On the Structure of Valiant's Complexity Classes

Peter Bürgisser[†]

Institut für Mathematik, Universität Zürich, Winterthurerstr. 190, CH-8057 Zürich, Switzerland,
`buerg@amath.unizh.ch`

In [26, 28] Valiant developed an algebraic analogue of the theory of NP-completeness for computations of polynomials over a field. We further develop this theory in the spirit of structural complexity and obtain analogues of well-known results by Baker, Gill, and Solovay [1], Ladner [18], and Schöning [23, 24].

We show that if Valiant's hypothesis is true, then there is a p-definable family, which is neither p-computable nor VNP-complete. More generally, we define the posets of p-degrees and c-degrees of p-definable families and prove that any countable poset can be embedded in either of them, provided Valiant's hypothesis is true. Moreover, we establish the existence of minimal pairs for VP in VNP.

Over finite fields, we give a *specific* example of a family of polynomials which is neither VNP-complete nor p-computable, provided the polynomial hierarchy does not collapse.

We define relativized complexity classes VP^h and VNP^h and construct complete families in these classes. Moreover, we prove that there is a p-family h satisfying VP^h = VNP^h.

Keywords: Structural complexity, Algebraic theories of NP-completeness, diagonalization, Poset of degrees.

1 Introduction

One of the most important developments in theoretical computer science is the concept of NP-completeness. Recently, initiated by a paper by Blum, Shub, and Smale [6] (BSS-model), there has been a growing interest in investigating such concepts over general algebraic structures, with the purpose of classifying the complexity of continous problems. But already ten years earlier, Valiant [26, 28] had developed a convincing analogue of the theory of NP-completeness in an entirely algebraic framework, in connection with his famous hardness result for the permanent [27]. In fact, the generating functions of many NP-complete graph problems turn out to be complete in Valiant's sense (cf. [7]). The major differences between the BSS-model and Valiant's model are the absence of uniformity conditions in the latter, and the fact that only straight-line computations are considered (no branching). Both structured models are adapted to the framework of polynomial computations, and we believe that they will be useful for classifing the intrinsic

---

complexity of problems in numerical analysis and in computer algebra (compare Smale [25], Heintz and Morgenstern [16]).

Our goal is to further develop Valiant's approach along the lines of discrete structural complexity theory.

We show that if Valiant's hypothesis is true, then, over any field, there is a p-definable family which is neither p-computable nor VNP-complete. A similar result due to Ladner [18] in the classical P-NP-setting is well-known. Ladner's proof is a diagonalization argument based on an effective enumeration of all polynomial time Turing machines. However, over uncountable structures, this approach causes problems. Malajovich and Meer [20] carried over Ladner's theorem to the setting of the BSS-model over the complex numbers by employing a transfer principle due to Blum et al. [5], which allows a reduction to the countable field of algebraic numbers. The corresponding question over the reals is still open, but it is known to be true under a nonuniformity assumption, cf. Ben-David et al. [3]. One of the reasons our proof works over any field is the nonuniformity of Valiant's model. (For a detailed treatment of these questions in a general model-theoretic context see Chapuis and Koiran [11].)

In [23] Schöning found a powerful and uniform technique for proving the existence of certain "diagonal" recursive sets. We develop a similar technique adapted to Valiant's setting. In this framework, the essence of enumeration and diagonalization arguments can be neatly captured by our notion of a σ-*limit set*, which serves as a substitute for the recursively presentable classes in Schöning's approach.

In Sect. 3 on page 76 we formalize this in a general abstract setting by studying certain compatible quasi-orders on the set $\Omega^{\mathbb{N}}$ of families in a quasi-ordered set $(\Omega, \leq)$, and by proving an abstract diagonalization theorem. Based on this theorem, we proceed in Sect. 4 on page 78 by providing an elegant proof that any countable poset can be embedded in the poset of degrees corresponding to a compatible quasi-order. This is applied in Sect. 5 on page 81 in Valiant's setting to an analogue of the polynomial Turing reduction (c-reduction), as well as to the p-projection. A similar result in the classical P-NP-setting for polynomial Turing or polynomial many-one degrees was stated by Ladner [18]; however, he presented a proof in a special case only. We further remark that the existence of minimal pairs for VP in VNP can be easily guaranteed by our approach. (See Landweber et al. [19] and Schöning [24] for corresponding results in the classical P-NP setting.)

A striking discovery is that we can describe *specific* families of polynomials which are neither VNP-complete nor p-computable. In fact, the family of cut enumerators over a finite field of characteristic p has this property, provided $\mathrm{Mod}_p\mathrm{NP}$ is not contained in P/poly. (The latter condition is satisfied if the polynomial hierarchy does not collapse at the second level.) In the classical, as well as in the BSS-setting, only artificial problems are known to have such properties. This is discussed in Section 6 on page 83.

Finally, in Sect. 7 on page 87, we define relative versions VP^h and VNP^h of Valiant's complexity classes with respect to a p-family h. For these, we have obtained some results in the spirit of Baker et al. [1]. (We remark that Emerson [13] has transfered such results to the BSS-model.) Over infinite fields, we can construct VP^h-complete and VNP^h-complete families with respect to p-projection. In particular, this gives a proof for the existence of VNP-complete families, which is independent of Valiant's intricate reduction for the permanent. Moreover, we can construct a p-family h satisfying $\mathrm{VP}^h = \mathrm{VNP}^h$. We do not know whether there exists a p-family h such that $\mathrm{VP}^h \neq \mathrm{VNP}^h$.

2   Valiant's Model

We briefly recall the main features of Valiant's algebraic model. For detailed expositions see von zur Gathen [15] and [9, Chap. 21].

In this section $\Omega := k[X_1, X_2, \ldots]$ denotes the polynomial ring over a fixed field k in countably many variables X_i. A *p-family* over k is a sequence $f = (f_n) \in \Omega^{\mathbb{N}}$ of multivariate polynomials such that the number of variables as well as the degree of f_n are polynomially bounded (*p-bounded*) functions of n. An example of a *p*-family is the permanent family $\mathrm{PER} = (\mathrm{PER}_n)$, where PER_n is the permanent of an n by n matrix with distinct indeterminate entries.

Let $L(f_n)$ denote the total complexity of f_n, that is, the minimum number of arithmetic operations $+, -, *$ sufficient to compute f_n from the variables X_i and constants in k by a straight-line program. We call a *p*-family f *p-computable* iff $n \mapsto L(f_n)$ is *p*-bounded. The *p*-computable families constitute the complexity class VP. We remark that the restriction to *p*-bounded degrees is a severe one: although X^{2^n} can be computed with only n multiplications, the corresponding sequence is not considered to be *p*-computable, as the degrees grow exponentially.

A *p*-family $f = (f_n)$ is called *p-definable* iff there exists a *p*-computable family $g = (g_n)$ with $g_n \in k[X_1, \ldots, X_{u(n)}]$ such that for all n

$$f_n(X_1, \ldots, X_{v(n)}) = \sum_{e \in \{0,1\}^{u(n)-v(n)}} g_n(X_1, \ldots, X_{v(n)}, e_{v(n)+1}, \ldots, e_{u(n)}) \ . \tag{1}$$

The set of *p*-definable families form the complexity class VNP. The class VP is obviously contained in VNP, and *Valiant's hypothesis* claims that this inclusion is strict. We can consider this as an algebraic counterpart of the well-known hypothesis $P \neq NP$ due to Cook [12]. Let us mention the following recent result due to the author, which reveals a close connection between these two hypotheses.

Theorem 2.1 ([8]) *If Valiant's hypothesis were false over the field k, then the nonuniform versions of the complexity classes* NC, P, NP, *and* PH *would be equal. In particular, the polynomial hierarchy would collapse to the second level. Hereby, we assume that k is finite or of characteristic zero; in the second case we assume a generalized Riemann hypothesis.*

We shall now define a quasi-order $\leq_p$ called *p-projection* on the set $\Omega^{\mathbb{N}}$ of families in Ω. A polynomial f_n is said to be a *projection* of a polynomial $g_m \in k[X_1, \ldots, X_u]$, for short $f_n \leq g_m$, iff

$$f_n(X_1, \ldots, X_{v(n)}) = g_m(a_1, \ldots, a_u) \tag{2}$$

for some $a_i \in k \cup \{X_1, \ldots, X_{v(n)}\}$. That is, f_n can be derived from g_m through substitution by indeterminates and constants. Let us call a function $t: \mathbb{N} \to \mathbb{N}$ *p-bounded from above and below* iff there exists some $c > 0$ such that $n^{1/c} - c \leq t(n) \leq n^c + c$ for all n. We call a *p*-family $f = (f_n)$ a *p-projection* of $g = (g_m)$, in symbols $f \leq_p g$, iff there exists a function $t: \mathbb{N} \to \mathbb{N}$ which is *p*-bounded from above and below such that

$$\exists n_0 \ \forall n \geq n_0 : f_n \leq g_{t(n)} \ . \tag{3}$$

We remark that our definition of $\leq_p$ differs slightly from the one given in [26]. On the one hand, we require the relation $f_n \leq g_{t(n)}$ to hold for sufficiently large n only. In turn, in order to guarantee transitivity of $\leq_p$, we have to make sure that $t(n) \to \infty$ as $n \to \infty$. For our purposes, it is convenient (but not essential) to achieve this by requiring that t is growing at least polynomially.

Finally, a *p*-family $g \in$ VNP is called *VNP-complete* (with regard to *p*-projection) iff any $f \in$ VNP is a *p*-projection of g.

