Logical Definability of NP Optimization Problems*

Phokion G. Kolaitis   Madhukar N. Thakur†

UCSC-CRL-93-10‡

Computer and Information Sciences
University of California Santa Cruz
Santa Cruz, CA 95064

Abstract: We investigate here NP optimization problems from a logical definability standpoint. We show that the class of optimization problems whose optimum is definable using first-order formulae coincides with the class of polynomially bounded NP optimization problems on finite structures. After this, we analyze the relative expressive power of various classes of optimization problems that arise in this framework. Some of our results show that logical definability has different implications for NP maximization problems than it has for NP minimization problems, in terms of both expressive power and approximation properties.

*To appear in *Information and Computation.* Research partially supported by NSF Grants CCR-8905038 and CCR-9108631.

†e-mail addresses: kolaitis@cse.ucsc.edu, thakur@cse.ucsc.edu

‡supersedes Technical report UCSC-CRL-90-48

1 Introduction and Summary of Results

It is well known that optimization problems had a major influence on the development of the theory of NP-completeness. As a matter of fact, many natural NP-complete problems are decision problems that are derived from an optimization problem by imposing a bound on the objective function ([GJ79]). In spite of this close connection, NP-completeness advanced along a strikingly different path than that of optimization theory. Non-deterministic Turing machines with polynomial-time bounds provide a fairly robust computational model for decision problems. This, in turn, made it possible to develop a rich structural complexity theory based on polynomial time reductions and to obtain various classifications of NP problems. There have been also several attempts to classify optimization problems and to study their structural properties. Some notable contributions include [OM90, Kre88, Wag86, PM81, ADP80, Joh74] (cf. also [BJY89] for a comprehensive survey of results in this area). Nevertheless, the absence of robust computational models for optimization problems has hindered the development of a structural optimization theory that is on a par with structural complexity theory. In particular, the approximation properties of optimization problems remain as one of the most persistent puzzles of optimization theory. Although all known natural NP-complete problems are polynomially isomorphic [BH77], their optimization counterparts may have dramatically different approximation properties, from possessing polynomial-time approximation schemes to being non-approximable within a constant factor (assuming P$\neq$NP).

Papadimitriou and Yannakakis [PY91] brought a fresh perspective to approximation theory by focusing on the logical definability of optimization problems. Their main motivation came from Fagin's [Fag74] characterization of NP in terms of definability in second-order logic on finite structures. An *existential second-order formula* is an expression of the form $(\exists \mathbf{S})\phi(\mathbf{S})$, where $\mathbf{S}$ is a sequence of predicates and $\phi(\mathbf{S})$ is a first-order formula. Fagin's theorem [Fag74] asserts that if C is a class of finite structures that is closed under isomorphisms, then C is NP-computable if and only if it is definable by an existential second-order formula. Moreover, it is well known that every such formula is equivalent to a formula in *Skolem normal form* (cf. [End72]), i.e., to one of the form $(\exists \mathbf{S})(\forall \mathbf{x})(\exists \mathbf{y})\psi(\mathbf{x},\mathbf{y},\mathbf{S})$, where ψ is a quantifier-free formula and $\mathbf{x},\mathbf{y}$ are finite sequences of variables. Thus, a class C of finite structures is NP-computable if and only if there is a formula $(\exists \mathbf{S})(\forall \mathbf{x})(\exists \mathbf{y})\psi(\mathbf{x},\mathbf{y},\mathbf{S})$, with ψ quantifier-free, such that for every finite structure $\mathbf{A}$ we have that

$$\mathbf{A} \in C \Longleftrightarrow \mathbf{A} \models (\exists \mathbf{S})(\forall \mathbf{x})(\exists \mathbf{y})\psi(\mathbf{x},\mathbf{y},\mathbf{S}).$$

Papadimitriou and Yannakakis [PY91] introduced the class MAX NP of maximization problems whose optimum can be defined as

$$\max_{\mathbf{S}} |\{\mathbf{x} : (\mathbf{A},\mathbf{S}) \models (\exists \mathbf{y})\psi(\mathbf{x},\mathbf{y},\mathbf{S})\}|,$$

where ψ is quantifier-free. Intuitively, in an NP decision problem one seeks predicates $\mathbf{S}$ witnessing some existential second-order sentence $(\exists \mathbf{S})(\forall \mathbf{x})(\exists \mathbf{y})\psi(\mathbf{x},\mathbf{y},\mathbf{S})$, while in the corresponding maximization problem in MAX NP one seeks predicates $\mathbf{S}$ that maximize the number of tuples $\mathbf{x}$ satisfying the existential first-order sentence $(\exists\mathbf{y})\psi(\mathbf{x},\mathbf{y},\mathbf{S})$. The canonical example of a problem in MAX NP is provided by MAX SAT, which asks for the maximum number of clauses that can be satisfied in a given Boolean formula.

Papadimitriou and Yannakakis [PY91] showed that for every optimization problem in MAX NP is *constant-approximable*, i.e., there is a polynomial time algorithm that approximates the optimum value of the problem within a constant factor. They also considered the subclass MAX SNP of MAX NP consisting of those maximization problems that are defined by quantifier-free formulae, i.e., the optimum of such problems can be defined as

$$\max_{\mathbf{S}} |\{\mathbf{x} : (\mathbf{A},\mathbf{S}) \models \psi(\mathbf{x},\mathbf{S})\}|,$$

where ψ is quantifier-free. They demonstrated that MAX SNP contains several natural maximization problems, such as MAX 3SAT, that are complete for MAX SNP via a certain reduction that preserves approximability. These results on the one hand revealed that the logical definability of an optimization problem may impact on its approximation properties and on the other provided supporting evidence for the conjecture that certain constant-approximable problems, such as MAX SAT and MAX 3SAT, do not have a *polynomial time approximation scheme*.[1] Recently, Arora et. al. [ALM$^+$92] confirmed this conjecture by establishing that, unless P = NP, no MAX SNP-complete problem has a polynomial time approximation scheme.

The expressive power of the class MAX NP was investigated by Panconesi and Ranjan [PR90], where it was established that MAX CLIQUE does not belong to this class (the proof of this result is actually due to D. Kozen). Moreover, Panconesi and Ranjan [PR90] proved that certain polynomial-time optimization problems are not in MAX NP. In an attempt to find a syntactic class of optimization problems containing MAX CLIQUE, they introduced the class MAX Π_1 of maximization problems whose optimum can be defined as

$$\max_{\mathbf{S}} |\{\mathbf{w} : (\mathbf{A},\mathbf{S}) \models (\forall\mathbf{x})\psi(\mathbf{w},\mathbf{x},\mathbf{S})\}|,$$

where ψ is a quantifier-free formula and $\mathbf{w},\mathbf{x}$ are sequences of first-order variables. With regard to approximation properties, Panconesi and Ranjan [PR90] showed that MAX Π_1 contains optimization problems that are not constant-approximable, unless P=NP. In addition, Panconesi and Ranjan [PR90] introduced and studied the class RMAX, a syntactic subclass of MAX Π_1 for which MAX CLIQUE is complete via an approximation preserving reduction. More recently, Arora and Safra [AS92] showed that, unless P=NP, MAX CLIQUE is not a constant-approximable problem.

What other classes of optimization problems can be obtained using the logical definability perspective and what is the exact expressive power of this framework? We

[1] An optimization problem is said to have a polynomial time approximation scheme, if for every constant $\epsilon \geq 0$ there is a polynomial time algorithm that approximates the optimum within a factor of $(1-\epsilon)$ for a maximization problem and $(1+\epsilon)$ for a minimization problem.

address these questions here by examining the class of all maximization problems whose optimum is definable using first-order formulae, i.e., it is given as

$$\max_{\mathbf{S}} |\{\mathbf{w} : (\mathbf{A}, \mathbf{S}) \models \phi(\mathbf{w}, \mathbf{S})\}|,$$

where $\phi(\mathbf{w}, \mathbf{S})$ is an arbitrary first-order formula with free variables from the sequence $\mathbf{w}$ and $\mathbf{S}$ is a sequence of predicate variables. We show first that this class coincides with the collection of *polynomially bounded* NP-maximization problems on finite structures, namely, the NP-maximization problems on finite structures whose optimum value is less than or equal to a polynomial of the input size. We classify next these problems according to the quantifier complexity of the first-order formulae used and we show that they form a proper hierarchy with exactly four levels:

$$\text{MAX } \Sigma_0 \subset \text{MAX } \Sigma_1 \subset \text{MAX } \Pi_1 \subset \text{MAX } \Pi_2,$$

where MAX Σ_0 = MAX SNP is obtained using quantifier-free formulae, MAX Σ_1 = MAX NP is obtained using existential formulae, MAX Π_1 is obtained using universal formulae, and finally MAX Π_2 is obtained using universal-existential formulae. In particular, MAX Π_2 can capture every polynomially bounded NP-maximization problem on finite structures. The above containments are strict and there are natural maximization problems witnessing the separation of the four classes. More specifically, we prove that MAX CONNECTED COMPONENT is in MAX Π_2, but not in MAX Π_1, while MAX SAT separates MAX Σ_1 from MAX Σ_0. As mentioned above, MAX CLIQUE is in MAX Π_1, but not in MAX Σ_1 (cf. [PR90]).

We focus next on the logical definability of NP-minimization problems. Panconesi and Ranjan [PR90] concentrated on maximization problems only, while Papadimitriou and Yannakakis [PY91] examined approximation properties of certain minimization problems by reducing them to maximization problems. At first sight, one may expect that hierarchy and proper containment results about classes of maximization problems should translate directly to analogous results about classes of minimization problems definable by similar formulae. It turns out, however, that this is *not* the case. Moreover, maximization and minimization problems defined by similar first-order formulae may have strikingly different approximation properties.

We show that the collection of polynomially bounded NP-minimization problems on finite structures coincides with the class of minimization problems whose optimum is defined using an existential-universal (Σ_2) first-order formula. After this, we establish that the polynomially bounded NP-minimization problems form a proper hierarchy with exactly two levels:

$$\text{MIN } \Sigma_0 = \text{MIN } \Sigma_1 \subset \text{MIN } \Pi_1 = \text{MIN } \Sigma_2.$$

We also show that MIN CHROMATIC NUMBER witnesses the separation between the two levels, namely, this problem is in MIN Π_1, but not in MIN Σ_1,

Recall that Papadimitriou and Yannakakis [PY91] showed that every maximization problem in MAX Σ_1 = MAX NP (and, hence, every problem in MAX Σ_0 = MAX SNP) is constant-approximable. In contrast, we prove here that MIN Σ_0 contains natural minimization problems, such as MIN 3NON-TAUTOLOGY, that are not constant-approximable, unless P=NP. Since the quantifier pattern of minimization problems does not have an impact on the approximation properties of the problems, we seek other syntactic properties that may have such an impact. To this effect, we introduce a natural subclass of MIN Σ_1 that is a syntactic dual of the class RMAX in [PR90]. This subclass of MIN Σ_1 contains MIN VERTEX COVER and has the property that every minimization problem in it is constant-approximable.

2 Preliminaries

This section contains the basic definitions and a minimum amount of the necessary background material.

Definition 2.1: An NP *optimization problem* is a tuple $\mathcal{Q} = (\mathcal{I}_\mathcal{Q}, \mathcal{F}_\mathcal{Q}, f_\mathcal{Q}, opt_\mathcal{Q})$ such that

- $\mathcal{I}_\mathcal{Q}$ is the set of input instances. It is assumed that $\mathcal{I}_\mathcal{Q}$ can be recognized in polynomial time.
- $\mathcal{F}_\mathcal{Q}(I)$ is the set of feasible solutions for the input I.
- $f_\mathcal{Q}$ is a polynomial time computable function, called the *objective function*. It takes positive integer values and is defined on pairs (I, T), where I is an input instance and T is a feasible solution of I.
- $opt_\mathcal{Q}$ is one of the two functions defined below with domain the set $\mathcal{I}_\mathcal{Q}$ of input instances and positive integers as values:

  $$opt_\mathcal{Q}(I) = \max_T f_\mathcal{Q}(I,T) \quad \text{or} \quad opt_\mathcal{Q}(I) = \min_T f_\mathcal{Q}(I,T).$$

  In the former case, we say $\mathcal{Q}$ is a *maximization* problem and in the latter case we say $\mathcal{Q}$ is a *minimization* problem.
- The following decision problem is in NP : Given $I \in \mathcal{I}_\mathcal{Q}$ and an integer k, does there exist a feasible solution $T \in \mathcal{F}_\mathcal{Q}(\mathcal{I})$ such that $f_\mathcal{Q}(I,T) \geq k$, for a maximization problem $\mathcal{Q}$ (or, $f_\mathcal{Q}(I,T) \leq k$, for a minimization problem $\mathcal{Q}$)?

