO}_{\text{NLO}} \cap \text{PBO} = \text{NLO} \cap \text{PBO}$. This obviously contradicts our assumption because $\text{LO}_{\text{NLO}} \subseteq \text{LSAS}_{\text{NLO}} \subseteq \text{APXL}_{\text{NLO}} \subseteq \text{NLO}$. Hence, the desired separation between L and NL follows immediately.

(2) We shall show separately two directions of ‘‘iff.’’

(If-part) To show that $\text{NC}^1 \neq \text{L}$, we assume that the following four classes are mutually different: $\text{NC}^1_{\text{ONLO}}$, $\text{NC}^1_{\text{AS}_{\text{NLO}}}$, $\text{APXNC}^1_{\text{NLO}}$, and LO_{NLO} . Toward a contradiction, we assume that $\text{NC}^1 = \text{L}$. Proposition 6.4(2) implies that $\text{NC}^1_{\text{ONLO}} = \text{LO}_{\text{NLO}}$. This is a clear contradiction; therefore, we obtain $\text{NC}^1 \neq \text{L}$, as requested.

(Only If-part) Assume that $\text{NC}^1 \neq \text{L}$. First, notice that, by Proposition 6.4(2), we obtain $\text{NC}^1_{\text{ONLO}} \cap \text{PBO} \neq \text{LO}_{\text{NLO}} \cap \text{PBO}$. Since $\text{NC}^1_{\text{ONLO}} \subseteq \text{NC}^1_{\text{AS}_{\text{NLO}}} \subseteq \text{APXNC}^1_{\text{NLO}}$, it suffices to verify four inequalities: (a) $\text{NC}^1_{\text{ONLO}} \cap \text{PBO} \neq \text{NC}^1_{\text{AS}_{\text{NLO}}} \cap \text{PBO}$, (b) $\text{NC}^1_{\text{AS}_{\text{NLO}}} \cap \text{PBO} \neq \text{APXNC}^1_{\text{NLO}} \cap \text{PBO}$, (c) $\text{LO}_{\text{NLO}} \cap \text{PBO} \not\subseteq \text{APXNC}^1_{\text{NLO}} \cap \text{PBO}$, where (c) implies $\text{LO}_{\text{NLO}} \neq \text{LO}_{\text{NLO}}$ and also $\text{NC}^1_{\text{AS}_{\text{NLO}}} \neq \text{LO}_{\text{NLO}}$.

(a) To lead to a contradiction, let us assume that $\text{NC}^1_{\text{ONLO}} \cap \text{PBO} = \text{NC}^1_{\text{AS}_{\text{NLO}}} \cap \text{PBO}$. We express MAX UAPP-VERTEX given in Section 5 as $(I, \text{SOL}, m, \text{MAX})$. Since MAX UAPP-VERTEX belongs to $\text{NC}^1_{\text{AS}_{\text{NLO}}} \cap \text{PBO}$ by Proposition 5.2(1), our assumption places MAX UAPP-VERTEX in $\text{NC}^1_{\text{ONLO}} \cap \text{PBO}$.

Take an NC^1 circuit C computing optimal solutions of MAX UAPP-VERTEX . Consider the set L composed of (G, s, w, ℓ) such that $(G, s, w) \in I$, $\ell \in \mathbb{N}^+$, and there exists a vertex $t \in V$ for which $w(t) \geq \ell$ and s and t are connected in G . This problem L falls into NC^1 , because it can be computed by the following NC^1 -circuit: on input $x = (G, s, w, \ell)$, we first compute $y = C(G, s, w)$, extract a vertex $t \in V$ from y satisfying $w(t) = m^*(x)$, and check whether $w(t) \geq \ell$. It therefore suffices to verify that $\text{USTCON} \leq_m^{\text{NC}^1} L$,

because this implies that $\text{USTCON} \in \text{NC}^1$, a contradiction. From any instance (G, s, t) to USTCON , we define (G, s, w, k) as follows. Define $w(t) = 1 + \log \log |V|$ and $w(v) = \log \log |V|$ for all $v \in V - \{t\}$. Let $k = w(t)$. Clearly, (G, s, w, k) can be computed from (G, s, t) using an appropriate NC^1 -circuit. It is not difficult to show that $(G, s, t) \in \text{USTCON}$ iff $(G, s, w, k) \in L$. Hence, $\text{USTCON} \leq_m^{\text{NC}^1} L$ follows.

(b) Since USTCON in L , it follows from our assumption that $\text{USTCON} \notin \text{NC}^1$. From USTCON , we construct the maximization problem $P = (I, \text{SOL}, m, \text{MAX})$ defined as in the proof (i) of Proposition 6.4(2). This problem P actually belongs to $\text{APXNC}_{\text{NLO}}^1 \cap \text{PBO}$, because there is an NC^1 algorithm computing 2-approximate solutions to P . For such a solution y on an instance x , we obtain $m(x, y)/m^*(x) \leq 2$, which means $m(x, y) \leq 2$. If $P \in \text{NC}^1\text{AS}_{\text{NLO}}$, then, by setting $r = 3/2$, we obtain r -approximate solutions z on input x using an NC^1 -circuit. If $x \in \text{USTCON}$, then, since $R(x, z) \leq 3/2$, $m(x, z) = 2$ follows; thus, z must be a s - t path. If $x \notin \text{USTCON}$, then $m(x, z) = 1$ and $z \neq t$. This implies that $\text{USTCON} \in \text{NC}^1$, a contradiction. Therefore, P is not in $\text{NC}^1\text{AS}_{\text{NLO}}$.