In [26] Valiant obtained the remarkable result that the permanent family (if char $k \neq 2$) and the family of Hamilton cycle polynomials are VNP-complete. It turns out that the generating functions of several NP-complete graph problems like Clique, factors, Hamilton cycles in planar graphs etc. are VNP-complete as well (cf. [7]).

3 An Abstract Diagonalization Theorem

Let a quasi-ordered set $(\Omega, \leq)$ be fixed. Elements of the set $\Omega^{\mathbb{N}}$ of sequences in Ω will be called *families* in the sequel. We may formally define a quasi-order $\leq_p$ (the abstract *p*-projection) on the set $\Omega^{\mathbb{N}}$ of families as in (3 on the page before). Two families f and g are said to be *p-equivalent* iff $f \leq_p g$ and $g \leq_p f$. We call the equivalence classes *p-degrees* and denote by $\mathcal{D}_p$ the poset of all *p*-degrees with the partial order induced by $\leq_p$. $f <_p g$ shall mean that $f \leq_p g$ but not $g \leq_p f$. The *join* $f \cup g$ of two families $f, g \in \Omega^{\mathbb{N}}$ is defined as

$$f \cup g := (f_0, g_0, f_1, g_1, f_2, g_2, \ldots) \ . \tag{4}$$

It is easy to see that the join of two *p*-degrees is well-defined and that it is the smallest upper bound of these *p*-degrees in $\mathcal{D}_p$. The poset $\mathcal{D}_p$ of *p*-degrees is thus a join-semilattice.

Definition 3.1 By a *cylinder* in $\Omega^{\mathbb{N}}$ we shall understand a set of families of the form $F \times \Omega^{\mathbb{N}}$, where $F \subseteq \Omega^n$ for some $n \in \mathbb{N}$. A *limit of cylinders* is defined as a countable intersection of cylinders. By σ-*limit set* in $\Omega^{\mathbb{N}}$ we shall understand a countable union of limits of cylinders.

We remark that countable unions and finite intersections of σ-limit sets are again σ-limit sets. In Example 3.5 on page 78 at the end of this section, we will see that the σ-limit sets are not closed under the formation of complements and countable intersections. Note that if $F_\nu \subseteq \Omega^{n_\nu}$ for $n_\nu \in \mathbb{N}$, then the cartesian product $\prod_\nu F_\nu$ is a limit of cylinders.

Lemma 3.2 $\{h \mid h \leq_p g\}$ *and* $\{h \mid f \leq_p h\}$ *are* σ-*limit sets for all* $f, g \in \Omega^{\mathbb{N}}$.

Proof. It is convenient to use the abbreviation

$$I(c, n) := \{m \mid n^{1/c} - c \leq m \leq n^c + c\} \ . \tag{5}$$

Note that $h \leq_p g$ can be expressed by the following predicate

$$\exists c > 0 \ \exists n_0 \ \forall n \geq n_0 \ \exists m : m \in I(c, n) \ \wedge \ h_n \leq g_m \ , \tag{6}$$

where the quantification is over natural numbers. Thus if we set

$$U(c, n) := \{u \in \Omega \mid \exists m : m \in I(c, n) \ \wedge \ u \leq g_m\} \ , \tag{7}$$

then we can write

$$\{h \mid h \leq_p g\} = \bigcup_{c, n_0} \left(\Omega^{n_0} \times \prod_{n \geq n_0} U(c, n)\right) \ , \tag{8}$$

hence $\{h \mid h \leq_p g\}$ is a σ-limit set.

On the other hand, we may write $\{h \mid f \leq_p h\}$ as the countable union over all c, n_0 of the following limits of cylinders

$$\bigcap_{n \geq n_0} \bigcup_{m \in I(n, c)} \left(\Omega^m \times \{v \in \Omega \mid f_n \leq v\} \times \Omega^{\mathbb{N}}\right) \ . \tag{9}$$

Therefore, $\{h \mid f \leq_p h\}$ is a σ-limit set. $\qquad\square$

We remark that the order $\leq_p$ is compatible with the join in the following sense: for all σ-limit sets $\mathcal{F} \subseteq \Omega^{\mathbb{N}}$ the set $\{(f, g) \mid f \cup g \in \mathcal{F}\}$ is a σ-limit subset of $(\Omega \times \Omega)^{\mathbb{N}}$ via the identification $\Omega^{\mathbb{N}} \times \Omega^{\mathbb{N}} \to (\Omega \times \Omega)^{\overline{\mathbb{N}}}$ sending $((f_n), (g_n))$ to $((f_n, g_n))$.

We call a family (f_n) a *finite variation* of a family (g_n) iff $f_n = g_n$ for all but finitely many n. Note that if f is a finite variation of g, then f and g are in the same p-degree. Subsets of $\Omega^{\mathbb{N}}$ which are closed under finite variation capture asymptotic properties of families (f_n) for $n \to \infty$. (In probability theory one calls them tail events.)

The following abstract diagonalization theorem is inspired by Schöning's uniform diagonalization theorem [23] (see also Balcázar et al. [2]). We note that the σ-limit sets serve as a substitute for the recursively presentable classes appearing there.

Theorem 3.3 *Let $\mathcal{F}, \mathcal{G}$ be σ-limit sets of $\Omega^{\mathbb{N}}$ which are closed under finite variation. Moreover, let $f, g \in \Omega^{\mathbb{N}}$ such that $f \notin \mathcal{F}$ and $g \notin \mathcal{G}$. Then there exists $h \in \Omega^{\mathbb{N}}$ satisfying $h \leq_p f \cup g$ and $h \notin \mathcal{F} \cup \mathcal{G}$.*

Proof. As $\mathcal{F}$ and $\mathcal{G}$ are σ-limit sets, we have representations $\mathcal{F} = \bigcup_i \bigcap_j \mathcal{F}_{ij}$ and $\mathcal{G} = \bigcup_i \bigcap_j \mathcal{G}_{ij}$ where $\mathcal{F}_{ij}$ and $\mathcal{G}_{ij}$ are cylinders in $\Omega^{\mathbb{N}}$. Clearly, we may assume that $\mathcal{F}_{i0} \supseteq \mathcal{F}_{i1} \supseteq \ldots$ and $\mathcal{G}_{i0} \supseteq \mathcal{G}_{i1} \supseteq \ldots$. We denote by π_n the projection $\Omega^{\mathbb{N}} \to \Omega^n, (f_\nu) \mapsto (f_0, \ldots, f_{n-1})$.

By induction, we are going to construct an infinite sequence $b_0 := 0 < a_1 < b_1 < a_2 < b_2 < \ldots$ of natural numbers such that the corresponding "mixture" h of the families f and g defined by

$$h_\nu := \begin{cases} f_\nu & \text{if } b_{s-1} \leq \nu < a_s \text{ for some } s \\ g_\nu & \text{if } a_s \leq \nu < b_s \text{ for some } s \end{cases} \tag{10}$$

satisfies for all i

$$\pi_{a_i}(h) \notin \pi_{a_i}(\mathcal{F}_{ia_i}), \quad \pi_{b_i}(h) \notin \pi_{b_i}(\mathcal{G}_{ib_i}) . \tag{11}$$

Let us first show that the resulting h fulfills the requirements of the theorem. It is clear that $h \leq_p f \cup g$. Assume by contradiction that $h \in \mathcal{F}$. Then there exists some i such that $h \in \mathcal{F}_{ij}$ for all j. Choosing $j = a_i$, we get a contradiction to (11). Analogously, one shows that $h \notin \mathcal{G}$.

Assume now that we have already constructed $0 < a_1 < b_1 < \ldots < a_{i-1} < b_{i-1}$. Then the elements h_ν for $\nu < b_{i-1}$ are already determined. Consider the following finite variation

$$\tilde{f} := (h_0, h_1, \ldots, h_{b_{i-1}-1}, f_{b_{i-1}}, f_{b_{i-1}+1}, \ldots) \tag{12}$$

of the family f. Since $f \notin \mathcal{F}$ and $\mathcal{F}$ is closed under finite variation, we have $\tilde{f} \notin \mathcal{F}$. Therefore, $\tilde{f} \notin \mathcal{F}_{ij}$ for some j. As $\mathcal{F}_{ij}$ is a cylinder, there exists N such that $\mathcal{F}_{ij} = \pi_n^{-1}(\pi_n(\mathcal{F}_{ij}))$ for all $n \geq N$. Now choose $a_i := \max\{b_{i-1} + 1, j, N\}$. Then we have

$$\pi_{a_i}^{-1}(\pi_{a_i}(\mathcal{F}_{ia_i})) \subseteq \pi_{a_i}^{-1}(\pi_{a_i}(\mathcal{F}_{ij})) = \mathcal{F}_{ij} . \tag{13}$$

Therefore, $\pi_{a_i}(\tilde{f}) \notin \pi_{a_i}(\mathcal{F}_{ia_i})$. The corresponding extension of the sequence h up to index $a_i - 1$ therefore satisfies the desired property. The index b_i can be found similarly by considering the finite variation $\tilde{g} = (h_0, \ldots, h_{a_i-1}, g_{a_i}, g_{a_i+1}, \ldots)$ of the family g. $\qquad\square$

By induction one can easily generalize Thm. 3.3 to an arbitrary finite number of σ-limit sets.