The above definition is due to [PR90] and is broad enough to encompass every known optimization problem arising in NP-completeness. We now restrict attention to *polynomially bounded* NP optimization problems [BJY89, LM81]. These are NP optimization problems in which the optimum value of the objective function on an instance is bounded by a polynomial in the length of that instance.

Definition 2.2: An NP optimization problem $\mathcal{Q}$ is said to be *polynomially bounded* if there is a polynomial p such that

$$opt_{\mathcal{Q}}(I) \leq p(|I|) \text{ for all } I \in \mathcal{I}_{\mathcal{Q}},$$

where $|I|$ is the length of the input I. Let MAX $\mathcal{PB}$ (MIN $\mathcal{PB}$) denote the class of all polynomially bounded NP maximization (minimization) problems.

Examples of polynomially bounded NP optimization problems are MAX CLIQUE, TRAVELING SALESMAN problem with weights 1 or 2, MIN CHROMATIC NUMBER, and MIN VERTEX COVER. On the other hand, INTEGER PROGRAMMING and the unrestricted version of the TRAVELING SALESMAN problem are examples of NP optimization problems that are not polynomially bounded. Indeed, it is possible to have an instance of the TRAVELING SALESMAN problem of size n in which the shortest tour has length 2^n, because in this problem inter-city distances are encoded in binary notation.

Since in the sequel we will study optimization problems from the perspective of logical definability, we review briefly some basic concepts from mathematical logic and introduce the notation that we will use here. We refer the reader to Enderton [End72] or to any other standard textbook of mathematical logic for a more detailed exposition.

Definition 2.3: A *vocabulary* (also known as a *similarity type*) $\sigma = \{\tilde{R}_1, \cdots, \tilde{R}_k\}$ is a finite set of predicate symbols. Each predicate symbol $\tilde{R}_i$ has a positive integer r_i as its designated arity. A *structure* $\mathbf{A} = (A, R_1, \cdots R_k)$ over the vocabulary σ consists of a set A, called the universe of $\mathbf{A}$, and relations $R_1, \cdots, R_k$ of arities $r_1, \cdots r_k$ on A, i.e., subsets of the Cartesian products $A^{r_1}, \ldots, A^{r_k}$ respectively. A *finite* structure is a structure whose universe is a finite set. The *size* $|\mathbf{A}|$ of a finite structure $\mathbf{A}$ is the cardinality of its universe.

For example, a graph is a structure $G = (V, E)$ over a vocabulary with a single binary predicate $\tilde{E}$. The universe of this structure is the set V of the vertices of the graph, while E is the set of the graph edges. In most cases an NP decision problem can either be described directly as a problem on finite structures or it can be easily encoded by such a problem. For example, CLIQUE and VERTEX COVER are problems about finite graphs, while an instance I of SATISFIABILITY can be identified with a finite structure $\mathbf{A}(I) = (X, C, P, N)$ over a vocabulary with one unary and two binary predicate symbols such that the universe X is the set of variables and clauses of I, the unary relation $C(x)$ expresses that x is a clause, and the binary relations $P(c, v)$ and $N(c, v)$ express respectively that a variable v occurs positively or negatively in a clause c.

We assume that the reader is familiar with the definition of the syntax and semantics of first-order logic over a vocabulary σ. Intuitively, the *formulae* of first-order logic over σ are built from the predicate symbols of σ, a special binary symbol $=$, and variables

$v_1, v_2, \ldots$ using the logical connectives $\wedge, \vee, \neg, \rightarrow$ and the quantifiers $(\exists v_i)$ and $(\forall v_i)$, $i \geq 1$. Every formula ϕ of first-order logic can be given semantics on structures over the vocabulary σ. The predicate symbols of σ are interpreted by the corresponding relations of the structure, the special binary symbol $=$ is always interpreted as equality on the universe of $\mathbf{A}$, while the variables v_i in the quantifiers $(\exists v_i)$ and $(\forall v_i)$, $i \geq 1$, are interpreted as ranging over elements of the universe of the structure. The formula ϕ becomes *true* or *false* on a structure $\mathbf{A}$ whenever a tuple of elements from the universe of the structure is assigned as interpretation to the sequence of the *free variables* of the formula, i.e., to those variables v_i that do not always occur within the scope of a quantifier $(\exists v_i)$ or $(\forall v_i)$ in the formula (cf. [End72] for the precise definitions).

Let $\mathbf{w}$ be a finite sequences of variables. We shall write $\phi(\mathbf{w})$ to indicate that $\mathbf{w}$ is the sequence of the free variables of the formula ϕ. Finally, if $\mathbf{A}$ is a structure over the vocabulary σ, then

$$\{\mathbf{w} : \mathbf{A} \models \phi(\mathbf{w})\}$$

is the set of all tuples from the universe A of $\mathbf{A}$ for which the formula ϕ becomes *true* (equivalently, the set of all tuples from A that *satisfy* ϕ). For example, if $G = (V, E)$ is a graph, then

$$\{w : G \models (\exists y)(\exists z)(\tilde{E}(w, y) \wedge \tilde{E}(w, z) \wedge \neg(y = z))\}$$

is the set of all vertices of degree at least 2, i.e., the set of all vertices with at least two distinct neighbors.

Notice that in order to simplify matters in the above expressions we mixed syntax with semantics by using the same notation for both a sequence of variables and a tuple of elements from the universe of the structure interpreting these variables. By the same token, from now on we shall take the liberty to use the same notation for both predicate symbols and relations on a structure interpreting these symbols. We trust that the reader is able to tell the difference from the context.

Let $\mathbf{S} = (S_1, \ldots, S_m)$ be a sequence of predicate symbols of arities $s_1, \ldots, s_m$ not in the vocabulary σ. We write $\phi(\mathbf{w}, \mathbf{S})$ to denote a formula of first-order logic over the vocabulary $\sigma \cup \{S_1, \ldots, S_m\}$ having $\mathbf{w}$ as its free variables. If $\mathbf{A} = (A, R_1, \ldots, R_k)$ is a structure over the vocabulary σ and $S_1, \ldots, S_m$ are relations on A of arities $s_1, \ldots, s_m$ respectively, then we write $(\mathbf{A}, \mathbf{S})$ to denote the *expanded* structure $(A, R_1, \ldots, R_k, S_1, \ldots, S_m)$. Thus,

$$\{\mathbf{w} : (\mathbf{A}, \mathbf{S}) \models \phi(\mathbf{w}, \mathbf{S})\}$$

denotes the set of all tuples from A for which the formula $\phi(\mathbf{w}, \mathbf{S})$ becomes true on the expanded structure $(\mathbf{A}, \mathbf{S})$. For example, if $G = (V, E)$ is a graph and S is a subset of V, then

$$\{x : (G, S) \models (\forall y)(E(x, y) \rightarrow S(y))\}$$

denotes the set of all vertices with the property that all their neighbors are in S.

It is well known that every formula of first-order logic is equivalent to a formula in *prenex normal form*, i.e., to a formula in which all quantifiers are to the left of all other symbols (cf. [End72, pages 150-151]). We write $\Sigma_n, n \geq 1$, for the class of first-order formulae in prenex normal form that have $n-1$ alternations of quantifiers and start with a block of existential quantifiers. For example, Σ_1 is the collection of existential formulae, while Σ_2 is the class of existential-universal formulae. Similarly, we write Π_n, $n \geq 1$, for the class of first-order formulae in prenex normal form with $n-1$ alternations of quantifiers, starting with a block of universal quantifiers. Thus, a Π_1 formula has universal quantifiers only, while Π_2 is the collection of universal-existential formulae. The class of quantifier-free formulae is denoted by Σ_0 or by Π_0.

From now on we assume that the instances of an optimization problem are given as finite structures over some vocabulary σ. We introduce next a framework for classifying optimization problems on finite structures in terms of their definability in first-order logic.

Definition 2.4: Let σ be a vocabulary and let $\mathcal{Q}$ be a maximization problem with finite structures $\mathbf{A}$ over σ as instances. We say that $\mathcal{Q}$ is in the class MAX Π_n, $n \geq 0$, if there is a Π_n formula $\phi(\mathbf{w}, \mathbf{S})$ with predicate symbols among those in σ and $\mathbf{S}$ such that for every instance $\mathbf{A}$ of $\mathcal{Q}$ we have that

$$opt_{\mathcal{Q}}(\mathbf{A}) = \max_{\mathbf{S}} |\{\mathbf{w} : (\mathbf{A}, \mathbf{S}) \models \phi(\mathbf{w}, \mathbf{S})\}|.$$

Similarly, we say that $\mathcal{Q}$ is in the class MAX Σ_n, $n \geq 0$, if its optimum is definable as above using a Σ_n formula $\phi(\mathbf{w}, \mathbf{S})$.

The classes MIN Π_n and MIN Σ_n, $n \geq 0$, of minimization problems are defined in an analogous way, with min in place of max. In particular, a minimization problem $\mathcal{Q}$ is in the class MIN Π_n, $n \geq 0$, if there is a Π_n formula $\phi(\mathbf{w}, \mathbf{S})$ with predicate symbols among those in σ and $\mathbf{S}$ such that for every instance $\mathbf{A}$ of $\mathcal{Q}$ we have that

$$opt_{\mathcal{Q}}(\mathbf{A}) = \min_{\mathbf{S}} |\{\mathbf{w} : (\mathbf{A}, \mathbf{S}) \models \phi(\mathbf{w}, \mathbf{S})\}|.$$

The classes MAX Σ_0 and MAX Σ_1 were introduced and studied by Papadimitriou and Yannakakis [PY91] under the names MAX SNP and MAX NP respectively, while the class MAX Π_1 was introduced by Panconesi and Ranjan [PR90]. We have chosen to use different names for MAX SNP and MAX NP here, because we are interested in having a uniform notation and terminology for all the classes of optimization problems obtained using first-order formulae. Moreover, the notation Π_n and Σ_n is consistent with the notation Π_n^p and Σ_n^p used for the polynomial hierarchy [Sto76]. We now give examples of natural problems in some of these classes.

• MAX 3SAT is a problem in the class MAX Σ_0 (cf. [PY91]). This problem asks for the maximum number of clauses that can be satisfied in a given Boolean formula in conjunctive normal form (CNF) with three literals per clause. We view every instance

I of MAX 3SAT as a finite structure $\mathbf{A}(I)$ over a vocabulary consisting of four ternary predicate symbols C_0, C_1, C_2, C_3. Under this encoding, the universe of the structure $A(I)$ is the set of variables of the formula, while each relation $C_i(w_1, w_2, w_3)$ is true if and only if $\{w_1, w_2, w_3\}$ is a clause with $w_1, \cdots, w_i$ appearing as negative literals and $w_{i+1}, \cdots, w_3$ appearing as positive literals, $0 \leq i \leq 3$. The optimum of 3SAT is given by

$$opt_{\text{MAX 3SAT}}(\mathbf{A}(I)) = \max_S |\{(w_1, w_2, w_3) : (\mathbf{A}(I), S) \models \phi(w_1, w_2, w_3, S)\}|,$$

where $\phi(w_1, w_2, w_3, S)$ is the formula

$$C_0(w_1, w_2, w_3) \wedge (S(w_1) \vee S(w_2) \vee S(w_3)) \vee C_1(w_1, w_2, w_3) \wedge (\neg S(w_1) \vee S(w_2) \vee S(w_3)) \vee$$

$$C_2(w_1, w_2, w_3) \wedge (\neg S(w_1) \vee \neg S(w_2) \vee S(w_3)) \vee C_3(w_1, w_2, w_3) \wedge (\neg S(w_1) \vee \neg S(w_2) \vee \neg S(w_3)).$$

• MAX SAT is a problem in the class MAX Σ_1 (cf. [PY91]). Under the encoding of SATISFIABILITY given in Section 2, if $\mathbf{A}(I)$ is the finite structure associated with an instance I of MAX SAT, then we have

$$opt_{\text{MAX SAT}}(\mathbf{A}(I)) = \max_S |\{w : (\mathbf{A}(I), S) \models (\exists y)[C(w) \wedge$$
$$((P(w, y) \wedge S(y)) \vee (N(w, y) \wedge \neg S(y)))]\}|.$$

Intuitively, in the above formulae for MAX SAT and MAX 3SAT the predicate symbol S encodes a truth assignment, i.e., it consists of the Boolean variables that are set to TRUE.