(c) Assuming that $\text{LO}_{\text{NLO}} \cap \text{PBO} \subseteq \text{APXNC}_{\text{NLO}}^1 \cap \text{PBO}$, we plan to derive that $\text{DSTCON} \in \text{NC}^1$. Consider MAX UB-VERTEX and set it as $(I_0, \text{SOL}_0, m_0, \text{MAX})$. Proposition 4.3 shows its $\leq_{\text{EX}}^{\text{NC}^1}$ -completeness for $\text{LO}_{\text{NLO}} \cap \text{PBO}$. From our assumption, we derive that this problem is also in $\text{APXNC}_{\text{NLO}}^1$. With a suitable constant $r > 1$, take an NC^1 circuit C producing r -approximate solutions to MAX UB-VERTEX . Let us construct another NC^1 circuit solving DSTCON . Given any instance (G, s, t) to DSTCON with $G = (V, E)$, we define a weight function w as $w(t) = r+1$ and $w(v) = 1$ for all $v \in V - \{t\}$. Here, we consider $x = (G, s, w)$. We run C to find an r -approximate solution, say, $u \in V$. If $(G, s, t) \in \text{DSTCON}$, then, since $\frac{m_0^*(x)}{m_0(x, u)} \leq r$, it follows that $m_0(x, u) \geq \frac{m_0^*(x)}{r} = 1 + \frac{1}{r} > 1$. Since $w(u) > 1$, u must be t . If $(G, s, t) \notin \text{DSTCON}$, then we can conclude that u cannot be t . Hence, DSTCON must be in NC^1 , as requested.

(3) Note that $\text{LO}_{\text{NPO}} \subseteq \text{PO}_{\text{NPO}}$ and $\text{LO}_{\text{NPO}} \subseteq \text{LSAS}_{\text{NPO}} \subseteq \text{APXL}_{\text{NPO}}$. In Proposition 6.4(3), we have already proven that $L \neq P$ iff $\text{LO}_{\text{NPO}} \neq \text{PO}_{\text{NPO}}$. It still remains to show that $L \neq P$ implies (a) $\text{PO}_{\text{NPO}} \not\subseteq \text{APXL}_{\text{NPO}}$, (b) $\text{LO}_{\text{NPO}} \neq \text{LSAS}_{\text{NPO}}$, and (c) $\text{LSAS}_{\text{NPO}} \neq \text{APXL}_{\text{NPO}}$. Note that $\text{PO}_{\text{NPO}} \not\subseteq \text{APXL}_{\text{NPO}}$ yields both $\text{APXL}_{\text{NPO}} \neq \text{PO}_{\text{NPO}}$ and $\text{LSAS}_{\text{NPO}} \neq \text{PO}_{\text{NPO}}$.

(a) For this claim, it suffices to prove the following statement: (*) $\text{PO}_{\text{NPO}} \subseteq \text{APXL}_{\text{NPO}}$ implies $L = P$. Let us assume that $\text{PO}_{\text{NPO}} \subseteq \text{APXL}_{\text{NPO}}$. Our goal is to show that CVP belongs to L since $L = P$ follows from the $\leq_m^{\text{NC}^1}$ -completeness of CVP for P . Consider MINCVP , which is closely related to the decision problem CVP , defined in Appendix. Notice that MINCVP is $\leq_{\text{EX}}^{\text{NC}^1}$ -complete for PO_{NPO} by Lemma 7.1. Let $\text{MINCVP} = (I, \text{SOL}, m, \text{MIN})$. Since $\text{MINCVP} \in \text{APXL}_{\text{NPO}}$ by our assumption, we can take a function $h \in \text{FL}$ and a constant $\gamma \geq 1$ such that $R(x, h(x)) \leq \gamma$ for all $x \in I$, where R indicates the performance ratio for MINCVP . Take an integer $k \geq 1$ for which $\gamma < 2^k$.

Let $z = \langle C_x \rangle$ be an instance of CVP . This implies that $\langle C_x \rangle \in \text{CVP}$ iff $1 \in \text{SOL}^*(\langle C_x, 1^n, 1 \rangle)$, where $n = |x|$. We define C'_x as follows. For the outcome z of C_x , C'_x outputs 0^k if $z = 0$; 10^{k-1} otherwise. The construction of C'_x can be done using log space. Given any outcome z' of C'_x , $\text{rep}(z')$ is either 0 or 2^k . Let $w = \langle C'_x, 1^n, 1^k \rangle$, an instance to MINCVP . It holds that, since $R(w, h(w)) \leq \gamma$, C_x outputs 0 (resp., 1) iff $h(w) = 0^k$ (resp., 10^{k-1}). Hence, we can determine whether $\langle C_x \rangle \in \text{CVP}$ simply by constructing $w = \langle C'_x, 1^n, 1^k \rangle$ and checking whether $h(w) = 10^{k-1}$. This implies that CVP belongs to L , as requested.

(b) In this case, we define $\text{MINASCVP} = (I, \text{SOL}, m, \text{MIN})$ as follows. Define $I = \{(C, x) \mid C : \text{circuit with } |x| \text{ inputs and 1 output, } x \in \{0, 1\}^*\}$ and $\text{SOL}(C, x) = \{y1^{|x|} \mid y \geq C(x)\}$ for any $(C, x) \in I$. Given any $(C, x) \in I$ and any $y \in \text{SOL}(C, x)$, define $m((C, x), y) = 2^{|x|+1} + C(x)$. Note that $\max_{y \in \text{SOL}(C, x)} \{m((C, x), y)\} \leq (1 + \frac{1}{2^{n+1}}) \min_{y \in \text{SOL}(C, x)} \{m((C, x), y)\}$. Consider the following algorithm M : on input $((C, x), r)$ with $r \in \mathbb{Q}^{>1}$, if $n+1 > \log(\frac{1}{r-1})$, then M outputs $1^{|x|+1}$; otherwise, M computes $C(x)$ by brute force and outputs $C(x)1^{|x|}$. Since $C(x)$ can be computed deterministically within $O(n^2)$ steps, M requires the total tape space of $O(\log n) + O(n^2)$, which equals $O(\log n) + O(\log^2(\frac{1}{r-1}))$. I also follows from $1 + 2^{-(n+1)} < r$ that M produces an r -approximate solution to x . Hence, MINASCVP belongs to LSAS_{NPO} .

Next, we claim that $\text{MINASCVP} \in \text{LO}_{\text{NPO}}$ implies $\text{CVP} \in L$. Assume that $\text{MINASCVP} \in \text{LO}_{\text{NPO}}$ and take a log-space deterministic Turing machine N that solves MINASCVP . The following log-space algorithm then solves CVP . On input $z = (C, x)$, run N on z and output y if N 's output is of the form $y1^{|x|}$. It is easy to verify that this algorithm is correct. Thus, CVP is in L . This contradicts our assumption that $L \neq P$ since CVP is P -complete. Therefore, MINASCVP does not belong to LO_{NPO} .