Corollary 3.4 *Let $\mathcal{F}_1, \ldots, \mathcal{F}_s$ be σ-limit sets of $\Omega^{\mathbb{N}}$ which are closed under finite variation and let $f_i \in \Omega^{\mathbb{N}} \setminus \mathcal{F}_i$ for $i = 1, \ldots, s$. Then there exists $h \in \Omega^{\mathbb{N}}$ satisfying $h \leq_p f_1 \cup \ldots \cup f_s$ and $h \notin \mathcal{F}_i$ for all i.*

Example 3.5 Consider $\Omega = \{0,1\}$ with the natural order $\leq$. We define the support of a family $h \in \Omega^{\mathbb{N}}$ as $\{\nu \mid h_\nu \neq 0\}$. The families with finite support form a σ-limit set $\mathcal{G}$. We claim that the complement $\mathcal{F}$ of $\mathcal{G}$ is not a σ-limit set. In fact, otherwise, Thm. 3.3 on the page before with the constant families $f = (0) \notin \mathcal{F}$ and $g = (1) \notin \mathcal{G}$ would imply the existence of a family $h \notin \mathcal{F} \cup \mathcal{G}$, which is absurd. This example also shows that the σ-limit sets are not closed under the formation of countable intersections: we have $\mathcal{F} = \cap_n \mathcal{F}_n$, where $\mathcal{F}_n$ denotes the σ-limit set $\mathcal{F}_n := \cup_{m \geq n}(\Omega^m \times \{1\} \times \Omega^{\mathbb{N}})$.

4   An Abstract Embedding Theorem

Again let a quasi-ordered set $(\Omega, \leq)$ be fixed and denote by $\leq_p$ the corresponding abstract p-projection. We extend our discussion to any quasi-order on $\Omega^{\mathbb{N}}$ which satisfies certain compatibility conditions.

Definition 4.1 A quasi-order $\leq_c$ of $\Omega^{\mathbb{N}}$ is called *compatible*, iff the following conditions are satisfied:

  (a) $\forall f, g : f \leq_p g \Rightarrow f \leq_c g$.

  (b) $\forall f, g, h : f \leq_c h, g \leq_c h \Rightarrow f \cup g \leq_c h$.

  (c) $\{h \mid h \leq_c g\}$ and $\{h \mid f \leq_c h\}$ are σ-limit sets for all $f, g \in \Omega^{\mathbb{N}}$.

Observe that $\leq_p$ is a compatible quasi-ordering by Lemma 3.2 on page 76.

In the sequel, let a compatible quasi-order $\leq_c$ on $\Omega^{\mathbb{N}}$ be fixed. We call two families f and g *c-equivalent* iff $f \leq_c g$ and $g \leq_c f$. The corresponding equivalence classes are a union of certain p-degrees and called *c-degrees*. $f <_c g$ shall mean that $f \leq_c g$, but not $g \leq_c f$. We say that $f <_p g$ holds *strongly* iff $f \leq_p g$ and $f <_c g$.

Let $(X, \subseteq)$ be a poset. A map $\varphi : X \to \Omega^{\mathbb{N}}$ is called an *embedding* of X in $\Omega^{\mathbb{N}}$ (with respect to $\leq_c$) if $x \subseteq y$ implies $\varphi(x) \leq_c \varphi(y)$ and vice versa. We call φ a *strong embedding* iff φ is an embedding and $x \subseteq y$ implies even $\varphi(x) \leq_p \varphi(y)$.

The goal of this section is to prove the following abstract embedding theorem.

Theorem 4.2 *For any countable poset $(X, \subseteq)$ and elements $f, g \in \Omega^{\mathbb{N}}$ with $f <_c g$ there is an embedding $X \to \{h \mid f <_c h <_c g\}$. If additionally $f \leq_p g$, then there is a strong embedding $X \to \{h \mid f <_p h <_p g\}$.*

The proof will be based on a sophisticated application of our abstract diagonalization theorem 3.3 on the page before. In the next lemma, we settle the special case where X consists of one point only.

Lemma 4.3 *For $f, g \in \Omega^{\mathbb{N}}$ with $f <_c g$ there exists $h \in \Omega^{\mathbb{N}}$ such that $f <_c h <_c g$. If also $f \leq_p g$, then we may additionally achieve that $f \leq_p h \leq_p g$.*

Proof. $\mathcal{F} := \{h \mid g \leq_c h \cup f\}$ and $\mathcal{G} := \{h \mid h \leq_c f\}$ are σ-limit sets, as $\leq_c$ is compatible. (Observe that $\mathcal{H} := \{h' \mid g \leq_c h'\}$ and thus $\mathcal{F} = \{h \mid h \cup f \in \mathcal{H}\}$ is a σ-limit set by the remark following Lemma 3.2 on page 76.) Moreover, $\mathcal{F}$ and $\mathcal{G}$ are closed under finite variation (use property (a) in Def. 4.1). By our assumption $f <_c g$ we have $f \notin \mathcal{F}$ and $g \notin \mathcal{G}$.

Theorem 3.3 on the preceding page implies the existence of some $h' \in \Omega^{\mathbb{N}}$ satisfying $h' \leq_p f \cup g$ and $h' \notin \mathcal{F} \cup \mathcal{G}$. Now put $h := h' \cup f$. Using property (b) in Def. 4.1 we conclude from $f \leq_c g$ that $f \leq_p h \leq_c g$. On the other hand, since $h' \notin \mathcal{F} \cup \mathcal{G}$, we have the strict inequalities $f <_c h <_c g$. Of course, if additionally $f \leq_p g$, then we even get $h \leq_p g$. $\square$

We need two further auxiliary results.

Lemma 4.4 *Any countable poset* $(X, \subseteq)$ *can be embedded in a countable lattice.*

Proof. For $y \in X$ denote by $X_y := \{x \in X \mid x \subseteq y\}$ the initial segment of y. Let $\mathcal{A}$ denote the boolean subalgebra generated by all initial segments. $\mathcal{A}$ is a countable lattice with respect to inclusion and the map $X \to \mathcal{A}, y \mapsto X_y$ obviously defines an order isomorphism. $\qquad\square$

Lemma 4.5 *Let* $(X, \subseteq)$ *be a countable lattice. Then there exists an enumeration* $x_0, x_1, x_2, \ldots$ *of* X *such that each* $X_n := \{x_0, \ldots, x_n\}$ *is closed under taking meets: that is,* $x \cap y \in X_n$ *for all* $x, y \in X_n$.

Proof. We may assume that X is infinite. Let $\nu \colon X \to \mathbb{N}$ be any enumeration of X. We proceed recursively: set $x_0 := \nu^{-1}(0)$. Assume now that $x_0, \ldots, x_{n-1}$ are already constructed. Let $z_n \in X \setminus X_{n-1}$ such that $\nu(z_n)$ is minimal and consider the finite set

$$M_n := \{z_n \cap x \mid x \in X_{n-1}, \ z_n \cap x \notin X_{n-1}\} \ . \tag{14}$$

If M_n is empty, we put $x_n := z_n$. Then X_n is obviously closed under taking meets.

Otherwise, let $z_n \cap x$ be a minimal element of M_n w.r.t. $\subseteq$ and define $x_n := z_n \cap x$. To show that X_n is closed under the formation of meets, let $a \in X_{n-1}$. We have $x_n \cap a = (z_n \cap x) \cap a = z_n \cap x'$, where $x' := x \cap a$ is in X_{n-1} by the induction hypothesis. By the minimality in M_n we see that in fact $x_n \cap a \in X_n$.

It remains to show that $x_0, x_1, \ldots$ exhaust all elements of X. If this were not the case, take $z \in X \setminus \{x_0, x_1, \ldots\}$ with a minimum value of ν. Hence there exists n_0 such that

$$\{y \mid \nu(y) < \nu(z)\} \subseteq \{x_0, \ldots, x_{n_0}\} \ . \tag{15}$$

Therefore, we have $z_n = z$ and $M_n \neq \emptyset$ for all $n > n_0$. It is easy to check that $M_{n+1} \subseteq M_n \setminus \{x_n\}$ for $n > n_0$. Thus we would obtain a infinite strictly descending chain of subsets in the finite set M_{n_0}, which is absurd. $\qquad\square$

Proof. (of Thm. 4.2 on the preceding page) Let $(X, \subseteq)$ be a countable poset and assume that $f <_p g$ holds strongly. (The case where we only know that $f <_c g$ can be settled similarly.)

By the Lemmas 4.4 and 4.5 we may assume that $(X, \subseteq)$ is a lattice and that $x_0, x_1, \ldots$ is an enumeration of X such that each $X_n = \{x_0, \ldots, x_n\}$ is closed under the formation of meets.

By induction on n, we shall construct maps $\varphi_n \colon X_n \to \Omega^{\mathbb{N}}$ satisfying $\varphi_n \mid_{X_{n-1}} = \varphi_{n-1}$ and such that the following two properties are satisfied:

(a) $f <_p \bigcap_{x \in X_n} \varphi_n(x)$ strongly , $\bigcup_{x \in X_n} \varphi_n(x) <_p g$ strongly .

(b) For all $x, y, y_1, \ldots, y_s \in X_n$ we have

$$x \subseteq y \Rightarrow \varphi_n(x) \leq_p \varphi_n(y), \ \varphi_n(x) \leq_c \varphi_n(y_1) \cup \ldots \cup \varphi_n(y_s) \Rightarrow x \subseteq y_1 \cup \ldots \cup y_s \ . \tag{16}$$

The induction start where $n = 0$ is guaranteed by Lemma 4.3 on page 78. Now let $n > 0$ and assume that $x_0, \ldots, x_{n-1}$ satisfying the claim are already constructed. To simplify notation we write $\varphi := \varphi_{n-1}$, $A := X_{n-1}$, and $z := x_n$. Let $a_1, \ldots, a_p$ denote the maximal elements of A which are smaller than z and $b_1, \ldots, b_q$ be the minimal elements of A which are bigger than z. Thus for all $x, y \in A$ the relation $x \subseteq z$ implies $x \subseteq a_i$ for some i and $z \subseteq y$ implies $b_j \subseteq y$ for some j.