• MAX CLIQUE is in the class MAX Π_1 (cf. [PR90]). Indeed, for MAX CLIQUE we have that

$$opt_{\text{MAX CLIQUE}}(G) = \max_S |\{w : (G, S) \models S(w) \wedge$$
$$(\forall y_1)(\forall y_2)[(S(y_1) \wedge S(y_2) \wedge (y_1 \neq y_2)) \rightarrow E(y_1, y_2)]\ \}|.$$

3 Polynomially Bounded NP Maximization Problems

In this section we investigate the relative expressive power of the classes MAX Π_n and MAX Σ_n, $n \geq 0$, and establish their basic relationship to the class MAX $\mathcal{PB}$ of polynomially bounded NP maximization problems.

Theorem 1: Let σ be a vocabulary and let $\mathcal{Q}$ be a maximization problem with finite structures $\mathbf{A}$ over σ as instances. Then $\mathcal{Q}$ is a polynomially bounded NP maximization problem if and only if there is a first-order formula $\phi(\mathbf{w}, \mathbf{S})$ with predicate symbols among those in σ and the sequence $\mathbf{S}$ such that for every instance $\mathbf{A}$ of $\mathcal{Q}$

$$opt_{\mathcal{Q}}(\mathbf{A}) = \max_{\mathbf{S}} |\{\mathbf{w} : (\mathbf{A}, \mathbf{S}) \models \phi(\mathbf{w}, \mathbf{S})\}|.$$

Moreover, $\phi(\mathbf{w}, \mathbf{S})$ can always be taken to be a Π_2 formula and, consequently,

$$\text{MAX } \mathcal{PB} = \text{MAX } \Pi_2 = \text{MAX } \Pi_n,\ n > 2.$$

Proof: It is clear that if a maximization problem $\mathcal{Q}$ is in the class MAX Π_n for some $n \geq 0$, then $\mathcal{Q}$ is a polynomially bounded NP maximization problem, since for any finite structure $\mathbf{A}$ there are polynomially many distinct tuples from $\mathbf{A}$ satisfying a given first-order formula.

For the other direction, assume that $\mathcal{Q}$ is a polynomially bounded NP maximization problem with finite structures $\mathbf{A}$ over the vocabulary σ as instances. Let m be a positive integer such that for any instance $\mathbf{A}$ we have that $\mathrm{opt}_{\mathcal{Q}}(\mathbf{A}) \leq |\mathbf{A}|^m$, where $|\mathbf{A}|$ is the size of the structure $\mathbf{A}$.

Consider now the following decision problem Q: Given a finite structure $\mathbf{A}$ over σ and a m-ary relation W on the universe A of $\mathbf{A}$, is there a feasible solution T for $\mathbf{A}$ such that $f_{\mathcal{Q}}(\mathbf{A}, T) \geq |W|$? Here, $f_{\mathcal{Q}}$ is the objective function of $\mathcal{Q}$ and $|W|$ is the cardinality of the m-ary relation W. Since $\mathcal{Q}$ is an NP optimization problem, we have that Q is a problem in NP. Moreover, Q can be viewed as an NP decision problem whose instances are finite structures over the vocabulary $\sigma \cup \{W\}$. By Fagin's [Fag74] characterization of NP in terms of definability in second-order logic, there is an existential second-order formula $(\exists \mathbf{S}^*)\psi(\mathbf{S}^*, W)$ such that the expanded structure $(\mathbf{A}, W)$ is a YES instance of Q if and only if $(\mathbf{A}, W) \models (\exists \mathbf{S}^*)\psi(\mathbf{S}^*, \mathbf{W})$. Since the maximization problem $\mathcal{Q}$ is bounded by $|\mathbf{A}|^m$, we have that

$$\mathrm{opt}_{\mathcal{Q}}(\mathbf{A}) = \max_{\mathbf{S}^*, W}\{|W| : (\mathbf{A}, \mathbf{S}^*, W) \models \psi(\mathbf{S}^*, W)\}$$

or, equivalently,

$$\mathrm{opt}_{\mathcal{Q}}(\mathbf{A}) = \max_{\mathbf{S}^*, W} |\{\mathbf{w} : (\mathbf{A}, \mathbf{S}^*, W) \models \psi(\mathbf{S}^*, W) \wedge W(\mathbf{w})\}|.$$

Let $\mathbf{S}$ denote the sequence $(\mathbf{S}^*, W)$ and let $\phi(\mathbf{w}, \mathbf{S})$ be the formula $\psi(\mathbf{S}^*, W) \wedge W(\mathbf{w})$. It follows that

$$\mathrm{opt}_{\mathcal{Q}}(\mathbf{A}) = \max_{\mathbf{S}} |\{\mathbf{w} : (\mathbf{A}, \mathbf{S}) \models \phi(\mathbf{w}, \mathbf{S})\}|.$$

Moreover, $\phi(\mathbf{w}, \mathbf{S})$ can be chosen to be a Π_2 formula, because Fagin's characterization of NP [Fag74] holds with a Π_2 formula $\psi(\mathbf{w}, \mathbf{S}^*)$. $\square$

Theorem 1 shows that MAX Π_2 is the entire class MAX $\mathcal{PB}$ of polynomially bounded NP maximization problems. In particular, it shows that MAX $\Sigma_2 \subseteq$ MAX Π_2. By restricting the quantifier prefix $\exists^*\forall^*$ of Σ_2 formulae, we obtain the class MAX Π_1 of [PR90], and the classes MAX Σ_1 = MAX NP and MAX Σ_0 = MAX SNP of [PY91]. It is clear that we have the following containments between these classes:

$$\begin{array}{ccccc} & \rotatebox{45}{$\subsetneq$} & \text{MAX } \Sigma_1 & \rotatebox{-45}{$\subsetneq$} & \\ \text{MAX } \Sigma_0 & & & & \text{MAX } \Sigma_2 \subseteq \quad \text{MAX } \Pi_2 = \text{MAX } \mathcal{PB}. \\ & \rotatebox{-45}{$\subsetneq$} & \text{MAX } \Pi_1 & \rotatebox{45}{$\subsetneq$} & \end{array}$$

We saw before examples of natural problems in the classes MAX Σ_0, MAX Σ_1, and MAX Π_1. We give next an example of a problem in the class MAX Π_2 that will be of particular interest to us in the sequel.

• MAX CONNECTED COMPONENT (MCC): Given an undirected graph G, find the size of the largest connected component in G.

Notice that actually MCC is an optimization problem on graphs that can be solved in polynomial time. Although Theorem 1 implies that MCC is in the class MAX Π_2, it is not obvious how to establish this directly. In what follows we produce a Π_2 formula ϕ that defines MCC in our framework.

In addition to a binary relation symbol E for the edges of the graph, the formula ϕ will involve the relation symbols $C, E, P, \leq, Z$. The intuition behind these is as follows: C is a unary relation symbol that represents the vertices of a connected component; $\leq$ is a binary relation that will vary over total orders on the vertices of the graph; P is a ternary relation symbol such that $P(x, y, k)$ indicates that the shortest path from x to y is of length k, where the integer k is encoded by the k^{th} element of the total order $\leq$; finally, Z is a unary predicate representing the smallest element (zero) of the total order $\leq$.

Let $\phi_1(\leq)$ be a formula asserting that $\leq$ is a total order and let $\phi_2(\mathcal{Z})$ be a formula asserting that Z is a singleton set containing the smallest element of $\leq$. Let also $pred(x, y)$ be a formula asserting that y is the predecessor of x under the above order. We leave it to the reader to verify that $\phi_1(\leq)$ and $pred(x, y)$ can be expressed as Π_1 formulae, while $\phi_2(Z)$ can be written as a conjunction of Π_1 and Σ_1 formulae. We are now ready to demonstrate that MCC is in the class MAX Π_2. Indeed, its optimum value on a graph G is given as

$$\begin{aligned} opt_{\mathrm{MCC}}(G) = \max_{(C,P,\leq,Z)} |\{w : (G, C, P, \leq, Z) \models\ & C(w) \wedge \phi_1(\leq) \wedge \phi_2(Z) \wedge \\ & (\forall x)(\forall y)((C(x) \wedge C(y)) \rightarrow (\exists z)P(x, y, z)) \wedge \\ & (\forall x)(\forall y)(\forall v)(\forall v')[(P(x, y, v) \wedge \neg Z(v) \wedge pred(v, v')) \rightarrow \\ & \qquad ((\exists z)P(x, z, v') \wedge E(z, y))] \wedge \\ & (\forall x)(\forall y)(\forall v)((P(x, y, v) \wedge Z(v)) \rightarrow (x = y))\ \}|. \end{aligned}$$

It is well known that the classes of Σ_1 and Π_1 formulae have incomparable expressive power on finite structures, while the class of Σ_2 formulae has strictly higher expressive power than the class of Π_1 formulae (cf. [CH82]). One might expect that similar results hold for the corresponding classes of maximization problems, but it turns out that this is not the case. The next result delineates the relationship between the classes of maximization problems and establishes that the polynomially bounded NP maximization problems form a hierarchy with exactly four distinct levels.

Theorem 2: The class MAX Σ_2 is contained in the class MAX Π_1. As a result,

$$\text{MAX } \Sigma_0 \subseteq \text{MAX } \Sigma_1 \subseteq \text{MAX } \Sigma_2 = \text{MAX } \Pi_1 \subseteq \text{MAX } \Pi_2.$$

Moreover, this sequence of containments is strict. In particular,

- MAX CONNECTED COMPONENT is in MAX Π_2, but not in MAX Π_1.
- MAX CLIQUE is in MAX Π_1, but not in MAX Σ_1 (cf. [PR90]).
- MAX SAT is in MAX Σ_1, but not in MAX Σ_0.

Proof: We give this proof in four parts.

Part A: In this part we prove that MAX Σ_2 is contained in the class MAX Π_1. Let $\mathcal{Q}$ be a MAX Σ_2 problem and let $\mathbf{A}$ be a finite structure that is an instance of $\mathcal{Q}$. Thus,

$$opt_{\mathcal{Q}}(\mathbf{A}) = \max_{\mathbf{S}} |\{\mathbf{w} : (\mathbf{A}, \mathbf{S}) \models (\exists \mathbf{x})(\forall \mathbf{y})\psi(\mathbf{w}, \mathbf{x}, \mathbf{y}, \mathbf{S})\}|,$$

where ψ is quantifier-free. Consider now the set

$$U(\mathbf{S}) = \{\mathbf{w} : (\mathbf{A}, \mathbf{S}) \models (\exists \mathbf{x})(\forall \mathbf{y})\psi(\mathbf{w}, \mathbf{x}, \mathbf{y}, \mathbf{S})\}$$

and notice that $opt_{\mathcal{Q}}(\mathbf{A}) = \max_{\mathbf{S}} |U(\mathbf{S})|$.

If $\mathbf{x}^*$ and $\mathbf{w}$ are tuples from the universe of $\mathbf{A}$ such that $(\mathbf{A}, \mathbf{S}) \models (\forall \mathbf{y})\psi(\mathbf{w}, \mathbf{x}^*, \mathbf{y}, \mathbf{S})$, then we say that $\mathbf{x}^*$ is a *witness of* $\mathbf{w}$ *relative to* $\mathbf{S}$. We now introduce an auxiliary predicate symbol R and define

$$\begin{aligned} V(\mathbf{S}, R) = \{(\mathbf{w}, \mathbf{x}^*) : (\mathbf{A}, \mathbf{S}, R) \models (\forall \mathbf{y})\psi(\mathbf{w}, \mathbf{x}^*, \mathbf{y}, \mathbf{S}) \wedge R(\mathbf{w}, \mathbf{x}^*) \wedge \\ (\forall \mathbf{x}_1)(\forall \mathbf{x}_2)((R(\mathbf{w}, \mathbf{x}_1) \wedge R(\mathbf{w}, \mathbf{x}_2)) \rightarrow \mathbf{x}_1 = \mathbf{x}_2)\} \end{aligned}$$

Intuitively, a pair $(\mathbf{w}, \mathbf{x}^*)$ is in the set $V(\mathbf{S}, R)$ if $\mathbf{x}^*$ is a witness of $\mathbf{w}$ relative to $\mathbf{S}$ and $\mathbf{x}^*$ is the only tuple $\mathbf{x}$ such that the pair $(\mathbf{w}, \mathbf{x})$ is in R. It is now easy to verify that for every fixed sequence $\mathbf{S}$ of relations we have that

$$|U(\mathbf{S})| = \max_{R} |V(\mathbf{S}, R)|$$

and, as a result,

$$opt_{\mathcal{Q}}(\mathbf{A}) = \max_{\mathbf{S}} |U(\mathbf{S})| = \max_{\mathbf{S}, R} |V(\mathbf{S}, R)|.$$

Since $V(\mathbf{S}, R)$ is defined using a Π_1 formula, it follows that $\mathcal{Q} \in$ MAX Π_1 and, consequently, the class MAX Σ_2 is contained in the class MAX Π_1.