(c) Toward a contradiction, we assume that $\text{LSAS}_{\text{NPO}} = \text{APXL}_{\text{NPO}}$. Here, we define a new problem $\text{MIN1CVP} = (I, \text{SOL}, m, \text{MIN})$ by setting $I = \{(C, x) \mid C : \text{circuit with } |x| \text{ inputs and 1 output, } x \in \{0, 1\}^+\}$, $\text{SOL}(C, x) = \{y1^{|x|} \mid y \geq C(x)\}$, and $m((C, x), y1^{|x|}) = \text{rep}_{(+)}(y1^{|x|})$ for any $(C, x) \in I$ and any $y1^{|x|} \in \text{SOL}(C, x)$.

To claim that MIN1CVP belongs to APXL_{NPO} , consider the following log-space deterministic Turing machine M : on each input $(C, x) \in I$, M produces $1^{|x|+1}$ on its output tape. Since $\text{rep}_{(+)}(M(C, x)) = 2^{|x|+2}$ and $\text{rep}_{(+)}(y1^{|x|}) \in \{2^{|x|+2}, 2^{|x|+1}\}$, $M(C, x)$ is a 2-approximate solution to (C, x) . Thus, MIN1CVP is in APXL_{NPO} . By our assumption, MIN1CVP is also in LSAS_{NPO} . There is a log-space deterministic Turing machine N that produces 3/2-approximate solutions to MIN1CVP . From this N , we define N' as follows. On input $z = (C, x)$, N' first runs N , and outputs 1 if $N(C, x)$ equals $1^{|x|+1}$, and outputs 0 otherwise. Clearly, N' solves CVP . Hence, CVP belongs to L , a contradiction against $\text{L} \neq \text{P}$ because CVP is \leq_m^{L} -complete for P . \square

To close this section, we wish to demonstrate further separations with *no unproven assumption*. Notice that AC^0 is known to be *properly* included inside NC^1 because the parity function, which is in NC^1 , requires at least non-uniform constant-depth circuits of super-polynomial size [1, 15]. The following proof relies on this fact.

Theorem 6.6 $\text{AC}^0\text{ONLO} \neq \text{AC}^0\text{ASNLO} \neq \text{APXAC}_{\text{NLO}}^0 \neq \text{NC}^1\text{ONLO}$.

Concerning NLO and NPO , most optimization and approximation classes enjoy the following *upward separation property*: if $\mathcal{D}_{\text{NLO}}^{(1)} \neq \mathcal{D}_{\text{NLO}}^{(2)}$, then $\mathcal{D}_{\text{NPO}}^{(1)} \neq \mathcal{D}_{\text{NPO}}^{(2)}$. Theorem 6.6 thus yields the separations $\text{AC}^0\text{ONPO} \neq \text{AC}^0\text{ASNPO} \neq \text{APXAC}_{\text{NPO}}^0 \neq \text{NC}^1\text{ONPO}$.

To derive Theorem 6.6, we shall use the *parity function* π , which is defined as $\pi(x_1, x_2, \dots, x_n) = x_1 \oplus x_2 \oplus \dots \oplus x_n$, where each variable x_i takes a binary value. The fact that the parity function is out of AC^0 was first claimed by Furst, Saxe, and Sipser [15] as well as Ajtai [1]. Using this π , we further define $\pi^*(x_{11}, \dots, x_{1n}, x_{21}, \dots, x_{2n}, \dots, x_{n1}, \dots, x_{nn})$ to be the n -bit string $\pi(x_{11}, \dots, x_{1n})\pi(x_{21}, \dots, x_{2n}) \dots \pi(x_{n1}, \dots, x_{nn})$. It is not difficult to show that π^* is in FNC^1 but not in FAC^0 because π resides in the difference $\text{NC}^1 - \text{AC}^0$.

Proof of Theorem 6.6. The subsequent proof consists of three parts: (1) $\text{NC}^1\text{ONLO} \not\subseteq \text{APXAC}_{\text{NLO}}^0$, (2) $\text{AC}^0\text{ASNLO} \neq \text{APXAC}_{\text{NLO}}^0$, and (3) $\text{AC}^0\text{ONLO} \neq \text{AC}^0\text{ASNLO}$.

(1) For this separation, we consider a minimization problem $\text{MIN M-PARITY} = (I, \text{SOL}, m, \text{MIN})$ defined as follows. Let $I = \bigcup_{n \in \mathbb{N}^+} \{0, 1\}^{n^2}$ and let $\text{SOL}(x) = \{y \in \{0, 1\}^{|x|} \mid \text{rep}_+(y) \geq \text{rep}_+(\pi^*(x))\}$ for each $x \in I$ with $|x| = n^2$. Let $m(x, y) = \text{rep}_+(y)$ for all $y \in \text{SOL}(x)$. Clearly, I is in L , $I \circ \text{SOL}$ is in auxL , and m is in auxFL . This indicates that MIN M-PARITY is an NLO problem. Next, we shall argue that $\text{MIN M-PARITY} \in \text{NC}^1\text{ONLO}$. Note that, for every $x \in I$, $\text{SOL}^*(x) = \{\pi^*(x)\}$ holds. Since π^* is in FNC^1 , it follows that MIN M-PARITY is NC^1 -solvable.

Let us prove that $\text{MIN M-PARITY} \notin \text{APXAC}_{\text{NLO}}^0$. To lead to a contradiction, we assume otherwise; that is, MIN M-PARITY is in $\text{APXAC}_{\text{NLO}}^0$. There exist a constant $\gamma > 1$ and a uniform family $\{C_n\}_{n \in \mathbb{N}^+}$ of AC^0 -circuits such that, for every $x \in I$, $C_{|x|}(x)$ computes a string y in $\text{SOL}(x)$ satisfying that (*) $\text{rep}_+(y)/\gamma \leq \text{rep}_+(\pi^*(x)) \leq \text{rep}_+(y)$. Take any number n satisfying $2^n > \gamma$ and any string $x \in \{0, 1\}^{n^2}$. Consider x^n and define $y = C_{n^2}(x^n)$. If $\pi(x) = 1$, then we obtain $\text{rep}_+(\pi^*(x^n)) = 2^{n+1}$. Since $|y| = n$ and $y \in \text{SOL}(x)$, it must hold that $\text{rep}_+(y) = 2^{n+1}$, implying $y = 1^n$. By contrast, if $\pi(x) = 0$, then we obtain $\text{rep}_+(\pi^*(x^n)) = 1$. By Condition (*), it follows that $1 \leq \text{rep}_+(y) \leq \gamma$. Since $\gamma < 2^n$, y must have the form $0z$ for a certain string z in $\{0, 1\}^{|x|^{-1}}$. As a consequence, $\pi(x)$ equals the first bit of y . This gives an AC^0 -circuit that computes π . This is a contradiction against the fact that π is not in AC^0 . Therefore, MIN M-PARITY is not in $\text{APXAC}_{\text{NLO}}^0$.