We are going to distinguish several cases.

Case 1: $p \geq 1, q \geq 1$.

We set $u := \varphi(a_1) \cup \ldots \cup \varphi(a_p)$ and $o := \varphi(b_1 \cap \ldots \cap b_q)$ (note that this is well defined, as A is closed under taking meets). Clearly $u \leq_p o$. For $x \in A$ and $y = (y_1, \ldots, y_s) \in A^s$ we define the set

$$\mathcal{G}_y := \{h \mid h \leq_c \varphi(y_1) \cup \ldots \cup \varphi(y_s)\} \tag{17}$$

if $z \nsubseteq y_1 \cup \ldots \cup y_s$, $s \geq 1$, and we define

$$\mathcal{F}_{x,y} := \{h \mid \varphi(x) \leq_c h \cup u \cup \varphi(y_1) \cup \ldots \cup \varphi(y_s)\} \tag{18}$$

if $x \nsubseteq z \cup y_1 \cup \ldots \cup y_s$, $s \geq 0$. All these sets $\mathcal{G}_y, \mathcal{F}_{x,y}$ are σ-limits and closed under finite variation. There is at most a finite number of them.

We claim that u lies in none of the sets $\mathcal{F}_{x,y}$. Otherwise, we would have $\varphi(x) \leq_c u \cup \varphi(y_1) \cup \ldots \cup \varphi(y_s)$ which implied by (b) that $x \subseteq a_1 \cup \ldots \cup a_p \cup y_1 \cup \ldots \cup y_s \subseteq z \cup y_1 \cup \ldots \cup y_s$, contradicting our assumption. In the same way one sees that o lies in none of the sets $\mathcal{G}_y$.

By Cor. 3.4 on page 77 of the abstract diagonalization theorem, there is some h which lies in none of the sets $\mathcal{G}_y$ and $\mathcal{F}_{x,y}$ and such that $h \leq_p u \cup o \leq_p o$. We extend now the map $\varphi = \varphi_{n-1}$ to X_n by setting $\varphi_n(z) := h \cup u$. Then condition (a) of the inductive claim is obviously satisfied. Moreover, we have

$$\varphi(a_i) \leq_p u \leq_p \varphi_n(z) \leq_p o \leq \varphi(b_j) \tag{19}$$

for all i, j, which, together with the inductive hypothesis, shows that $x \subseteq y$ implies $\varphi_n(x) \leq_p \varphi_n(y)$ for all $x, y \in X_n$. To prove the second part of the claim (b) assume that $\varphi_n(z) \leq_c \varphi(y_1) \cup \ldots \cup \varphi(y_s)$. If we had $z \nsubseteq y_1 \cup \ldots \cup y_s$, then we would obtain the contradiction $h \in \mathcal{G}_y$, Similarly, $\varphi(x) \leq_c \varphi_n(z) \cup \varphi(y_1) \cup \ldots \cup \varphi(y_s)$ implies $x \subseteq z \cup y_1 \cup \ldots \cup y_s$, since $h \notin \mathcal{F}_{x,y}$. This proves part (b) of the claim.

Case 2: $p \geq 1, q = 0$.

By Lemma 4.3 on page 78 there exists g' such that $\cup_{x \in A}\varphi(x) <_p g' <_p g$ holds strongly. We set $u := \varphi(a_1) \cup \ldots \cup \varphi(a_p)$ but now we define $o := g'$. For the sets $\mathcal{G}_y$ and $\mathcal{F}_{x,y}$ introduced as above we have $u \notin \mathcal{F}_{x,y}$ and $o \notin \mathcal{G}_y$. By Cor. 3.4 on page 77 there is some h lying in none of the $\mathcal{G}_y, \mathcal{F}_{x,y}$ and such that $h \leq_p u \cup o \leq_p o$, and we define $\varphi_n(z) := h \cup u$. Then condition (a) of the claim remains valid, as $\cup_{x \in X_n} \leq_p g'$ and $g' <_p g$ strongly holds. The remaining conditions can be checked as in Case 1.

Case 3: $p = 0, q \geq 1$.

By Lemma 4.3 on page 78 there exists f' such that $f <_p f' <_p \cap_{x \in A}$ holds strongly. We take $u := f'$ and $o := \varphi(b_1 \cap \ldots \cap b_q)$. The sets $\mathcal{G}_y$ and $\mathcal{F}_{x,y}$ are defined as before, but using the element $u = f'$. We proceed now similarly as before.

Case 4: $p = 0, q = 0$.

By Lemma 4.3 on page 78 there exist f', g' satisfying $f <_p f' <_p \cap_{x \in A}$ and $\cup_{x \in A} \varphi(x) <_p g' <_p g$ strongly. Take $u := f'$, $o := g'$, and proceed similarly as before. $\square$

5 Structure of Valiant's Complexity Classes

In this section, we apply our previous results to the setting of Valiant. Let $\Omega := k[X_1, X_2, \ldots]$ denote the polynomial ring over a fixed field k in countably many variables X_i and consider the projection $\leq$, which is a quasi-order on Ω. (Recall that $f \leq g$ iff f can be obtained from g by a substitution of its variables by variables or constants in k.) The corresponding quasi-order $\leq_p$ on $\Omega^{\mathbb{N}}$ is the usual p-projection.

To avoid confusions, we remark that in the future symbols like $f, g, h, \ldots$ will be used to denote either polynomials or sequences of polynomials; it will always be clear from the context what is meant.

We introduce the concept of oracle computations. Let a polynomial $g \in k[X_1, \ldots, X_s]$ be given. We consider straight-line programs which, beside the usual arithmetic operations, have the ability to evaluate the "oracle polynomial" g at previously computed values at unit cost. This can easily be formalized by considering straight-line programs Γ of type $\{+, -, *, o\}$, where the symbol o stands for the oracle operation of arity s.

Definition 5.1 The *oracle complexity* $L^g(f_1, \ldots, f_t)$ of a set of polynomials $f_1, \ldots, f_t \in \Omega$ with respect to the oracle polynomial g is the minimum number of arithmetic operations $+, -, *$ and evaluations of g (at previously computed values) that are sufficient to compute the f_j from the indeterminates X_i and constants in k.

We introduce next the notion of *c-reduction*, which can be seen as an analogue of the polynomial Turing reduction for Valiant's setting. (*c* is an acronym for computation.) One might also interpret the p-projection as an analogue of the polynomial many-one reduction, however, the p-projection is much finer.

Definition 5.2 Let $f = (f_n)$, $g = (g_n) \in \Omega^{\mathbb{N}}$. We call f a *c-reduction* (or polynomial oracle reduction) of g, shortly $f \leq_c g$, iff there is a p-bounded function $t : \mathbb{N} \to \mathbb{N}$ such that the map $n \mapsto L^{g_{t(n)}}(f_n)$ is p-bounded.

It easy to check that $\leq_c$ is a quasi-order of $\Omega^{\mathbb{N}}$. Note that for a p-family f we have $f \leq_c 0$ iff f is p-computable.

Lemma 5.3 *The c-reduction $\leq_c$ is a compatible quasi-order on $\Omega^{\mathbb{N}}$.*

Proof. The verification of conditions (a) and (b) of Def. 4.1 on page 78 is straightforward. Condition (c) will be shown similarly as in the proof of Lemma 3.2 on page 76. We can express $h \leq_c g$ by the following predicate

$$\exists c \ \forall n \ \exists m : m \leq n^c + c \ \wedge \ L^{g_m}(h_n) \leq n^c + c \ . \tag{20}$$

If we write

$$V(c, n) := \{u \in \Omega \mid \exists m \leq n^c + c : L^{g_m}(u) \leq n^c + c\} \ , \tag{21}$$

then we have $\{h \mid h \leq_c g\} = \bigcup_c \prod_n V(c, n)$, which shows that this is a σ-limit set.

On the other hand, we may write $\{h \mid f \leq_c h\}$ as the countable union over all c of the following limits of cylinders

$$\bigcap_n \bigcup_{m \leq n^c + c} \left(\Omega^m \times \{v \in \Omega \mid L^v(f_n) \leq n^c + c\} \times \Omega^{\mathbb{N}} \right) . \tag{22}$$

$\square$

Corollary 5.4 *The set of p-families as well as the classes* VP *and* VNP *are* σ-*limit sets.*

Proof. We leave it to the reader to check that the set $\mathcal{P}$ of p-families is a σ-limit. Let g be VNP-complete w.r.t. p-projection. We have $\text{VP} = \{f \in \Omega^{\mathbb{N}} \mid f \leq_c 0\} \cap \mathcal{P}$ and $\text{VNP} = \{f \in \Omega^{\mathbb{N}} \mid f \leq_p g\}$. Thus we may conclude from Lemma 5.3 on the page before and the fact that the p-projection is compatible, that both of these sets are σ-limits. $\square$

Let us call a p-degree or a c-degree *p-definable* iff it contains a p-definable family. Note that a p-definable p-degree consists of p-definable families only, whereas a p-definable c-degree might also contain families which are not in VNP. This is because $f \leq_c g$ and $g \in \text{VNP}$ might not imply that $f \in \text{VNP}$. We denote by $\mathcal{PD}_p$ the set of p-degrees of p-definable families and by $\mathcal{PD}_c$ the set of c-degrees of p-definable families.