Part B: We showed earlier that MCC is in the class MAX Π_2. In this part of the proof we show that MCC is not in the class MAX Π_1. Towards a contradiction, assume that the optimum of MCC is given by

$$opt_{\mathrm{MCC}}(G) = \max_{\mathbf{S}} |\{\mathbf{w} : (G, \mathbf{S}) \models (\forall \mathbf{y})\psi(\mathbf{w}, \mathbf{y}, \mathbf{S})\}|,$$

where ψ is quantifier-free and $\mathbf{w}$ ranges over tuples of arity m.

Let G be a graph that is a path with vertices $\{a_1, \cdots, a_n\}$, for some $n > 8m+1$, and edges $\{a_i, a_{i+1}\}, 1 \le i \le n-1$. Consider the subgraphs H_i, $1 \le i \le \lfloor n/2 \rfloor$, obtained from G by deleting a_i and all edges incident to it. Assume that the maximum value in the above expression occurs at $\mathbf{S} = \mathbf{S}^*$ and let $\mathbf{S}_i^*$ be the restriction of $\mathbf{S}^*$ to the vertex set $\{a_1, \cdots, a_{i-1}, a_{i+1}, \cdots, a_n\}$ of H_i. Since $opt_{\mathrm{MCC}}(H_i) = n - i$, we have that

$$|\{\mathbf{w} : (H_i, \mathbf{S_i^*}) \models (\forall \mathbf{y})\psi(\mathbf{w}, \mathbf{y}, \mathbf{S}_i^*)\}| \le n - i.$$

We now claim that each a_i, $1 \le i \le n$, occurs in at least i tuples in the set $\{\mathbf{w} : (G, \mathbf{S}^*) \models (\forall \mathbf{y})\psi(\mathbf{w}, \mathbf{y}, \mathbf{S}^*)\}$. Indeed, otherwise we would have that

$$|\{\mathbf{w} \in {H_i}^m : (G, \mathbf{S}^*) \models (\forall \mathbf{y})\psi(\mathbf{w}, \mathbf{y}, \mathbf{S}^*)\}| > n - i.$$

Since universal formulae are preserved under substructures, if $\mathbf{b}$ is an m-tuple from H_i such that $(G, \mathbf{S}^*) \models (\forall \mathbf{y})\psi(\mathbf{b}, \mathbf{y}, \mathbf{S}^*)$, then $(H_i, \mathbf{S_i^*}) \models (\forall \mathbf{y})\psi(\mathbf{b}, \mathbf{y}, \mathbf{S}_i^*)$. Thus, $|\{\mathbf{w} : (H_i, \mathbf{S_i^*}) \models (\forall \mathbf{y})\psi(\mathbf{w}, \mathbf{y}, \mathbf{S}_i^*)\}| > n - i$, which is a contradiction. Therefore, each a_i occurs in at least i tuples in the set $\{\mathbf{w} : (G, \mathbf{S}^*) \models (\forall \mathbf{y})\psi(\mathbf{w}, \mathbf{y}, \mathbf{S}^*)\}$. As a result, the total number of occurrences of all a_i's in this set is at least $(\sum_{i=1}^{i=\lfloor n/2 \rfloor} i) > nm$, since $n > 8m + 1$. On the other hand, since $\mathbf{w}$ ranges over tuples of arity m and $|\{\mathbf{w} : (G, \mathbf{S}^*) \models (\forall \mathbf{y})\psi(\mathbf{w}, \mathbf{y}, \mathbf{S}^*)\}| = n$, the total number of occurrences of all a_i's in this set is at most nm. Thus, we have arrived at a contradiction.

Part C: As mentioned in the Introduction, D. Kozen showed that MAX CLIQUE is in the class MAX Π_1, but not in the class MAX Σ_1 (cf. [PR90]).

Part D: We have seen before that MAX SAT is in the class MAX Σ_1. In this part of the proof we show that MAX SAT is not in the class MAX Σ_0. Let I be an instance of SAT and let $\mathbf{A}(I) = (X, C, P, N)$ be its encoding as a finite structure. Recall that X consists of the variables and the clauses of I, while the unary relation C consists of the clauses of I. Also recall that $(c, v) \in P$ (respectively, $(c, v) \in N$) if and only if the variable v occurs positively (respectively, negatively) in the clause c. Towards a contradiction, assume that MAX SAT is in the class MAX Σ_0. Therefore, there is a quantifier-free formula $\psi(\mathbf{w}, \mathbf{S})$ such that for every finite structure $\mathbf{A}(I)$ encoding an instance I of MAX SAT we have that

$$opt_{\text{MAX SAT}}(\mathbf{A}(I)) = \max_{\mathbf{S}} |\{\mathbf{w} : (\mathbf{A}(I), \mathbf{S}) \models \psi(\mathbf{w}, \mathbf{S})\}|,$$

where $\mathbf{w}$ ranges over m-tuples $(w_1, w_2, \cdots, w_m)$ and $\mathbf{S} = (S_1, \cdots, S_t)$. We distinguish two cases and show that in either case we arrive at a contradiction.

Case 1: Assume that for every structure $\mathbf{A}(I)$ encoding an instance I the maximum number of clauses satisfiable is given by

$$opt_{\text{MAX SAT}}(\mathbf{A}(I)) = \max_{\mathbf{S}} |\{(\underbrace{w, \cdots, w}_{m}) : (\mathbf{A}(I), \mathbf{S}) \models \psi(\underbrace{w, \cdots, w}_{m}, \mathbf{S})\}|.$$

Let $\psi'(w, \mathbf{S})$ be the formula obtained from ψ by replacing each occurrence of every variable by w. It is clear that

$$opt_{\text{MAX SAT}}(\mathbf{A}(I)) = \max_{\mathbf{S}} |\{w : (\mathbf{A}(I), \mathbf{S}) \models \psi'(w, \mathbf{S})\}|.$$

Since ψ is a quantifier-free formula, ψ' is also a quantifier-free formula whose only variable is w. As a result, in $\psi'(w, \mathbf{S})$ the only occurrences of the predicate symbols C, P, N and $S_1, \cdots, S_t$ in $\mathbf{S}$ are amongst the following:

$$\begin{gathered} C(w), \neg C(w), P(w,w), \neg P(w,w), N(w,w), \neg N(w,w), \\ S_l(\underbrace{w, \cdots, w}_{\alpha[l]}), \neg S_l(\underbrace{w, \cdots, w}_{\alpha[l]}), 1 \leq l \leq t, \end{gathered}$$

where $\alpha[l]$ is the arity of S_l. For every instance I encoded by a finite structure $\mathbf{A}(I) = (X, C, P, N)$, it is the case that $\mathbf{A}(I) \not\models P(x, x)$ and $\mathbf{A}(I) \not\models N(x, x)$, for all $x \in X$, because the first arguments of P, N refer to a clause, the second to a variable and the variables are different from the clauses. Let $\psi''(w, \mathbf{S})$ be the formula obtained from $\psi'(w, \mathbf{S})$ by replacing each occurrence of $P(w, w)$, $N(w, w)$ by the logical constant $FALSE$, and each occurrence of $\neg P(w, w)$, $\neg N(w, w)$ by the logical constant $TRUE$. Then we have that for every instance I

$$opt_{\text{MAX SAT}}(\mathbf{A}(I)) = \max_{\mathbf{S}} |\{w : (\mathbf{A}(I), \mathbf{S}) \models \psi''(w, \mathbf{S})\}.$$

Let I_1, I_2 be two instances of MAX SAT, each having the same number of variables and the same number of clauses, but differing in the maximum number of satisfiable clauses. Without loss of generality, we can find structures $\mathbf{A}(I_1) = (X_1, C_1, P_1, N_1)$ and $\mathbf{A}(I_2) = (X_2, C_2, P_2, N_2)$ encoding I_1, I_2 respectively, such that $X_1 = X_2$ and $C_1 = C_2$. Since $\psi''(w, \mathbf{S})$ does not have any occurrences of the symbols P and N, we have

$$\{w : (\mathbf{A}(I_1), \mathbf{S}) \models \psi''(w, \mathbf{S})\} = \{w : (\mathbf{A}(I_2), \mathbf{S}) \models \psi''(w, \mathbf{S})\}.$$

for all values of $\mathbf{S}$. Therefore,

$$opt_{\text{MAX SAT}}(\mathbf{A}(I_1)) = opt_{\text{MAX SAT}}(\mathbf{A}(I_2)),$$

which is a contradiction.

Case 2: Assume that there is some instance I_1 such that its encoding by the structure $\mathbf{A}(I_1) = (X_1, C_1, P_1, N_1)$ satisfies

$$opt_{\text{MAX SAT}}(\mathbf{A}(I_1)) \neq \max_{\mathbf{S}} |\{(\underbrace{w, \cdots, w}_{m}) : (\mathbf{A}(I_1), \mathbf{S}) \models \psi(\underbrace{w, \cdots, w}_{m}, \mathbf{S})\}|.$$

For simplicity, we write $\mathbf{A}_1$ for the structure $\mathbf{A}(I_1)$.

Let $\mathbf{S}^*$ be a sequence $(S_1^*, S_2^*, \cdots, S_t^*)$ of predicates that realizes $opt_{\text{MAX SAT}}(\mathbf{A}_1)$, i.e.,

$$opt_{\text{MAX SAT}}(\mathbf{A}_1) = |\{(w_1, \cdots, w_m) : (\mathbf{A}_1, \mathbf{S}^*) \models \psi(w_1, \cdots, w_m, \mathbf{S}^*)\}|.$$

Let $x_1^1, x_2^1, \cdots, x_n^1$ be the variables and the clauses of I_1, i.e., $X_1 = \{x_1^1, x_2^1, \ldots, x_n^1\}$. We now construct $n-1$ additional structures, $\mathbf{A}_2, \cdots, \mathbf{A}_n$, where $\mathbf{A}_i = (X_i, C_i, P_i, N_i)$ with $X_i = \{x_1^i, x_2^i, \cdots, x_n^i\}, 2 \leq i \leq n$, such that they are all isomorphic to $\mathbf{A}_1$ via the mapping x_u^i to x_u^1, for $1 \leq i, u \leq n$.

We define next a structure $\mathbf{A} = (X, C, P, N)$ as follows:

$$\begin{aligned} X &= \bigcup_i^n X_i, \quad C = \bigcup_i^n C_i, \\ P &= \{(x_u^i, x_v^j) : P_1(x_u^1, x_v^1), 1 \leq u, v, i, j \leq n\}, \\ N &= \{(x_u^i, x_v^j) : N_1(x_u^1, x_v^1), 1 \leq u, v, i, j \leq n\}. \end{aligned}$$

It can be seen that $\mathbf{A}$ encodes an instance of MAX SAT. Also, observe that $|C| = n|C_1| \leq n(n-1)$, as the universe of the structure $\mathbf{A}_1$ has at least one variable. Therefore, $opt_{\text{MAX SAT}}(\mathbf{A}) \leq n(n-1)$. We will arrive at a contradiction by showing that $opt_{\text{MAX SAT}}(\mathbf{A}) \geq n^2$.

For $1 \leq l \leq t$, let

$$\begin{aligned} \mathcal{S}_l^* = \{(x_{u_1}^{i_1}, x_{u_2}^{i_2}, \cdots, x_{u_{\alpha[l]}}^{i_{\alpha[l]}}) : \quad & S_l^*(x_{u_1}^1, x_{u_2}^1, \cdots, x_{u_{\alpha[l]}}^1), \quad \text{where} \\ & 1 \leq i_1, \cdots, i_{\alpha[l]} \leq n \text{ and } 1 \leq u_1, \cdots, u_{\alpha[l]} \leq n\}, \end{aligned}$$

and let $\mathcal{S}^*$ denote the sequence $(\mathcal{S}_1^*, \mathcal{S}_2^*, \cdots, \mathcal{S}_t^*)$. We will show that $|V| \geq n^2$, where

$$V = \{(w_1, \cdots, w_m) : (\mathbf{A}, \mathcal{S}^*) \models \psi(w_1, \cdots, w_m, \mathcal{S}^*)\}.$$

Let

$$V_1 = \{(w_1, \cdots, w_m) : (\mathbf{A}_1, \mathbf{S}^*) \models \psi(w_1, \cdots, w_m, \mathbf{S}^*)\}$$

From the hypothesis of Case 2, it follows that

$$V_1 \neq \{(w, \cdots, w) : (\mathbf{A}_1, \mathbf{S}^*) \models \psi(w, \cdots, w, \mathbf{S}^*)\}.$$

Indeed, otherwise we would have

$$\begin{aligned} opt_{\text{MAX SAT}}(\mathbf{A}_1) &= \max_{\mathbf{S}} |\{(w_1, \cdots, w_m) : (\mathbf{A}_1, \mathbf{S}) \models \psi(w_1, \cdots, w_m, \mathbf{S})\}| \\ &\geq \max_{\mathbf{S}} |\{(w, \cdots, w) : (\mathbf{A}_1, \mathbf{S}) \models \psi(w, \cdots, w, \mathbf{S})\}| \\ &\geq |\{(w, \cdots, w) : (\mathbf{A}_1, \mathbf{S}^*) \models \psi(w, \cdots, w, \mathbf{S}^*)\}| \\ &= |V_1| = opt_{\text{MAX SAT}}(\mathbf{A}_1). \end{aligned}$$

Thus, $opt_{\text{MAX SAT}}(\mathbf{A}_1) = \max_{\mathbf{S}} |\{(w, \cdots, w) : (\mathbf{A}_1, \mathbf{S}) \models \psi(w, \cdots, w, \mathbf{S})\}|$, which contradicts the hypothesis of Case 2.