(2) Based on the parity function π , we define a simple minimization problem, MIN 1-PARITY , which is somewhat similar to MIN M-PARITY defined in the proof of (1). MIN 1-PARITY takes an instance $x \in \{0, 1\}^n$ and finds a solution $y \in \{0, 1\}$ such that $\text{rep}(1y) \geq \text{rep}(1\pi(x))$ with the measure function $m(x, y) = \text{rep}(1y)$.

Here, we claim that MIN 1-PARITY is in $\text{APXAC}_{\text{NLO}}^0$. Consider the following AC^0 -circuit C : on input x , output $y_1 = 1$. Since $\text{rep}(1y_1) = 3$ and $\text{rep}(1\pi(x)) \in \{2, 3\}$, it follows that $\text{rep}(1y_1)/2 \leq \text{rep}(1\pi(x)) \leq \text{rep}(1y_1)$. Thus, C is a 2-approximate algorithm for MIN 1-PARITY . We then conclude that MIN 1-PARITY belongs to $\text{APXAC}_{\text{NLO}}^0$.

Assume that $\text{AC}^0\text{ASNLO} = \text{APXAC}_{\text{NLO}}^0$. This means that MIN 1-PARITY is in AC^0ASNLO . Take a uniform family $\{C_n\}_{n \in \mathbb{N}}$ of AC^0 -circuits such that, for any $x \in \{0, 1\}^*$, $C_{|x|}(x)$ outputs a 5/4-approximate solution y . For simplicity, write y for $C_{|x|}(x)$. Consider the case of $\pi(x) = 0$. If $y = 1$, then $\text{rep}(1y) = 3$. This contradicts the inequalities: $(4/5)\text{rep}(1y) \leq \text{rep}(1\pi(x)) \leq \text{rep}(1y)$. This implies that y is 0. In the case of $\pi(x) = 0$, when $y = 0$, it does not hold that $(4/5)\text{rep}(1y) \leq \text{rep}(1\pi(x)) \leq \text{rep}(1y)$. Hence, y must be 1. Therefore, $\{C_n\}_{n \in \mathbb{N}}$ computes π correctly. This implies that π is actually in AC^0 , a contradiction against $\pi \notin \text{AC}^0$. Therefore, $\text{AC}^0\text{ASNLO} \neq \text{APXAC}_{\text{NLO}}^0$ follows.

(3) We define another minimization problem MIN BIT-PARITY, which belongs to AC^0AS_{NLO} by setting $I = \bigcup_{n \in \mathbb{N}^+} \{0, 1\}^n$, $SOL(x) = \{y \in \{0, 1\}^{|x|} \mid rep(y) \geq rep(1^{|x|}\pi(x))\}$ for each $x \in I$, and $m(x, y) = \max\{1, rep(y)\}$ for $y \in SOL(x)$.

Consider a circuit C_n that takes input (x, r) with $n = |x|$ and, if $r \geq 1 + \frac{1}{2^{n+1}-2}$, then produces 1^{n0} , and if $1 < r < 1 + \frac{1}{2^{n+1}-2}$, then computes $\pi(x)$ exactly and outputs $1^n\pi(x)$. Note that the value $\pi(x)$ can be exactly computed by an appropriately chosen circuit, say, D of size $n^{O(1)}$ and depth $O(\log n)$. When $1 < r < 1 + \frac{1}{2^{n+1}-2}$, since this is equivalent to $n < \log(\frac{2r-1}{r-1}) - 1$ and r is treated as a constant, D has depth $O(\log \log(\frac{2r-1}{r-1}))$, which is bounded by a certain constant. Thus, C_n is an AC^0 -circuit parameterized by r . Let $y = C_n(x, r)$. In the case of $r \geq 1 + \frac{1}{2^{n+1}-2}$, we obtain $rep(y) = 2^{n+1} - 1$. When $\pi(x) = 0$, since $rep(1^n\pi(x)) = 2^{n+1} - 2$, it follows that $rep(y)/r = 2^{n+1} - 2 \leq rep(1^n\pi(x)) \leq rep(y)$. In contrast, when $\pi(x) = 1$, clearly we obtain $rep(y)/r < rep(1^n\pi(x)) \leq rep(y)$. This implies that C_n computes r -approximate solutions.

Toward a contradiction, we assume that $AC^0ONLO = AC^0AS_{NLO}$. This implies that MIN BIT-PARITY is in AC^0ONLO . Let $MIN\ BIT-PARITY = (I, SOL, m, MIN)$. Next, we argue that $\pi \in AC^0$, leading to a clear contradiction. Since $MIN\ BIT-PARITY \in AC^0AS$, there is a uniform family $\{C'_n\}_{n \in \mathbb{N}}$ of AC^0 -circuits solving MIN BIT-PARITY. Since $SOL(x) = \{1^{|x|}\pi(x)\}$, it follows that $C'_n(x) = 1^{|x|}\pi(x)$. From C'_n , we can design another AC^0 -circuit that computes $\pi(x)$. Thus, we obtain the desired conclusion of $\pi \in AC^0$. This yields the separation between AC^0ONLO and AC^0AS_{NLO} . \square

7 Discussions and Future Research Directions

We have refined the existing framework of combinatorial optimization problems in hope of providing a useful means of classifying various optimization problems lying inside the P-solvable class PO_{NPO} ; in particular, we have been focused on NL optimization problems or NLO problems, following early work of Álvarez and Jenner [4, 5] and Tantau [39]. In a course of our exploring study on such optimization problems, unfortunately, we have left numerous fundamental issues unsettled. For an advance of our knowledge, they definitely require reasonable answers and explanations. As a quick guide to those pending issues, we want to shed clear light on some of the challenging issues arising naturally in the course of our study.