Remark 5.5   1. The poset $\mathcal{PD}_p$ has a unique maximal p-degree which consists of the VNP-complete families with respect to p-projection. Any family (f_n) of constants (i.e., $f_n \in k$ for all n) constitutes a minimal p-degree in $\mathcal{PD}_p$, and these are all the minimal p-degrees in $\mathcal{PD}_p$. (Hence $\mathcal{PD}_p$ has at least the cardinality of the continuum.)

2. The poset $\mathcal{PD}_c$ has a unique maximal c-degree which consists of the VNP-complete families with respect to c-reduction. The complexity class VP forms the unique minimal c-degree in $\mathcal{PD}_c$. Valiant's hypothesis "VNP $\neq$ VP" means that $\mathcal{PD}_c$ consists of more than one element.

The main result of this section is analogous to that of Ladner's work [18]. It follows now easily from our abstract embedding theorem 4.2 on page 78.

Theorem 5.6 *Any countable poset can be embedded in the poset* $\mathcal{PD}_p$. *If Valiant's hypothesis is true, then any countable poset can be embedded in the poset* $\mathcal{PD}_c$.

Note that the result on $\mathcal{PD}_p$ is unconditional due to Remark 5.5.1.

Corollary 5.7 *If Valiant's hypothesis is true, then there is a p-definable family which is neither p-computable nor* VNP-*complete with respect to c-reduction.*

We finally show that an analogue of Schöning's general minimal pair theorem [24] holds in Valiant's setting. We call a pair of families $\varphi, \psi \in \Omega^{\mathbb{N}}$ a *minimal pair* for VP iff φ and ψ are not contained in VP and

$$\forall h \in \Omega^{\mathbb{N}} : h \leq_c \varphi \ \wedge \ h \leq_c \psi \Longrightarrow h \in \text{VP} . \tag{23}$$

Theorem 5.8 *Assume that* $\mathcal{F} \subseteq \Omega^{\mathbb{N}}$ *is a* σ-*limit set containing* VP *which is closed under finite variation, and let* $f, g \in \Omega^{\mathbb{N}} \setminus \mathcal{F}$. *Then there exist* $\varphi, \psi \in \Omega^{\mathbb{N}} \setminus \mathcal{F}$ *such that* $\varphi \leq_p f$, $\psi \leq_p g$, *and such that* φ, ψ *is a minimal pair for* VP.

Proof. Let $\mathcal{F} = \bigcup_i \bigcap_j \mathcal{F}_{ij}$ with cylinders $\mathcal{F}_{ij}$ satisfying $\mathcal{F}_{ij} \supseteq \mathcal{F}_{ij+1}$. By induction, we will construct a sequence $0 = a_{10} < a_{11} < a_{12} < a_{13} < a_{14} < a_{15} < a_{20} < \ldots < a_{25} < \ldots$ of natural numbers satisfying the requirements below. We define families φ and ψ corresponding to the sequence $(a_{ij})_{1 \leq i, 0 \leq j \leq 5}$, by setting

$$\varphi_\nu := \begin{cases} f_\nu & \text{if } \exists i : a_{i0} \leq \nu < a_{i1} \\ 0 & \text{otherwise} \end{cases} \qquad \psi_\nu := \begin{cases} g_\nu & \text{if } \exists i : a_{i3} \leq \nu < a_{i4} \\ 0 & \text{otherwise.} \end{cases} \tag{24}$$

The requirements are:

$$
\begin{array}{llll}
(0) & \pi_{a_{i1}}(\varphi) \notin \pi_{a_{i1}}(\mathcal{F}_{ia_{i1}}) & (3) & \pi_{a_{i4}}(\psi) \notin \pi_{a_{i4}}(\mathcal{F}_{ia_{i4}}) \\
(1) & \max_{m \leq a_{i1}} L(f_m) \leq a_{i2} & (4) & \max_{m \leq a_{i4}} L(g_m) \leq a_{i5} \\
(2) & 2^{a_{i2}} \leq a_{i3} & (5) & 2^{a_{i5}} \leq a_{i+10} \ .
\end{array}
\tag{25}
$$

As in the proof of the abstract diagonalization theorem 3.3 on page 77, one can show that it is possible to construct a sequence (a_{ij}) satisfying all these requirements. (Only conditions (0) and (3) require some attention.)

Let us show that φ, ψ have the desired properties. It is clear that $\varphi \leq_p f$ and $\psi \leq_p g$. Moreover, we have $\varphi, \psi \notin \mathcal{F}$ due to conditions (0) and (3). It remains to prove that φ, ψ is a minimal pair. So let us assume that $h \leq_c \varphi$ and $h \leq_c \psi$ for some $h \in \Omega^{\mathbb{N}}$. Then there exist p-bounded functions $u, v, w : \mathbb{N} \to \mathbb{N}$ satisfying

$$L^{\varphi_{u(n)}}(h_n) \leq w(n) \ , \quad L^{\psi_{v(n)}}(h_n) \leq w(n) \ . \tag{26}$$

It suffices to verify that $L(h_n) \leq nw(n)$ for sufficiently large n.

We are going to distinguish two cases. Suppose first that $a_{i2} \leq n < a_{i5}$. We may assume that $a_{j0} \leq u(n) < a_{j1}$ for some j, since otherwise $\varphi_{u(n)} = 0$ and we are done. Thus $\varphi_{u(n)} = f_{u(n)}$. For sufficiently large n we have by condition (5) that $u(n) \leq 2^n < 2^{a_{i5}} \leq a_{i+10}$. This implies that $j \leq i$, hence $u(n) < a_{i1}$. Therefore, using condition (1), we have $L(f_{u(n)}) \leq a_{i2} \leq n$. We conclude that indeed

$$L(h_n) \leq L^{\varphi_{u(n)}}(h_n) L(\varphi_{u(n)}) \leq nw(n) \ . \tag{27}$$

The discussion of the other case where $a_{i5} \leq n < a_{i+12}$ is similar and left to the reader. $\qquad \square$

By applying the theorem to $\mathcal{F} = \mathrm{VP}$ and choosing $f = g$ to be VNP-complete we obtain the following corollary. (Note that VP is a σ-limit set by Corollary 5.4 on the page before.)

Corollary 5.9 *There exists a minimal pair φ, ψ for* VP *in* VNP, *provided* VP $\neq$ VNP.

6   A Specific Family neither Complete nor p-Computable

For $1 \leq i < j \leq n$ let X_{ij} be distinct indeterminates and set $X_{ji} := X_{ij}$. Moreover, let q be a power of the prime p. The *cut enumerator* Cut_n^q is the following multivariate polynomial over the finite field $\mathbb{F}_q$

$$\mathrm{Cut}_n^q := \sum_S \prod_{i \in A, j \in B} X_{ij}^{q-1} \ , \tag{28}$$

where the sum is over all cuts $S = \{A, B\}$ of the complete graph K_n on the set of nodes $\underline{n} := \{1, 2, \ldots, n\}$. (A cut of a graph is a partition of its set of nodes into two nonempty subsets.) It is easy to see that $\mathrm{Cut}^q := (\mathrm{Cut}_n^q)$ is a p-definable family.

To motivate this definition, consider a complete graph $K_n = (\underline{n}, E_n)$ endowed with a weight function $w\colon E_n \to \mathbb{N}$. We define the weight $w(S)$ of a cut $S = \{A, B\}$ as the sum of the weights of all edges separated by S. Let $c(s)$ denote the number of cuts of weight s. (Notice that the w_{ij} are interpreted here as additive weights, whereas the X_{ij} above are viewed as multiplicative weights.) Under the substitution $X_{ij} \mapsto x_{ij} := T^{w_{ij}}$, T being a formal variable, the cut polynomial Cut_n^q becomes

$$\mathrm{Cut}_n^q(x) = \sum_S T^{(q-1)w(S)} = \sum_s (c(s) \bmod p)\, T^{(q-1)s} \ , \tag{29}$$

which can be interpreted as the generating function of the sequence $(c(s) \bmod p)_s$.

The main result of this section states that Cut^q is an explicit example of a p-family, which is neither p-computable nor complete in VNP. For the definition of the complexity classes $\mathrm{Mod}_p\mathrm{NP}/\mathrm{poly}$ see below.

Theorem 6.1 *The family of cut enumerators Cut^q over a finite field $\mathbb{F}_q$ is neither p-computable nor VNP-complete with respect to c-reduction, provided $\mathrm{Mod}_p\mathrm{NP} \not\subseteq \mathrm{P}/\mathrm{poly}$. The latter condition is satisfied if the polynomial hierarchy does not collapse at the second level.*

The proof of this theorem requires two auxiliary results. The first of them states that the cut polynomial $\mathrm{Cut}_n^q(x)$ can be evaluated over $\mathbb{F}_q$ by boolean circuits of polynomial size. The reader should be aware that this does not necessarily imply that the cut polynomial can also be evaluated by arithmetic circuits of p-bounded size (i.e., the p-computability of the family Cut^q).

Lemma 6.2 *To a symmetric matrix $x \in \mathbb{F}_q^{n \times n}$ we assign the graph $G(x)$ on the set of nodes $\underline{n}$ by requiring that $\{i, j\}$ is an edge iff $x_{ij} = 0$. Then we have*

$$\mathrm{Cut}_n^q(x) = 2^{N(x)-1} - 1 \bmod p \ , \tag{30}$$

where $N(x)$ is the number of connected components of $G(x)$. In particular, the value $\mathrm{Cut}_n^q(x)$ can be computed from a symmetric $x \in \mathbb{F}_q^{n \times n}$ in polynomial time by a Turing machine.