We now know that there is a tuple $\mathbf{e}$ in V_1 with at least two distinct components x_p^1 and x_q^1. For every i, j with $1 \leq i, j \leq n$, let $\mathbf{e}_{i,j}$ be obtained from $\mathbf{e}$ by replacing every occurrence of x_p^1 by x_p^i and every occurrence of x_q^1 by x_q^j. Also, let $\mathbf{A}_{i,j}$ denote the substructure of $\mathbf{A}$ with universe

$$\{x_1^1, \cdots, x_{p-1}^1, x_p^i, x_{p+1}^1, \cdots, x_{q-1}^1, x_q^j, x_{q+1}^1, \cdots, x_n^1\}.$$

It is clear that $\mathbf{A}_{i,j}$ is isomorphic to $\mathbf{A}_1$. Moreover, the restriction of $\mathcal{S}^*$ to the above set is a sequence of predicates isomorphic to $\mathbf{S}^*$, where the isomorphism maps x_p^i to x_p^1, maps x_q^i to x_q^1, and is the identity on the rest of the elements. Let $\mathcal{S}_{i,j}^*$ denote the restriction of $\mathcal{S}^*$ to the universe of $\mathbf{A}_{i,j}$ and observe that $(\mathbf{A}_{i,j}, \mathcal{S}_{i,j}^*) \models \psi(\mathbf{e}_{i,j}, \mathcal{S}_{i,j}^*)$ for $1 \leq i, j \leq n$. Since Σ_0 sentences are preserved under extensions, it is also true that $(\mathbf{A}, \mathcal{S}^*) \models \psi(\mathbf{e}_{i,j}, \mathcal{S}^*)$ for $1 \leq i, j \leq n$. As there are n^2 distinct such elements $\mathbf{e}_{i,j}$, we have that $|V| \geq n^2$. It follows that $opt_{\mathrm{MAX\ SAT}}(\mathbf{A}) \geq n^2$, which is a contradiction. The proof that MAX SAT is not in the class MAX Σ_0 is now complete. $\Box$

4 Polynomially Bounded NP Minimization Problems

The logical definability of NP minimization problems has not been explored in the literature so far. We undertake this investigation here by studying the classes MIN Σ_n and MIN Π_n, $n \geq 0$, of minimization problems that are definable using first-order formulae. Our findings for the expressive power and the relations between these classes unveil a strikingly different picture from the one for the corresponding maximization classes.

We begin by presenting examples of natural minimization problems in the classes MIN Σ_0, MIN Σ_1, and MIN Σ_2.

• MIN 3NON-TAUTOLOGY (3NT): Given a Boolean formula in disjunctive normal form with three literals per disjunct (3DNF), find the minimum number of satisfiable disjuncts.

MIN 3NON-TAUTOLOGY is an optimization problem in the class MIN Σ_0 that arises from the NP-complete problem NON-TAUTOLOGY of 3DNF formulae [GJ79]: Given a Boolean formula in 3DNF, is there a truth assignment that makes this formula false?

We view every instance I of MIN 3NT as a finite structure $\mathbf{A}(I)$ with four ternary predicates D_0, D_1, D_2, D_3, where $D_i(w_1, w_2, w_3)$ is true if and only if the set $\{w_1, w_2, w_3\}$ is a disjunct with $w_1, \cdots, w_i$ appearing as negative literals and $w_{i+1}, \cdots, w_3$ appearing as positive literals, $0 \leq i \leq 3$. The optimum of 3NT is given by

$$opt_{\mathrm{3NT}}(I) = \min_S |\{(w_1, w_2, w_3) : (\mathbf{A}(I), S) \models \phi(w_1, w_2, w_3, S)\}|,$$

where $\phi(w_1, w_2, w_3, S)$ is the following quantifier-free formula asserting that (w_1, w_2, w_3) is a disjunct of the 3DNF formula encoded by $\mathbf{A}(I)$ and that S is a truth assignment that satisfies this disjunct.

$$(D_0(w_1, w_2, w_3) \wedge S(w_1) \wedge S(w_2) \wedge S(w_3)) \vee (D_1(w_1, w_2, w_3) \wedge \neg S(w_1) \wedge S(w_2) \wedge S(w_3)) \vee$$

$(D_2(w_1, w_2, w_3) \wedge \neg S(w_1) \wedge \neg S(w_2) \wedge S(w_3)) \vee (D_3(w_1, w_2, w_3) \wedge \neg S(w_1) \wedge \neg S(w_2) \wedge \neg S(w_3))$.

• MIN VERTEX COVER: Given a graph $G = (V, E)$, find the smallest cardinality of a *vertex cover*, i.e., a subset S of the vertices such that every edge of G is adjacent to at least one vertex in S.

It is easy to see that MIN VERTEX COVER is in the class MIN Σ_1. Indeed, on any graph G the optimum is given by

$$\begin{aligned}
opt_{\text{MIN VC}}(G) &= \min_S \{|S| : (G, S) \models (\forall y_1)(\forall y_2)\,[\,E(y_1, y_2) \rightarrow (S(y_1) \vee S(y_2))\,]\,\} \\
&= \min_S |\{w : (G, S) \models \;[\,(\forall y_1)(\forall y_2)[E(y_1, y_2) \rightarrow (S(y_1) \vee S(y_2))]\,] \;\rightarrow S(w)\}| \\
&= \min_S |\{w : (G, S) \models (\exists y_1)(\exists y_2)[(E(y_1, y_2) \wedge \neg S(y_1) \wedge \neg S(y_2)) \vee S(w)]\}|\,.
\end{aligned}$$

• MIN CHROMATIC NUMBER: Given a graph $G = (V, E)$, find the minimum number of colors that can be assigned to the vertices of G so that no two adjacent vertices are of the same color.

MIN CHROMATIC NUMBER is an optimization problem that plays am important role in both graph theory and complexity theory (cf. [GJ79]). We now show that it is in the class MIN Σ_2. Consider first the following Π_2 sentence $\psi(S)$ asserting that the binary predicate S is a coloring of a graph G:

$$\begin{aligned}
\psi(S) \;\equiv\; & (\forall x)(\exists c)S(x, c) \;\wedge\; (\forall x)(\forall c_1)(\forall c_2)[S(x, c_1) \wedge S(x, c_2) \rightarrow (c_1 = c_2)] \\
& \wedge (\forall x)(\forall y)(\forall c_1)(\forall c_2)[E(x, y) \wedge S(x, c_1) \wedge S(y, c_2) \rightarrow (c_1 \neq c_2)].
\end{aligned}$$

It now follows that for every graph $G = (V, E)$

$$opt_{\text{CHROMATIC NUMBER}}(G) = \min_S |\{c : (G, S) \models \psi(S) \rightarrow (\exists x)S(x, c)\}|.$$

It is clear that the expression $\psi(S) \rightarrow (\exists x)S(x, c)$ is equivalent to a Σ_2 formula and, consequently, MIN CHROMATIC NUMBER is in the class MIN Σ_2.

Our next result shows that the class MIN Σ_2 contains all polynomially bounded NP minimization problems. It should be compared with Theorem 1 in Section 3.

Theorem 3: Let σ be a vocabulary and let $\mathcal{Q}$ be an NP minimization problem with finite structures $\mathbf{A}$ over σ as instances. Then $\mathcal{Q}$ is a polynomially bounded NP minimization problem if and only if there is a first-order formula $\phi(\mathbf{w}, \mathbf{S})$ with predicate symbols among those in σ and $\mathbf{S}$ such that for every instance $\mathbf{A}$ of $\mathcal{Q}$

$$opt_{\mathcal{Q}}(\mathbf{A}) = \min_{\mathbf{S}} |\{\mathbf{w} : (\mathbf{A}, \mathbf{S}) \models \phi(\mathbf{w}, \mathbf{S})\}|.$$

Moreover, $\phi(\mathbf{w}, \mathbf{S})$ can always be taken to be a Σ_2 formula and, consequently,

$$\text{MIN } \mathcal{PB} = \text{MIN } \Sigma_2 = \text{MIN } \Sigma_n, \;\; n > 2.$$

Proof: Following the same arguments as in Theorem 1, we can show that if $\mathcal{Q}$ is a polynomially bounded NP minimization problem, then there is a Π_2 formula $\psi(\mathbf{S}^*, W)$ such that

$$opt_{\mathcal{Q}}(\mathbf{A}) = \min_{\mathbf{S}^*, W}\{|W| : (\mathbf{A}, \mathbf{S}^*, W) \models \psi(\mathbf{S}^*, W)\}$$

It follows that

$$\begin{aligned} opt_{\mathcal{Q}}(\mathbf{A}) &= \min_{\mathbf{S}^*, W} |\{\mathbf{w} : (\mathbf{A}, \mathbf{S}^*, W) \models \psi(\mathbf{S}^*, W) \rightarrow W(\mathbf{w})\}| \\ &= \min_{\mathbf{S}^*, W} |\{\mathbf{w} : (\mathbf{A}, \mathbf{S}^*, W) \models \neg\psi(\mathbf{S}^*, W) \vee W(\mathbf{w})\}| \,. \end{aligned}$$

Let $\mathbf{S}$ denote the sequence $(\mathbf{S}^*, W)$ and let $\phi(\mathbf{w}, \mathbf{S})$ be a Σ_2 formula that is logically equivalent to

$$\neg\psi(\mathbf{S}^*, W) \vee W(\mathbf{w}).$$

We can now conclude that

$$opt_{\mathcal{Q}}(\mathbf{A}) = \min_{\mathbf{S}} |\{\mathbf{w} : (\mathbf{A}, \mathbf{S}) \models \phi(\mathbf{w}, \mathbf{S})\}|.$$

□

Remark 1: Notice that, unlike the case of maximization problems, if

$$opt_{\mathcal{Q}}(\mathbf{A}) = \min_{\mathbf{S}^*, W}\{|W| : (\mathbf{A}, \mathbf{S}^*, W) \models \psi(\mathbf{S}^*, W)\},$$

then it is not true that

$$opt_{\mathcal{Q}}(\mathbf{A}) = \min_{\mathbf{S}^*, W} |\{\mathbf{w} : (\mathbf{A}, \mathbf{S}^*, W)\psi(\mathbf{S}^*, W) \wedge W(\mathbf{w})\}|,$$

because the minimum cardinality of the above set is zero, which occurs when W is empty. Instead, as we saw above

$$opt_{\mathcal{Q}}(\mathbf{A}) = \min_{\mathbf{S}^*, W} |\{\mathbf{w} : (\mathbf{A}, \mathbf{S}^*, W) \models \neg\psi(\mathbf{S}^*, W) \vee W(\mathbf{w})\}|$$

This explains the "dual" behavior in logical definability between maximization and minimization problems, viz. MAX $\mathcal{PB}$ = MAX Π_2, while MIN $\mathcal{PB}$ = MIN Σ_2.

The above theorem 3 implies also that the class MIN Π_2 is contained in the class MIN Σ_2. Thus, at this point we have the following picture for the relationship between the classes MIN Σ_i and MIN Π_i, $0 \leq i \leq 2$.

$$\begin{array}{ccccccc} & & \text{MIN }\Sigma_1 & & & & \\ \text{MIN }\Sigma_0 & \begin{array}{c}\subsetneq\\ \subsetneq\end{array} & & \begin{array}{c}\subsetneq\\ \subsetneq\end{array} & \text{MIN }\Pi_2 \subseteq & & \text{MIN }\Sigma_2 = \text{MIN }\mathcal{PB} \\ & & \text{MIN }\Pi_1 & & & & \end{array}$$

In what follows we establish that the above picture can be simplified considerably. More specifically, we will show that the class MIN $\mathcal{PB}$ coincides with the class MIN Π_1, while the class MIN Σ_1 collapses to the class MIN Σ_0. In particular, MIN VERTEX COVER will turn out to be a member of the class MIN Σ_0. These results are rather surprising, especially when compared with Theorem 2 for the maximization classes, which asserts that MAX Π_1 is a proper subclass of MAX $\mathcal{PB}$ and that MAX Σ_0 is a proper subclass of MAX Σ_1.