1. [Class Separations and Inclusions] We have introduced numerous classes of refined optimization problems in Section 2.2 but a few classes are successfully proven to be different (cf. Theorem 6.6). One of our ultimate goals is to give proofs for separations of other important optimization and approximation classes. As seen in Section 6, those separations are closely tied up to long-standing open questions regarding the classes of their associated decision problems, and thus the separations of refined classes immediately lead to definitive answers to such open questions. Beside the separation issues, there are important unsettled questions concerning the inclusion relationships among newly refined classes. For example, contrary to the well-known inclusion $NC^1 \subseteq L$, we suspect that $APXNC^1_{NLO} \not\subseteq LONLO$, $NC^1AS_{NLO} \not\subseteq LONLO$, and $APXNC^1 \not\subseteq LSAS_{NLO}$.
2. [Completeness Issues] (1) In this paper, a number of optimization problems have been demonstrated to be complete for certain optimization and approximation classes within NLO. It is imperative to find more natural and useful optimization problems and prove them to be complete for target classes, in particular, $LSAS_{NLO} \cap PBO$ and $NC^1AS_{NLO} \cap PBO$. (2) In Theorem 3.3 and Corollary 3.4, we have shown that both MaxNL and MinNL contain $\leq_{sAP}^{NC^1}$ -complete problems. Unlike NPO, we do not know whether NLO ($= \text{MaxNL} \cup \text{MinNL}$) contains $\leq_{sAP}^{NC^1}$ -complete problems unless auxFL is closed under division (cf. Proposition 3.5. Can we remove this closure property? Similarly, are there any complete problem in $APXL_{NLO}$ without any assumption? Moreover, several optimization problems discussed in this paper have not proven to be complete. Those pending problems include MAX λ -DFA, MAX B-VERTEX, and MIN UPATH-WEIGHT. (3) There is a large collection of works presenting L-complete and NL-complete decision problems [3, 12, 27, 26]. In most cases, it is relatively straightforward to turn those problems into their associated NLO problems. Which of them are actually complete for refined optimization and approximation classes under suitably chosen reductions?
3. [Polynomially-Bounded Problems] The relationship between NL and NLO appears to be quite different from the relationship between NP and NPO. Such difference comes primarily from an architectural limitation of log-space (auxiliary) Turing machines. In particular, polynomially-bounded optimization problems are quite special for log-space computation. We have proven that $NC^1ONLO = LONLO$ iff

$\text{NC}^1\text{O}_{\text{NLO}} \cap \text{PBO} = \text{LO}_{\text{NLO}} \cap \text{PBO}$ (Proposition 6.4(2)) but we do not know whether (i) $\text{LO}_{\text{NLO}} = \text{NLO}$ iff $\text{LO}_{\text{NLO}} \cap \text{PBO} = \text{NLO} \cap \text{PBO}$ and (ii) $\text{NC}^1\text{O}_{\text{NLO}} = \text{NLO}$ iff $\text{NC}^1\text{O}_{\text{NLO}} \cap \text{PBO} = \text{NLO} \cap \text{PBO}$? Prove or disprove these equivalences (i)–(ii).

4. [Further Refinement of Approximability] Beyond APXL_{NLO}, Tantau [39] discussed classes of NLO problems whose $n^{O(1)}$ -approximate (as well as $2^{n^{O(1)}}$ -approximate) solutions are computed by log-space deterministic Turing machines. We briefly call them poly-APXL_{NLO} and exp-APXL_{NLO}. Those two classes fill a seemingly wide gap between APXL_{NLO} and NLO. For example, we have already mentioned in Section 3.2 and 4.2 that MIN UPATH-WEIGHT and MAX B-VERTEX belong to poly-APXL_{NLO}. Find more natural complete problems falling into these special approximation classes. These classes can be further expanded into a more general class denoted by $\ell(n)$ -APXL_{NLO}, where $\ell(n)$ refers to $\ell(n)$ -approximate solutions. Explore its features by finding natural problems inside it.
5. [Heuristic and Probabilistic Approaches] This paper has aimed at cultivating a basic theory of NLO problems with a major focus on complete problems. As a consequence, this paper has neglected any practical heuristic approach toward the NLO problems. The next possible step is to find such heuristic approaches to real-life NLO problems. Study also the power of probabilistic algorithms to solve NLO problems.
6. [Help by Advice] In 1980s, Karp and Lipton [29] studied a notion of advice, which is an external source used to enhance any underlying Turing machines, and they introduced two major complexity classes, P/log and P/poly, where “log” and “poly” refer to advice sizes of $O(\log n)$ and $n^{O(1)}$, respectively. In analogy with these advised classes, it may be possible to formulate NLO/log and NLO/poly containing advised optimization problems. Explore the properties of those advised classes and show the separations among them.
7. [Fixed Parameter Complexity] Recall that optimization problems in LSAS_{NLO} and $\text{NC}^1\text{AS}_{\text{NLO}}$ admit certain approximation schemes, which are algorithms whose resources (i.e., tape space or circuit size) described in terms of *fixed parameters*. NLO problems whose log-space approximation schemes have the performance ratios of, e.g., $k + \log n$ and $2^k \log n$, are treated equally inside LSAS_{NLO}, but those schemes behave differently in practice. Therefore, it is more practical to specify those fixed parameters and refine those classes.
8. [More Classes Inside PO_{NPO}] (1: Classes inside NLO) If the number of solutions of a given problem is relatively small, can we use this fewness information to find an optimal solution much more efficiently? Similarly to FewP (few polynomial time), we may define its log-space optimization counterpart, denoted by FewLO, which is a class of NLO problems whose sets of feasible solutions have sizes bounded from above by certain polynomials in $|x|$. Study the computational complexity of optimization problems inside FewLO.
(2: Classes between PO_{NPO} and NLO) A focal point of the current paper is optimization problems inside NLO. It still remains to fill the gap between PO_{NPO} and NLO. Uniform circuit families, such as NC^k , AC^k , SAC^k , and TC^k for $k \geq 1$, may be used to form optimization/approximation classes beyond NLO. For example, SAC^1 -circuits may introduce optimization and approximation classes, say, $\text{SAC}^1\text{O}_{\text{NPO}}$ and $\text{APXSAC}^1_{\text{NPO}}$. Find complete problems for those new classes lying above NLO. Concerning to the power of SAC^1 , it is also important to determine whether NLO coincides with $\text{SAC}^1\text{O}_{\text{NLO}}$, expanding Lemma 6.2(1).
9. [Other Categories of Optimization Problems] Lately, Yamakami [41] studied OptCFL, which is a CFL analogue of Krentel’s OptP [32] and Álvarez and Jenner’s OptL [4, 5], where CFL stands for *context-free languages*. This fact suggests a possibility of introducing CFL-based optimization problems under the current framework of refined optimization problems. Cultivate a theory based on those weak optimization problems.