Proof. For any nonzero $\lambda \in \mathbb{F}_q$ we have $\lambda^{q-1} = 1$ by Fermat's theorem. Therefore, a partition $\{A, B\}$ of $\underline{n}$ contributes to $\mathrm{Cut}_n^q(x)$ either zero or one. The contribution is one iff $x_{ij} \neq 0$ for all $i \in A$, $j \in B$, which is the case iff none of the nodes of A is connected with any node in B in the graph $G(x)$. This in turn means that A and B are both a union of certain connected components of the graph $G(x)$. The number of such partitions clearly equals $2^{N(x)-1} - 1$, where $N(x)$ is the number of connected components of $G(x)$. This proves the lemma. $\qquad\square$

It is now time to recall a few facts from discrete complexity theory. For a prime number p the class $\mathrm{Mod}_p\mathrm{NP}$ is defined as the set of languages $\{x \in \{0, 1\}^* \mid \varphi(x) \equiv 1 \bmod p\}$, where $\varphi\colon \{0, 1\}^* \to \mathbb{N}$ is a function in #P (cf. Cai and Hemachandra [10]). This generalizes the class parity polynomial time $\oplus P$ introduced by Papadimitriou and Zachos [22]. We remark that if $\varphi\colon \{0, 1\}^* \to \mathbb{N}$ is #P-complete with respect to parsimonious reductions, then the corresponding language $\{x \mid \varphi(x) \equiv 1 \bmod p\}$ is $\mathrm{Mod}_p\mathrm{NP}$-complete (with respect to polynomial many-one reductions). We denote by $\mathcal{C}/\mathrm{poly}$ the nonuniform version of the complexity class $\mathcal{C}$, cf. Karp and Lipton [17].

The counting problem #Cut is the following: given a complete graph K_n with a weight function $w\colon E_n \to \mathbb{N}$ and $s \in \mathbb{N}$, what is the number of cuts of weight s? Hereby, we assume the edge weights

to be encoded in unary. The related decision problem Mod_pCUT just asks for the residue class modulo p of the number of cuts of weight s. This problem is clearly in the class Mod_pNP.

It is well known that the computation of a cut of maximal weight of a given graph is NP-hard. By a straightforward modification of the proof of this fact given in Papadimitriou [21, p. 191], one can strengthen this as follows. We will provide the detailed proof of this claim at the end of this section.

Lemma 6.3 *#CUT is #P-complete with respect to parsimonious reductions. Thus Mod_pCUT is Mod_pNP-complete.*

Proof. (of Thm. 6.1 on the preceding page) Let L be a language in Mod_pNP, say $L = \{x \in \{0,1\}^* \mid \varphi(x) \equiv 1 \bmod p\}$, where $\varphi \colon \{0,1\}^* \to \mathbb{N}$ is in the class #P. In [8] it is shown that there exists a p-definable family (f_n) over $\mathbb{F}_p$ such that $f_n \in \mathbb{F}_p[X_1, \ldots, X_n]$ and

$$\forall n \; \forall x \in \{0,1\}^n : f_n(x) = \varphi(x) \bmod p \; . \tag{31}$$

Assume now that Cut^q is VNP-complete over $\mathbb{F}_q$ with respect to c-reduction. Then we have $(f_n) \leq_c \text{Cut}^q$, hence there is a p-bounded function $t \colon \mathbb{N} \to \mathbb{N}$ such that $L^{\text{Cut}^q_{t(n)}}(f_n)$ is p-bounded. Lemma 6.2 on the page before tells us that $\text{Cut}^q_{t(n)}$ can be evaluated over $\mathbb{F}_q$ by boolean circuits of p-bounded size in n. Hence, by simulating straight-line programs by boolean circuits, we can design for each n a boolean circuit C_n of p-bounded size in n, which computes $f_n(x)$ from $x \in \mathbb{F}_q^n$. This implies that the language L is contained in P/poly. We therefore arrive at the conclusion $\text{Mod}_p\text{NP} \subseteq \text{P/poly}$.

For fixed $m, n \geq 1$ consider a field extension $K = \mathbb{F}_q(\xi)$ of $\mathbb{F}_q$ of degree $(q-1)m\binom{n}{2}$. To an instance $w \colon E_n \to \mathbb{N}$ of #CUT satisfying $\max w \leq m$ we assign the symmetric matrix $x \in K^{n \times n}$ defined by $x_{ij} := \xi^{w_{ij}}$. Then we have by (29 on the preceding page)

$$\text{Cut}^q_n(x) = \sum_s (c(s) \bmod p)\xi^{(q-1)s} \; , \tag{32}$$

where $c(s)$ is the number of cuts in K_n of weight s. The coefficients $c(s) \bmod p$ are uniquely determined by $\text{Cut}^q_n(x)$ since the above summation is over $s < m\binom{n}{2}$.

Assume now that Cut^q is p-computable over $\mathbb{F}_q$. Hence for each n there is a straight-line program Γ_n of p-bounded size in n, which computes $\text{Cut}^q_n(X)$ from constants in $\mathbb{F}_q$ and the indeterminates X_{ij} in the polynomial ring $\mathbb{F}_q[X_{ij} \mid 1 \leq i, j \leq n]$. By the universal property of the polynomial ring, Γ_n will compute $\text{Cut}^q_n(x)$ in the $\mathbb{F}_q$-algebra K from the same constants and $x \in K^{n \times n}$. We may simulate this computation by a boolean circuit of p-bounded size, since the arithmetic operations in K can be simulated by p-bounded circuits. Here it is important to note that the degree of the field extension $K/\mathbb{F}_q$ is p-bounded, as m is assumed to be encoded in unary (see the definition of #CUT). In this way, we could solve the Mod_pCUT problem in nonuniform polynomial time. As Mod_pCUT is Mod_pNP-complete by Lemma 6.3, this would imply that $\text{Mod}_p\text{NP} \subseteq \text{P/poly}$.

It remains to show that $\text{Mod}_p\text{NP} \subseteq \text{P/poly}$ implies the collapse of the polynomial hierarchy at the second level. By a well-known result of Karp and Lipton [17] this collapse would be a consequence of the inclusion $\text{NP/poly} \subseteq \text{P/poly}$. Therefore, it is sufficient to prove that

$$\text{NP/poly} \subseteq \text{Mod}_p\text{NP/poly} \; . \tag{33}$$

This follows from a well known randomized reduction due to Valiant and Vazirani [29]. (For details see [8].) □

We remark that one can prove the absolute statement that Cut^q is not VNP-complete with respect to p-projection. It would be interesting to find out whether Cut^2, interpreted as family over the rationals, is VNP-complete.

We supply now the proof of Lemma 6.3 on the preceding page. Consider the auxiliary counting problem #NAESAT which is defined as follows: given a set of boolean variables and a set of clauses each consisting of exactly three literals, compute N, where $2N$ equals the number of truth assignments of the variables such that in none of the clauses all three literals have the same truth value. (Note that the latter number must always be even!)

Lemma 6.4 *There is a parsimonious reduction from #SAT to #NAESAT.*

Proof.   The reduction from CIRCUIT SAT to NAESAT given in Example 8.3 (p. 163) and Thm. 9.3 (p. 187) of Papadimitriou [21] is easily checked to be parsimonious. On the other hand, SAT can be parsimoniously reduced to CIRCUIT SAT in an obvious way.   □

To prove Lemma 6.3 on the page before it suffices now to show the next lemma.

Lemma 6.5 *There is a parsimonious reduction from #NAESAT to #CUT.*

Proof.   We slightly modify the reduction from NAESAT to MAX CUT from Papadimitriou [21, Thm. 9.5, p. 191] in order to make it parsimonious.

Let be given a set of variables $x_1, \ldots, x_n$ and a set of clauses $C_1, \ldots, C_m$ each consisting of exactly three literals. We may assume that in no clause all literals are equal since otherwise the formula is not satisfiable in the sense of NAESAT. Moreover, we may remove the clauses which contain a variable and its negation since these are always satisfiable in the sense of NAESAT. Let m_3 denote the number of clauses with three different literals and m_2 be the number of clauses in which two literals coincide. We have $m = m_2 + m_3$.

Let G be the complete graph having as nodes the variables x_i and its negations $\neg x_i$. We define the weight function of G as follows. The *horizontal edges* $\{x_i, \neg x_i\}$ have the weight $m + 1$. The remaining edges edges $e = \{u, v\}$ (the *nonhorizontal* ones) have as weight the number of clauses C_j in which both of the literals u and v appear. If we express this event by $e \subseteq C_j$, we may write for such nonhorizontal edges e

$$w(e) = |\{C_j \mid e \subseteq C_j\}| \ . \tag{34}$$

Finally, we put $s := (m+1)n + m_2 + 2m_3$. Note that the weight of each edge is at most $m + 1$. Thus we may encode the edge weights in unary.

Let S be a cut of G and denote by $\mathcal{E}_h$ the set of horizontal edges separated by S, and by $\mathcal{E}$ the set of nonhorizontal edges separated by S. We have

$$
\begin{aligned}
w(S) &= \sum_{e \in \mathcal{E}_h \cup \mathcal{E}} w(e) = (m+1)|\mathcal{E}_h| + \sum_{e \in \mathcal{E}} |\{(e, C_j) \mid e \subseteq C_j\}| \\
&= (m+1)|\mathcal{E}_h| + \sum_{j=1}^{m} |\{(e, C_j) \mid e \in \mathcal{E}, e \subseteq C_j\}| \\
&\leq (m+1)n + (m_2 + 2m_3) = s \ .
\end{aligned}
$$

Equality holds if and only if S separates all x_i from $\neg x_i$ and if S separates the literals of any clause. This is exactly the case if S defines a truth assignment in the sense of NAESAT. (The cut S separates the true literals from the false ones.) This proves that the number of satisfying truth assignments is exactly twice the number of cuts of weight s in G.   □

7 Relativized Complexity Classes

Our investigations here are inspired by the well-known results of Baker, Gill, and Solovay [1] on relativations of the classical P-NP question.