Example 1: Before stating and proving the next theorem, we illustrate an instance of it by showing that MIN VERTEX COVER is in the class MIN Σ_0. As we saw earlier, for every graph $G = (V, E)$ we have that

$$opt_{\text{MIN VC}}(G) = \min_S \ |U(S)|,$$

where

$$U(S) = \{w : (G, S) \models (\exists y_1)(\exists y_2)[(E(y_1, y_2) \wedge \neg S(y_1) \wedge \neg S(y_2)) \vee S(w)]\}.$$

Let

$$V(S) = \{(w, x) : (G, S) \models (w = x \wedge S(w)) \vee (E(w, x) \wedge \neg S(w) \wedge \neg S(x))\}.$$

We now claim that for every graph $G = (V, E)$ we have that

$$\min_S |U(S)| = \min_S |V(S)|.$$

Notice that if S^* is a minimum vertex cover for G, then $V(S^*) = \{(w, w) : S^*(w)\}$ and, as a result,

$$\min_S |U(S)| = |S^*| = |V(S^*) \geq \min_S |V(S)|.$$

For the other direction, let S' be a set of vertices such that $|V(S')| = \min_S |V(S)|$. We will show that we can add vertices to S' until it becomes a vertex cover of G without changing the cardinality of the set $V(S')$. Indeed, if (w_1, x_1) is a pair of vertices of G such that

$$(G, S') \models E(w_1, x_1) \wedge \neg S'(w_1) \wedge \neg S'(x_1),$$

we put $S'_1 = S' \cup \{w_1\}$. Then $|V(S'_1)| \leq |V(S')|$, because $V(S'_1)$ contains (w_1, w_1), but it does not contain (w_1, x_1) and, perhaps, other pairs of the form (w_1, x). On the other hand, the minimality property of S' yields that $|V(S')| \leq |V(S'_1)|$ and, consequently, $|V(S')| = |V(S'_1)|$. By repeating this process, we can find a vertex cover S'' of G such that $|V(S')| = |V(S'')|$. It follows that

$$\min_S |U(S)| \leq |V(S'')| = |V(S')| = \min_S |V(S)|$$

and, thus, $\min_S |U(S)| = \min_S |V(S)|$. Since $V(S)$ is defined using a quantifier-free formula, we conclude that MIN VERTEX COVER is in the class MIN Σ_0.

Notice that the quantifier-free formula that defines MIN VERTEX COVER has two free variables w and x, while the Σ_1 formula that defines it has a single free variable w. It turns out that this increase in arity is inevitable, i.e., there is no quantifier-free formula $\psi(w, \mathbf{S})$ with w as its only free variable such that on every graph $G = (V, E)$

$$opt_{\text{MIN VC}}(G) = \min_{\mathbf{S}} |\{w : (G, \mathbf{S}) \models \psi(w, \mathbf{S})\}.$$

Indeed, if such a formula existed, then on every graph $G = (V, E)$ we would have that

$$opt_{\text{MAX CLIQUE}}(G) = \max_{\mathbf{S}} |\{w : (G, \mathbf{S}) \models \neg\psi(w, \mathbf{S})\},$$

which would imply that MAX CLIQUE is in the class MAX Σ_0 and, a fortiori, in the class MAX Σ_1, contradicting Theorem 2.

We are now ready to state and prove the main result of this section.

Theorem 4: The class MIN Σ_1 is contained in the class MIN Σ_0 and the class MIN Σ_2 is contained in the class MIN Π_1. As a result,

$$\text{MIN } \Sigma_0 = \text{MIN } \Sigma_1 \ \subseteq \ \text{MIN } \Pi_1 = \text{MIN } \Sigma_2 = \text{MIN } \mathcal{PB}.$$

Moreover, MIN Σ_1 is a proper subclass of MIN Π_1. In particular, MIN CHROMATIC NUMBER is in MIN Π_1, but not in MIN Σ_1.

Proof: We give this proof in two parts.

Part A: In this part we show that MIN Σ_1 is a subclass of MIN Σ_0 and that MIN Σ_2 is a subclass of MIN Π_1.

Let $\mathcal{Q}$ be a problem in MIN Σ_1 with finite structures over a vocabulary σ as instances. Then there is a quantifier-free formula $\phi(\mathbf{w}, \mathbf{x}, \mathbf{S})$ with predicate symbols from $\sigma \cup \mathbf{S}$ such that for every finite structure $\mathbf{A}$ over the vocabulary σ

$$opt_{\mathcal{Q}}(\mathbf{A}) = \min_{\mathbf{S}} |\{\mathbf{w} : (\mathbf{A}, \mathbf{S}) \models (\exists \mathbf{x})\phi(\mathbf{w}, \mathbf{x}, \mathbf{S})\}|.$$

We can assume, without loss of generality, that the number of variables in the sequence $\mathbf{w}$ is the same as the number of variables in the sequence $\mathbf{x}$. Indeed, let $\mathbf{w}$ be the sequence $(w_1, \cdots, w_m)$ and $\mathbf{x}$ be the sequence $(x_1, \cdots, x_l)$. If $m > l$, we can increase the length of the sequence $\mathbf{x}$ by adding dummy variables $x_{l+1}, \cdots, x_m$. If $m < l$, we introduce new variables, $w_{m+1}, \cdots, w_l$ and express the optimum of $\mathcal{Q}$ as follows:

$$opt_{\mathcal{Q}}(\mathbf{A}) = \min_{\mathbf{S}} |\{(w_1, \cdots, w_m, w_{m+1}, \cdots, w_l) :$$
$$(\mathbf{A}, \mathbf{S}) \models (\exists \mathbf{x})\phi(w_1, \cdots, w_m, \mathbf{x}, \mathbf{S}) \wedge w_m = w_{m+1} = \cdots = w_l\}|.$$

In what follows, we will assume that the number of variables in the sequence $\mathbf{w}$ is the same as the number of variables in the sequence $\mathbf{x}$. Our goal is to find a quantifier-free formula ψ that defines $opt_{\mathcal{Q}}(\mathbf{A})$ on every structure $\mathbf{A}$ over σ. The idea is similar to the one used to construct the quantifier-free formula that defined MIN VERTEX COVER in the preceding Example 1, but the construction of ψ in the general case requires an auxiliary predicate symbol R that is different from all predicate symbols in $\mathbf{S}$.

Put

$$U(\mathbf{S}) = \{\mathbf{w} : (\mathbf{A}, \mathbf{S}) \models (\exists \mathbf{x})\phi(\mathbf{w}, \mathbf{x}, \mathbf{S})\}$$

and notice that

$$\begin{aligned} |U(\mathbf{S})| &= \min_R \{|R| : (\mathbf{A}, \mathbf{S}, R) \models (\forall \mathbf{w})((\exists \mathbf{x})\phi(\mathbf{w}, \mathbf{x}, \mathbf{S}) \rightarrow R(\mathbf{w}))\} \\ &= \min_R \{|R| : (\mathbf{A}, \mathbf{S}, R) \models (\forall \mathbf{w})(\forall \mathbf{x})(\neg\phi(\mathbf{w}, \mathbf{x}, \mathbf{S}) \vee R(\mathbf{w}))\}. \end{aligned}$$

Let

$$V(\mathbf{S}, R) = \{(\mathbf{w}, \mathbf{x}) : (\mathbf{A}, \mathbf{S}, R) \models [(\mathbf{w} = \mathbf{x}) \wedge R(\mathbf{w})] \vee [\phi(\mathbf{w}, \mathbf{x}, \mathbf{S}) \wedge \neg R(\mathbf{w})].$$

We now claim that for every graph $G = (V, E)$ and every sequence $\mathbf{S}$ of relations on V we have that

$$|U(\mathbf{S})| = \min_R |V(\mathbf{S}, R)|.$$

Notice first that if R^* is a relation such that $|U(\mathbf{S}) = |R^*|$, then

$$V(\mathbf{S}, R^*) = \{(\mathbf{w}, \mathbf{w}) : R^*(\mathbf{w})\}$$

and, as a result,

$$|U(\mathbf{S})| = |R^*| = |V(\mathbf{S}, R)| \geq \min_R |V(\mathbf{S}, R)|.$$

For the other direction, let R' be a relation such that $|V(\mathbf{S}, R')| = \min_R |V(\mathbf{S}, R')|$. If $(\mathbf{w_1}, \mathbf{x_1})$ is a pair of tuples from G such that $(G, \mathbf{S}, R') \models \phi(\mathbf{w_1}, \mathbf{x_1}, \mathbf{S}) \wedge \neg R(\mathbf{w_1})$, we put $R'_1 = R \cup \{\mathbf{w_1}\}$. Then $|V(\mathbf{S}, R'_1)| \leq |V(\mathbf{S}, R')|$, because $V(\mathbf{S}, R'_1)$ contains $(\mathbf{w_1}, \mathbf{w_1})$, but does not contain $(\mathbf{w_1}, \mathbf{x_1})$ and, perhaps, other pairs of the form $(\mathbf{w_1}, \mathbf{x})$. On the other hand, the minimality property of R' yields that $|V(\mathbf{S}, R')| \leq |V(\mathbf{S}, R'_1)|$ and, consequently, $|V(\mathbf{S}, R')| = |V(\mathbf{S}, R'_1)|$. By repeating this process, we can find a relation R'' on G such that $|V(\mathbf{S}, R')| = |V(\mathbf{S}, R'')|$ and

$$(\mathbf{A}, \mathbf{S}, R'') \models (\forall \mathbf{w})(\forall \mathbf{x})(\neg\phi(\mathbf{w}, \mathbf{x}, \mathbf{S}) \vee R(\mathbf{w}))\}.$$

It follows that for every sequence $\mathbf{S}$ of relations we have

$$|U(\mathbf{S})| \leq |V(\mathbf{S}, R'')| = |V(\mathbf{S}, R')| = \min_R |V(\mathbf{S}, R)|$$

and, hence, $|U(\mathbf{S})| = \min_R |V(\mathbf{S}, R)|$. Thus,

$$\begin{aligned} opt_{\mathcal{Q}}(\mathbf{A}) &= \min_{\mathbf{S}, R} |V(\mathbf{S}, R)| \\ &= \min_{\mathbf{S}, R} |\{(\mathbf{w}, \mathbf{x}) : (\mathbf{A}, \mathbf{S}, R) \models [(\mathbf{w} = \mathbf{x}) \wedge R(\mathbf{w})] \vee [\phi(\mathbf{w}, \mathbf{x}, \mathbf{S}) \wedge \neg R(\mathbf{w})], \end{aligned}$$

which establishes that $\mathcal{Q}$ is in the class MIN Σ_0.

A similar argument establishes that MIN Σ_2 is a subclass of MIN Π_1.

Part B: In this part of the proof we show that MIN CHROMATIC NUMBER is in the class MIN Π_1, but not in the class MIN Σ_0.

We have already seen that MIN CHROMATIC NUMBER is in the class MIN Σ_2 and hence, by what we proved above, it is in the class MIN Π_1. We now show that MIN CHROMATIC NUMBER is not in the class MIN Σ_0. Towards a contradiction, assume that there is a quantifier-free formula $\psi(\mathbf{w}, \mathbf{S})$ such that for every graph G

$$opt_{\text{CHROMATIC NUMBER}}(G) = \min_{\mathbf{S}} |\{\mathbf{w} : (G, \mathbf{S}) \models \psi(\mathbf{w}, \mathbf{S})\}|.$$

Let k be a positive integer, let H_1 be a graph with $opt_{\text{CHROMATIC NUMBER}}(H_1) = k$, and let H_2 be an isomorphic copy of H_1. If G is the disjoint union (direct sum) of H_1 and H_2, then it is clear that $opt_{\text{CHROMATIC NUMBER}}(G) = k$. Let $\mathbf{S}^*$ be a sequence of relations on G such that

$$|\{\mathbf{w} : (G, \mathbf{S}^*) \models \psi(\mathbf{w}, \mathbf{S}^*)\}| = k$$

and let $\mathbf{S}_1^*$ and $\mathbf{S}_2^*$ be the restrictions of $\mathbf{S}^*$ to the vertex sets of H_1 and H_2 respectively. If $\mathbf{b}$ is a tuple from H_i, $i = 1, 2$, such that $(H_i, \mathbf{S}_1^*) \models \psi(\mathbf{b})$, then it is also the case that $(G, \mathbf{S}^*) \models \psi(\mathbf{b})$, because quantifier-free formulae are preserved under extensions. Notice, however, that for $i = 1, 2$

$$|\{\mathbf{w} : (H_i, \mathbf{S}_i^*) \models \psi(\mathbf{w}, \mathcal{S}_i^*)\}| \geq k,$$

and, moreover, the sets $\{\mathbf{w} : (H_1, \mathbf{S}_1^*) \models \psi(\mathbf{w}, \mathcal{S}_1^*)\}$ and $\{\mathbf{w} : (H_2, \mathbf{S}_2^*) \models \psi(\mathbf{w}, \mathbf{S}_2^*)\}$ are disjoint. Therefore,

$$|\{\mathbf{w} : (G, \mathbf{S}^*) \models \psi(\mathbf{w}, \mathbf{S}^*)\}| \geq 2k,$$

which is a contradiction. Thus, MIN CHROMATIC NUMBER is not in the class MIN Σ_0. $\Box$

5 Approximation Properties of NP Minimization Problems

In this section, we focus on the approximation properties of the minimization classes and contrast them with those of the maximization classes.