Appendix: Proof of Proposition 3.2

Proposition 3.2 states the completeness of MIN WEIGHT-ST-CUT for PO_{NPO} . Here, we restate the proposition.

Proposition 3.2 (again). MIN ST-CUT (on weighted directed graphs) is $\leq_{\text{EX}}^{\text{NC}^1}$ -complete for PO_{NPO} .

Note that an unweighted version of MIN WEIGHT-ST-CUT is also known as the *edge connectivity problem*. Gabow [16] gave a sequential algorithm of finding the edge connectivity c of a (directed or undirected) graph of n vertices and m edges in time $O(cn \log(n^2/m))$.

For readability, we have left Proposition 3.2 unproven in Section 3.1. In this appendix, we shall give the proof of Proposition 3.2 for completeness. As noted in Section 3.1, this result is drawn from the P-completeness proof of Goldschlager et al. [17] for a decision version of the maximum s - t flow problem.

To prove the proposition, we first look into the decision problem CVP (Circuit Value Problem) into a minimization problem in PO_{NLO} . Given each binary string x , the notation $\text{rep}(x)$ denotes one plus the natural number represented in binary by x .

MINIMUM CIRCUIT VALUE PROBLEM (MINCVP):

- INSTANCE: $\langle C_x, 1^n, 1^k \rangle$ with a circuit C_x of n inputs and k outputs with AND, OR, and NOT gates for an input $x \in \{0, 1\}^n$ (which must be specified in the description of the circuit).
- SOLUTION: a binary string y of length k such that $\text{rep}(y) \geq \text{rep}(z)$, where z is the outcome of the circuit C_x .
- MEASURE: the number $\text{rep}(y)$.

We further demand a circuit to satisfy the following conditions: (i) it is monotone (i.e., using only AND and OR gates), (ii) each input has fan-out at most 1, (iii) each gate has fan-out at most 2, and (iv) the top (i.e., root) gate is an OR gate. With those conditions, we obtain the MINIMUM MONOTONE CIRCUIT VALUE 2 PROBLEM (MINMCV2) defined similarly to MINCVP. The decision version of MINMCV2 was shown to be \leq_m^L -complete for P [17].

Here, we establish the $\leq_{\text{EX}}^{\text{NC}^1}$ -completeness of MINCVP and MINMCV2.

Lemma 7.1 *MINCVP and MINMCV2 are $\leq_{\text{EX}}^{\text{NC}^1}$ -complete for PO_{NPO} .*

Proof. We first show that MINCVP is in PO_{NPO} . Given an instance $\langle C_x, 1^n, 1^k \rangle$ of MINCVP, it is possible to calculate the outcome of C_x by evaluating each gate of C_x one by one within polynomial time. Hence, MINCVP is in PO_{NPO} .

Next, we shall show the $\leq_{\text{EX}}^{\text{NC}^1}$ -hardness of MinCVP for PO_{NPO} . Let $P = (I, \text{SOL}, m, \text{goal})$ be any optimization problem in PO_{NPO} . Let x be any instance of P . Since P is P-solvable, we take a function $h \in \text{FP}$ satisfying $h(x) \in \text{SOL}^*(x)$ for all $x \in (I \circ \text{SOL})^\exists$. For each x , we pad extra k'_x zeros so that $k'_x \geq 0$ and $|h(x)10^{k'_x}| = k(n)$ for a certain absolute polynomial k . For convenience, we set $h'(x) = h(x)10^{k'_x}$. It is known that any function in FP can be computed by a certain NC^1 -uniform family of polynomial-size Boolean circuits (see, e.g., [14]).

Let us define an $\leq_{\text{EX}}^{\text{NC}^1}$ -reduction (f, g) from P to MinCVP as follows. Since $h' \in \text{FP}$, take an NC^1 -uniform family $\{C_n\}_{n \in \mathbb{N}^+}$ of polynomial-size Boolean circuits C_n of n inputs and $k(n)$ outputs that computes h' . Define $f(x) = \langle C_x, 1^{|x|}, 1^{k(|x|)} \rangle$, where C_x is obtained from $C_{|x|}$ by replacing $|x|$ input variables with $|x|$ constant bits $x \in \{0, 1\}^{|x|}$. Since $\{C_n\}_{n \in \mathbb{N}^+}$ is NC^1 -uniform, there is an NC^1 -circuit family to construct each C_x . Hence, f must be in FNC^1 . If y is of the form $z10^{k'}$, then we write \tilde{y} to denote z . Define $g(x, y) = \tilde{y}$. It is not difficult to show that if $y = C_x$ then $g(x, y) = \tilde{y} = h(x)$. Hence, P is $\leq_{\text{EX}}^{\text{NC}^1}$ -reducible to MinCVP. The transformation f is computed by a certain family of NC^1 -circuits.

The case for MINMCV2 is similar, however, by using restricted monotone circuits for h' . \square

Using Lemma 7.1, we want to derive Proposition 3.2.

Proof of Proposition 3.2. It is known that MIN ST-CUT is in PO_{NPO} . Thus, it suffices to reduce MINMCV2 to MIN ST-CUT by an appropriate $\leq_{\text{EX}}^{\text{NC}^1}$ -reduction (f, g) . For convenience, we set $\text{MINMCV2} = (I_1, \text{SOL}_1, m_1, \text{MIN})$. Let $\langle C_x, 1^n, 1^k \rangle$ be any instance of MinMCV2, where C_x is expressed as a sequence $(\alpha_n, \alpha_{n-1}, \dots, \alpha_0)$.