Relative versions of the complexity classes VP and VNP can be defined as follows.

Definition 7.1 Let h be a p-family. VP^h consists of all p-families f such that $f \leq_c h$. VNP^h is the set of all p-families $f = (f_n)$ which can be obtained from some $g = (g_n) \in VP^h$ in the sense of (1 on page 75). We call the families in VP^h and VNP^h p-computable and p-definable relative to h, respectively.

Note that VP^h and VNP^h specialize to VP and VNP, respectively, if h is p-computable. We remark that VP^h is closed under c-reduction and VNP^h is closed under p-projection.

Our first goal is to establish the existence of complete families for the complexity classes VP^h and VNP^h. In particular, this gives a proof for the existence of VNP-complete families, which is independent of Valiant's intricate reduction for the permanent. The idea is to use a generalization of the concept of generic computations (cf. [9, Chap. 9]). In order to avoid an exponential growth of degrees, we combine this with an auxiliary result on the computation of homogeneous components (Prop. 7.2), which works by evaluation and interpolation, and requires that k contains sufficiently many points. In the sequel, we will therefore assume that k is an infinite field.

It is useful to introduce the following auxiliary notion. Let $h, f_1, \ldots, f_t$ be polynomials over the field k. We define the *h-complexity* $L^h(f_1, \ldots, f_t)$ as the minimum number of multiplications and evaluations of h that are sufficient to compute all f_i from the indeterminates and constants in k (we do not allow divisions). Note that for $h = X_1 X_2$ this specializes to the multiplicative (or nonscalar) complexity. We further remark that if h is a projection of h', then we have $L^{h'} \leq L^h$ as well as $\mathcal{L}^{h'} \leq \mathcal{L}^h$.

The h-complexity may be characterized in a way similar to the multiplicative complexity. Let us define an *h-computation sequence* of length r on $X_1, \ldots, X_n$ as a sequence of polynomials $g_{-n}, g_{-n+1}, \ldots, g_r$ such that $g_{-n} = 1, g_{-n+1} = X_1, \ldots, g_0 = X_n$, and such that we have

$$g_\rho = h\left(\sum_{j=-n}^{\rho-1} \alpha_{\rho 1 j} g_j, \ldots, \sum_{j=-n}^{\rho-1} \alpha_{\rho s j} g_j\right) \tag{35}$$

for all $1 \leq \rho \leq r$ and some $\alpha_{\rho \sigma j}$ in k. We say that such a sequence *computes* $f_1, \ldots, f_t$ iff all f_i are contained in the k-linear hull of $g_{-n}, \ldots, g_r$.

In what follows, we will assume that $X_1 X_2$ is a projection of h, in which case we say that h *contains the multiplication*. Then it is not hard to see that the h-complexity r of $f_1, \ldots, f_t$ equals the minimum length of an h-computation sequence which computes all f_i. Moreover, the complexity L^h and the h-complexity r are polynomially related as follows: we have $r \leq L^h(f_1, \ldots, f_t) \leq 2s(n+1)(r+1) + sr^2$, when s is the number of variables of h.

Proposition 7.2 *Let f be a polynomial in $a_1, \ldots, a_m$ and $X_1, \ldots, X_n$ having degree at most $d \geq 1$ in the X-variables. We denote by $f^{(\delta)}$ the homogeneous part of f of degree δ with respect to the X-variables. Then we have*

$$\mathcal{L}^h(\{f^{(\delta)} \mid \delta \leq d\}) \leq (1 + d \deg h)\, \mathcal{L}^h(f) \;. \tag{36}$$

Proof. We will use the abbreviation $f^{\leq d} := \sum_{\delta \leq d} f^{(\delta)}$ and write $D := \deg h$. Let $(g_\rho)_{\rho \geq -n}$ be an h-computation sequence of length $r := \mathcal{L}^h(f)$ as in (35) which computes f. We define a related sequence $(u_\rho)_{\rho \geq -n}$ by setting $u_\rho := g_\rho$ for $-n \leq \rho \leq 0$, and for $\rho > 0$

$$u_\rho = h(v_{\rho 1}, \ldots, v_{\rho s})^{\leq d} \;, \tag{37}$$

where $v_{\rho\sigma} = \sum_{j=-n}^{\rho-1} \alpha_{\rho\sigma j} u_j$. It is easy to check that $u_\rho = g_\rho^{\leq d}$ for all ρ.

The homogeneous parts of f (w.r.t. X) are a k-linear combination of the homogeneous parts of the g_ρ. Therefore, it suffices to prove that all homogeneous parts of u_ρ up to degree d can be computed from the homogeneous parts of $u_{-n}, \ldots, u_{\rho-1}$ up to degree d by some k-linear operations and $1 + dD$ evaluations of h.

The polynomial $w_\rho := h(v_{\rho 1}, \ldots, v_{\rho s})$ has degree at most dD. By definition, $u_\rho^{(\delta)} = w_\rho^{(\delta)}$ for $\delta \leq d$. We have for $\lambda \in k$ that

$$\sum_{\delta \leq dD} \lambda^\delta w_\rho^{(\delta)} = w_\rho(\lambda X) = h\left(\sum_{\delta \leq d} \lambda^\delta v_{\rho 1}^{(\delta)}, \ldots, \sum_{\delta \leq d} \lambda^\delta v_{\rho s}^{(\delta)}\right) . \tag{38}$$

Hence we can compute $w_\rho(\lambda X)$ from the $v_{\rho\sigma}^{(\delta)}$ and thus from the $u_j^{(\delta)}$ for $j < \rho, \delta \leq d$ by k-linear operations and just one evaluation of h. We can thus compute the homogeneous parts of w_ρ as a k-linear combination of $w_\rho(\lambda X)$ for $1 + dD$ different values of $\lambda \in k$ (interpolation). $\qquad\square$

In the sequel, we will use the abbreviations $X^\mu := X_1^{\mu_1} \cdots X_n^{\mu_n}$ and $|\mu| := \sum \mu_i$ for $\mu \in \mathbb{N}^n$. Moreover, we set $\deg 0 := -\infty$ for the zero polynomial.

Definition 7.3 Let a polynomial $h \in k[X_1, \ldots, X_s]$ of degree D be given.

(a) We define the *generic h-computation* $(G_\rho)_{\rho \geq -n}$ on $X_1, \ldots, X_n$ over k recursively as follows: $G_{-n} := 1, G_{-n+1} := X_1, \ldots, G_0 := X_n$, and for all $\rho > 0$ we set

$$\begin{aligned} G_\rho \quad := \quad & h\left(\sum_{j=-n}^{\rho-1} a_{\rho 1 j} G_j, \ldots, \sum_{j=-n}^{\rho-1} a_{\rho s j} G_j\right) \\ & + b_\rho - h\left(\sum_{j=-n}^{\rho-1} a_{\rho 1 j} G_{j0}, \ldots, \sum_{j=-n}^{\rho-1} a_{\rho s j} G_{j0}\right) . \end{aligned}$$

Here the $a_{\rho\sigma j}$ and b_ρ denote different indeterminates and G_{j0} is the constant term of G_j with respect to the X-variables. We write $G_\rho = \sum_\mu G_{\rho\mu} X^\mu$, where $G_{\rho\mu}$ depends only on the a and b-variables.

(b) The n^{th} *generic polynomial computed relative to h* is defined as

$$C_n(h) := \sum_{|\mu| \leq n} \sum_{\rho=-n}^{n} c_\rho G_{\rho\mu} X^\mu . \tag{39}$$

Here the c_ρ denote additional indeterminates. Thus $C_n(h)$ is the sum of the X-homogeneous parts up to degree n of $\sum_{\rho=-n}^{n} c_\rho G_\rho$.

(c) The n^{th} *generic polynomial defined relative to h* is

$$D_n(h) := \sum_{v=0}^{n} d_v \sum_{e \in \{0,1\}^{n-v}} C_n(h)(a, b, c, X_1, \ldots, X_v, e_{v+1}, \ldots, e_n) , \tag{40}$$

where the d_v denote additional indeterminates. (Note that the $X_{v+1}, \ldots, X_n$ are substituted by 0 or 1.)

The following technical lemma summarizes some of the properties of h-generic computations as well as of the polynomials $C_n(h)$ and $D_n(h)$. Recall that $h \leq h'$ means that h is a projection of h'.

Lemma 7.4 *We have for $\rho > 0$ and $\mu \neq 0$:*

(a) $G_{\rho 0} = b_\rho$.

(b) $G_{\rho\mu}$ depends on at most $s\rho(n+1+\frac{\rho-1}{2}) + \rho$ variables.

(c) $\mathcal{L}^h(G_1, \ldots, G_\rho) \le s\rho(n+\rho-1) + 2\rho$.

(d) $\deg G_{\rho\mu} \le 1 + 2D\rho|\mu|$.

(e) $C_n(h)$ and $D_n(h)$ are polynomials in at most $2sn^3 + 5n + 2$ variables and have degree at most $2Dn^2 + n + 3$.

(f) $\mathcal{L}^h(C_n(h)) \le (1 + Dn)(2sn^2 + 4n)$.