Definition 5.1: [PS82] Let $\mathcal{Q} = (\mathcal{I}_{\mathcal{Q}}, \mathcal{F}_{\mathcal{Q}}, f_{\mathcal{Q}}, opt)$ be an NP optimization problem and let A be an algorithm which, given an instance $I \in \mathcal{I}_{\mathcal{Q}}$, returns a feasible solution $T \in \mathcal{F}_{\mathcal{Q}}$. We say that A is an *ϵ-approximation algorithm* for $\mathcal{Q}$ for some $\epsilon \geq 0$ if

$$\frac{|f_{\mathcal{Q}}(I, T) - opt(I)|}{opt_{\mathcal{Q}}(I)} \leq \epsilon$$

for all instances I. The feasible solution T is said to be an *ϵ-approximate solution* for the instance I. An NP optimization problem $\mathcal{Q}$ is *constant-approximable* if for some $\epsilon > 0$ there is a polynomial time ϵ-approximation algorithm for $\mathcal{Q}$. For maximization problems we also require that $\epsilon < 1$, otherwise all maximization problems would be trivially constant-approximable.

MAX 3SAT, MAX SAT, MIN VERTEX COVER, and TRAVELING SALESMAN with Δ-inequality are important examples of constant-approximable optimization problems. Papadimitriou and Yannakakis [PY91] proved that every problem in MAX Σ_1 (and, a fortiori, every problem in MAX Σ_0) is constant-approximable. In contrast to this, we show below that MIN Σ_0 contains natural problems that are not constant-approximable, unless P$\neq$NP. In fact, it turns out that an already familiar problem from the previous section has this property.

Theorem 5: MIN 3NON-TAUTOLOGY is not constant-approximable, unless P=NP.

Proof: Assume that for some $\epsilon > 0$ there is an ϵ-approximation algorithm A for MIN 3NT. We show below that A can be used to solve in polynomial time the NON-TAUTOLOGY problem of 3DNF formulae, a problem that is known to be NP-complete.

Given an instance ϕ of NON TAUTOLOGY of 3DNF formulae, we create in polynomial time an instance Φ of MIN 3NT as follows: Let x be a variable not occurring in ϕ and let $\overline{x}$ be its negated literal. The formula Φ is a disjunction of $x \vee \overline{x}$ and of n copies of every disjunct of ϕ, where $n > (1 + \epsilon)$.

If ϕ is a non-tautology, then $opt_{\mathrm{3NT}}(\Phi) = 1$, because every truth assignment satisfies exactly one of the disjuncts x and $\overline{x}$, and there is a truth assignment under which no disjuncts in any copy of ϕ are satisfied. If ϕ is a tautology, then there is no truth assignment that falsifies every disjunct in ϕ. Hence, in Φ at least one disjunct from each copy of ϕ is satisfied under every truth assignment. Therefore, $opt_{\mathrm{3NT}}(\Phi) \geq n + 1$.

It follows that the formula ϕ is a non-tautology if and only if the algorithm A on input Φ returns a value less than or equal to $(1 + \epsilon)$. Thus, we have exhibited a polynomial time algorithm for solving an NP-complete problem, which implies that P=NP. $\Box$

We now consider an approximation preserving reduction and in Theorem 6 we prove that MIN 3NT is a complete problem for the class MIN Σ_0 under this reduction.

Papadimitriou and Yannakakis [PY91] introduced a notion of L-reduction between optimization problems. Panconesi and Ranjan [PR90] generalized this to the notion of P-reduction. We use here a variant of these reductions introduced by Crescenzi and Panconesi [CP91].

Definition 5.2: [CP91] Let $\mathcal{Q}$ and $\mathcal{R}$ be two NP optimization problems. An *approximability preserving reduction* (or, A-*reduction*) from $\mathcal{Q}$ to $\mathcal{R}$ is a triple $\tau = (t_1, t_2, c)$ for which the following hold:

- t_1 and t_2 are polynomially computable functions with $t_1 : \mathcal{I}_\mathcal{Q} \rightarrow \mathcal{I}_\mathcal{R}$ and $t_2 : \mathcal{I}_\mathcal{R} \times \mathcal{F}_\mathcal{R} \rightarrow \mathcal{F}_\mathcal{Q}$.
- c is a function from non-negative rationals to non-negative rationals such that if T is an ϵ-approximate solution for an instance $t_1(I)$ of $\mathcal{R}$, then $t_2(I,T)$ is a $c(\epsilon)$-approximate solution for $\mathcal{Q}$.

If there is an A-reduction form $\mathcal{Q}$ to $\mathcal{R}$, then we say that $\mathcal{Q}$ is A-*reducible* to $\mathcal{R}$ and we write $\mathcal{Q} \leq_{\mathrm{A}} \mathcal{R}$,

The A-reduction defined above is a more relaxed reducibility than the L-reduction defined by Papadimitriou and Yannakakis [PY91]. In the latter the optimum solutions of the two problems $\mathcal{Q}$ and $\mathcal{R}$ are required to be within a constant factor of each other. Although this is the case with many optimization problems, a reduction may preserve approximability (within a constant factor of the optimal) without having this property.

The following propositions follow easily from the definitions.

Proposition 1: If $\mathcal{R}$ is constant-approximable and $\mathcal{Q} \leq_{\mathrm{A}} \mathcal{R}$, then $\mathcal{Q}$ is constant-approximable.

Proposition 2: A-reductions compose.

Definition 5.3: An NP optimization problem $\mathcal{Q}$ is *approximation complete* for a class $\mathcal{C}$ of optimization problems if $\mathcal{Q} \in \mathcal{C}$ and every problem $\mathcal{R} \in \mathcal{C}$ can be A-reduced to $\mathcal{Q}$.

With the necessary definitions behind us, we can now state and prove the following result.

Theorem 6: MIN 3NON-TAUTOLOGY is approximation complete for MIN Σ_0.

Proof: We have shown before that MIN 3NT is in MIN Σ_0. We now prove that every problem in MIN Σ_0 is A-reducible to it. Let $\mathcal{Q}$ be a problem in MIN Σ_0, let I be an instance of it, and let $\mathbf{A}(I)$ be a structure encoding I. Then there is a quantifier-free formula $\psi(\mathbf{w}, \mathbf{S})$ such that

$$opt_\mathcal{Q}(\mathbf{A}(I)) = \min_{\mathbf{S}} |\{\mathbf{w} : \mathbf{A}(I) \models \psi(\mathbf{w}, \mathbf{S})\}|.$$

Assume that the arity of $\mathbf{w}$ is k and that the size $|\mathbf{A}(I)|$ of $\mathbf{A}(I)$ is equal to n. Let $\{\mathbf{w}_1, \mathbf{w}_2, \cdots, \mathbf{w}_{n^k}\}$ be the possible values of $\mathbf{w}$ on $\mathbf{A}$. For every $\mathbf{w}_i$ we consider the Boolean circuit B_i, composed of gates AND, OR and NOT, that represents the formula $\psi(\mathbf{w}_i, \mathbf{S})$. The inputs to the circuit are of the form $S_i(\mathbf{w}'_i)$, where S_i is a predicate symbol from the sequence $\mathbf{S}$ of symbols and $\mathbf{w}'_i$ is a projection of $\mathbf{w}_i$ of arity the same as the arity of S_i.

Given an instance I of $\mathcal{Q}$, we construct an instance $t_1(I)$ of MIN 3NT. Corresponding to the output of every gate g in the circuit B_i, we have a variable g in $t_1(I)$. The other variables of $t_1(I)$ are the input variables of the circuit. The disjuncts of $t_1(I)$ are as follows. If g is the output of a NOT gate with input x, then we have $(g \wedge x)$ and $(\overline{g} \wedge \overline{x})$ as disjuncts. If g is the output of an AND gate with inputs x_1, x_2, then we have $(\overline{x}_1 \wedge \overline{x}_2 \wedge \overline{g})$, $(\overline{x}_1 \wedge x_2 \wedge \overline{g})$, $(x_1 \wedge \overline{x}_2 \wedge \overline{g})$, and $(x_1 \wedge x_2 \wedge g)$. If g is the output of an OR gate with inputs x_1, x_2, then we have $(\overline{x}_1 \wedge x_2 \wedge \overline{g})$, $(x_1 \wedge \overline{x}_2 \wedge \overline{g})$, $(x_1 \wedge x_2 \wedge \overline{g})$, and $(\overline{x}_1 \wedge \overline{x}_2 \wedge g)$ as disjuncts. Finally, if g is the output of the circuit B_i, then we have a disjunct (g).

Given any input to the circuit B_i, we can set the Boolean values of the intermediate gates such that every disjunct is falsified. The disjuncts are designed such that if g is the output of the AND gate with inputs x_1 and x_2, then setting g to $x_1 \wedge x_2$ will result in falsifying all the disjuncts corresponding to this gate. Similarly, for disjuncts corresponding to OR and NOT gates, if we set the output to the disjunction of the inputs or the negation of the input respectively, then all the disjuncts that correspond to the gate are falsified. Thus, if a truth assignment falsifies $\psi(\mathbf{w}_i, \mathbf{S})$, then we can falsify all the disjuncts corresponding to the circuit B_i. Moreover, if it satisfies $\psi(\mathbf{w}_i, \mathbf{S})$, then the minimum number of disjuncts (corresponding to B_i) satisfied is 1. Hence, $opt_{\mathcal{Q}}(I)$ is equal to the minimum number of satisfiable disjuncts in the instance $t_1(I)$ of 3NT.

In addition, it is straightforward to define the mapping t_2 such that, given an ϵ-approximate truth assignment to the instance $t_1(I)$, we obtain an ϵ-approximate solution to $\mathcal{Q}$. Thus, $\mathcal{Q} \leq_{\mathrm{A}}$ MIN 3NT. □

The preceding Theorem 5 reveals that the pattern of the quantifier prefix does not impact on the approximability of minimization problems, unlike the case of maximization problems. As a result, we have to seek other syntactic features that may imply good approximation properties. We introduce below classes of minimization problems defined by imposing restrictions on the quantifier-free part of formulae and we show that there are natural complete problems for these classes.

Definition 5.4: Let MIN $\mathrm{F}^+\Pi_1(k)$, $k \geq 2$, (F stands for *feasible*) be the class of all minimization problems $\mathcal{Q}$ whose optimum can be expressed as:

$$opt_{\mathcal{Q}}(\mathbf{A}) = \min_S\{|S| : (\mathbf{A}, S) \models (\forall \mathbf{y})\psi(\mathbf{y}, S)\},$$

or, equivalently,

$$opt_{\mathcal{Q}}(\mathbf{A}) = \min_S |\{\mathbf{w} : (\mathbf{A}, S) \models ((\forall \mathbf{y})\psi(\mathbf{y}, S)) \rightarrow S(\mathbf{w})\}|,$$

where S is a single predicate, $\psi(\mathbf{w}, S)$ is a quantifier-free CNF formula in which all occurrences of S are positive, and S occurs at most k times in each clause. We also let

$$\text{MIN } \mathrm{F}^+\Pi_1 = \bigcup_k \text{MIN } \mathrm{F}^+\Pi_1(k)$$

denote the union of these classes.

Notice that the second equation in the above definition shows that the class MIN $F^+\Pi_1$ is a subclass of MIN Σ_1. The canonical example of a problem in the class MIN $F^+\Pi_1(2)$ is MIN VERTEX COVER, since its optimum is given by

$$opt_{\text{MIN VC}}(G) = \min_S\{|S| : (G,S) \models (\forall y_1)(\forall y_2)(\neg E(y_1,y_2) \vee S(y_1) \vee S(y_2))\}.$$

The expressive power of the class $F^+\Pi_1(2)$ has been investigated in [KT91].

By generalizing the vertex cover problem to k-hypergraphs, $k \geq 2$, we can obtain the problem MIN k-HYPERVERTEX COVER. This is a typical example of a problem in the class MIN $F^+\Pi_1(k)$.