First, we consider the case where $k = 1$. The function f is of the form $f(\langle C_x, 1^n, 1^k \rangle) = \langle G, s, t, c \rangle$, where $G = (V, E)$ is defined as follows. Let $V = \{i \mid 0 \leq i \leq n\} \cup \{s, t\}$. Let $c(s, i) = b \cdot 2^i$ if α_i is an input $b \in \{0, 1\}$. If α_i is of the form $\text{AND}(j, k)$ or $\text{OR}(j, k)$, then let $c(j, i) = 2^j$ and $c(k, i) = 2^k$. Moreover, when α_i is $\text{AND}(j, k)$, let $c(i, t) = 2^j + 2^k - d \cdot 2^i$; when α_i is $\text{OR}(j, k)$, let $c(i, s) = 2^j + 2^k - d \cdot 2^i$, where d is the fan-out of α_i . Now, let $E = \{(i, j) \mid c(i, j) \text{ is defined}\}$. Note that $c(i, j) \leq 2^{n+1} \leq 2^{|V|}$ for any $(i, j) \in E$. For any s - t cut (S'_0, S'_1) , we set $g(\langle G, s, t, c \rangle, (S'_0, S'_1)) = 1$ if the cut capacity of (S'_0, S'_1) is odd; 0 otherwise. Note that g can be computed by certain NC^1 -circuits. Let $x = \langle C_x, 1^n, 1 \rangle$. If (S'_0, S'_1) is an s - t cut, then $m_1(x, g(x, (S'_0, S'_1))) = m_1^*(x) = 1$.

Here, we define a special s - t cut (S_0, S_1) inductively. The sink t and all vertices corresponding to inputs with 0 given to C_x are in S_1 . For a vertex i corresponding to a gate AND or OR, if this gate outputs 0, then the vertex i is in S_1 . The set S_0 is defined as $S_0 = V - S_1$. It is possible to verify that C_n outputs 1 iff the cut capacity of (S_0, S_1) is odd. Moreover, we can prove that (S_0, S_1) is minimal.

Let $k \geq 2$. Given C_x with n input bits and k output bits, we first make k copies $C_{x,1}, C_{x,2}, \dots, C_{x,k}$ of C_x and, for each $C_{x,i}$ ($i \in [k]$), we add the following new nodes: an input β having 0, two gates $AND(\beta, \gamma_1, \gamma_2, \dots, \gamma_{i-1}, \gamma_{i+1}, \dots, \gamma_k)$ and $OR(\gamma_i, \zeta)$, where γ_j corresponds to the j th output bit of C_x and ζ is the index corresponding to this AND gate and by renaming all vertices v of $C_{x,i}$ as $v^{(i)}$. Let $\tilde{C}_{x,i}$ denote the circuit obtained from $C_{x,i}$ by adding these extra gates. Clearly, the output of C_x coincides with the k -bit string $r_1 r_2 \dots r_n$, where r_i is the output of $C_{x,i}$. For each $\tilde{C}_{x,i}$, we define a graph $G_i = (V_i, E_i)$ and a weight function c_i as in the case of $k = 1$.

Finally, we add new source \tilde{s} and new sink \tilde{t} , and then define \tilde{c} as follows. We set $\tilde{c}(\tilde{s}, s_i) = 2^{n^2}$ and $\tilde{c}(t_i, \tilde{t}) = 2^{n^2}$ for every $i \in [k]$. Moreover, we define $\tilde{c}(\ell_1, \ell_2) = 2^{in} c_i(\ell'_1, \ell'_2)$ if ℓ'_1, ℓ'_2 are in the same $\tilde{C}_{x,i}$; 0 otherwise. A new graph $\tilde{G} = (\tilde{V}, \tilde{E})$ is defined similarly as before: let $\tilde{V} = \bigcup_{i \in [k]} V_i \cup \{\tilde{s}, \tilde{t}\}$ (provided that $V_i \cap V_j = \emptyset$ for any two distinct $i, j \in [k]$) and $\tilde{E} = \{(i, j) \mid \tilde{c}(i, j) \text{ is defined}\}$. Let $x = \langle C_x, 1^n, 1^k \rangle$. Note that $m_1(x, g(x, (\tilde{S}_0, \tilde{S}_1))) = m_1^*(x) = y_1 y_2 \dots y_k$. Hence, (f, g) reduces MINMCV2 to MIN WEIGHT-ST-CUT. Given a minimal $s^{(i)}-t^{(i)}$ cut $(S_{0,i}, S_{1,i})$ of G_i for each $i \in [k]$, define $\tilde{S}_0 = \{\tilde{s}\} \cup (\bigcup_{i \in [k]} S_{0,i})$ and $\tilde{S}_1 = \{\tilde{t}\} \cup (\bigcup_{i \in [k]} S_{1,i})$. The cut $(\tilde{S}_0, \tilde{S}_1)$ is a minimal $\tilde{s}-\tilde{t}$ cut of \tilde{G} . We define g as follows: the i th bit of $g(\langle \tilde{G}, \tilde{s}, \tilde{t}, \tilde{c} \rangle, (\tilde{S}_0, \tilde{S}_1))$ is 1 iff the $s^{(i)}-t^{(i)}$ cut $(S_{0,i}, S_{1,i})$ has odd capacity. This equivalence gives rise to the desired equivalence: C_x outputs $y_1 y_2 \dots y_k$ iff $g(x, (\tilde{S}_0, \tilde{S}_1)) = y_1 y_2 \dots y_k$. \square

References

- [1] M. Ajtai. Σ_1^1 -formula on finite structures. *Ann. Pure Appl. Logic* 24, 1–48, 1983.
- [2] E. Allender, D. Bruschi, and G. Pighizzini. The complexity of computing maximal word functions. *Comput. Complexity* 3 (1993) 368–391.
- [3] C. Àlvarez and R. Greenlaw. A compendium of problems complete for symmetric logarithmic space. *Comp. Complex.* 9 (2000) 123–145.
- [4] C. Àlvarez and B. Jenner. A very hard log-space counting class. *Theoret. Comput. Sci.* 107 (1993) 3–30.
- [5] C. Àlvarez and B. Jenner. A note on logspace optimization. *Comput. Complexity* 5 (1995) 155–166.
- [6] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi. *Complexity and Approximation: Combinatorial Optimization Problems and Their Approximability Properties*, Springer-Verlag, 2003.
- [7] D. A. M. Barrington, N. Immerman, and H. Straubing. On uniformity within NC^1 . *J. Comput. System. Sci.* 41 (1990) 274–306.
- [8] M. Beaudry and P. McKenzie. Circuits, matrices, and nonassociative computation. *J. Comput. System Sci.* 50 (1995) 441–455.
- [9] S. Buss. The Boolean formula value problem is in ALOGTIME. In the *Proc. of the 19th Symposium on Theory of Computing (STOC'87)*, pp.123–131, ACM Press, 1987.
- [10] A. K. Chandra, L. Stockmeyer, and U. Vishkin. Constant depth reducibility. *SIAM J. Comput.* 13 (1984) 423–439.
- [11] P. Crescenzi, V. Kann, R. Silverstri, and L. Trevisan. Structure in approximation classes. *SIAM J. Comput.*, 28 (1999) 1750–1782.
- [12] S. A. Cook and P. McKenzie. Problems complete for deterministic logarithmic space. *J. Algorithms* 8 (1987) 385–394.
- [13] C. Damm, M. Holzer, and P. Rossmanith. Expressing uniformity via oracle. *Theory Comput. Syst.* 30 (1997) 355–366.
- [14] D. Du and K. Ko. *Theory of Computational Complexity*. John Wiley & Sons, Inc., 2000.
- [15] M. Furst, J. Saxe, and M. Sipser. Parity, circuits and the polynomial-time hierarchy. *Math. Systems Theory* 17, 13–27, 1984.
- [16] H. N. Gabow. A matroid approach to finding edge connectivity and packing arborescences. In the *Proc. of the 23rd Annual Symposium on Theory of Computing (STOC'91)*, ACM Press, pp.112–122, 1991.
- [17] L. M. Goldschlager, R. A. Shaw, and J. Staples. The maximum flow problem is log space complete for P. *Theoret. Comput. Sci.* 21 (1982) 105–111.
- [18] J. Håstad. *Computational Limitations for Small-Depth Circuits*. Ph.D. dissertation, MIT Press, Cambridge, MA, 1986.

- [19] W. Hesse. Division is in uniform TC^0 . In the Proc. of the 28th International Colloquium on Automata, Languages and Programming (ICALP 2001), Lecture Notes in Computer Science, Springer, vol.2076, pp.104–114, 2001.
- [20] W. Hesse, E. Allender, D. A. M. Barrington: Uniform constant-depth threshold circuits for division and iterated multiplication. *J. Comput. Syst. Sci.* 65 (2002) 695–716. Corrigendum, *J. Comput. Syst. Sci.* 80 (2014) 496–497.
- [21] J. E. Hopcroft, R. Motowani, and J. D. Ullman. Introduction to Automata Theory, Languages, and Computation. Second Edition, Addison Wesley, 2001.
- [22] N. Immerman. Nondeterministic space is closed under complement. *SIAM J. Comput.* 17 (1988) 935–938.
- [23] B. Jenner. Knapsack problems for NL. *Inf. Process. Lett.* 54 (1995) 169–174.
- [24] B. Jenner, K. Lange, and P. McKenzie. Tree isomorphism and some other complete problems for deterministic logspace. #1059, DIRO, University of Montréal, 1997.
- [25] B. Jenner, P. McKenzie, and J. Torán. A note on the hardness of tree isomorphism. In the Proc. of the 13th Annual IEEE Conference on Computational Complexity (CCC’98), IEEE Computer Society, pp.101–105 (1998).
- [26] N. D. Jones. Space-bounded reducibility among combinatorial problems. *J. Comput. System Sci.* 11 (1975) 68–75.
- [27] N. D. Jones, Y. E. Lien, and W. T. Laaser. New problems complete for nondeterministic log space. *Math. Systems Theory* 10 (1976) 1–17.
- [28] D. R. Karger. Global min-cuts in RNC, and other ramifications of a simple min-cut algorithm. In the Proc. of the 4th ACM/SIGACT-SIAM Symposium on Discrete Algorithms (SODA’93), pp.21–30, 1993.
- [29] R. Karp and R. Lipton. Turing machines that take advice. *L’Enseignement Mathématique* 28, 191–210, 1982.
- [30] S. Khanna, R. Motwani, M. Sudan, and U. V. Vazirani. On syntactic versus computational views of approximability. *SIAM J. Comput.*, 28, 164–191, 1998.
- [31] J. Köbler, U. Schöning, and J. Torán. On counting and approximation. *Acta Informatica* 26 (1989) 363–379.
- [32] M. Krentel. The complexity of optimization problems. *J. Comput. System Sci.* 36 (1988) 490–509.
- [33] R. E. Ladner. The circuit value problem is log-space complete for P. *SIGACT News* 7 (1975) 18–20.
- [34] A. Nickelsen and T. Tantau. The complexity of finding paths in graphs with bounded independence number. *SIAM J. Comput.* (2005) 1176–1195.
- [35] C. H. Papadimitriou and M. Yannakakis. Optimization, approximation and complexity classes. *J. Comput. System Sci.*, 43, 425–440, 1991.
- [36] O. Reingold. Undirected connectivity in log-space. *J. ACM* 55 (2008) article 17.
- [37] W. L. Ruzzo. On uniform circuit complexity. *J. Comput. System Sci.* 21 (1981) 365–383.
- [38] R. Szelepcsényi. The method of forced enumeration for nondeterministic automata. *Acta Inform.* 26 (1988) 279–284.
- [39] T. Tantau. Logspace optimisation problems and their approximation properties. *Theory Comput. Syst.* 41 (2007) 327–350.
- [40] T. Yamakami. Optimization, randomized approximability, and constraint satisfaction problems. In the Proc. of the 22nd International Symposium on Algorithms and Computation (ISAAC 2011), Lecture Notes in Computer Science, Springer, vol.7074, pp.454–463, 2011. See also arXiv:1109.3651.
- [41] T. Yamakami. Structural complexity of multi-valued partial functions computed by nondeterministic push-down automata (extended abstract). In the Proc. of the 15th Italian Conference of Theoretical Computer Science (ICTCS 2014), CEUR Workshop Proceedings, vol.1231, pp.225–236, 2014. See also arXiv:1508.05814.
- [42] A. C. Yao. Separating the polynomial-time hierarchy by oracles. In the Proc. of the 26th IEEE Symposium on Foundations of Computer Science (FOCS’85), pp.1–10, 1985.