Definition 5.5: A *k-hypergraph* is a structure $H = (V, E)$ with $E \subseteq V^k$. A *hypervertex cover* is a set $S \subseteq V$ such that for every k-tuple $(v_1, \ldots, v_k)$ in E we have that S contains some v_i.

Notice that a 2-hypergraph can be viewed as an ordinary graph. Moreover, a hypervertex cover for a 2-hypergraph is a vertex cover in the usual sense of the term.

- The MIN k-HYPERVERTEX COVER problem is to find the cardinality of the smallest hypervertex cover in a k-hypergraph. Its optimum is expressed as:

$$opt_{\text{MIN }k-\text{HVC}}(G) = \min_S\{|S| : (G,S) \models (\forall y_1)\cdots(\forall y_k)(E(y_1,\cdots,y_k) \rightarrow S(y_1) \vee \cdots \vee S(y_k))\}.$$

The MIN VERTEX COVER problem has a rather straightforward polynomial time 1-approximation algorithm [GJ79] that is based on the idea of maximal matching. By generalizing the notion of maximal matching to hypergraphs, we can obtain a polynomial time k-approximation algorithm for the MIN k-HYPERVERTEX COVER problem.

Theorem 7: MIN k-HYPERVERTEX COVER is approximation complete for MIN $F^+\Pi_1(k)$, $k \geq 2$, under A-reductions. As a result, every problem in MIN $F^+\Pi_1$ is constant-approximable.

Proof: Let $\mathcal{Q}$ be a problem in MIN $F\Pi_1(k)$, let I be an instance of it, and let $\mathbf{A}(I)$ be a structure encoding I. Then there is a quantifier-free formula $\psi(\mathbf{y}, S)$ in CNF satisfying the conditions in definition 5.4 such that

$$opt_{\mathcal{Q}}(\mathbf{A}(I)) = \min_S\{|S| : (\mathbf{A}(I), S) \models (\forall \mathbf{y})\psi(\mathbf{y}, S)\}.$$

Assume that the arity of S is m, the arity of $\mathbf{y}$ is k, and the size $|\mathbf{A}(I)|$ of $\mathbf{A}$ is equal to n. Let $\{\mathbf{y}_1, \mathbf{y}_2, \cdots, \mathbf{y}_{n^k}\}$ be the possible values of $\mathbf{y}$ on $\mathbf{A}$. If we let ψ_i be the formula $\psi(\mathbf{y}_i, S)$, then

$$opt_{\mathcal{Q}}(\mathbf{A}(I)) = \min_S\{|S| : (\mathbf{A}(I), S) \models \bigwedge_i \psi_i\}.$$

Notice that $\bigwedge_i \psi_i$ is a CNF formula whose variables are of the form $S(\overline{\mathbf{y}})$, where $\overline{\mathbf{y}}$ is a sequence of length m. From the definition of MIN $\mathrm{F}^{+}\Pi_1(k)$ we know that S occurs at most k times in a clause of ψ. Without loss of generality, we can assume that S occurs exactly k times in each clause. Indeed, if S appears less than k times in a clause, then we can repeat one of its occurrences in that clause. Clauses with no occurrences of S depend only on the structure $\mathbf{A}(I)$ and are true independent of S, hence they can be neglected (if such disjuncts are falsified by $\mathbf{A}(I)$, then we do not have a feasible solution).

Given a structure $\mathbf{A}(I)$ encoding an instance I of a problem in MIN $\mathrm{F}\Pi_1(k)$, we construct an instance $G = (V, E)$ of the MIN k-HYPERVERTEX COVER problem as follows. The set V of vertices of G is the set of all m-tuples from the universe of $\mathbf{A}(I)$. Moreover, if $S(\overline{\mathbf{y}}_{i_1}), S(\overline{\mathbf{y}}_{i_2}), \cdots, S(\overline{\mathbf{y}}_{i_k})$ appear in the same clause in the CNF formula, then $\{\overline{\mathbf{y}}_{i_1}, \overline{\mathbf{y}}_{i_2}, \cdots, \overline{\mathbf{y}}_{i_k}\}$ is an edge in G.

Now observe that $S = \{\overline{\mathbf{y}}_{j_1}, \overline{\mathbf{y}}_{j_2}, \cdots, \overline{\mathbf{y}}_{j_t}\}$ is a hypervertex cover for G if and only if we have $(\mathbf{A}(I), S) \models (\forall \mathbf{y})\psi(\mathbf{y}, S)$.

It follows that $\mathcal{Q}$ is A-reducible to MIN k-HYPERVERTEX COVER and so MIN k-HYPERVERTEX COVER is complete for MIN $\mathrm{F}\Pi_1(k)$. $\square$

The good approximation properties of the class MIN $\mathrm{F}\Pi_1$ should be contrasted with those of the class RMAX introduced in [PR90]. This is a syntactic subclass of MAX Π_1 that is in some sense the "dual" of MIN $\mathrm{F}\Pi_1$. More formally, RMAX is the class of NP maximization problems with optimum definable as

$$opt_{\mathcal{Q}}(\mathbf{A}) = \max_S \{|S| : \mathbf{A} \models (\forall \mathbf{y})\psi(\mathbf{y}, S)\}$$

where S is a single predicate and ψ is a quantifier-free CNF formula in which all occurrences of S are negative. MAX CLIQUE is the canonical example of a problem in RMAX. As mentioned in the Introduction, Arora and Safra [AS92] showed that MAX CLIQUE is not constant-approximable, unless P=NP.

Remark 2: We now consider briefly the effect of taking the *A-closure* of the classes MAX Π_n and MAX Σ_n, i.e., all optimization problems that have an A-reduction to a problem in one of these classes. We have seen before that fine distinctions between NP-maximization problems can be made by focusing on their logical definability. It turns out, however, that some of the distinctions manifested in Theorem 2 disappear by passing to A-closures. Indeed, it can be shown that MAX Π_1 contains problems that are complete for the class MAX Π_2 via A-reductions, such as the MAX Number of Satisfiable Formulae (MAX NSF) problem of [PR90]. As a result, the A-closure of MAX Π_1 contains all polynomially bounded maximization problems. It should be pointed out that a similar situation holds with NP decision problems. For example, 3-COLORABILITY is expressible using a *strict* Σ^1_1 formula, i.e., an existential second-order formula whose first-order part has universal quantifiers only. It is known that NP problems definable by such formulae have certain special properties that are not shared by all NP problems, in particular their asymptotic probabilities obey a 0-1 law ([KV87]). On the other hand, the closure of strict Σ^1_1 formulae under polynomial reductions is the entire class of NP problems.

6 Concluding Remarks and Open Problems

In this paper we investigated NP optimization problems from the standpoint of logical definability and analyzed the relative expressive power of the various classes of NP optimization problems that arise in this framework. One of our findings is that logical definability has different implications for NP maximization problems than it has for NP minimization problems. The original motivation in [PY91] for pursuing the logical definability approach was to find syntactic classes of NP maximization problems with good approximation properties, such as MAX Σ_1, and to pinpoint natural complete problems for these classes. Since the class MIN Σ_1 contains problems that are not constant-approximable (modulo P$\neq$NP), it would be interesting to find syntactic subclasses of MIN Σ_1 that contain constant-approximable problems only. Theorem 7 shows that the class MIN $F^+\Pi_1$ is a first step in this direction.

The TRAVELING SALESMAN problem with possible distances 1 or 2 is an important example of a minimization problem that is constant-approximable. Papadimitriou and Yannakakis [PY90] have shown that every problem in the class MAX Σ_0 is L-reducible to the TRAVELING SALESMAN problem with possible distances 1 or 2. It is an open problem to identify a natural class of minimization problems for which the TRAVELING SALESMAN problem with distances $1, 2$ is complete.

Papadimitriou and Yannakakis [PY91] proved that MAX 3SAT and a host of other problems are complete for MAX Σ_0. Panconesi and Ranjan [PR90] introduced the problem MAX Number of Satisfiable Formulae (MAX NSF) and proved it complete for MAX Π_1. As mentioned earlier, it can be shown that this problem is also complete for the class MAX Π_2 = MAX $\mathcal{PB}$. It is not known, however, if MAX Σ_1 possesses complete problems. On the side of minimization, we proved here that MIN 3NT is complete for the class MIN Σ_0, which, by Theorem 4, is the same as the class MIN Σ_1. It would be interesting to investigate the existence of complete problems for the class MIN Π_1.

Acknowledgements: We are grateful to Christos H. Papadimitriou for several useful telephone conversations, to Alessandro Panconesi for several interesting e-mail exchanges, and to Moshe Y. Vardi for giving us feedback in person on some of the work reported here. Thanks are also due to Phil Long and Shankar Ramamoorthy for valuable comments on the details of the proofs. Finally, we wish to thank Ron Fagin for reading carefully an earlier draft of this paper and providing us with numerous comments that enhanced the readability of the paper.

References

[ADP80] G. Ausiello, A. D'Atri, and M. Protasi. Structure preserving reductions among convex optimization problems. *J. Computer and System Sciences*, 21:136–153, 1980.

[ALM+92] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and intractability of approximation problems. In *Proc. 33rd IEEE Symp. on Foundations of Computer Science*, 1992.

[AS92] S. Arora and S. Safra. Approximating clique is NP-complete. In *Proc. 33rd IEEE Symp. on Foundations of Computer Science*, 1992.

[BH77] L. Berman and J. Hartmanis. On isomorphisms and density of NP and other complete sets. *SIAM Journal for Computing*, 6(2):305 – 322, 1977.

[BJY89] D. Bruschi, D. Joseph, and P. Young. A structural overview of NP optimization problems. Technical Report 861, Department of Computer Science, University of Wisconsin, Madison, 1989.

[CH82] A. Chandra and D. Harel. Structure and complexity of relational queries. *Journal of Computer and System Sciences*, 25:99–128, 1982.

[CP91] P. Crescenzi and A. Panconesi. Completeness in approximation classes. *Information and Computation*, 93(2):241 – 262, 1991. First appeared in Proceedings of FCT'89.

[End72] H. Enderton. *A mathematical introduction to logic*. Academic Press, New York, 1972.

[Fag74] R. Fagin. Generalized first–order spectra and polynomial time recognizable sets. In R. Karp, editor, *Complexity of Computations*, pages 43–73. SIAM-AMS Proc. 7, 1974.

[GJ79] M. R. Garey and D. Johnson. *Computers and Intractability: A Guide to the theory of NP-completeness*. Freeman, San Francisco, 1979.

[Joh74] D. Johnson. Approximation algorithms for combinatorial problems. *J. Computer and System Sciences*, 9:256–278, 1974.

[Kre88] M. Krentel. The complexity of optimization problems. *J. Computer and System Sciences*, 36:490–509, 1988.

[KT91] Ph. G. Kolaitis and M. N. Thakur. Approximation properties of NP minimization classes. In *Proc. Structure in Complexity Theory, Sixth Annual Conference*, pages 353–366, 1991. Full paper to appear in *Journal of Computer and System Scienecs*.

[KV87] Ph. G. Kolaitis and M. Y. Vardi. The decision problem for the probabilities of higher-order properties. In *Proc. 19th ACM Symp. on Theory of Computing*, pages 425–435, 1987.

[LM81] E. Leggett and J. Moore. Optimization problems and the polynomial time hierarchy. *Theoretical Computer Science*, 15:279 – 289, 1981.

[OM90] P. Orponen and H. Manila. On approximation preserving reductions: Complete problems and robust measures. To be published; earlier version: TR, University of Helsiki, 1990.

[PM81] A. Paz and S. Moran. Non-deterministic polynomial optimization problems and their approximation. *Theoretical Computer Science*, 15:251–257, 1981.

[PR90] A. Panconesi and D. Ranjan. Quantifiers and approximation. In *Proceedings 22nd Annual ACM Symp. on Theory of Computing*. ACM, N. Y., 1990.

[PS82] C. Papadimitriou and K. Steiglitz. *Combinatorial Optimization - Algorithms and Complexity*. Prentice Hall, New Jersey, 1982.

[PY90] C. Papadimitriou and M. Yannakakis. The travelling salesman problem with distance one and two. Unpublished manuscript, 1990.

[PY91] C. Papadimitriou and M. Yannakakis. Optimization, approximation and complexity classes. *Journal of Computer and System Sciences*, 43(3):425 – 440, 1991. First appeared in the Proceedings 20th Annual ACM Symp. on Theory of Computing, 1988.

[Sto76] L. J. Stockmeyer. The polynomial time hierarchy. *Theoretical Computer Science*, 3:1–22, 1976.

[Wag86] K. Wagner. More complicated questions about maxima and minima, and some closure properties of NP. *Theoretical Computer Science*, 51:53–80, 1986.