(g) If we abbreviate the a, b, c-variables occuring in $C_n(h)$ by $Z_1, \ldots, Z_w$, we have

$$\{f \in k[X_1, \ldots, X_n] \mid \deg f \le n, \ \mathcal{L}^h(f) \le n\} \subseteq \{C_n(h)(z, X) \mid z \in k^w\} \ . \tag{41}$$

(h) $h \le C_n(h)$ if $\deg h \le n$. Moreover $C_n(h) \le D_n(h)$.

(i) For all $n \le n'$ and $h \le h'$ we have $C_n(h) \le C_{n'}(h')$ and $D_n(h) \le D_{n'}(h')$.

Proof. Claims (a), (b), and (c) follow by straightforward calculations.

We will prove Claim (d) by induction on ρ. We remark first that $\deg G_{\rho\mu} \le 0$ for $\rho \le 0$ and all μ. Let now $\rho > 0$ be fixed and put

$$V_\sigma := \sum_\mu V_{\sigma\mu} X^\mu := \sum_{j=-n}^{\rho-1} a_{\rho\sigma j} G_j \ . \tag{42}$$

By the induction hypothesis we have for all $\mu \ne 0$

$$\deg V_{\sigma\mu} \le 2 + 2D(\rho-1)|\mu| \ . \tag{43}$$

This estimate is also true for $\mu = 0$, since $\deg G_{j0} \le 1$. What we have to do is to prove the estimate

$$\forall \mu \ne 0 : \deg H_\mu \le 1 + 2D\rho|\mu| \tag{44}$$

for the polynomial $H := \sum_\mu H_\mu X^\mu := h(V_1, \ldots, V_s)$. Clearly, it is sufficient to verify this for the power products $h = X_1^{e_1} \cdots X_s^{e_s}$ of degree at most D. In this case, we have $H = V_1^{e_1} \cdots V_s^{e_s}$, and we obtain

$$H_\nu = \sum \prod_{\sigma=1}^{s} \prod_{\varepsilon=1}^{e_\sigma} V_{\sigma\mu_\sigma(\varepsilon)} \ , \tag{45}$$

where the sum runs over all systems of maps $\mu_\sigma : \{1, 2, \ldots, e_\sigma\} \to \mathbb{N}^n$, $1 \le \sigma \le s$ satisfying $\sum_\sigma \sum_\varepsilon \mu_\sigma(\varepsilon) = \nu$. Using (43) we conclude that for $\nu \ne 0$

$$\begin{aligned}
\deg H_\nu &\le \quad \sum_\sigma \sum_\varepsilon \left(2 + 2D(\rho-1)|\mu_\sigma(\varepsilon)|\right) \\
&\le \quad 2\sum_\sigma e_\sigma + 2D(\rho-1)\sum_\sigma \sum_\varepsilon |\mu_\sigma(\varepsilon)| \\
&\le \quad 2D + 2D(\rho-1)|\nu| \\
&\le \quad 1 + 2D\rho|\nu| \ ,
\end{aligned}$$

which proves Claim (d).

Claim (e) follows immediately from the Claims (b) and (d), whereas Claim (f) is a consequence of Claim (c) and Prop. 7.2 on page 87.

To show Claim (g) assume $f \in k[X_1, \ldots, X_n]$ such that $\mathcal{L}^h(f) \le n$. There is an h-computation sequence (g_ρ) of length n which computes f, say

$$g_\rho = h(\textstyle\sum_{j<\rho} \alpha_{\rho 1 j} g_j, \ldots, \sum_{j<\rho} \alpha_{\rho s j} g_j) \ , \tag{46}$$

and $f = \sum_{\rho=-n}^{n} \gamma_\rho g_\rho$, with $\alpha_{\rho\sigma j}, \gamma_\rho \in k$. Thus the substitution $a_{\rho\sigma j} \mapsto \alpha_{\rho\sigma j}$, $b_\rho \mapsto g_\rho(X = 0)$ sends G_ρ to g_ρ. If we additionally substitute the c_ρ by the γ_ρ, then the polynomial $C_n(h)$ is mapped to f, provided that $\deg f \le n$.

(h) $C_n(h) \le D_n(h)$ is obtained by substituting $d_n \mapsto 1$ and $d_v \mapsto 0$ if $v < n$. To show that $h \le C_n(h)$ consider the substitution φ which maps X_1 and c_1 to 1, sends the $a_{1,\sigma,-n+1}$ to X_σ for $1 \le \sigma \le s$, maps b_1 to $h(0, \ldots, 0)$, and sends all the remaining a-variables and c-variables to 0. All other variables shall remain invariant under φ. G_1 is mapped to $h = h(X_1, \ldots, X_s)$ under φ. We have $\deg_X G_1 \le n$ as $\deg h \le n$. From this it easily follows that $C_n(h)$ is mapped to h under φ.

Before proving Claim (i) it is useful to make the following general observation. Let $A := k[a_1, \ldots, a_m, X_1, \ldots, X_n]$ and consider a substitution (k-algebra morphism) $\varphi: A \to A$ which fixes the X-variables and such that $\varphi(a_i) \in k \cup \{a_1, \ldots, a_m\}$. Let $f^{(\delta)}$ denote the homogeneous part of $f \in A$ with respect to the X-variables. Then we have $\varphi(f^{(\delta)}) = \varphi(f)^{(\delta)}$. If $\sigma: A \to A$ is another substitution which leaves the a-variables invariant and such that $\sigma(X_i) \in k \cup \{X_1, \ldots, X_n\}$, then we have $\varphi(\sigma(f)) = \sigma(\varphi(f))$ for all $f \in A$.

(i) We show first that $C_n(h) \le C_{n'}(h)$ for $n \le n'$. Let (G'_ρ) denote the generic h'-computation on $X_1, \ldots, X_{n'}$, and (G_ρ) be the generic h-computation on $X_1, \ldots, X_n$. The substitution φ which maps $a_{\rho\sigma j}$ to 0 for all $-n' + n < j \le 0$ and which leaves all other a-variables and the b-variables invariant sends G'_ρ to G_ρ (up to a renaming of the variables). This can be proven by induction on $\rho > 0$. Note that $G'_\rho \mapsto G_\rho$ implies $G'_{\rho 0} \mapsto G_{\rho 0}$ by our general observation, as φ fixes the X-variables. By additionally requiring $c_\rho \mapsto 0$ for either $-n' + n < \rho \le 0$ or $\rho > n$ we get

$$\varphi(\textstyle\sum_{\rho=-n'}^{n'} c_\rho G'_\rho) = \sum_{\rho=-n}^{n} c_\rho G_\rho \ . \tag{47}$$

Since φ leaves the X-variables invariant, it commutes with taking homogeneous parts with respect to the X-variables. Thus $\varphi(C_{n'}(h')) = C_n(h)$.

A slight modification of the reasoning before yields a substitution φ which leaves the X-variables invariant and such that

$$\varphi\big(C_{n'}(h')(Z', \overline{X}_1 \ldots, \overline{X}_{n'-n}, X_1, \ldots, X_n)\big) = C_n(h)(Z, X_1, \ldots, X_n) \ , \tag{48}$$

where Z', Z stand for the corresponding a, b, c-variables. This implies by our general observation

$$\varphi\big(C_{n'}(h')(Z', \overline{X}_1 \ldots, \overline{X}_{n'-n}, X_1, \ldots, X_u, e_{u+1}, \ldots, e_n)\big)$$
$$= \quad C_n(h)(Z, X_1, \ldots, X_u, e_{u+1}, \ldots, e_n)$$

for $0 \le u \le n$ and $e_i \in \{0, 1\}$. If we extend φ by sending $d_{n'-n+u}$ to d_u for $0 \le u \le n$ and mapping the remaining d's to zero, we get $\varphi(D_{n'}(h')) = D_n(h)$. Hence $D_n(h) \le D_{n'}(h')$.

Assume now $h \le h'$ and let (G_ρ) and (G'_ρ) be the corresponding generic computations on $X_1, \ldots, X_n$. It is not hard to see that there exists a substitution which only changes the a-variables and that maps all G'_ρ to G_ρ. From this one concludes as above that $C_n(h) \le C_n(h')$ and $D_n(h) \le D_n(h')$. $\qquad\square$

We will interpret families of polynomials $(f_{m,n})$ with double indices as sequences of polynomials by enumerating pairs $(m,n) \in \mathbb{N}^2$ according to $(m,n) \mapsto m + (m+n)(m+n+1)/2$.

We can now state the first result of this section.

Theorem 7.5 *Let $h = (h_n)$ be a p-family such that any h_n contains the multiplication. Then the double-indexed families $(C_n(h_m))$ and $(D_n(h_m))$ are VP^h-complete, resp. VNP^h-complete with regard to p-projection.*

Proof. By Lemma 7.4 on page 88(e) both families $(C_n(h_m))$ and $(D_n(h_m))$ are p-families. Part (f) of that lemma implies that $(C_n(h_m))$ is p-computable relative to h.

Assume $(f_j) \in \text{VP}^h$, where f_j is a polynomial in $u(j)$ variables. By definition, there exists a p-bounded function $m: \mathbb{N} \to \mathbb{N}$ such that $\mathcal{L}^{h_{m(j)}}(f_j)$ is p-bounded in j. Let $n(j)$ denote the maximum of $u(j)$, $\deg f_j$, and $\mathcal{L}^{h_{m(j)}}(f_j)$. It is clear that $n(j)$ is p-bounded in j. We denote the a,b,c-variables in $C_{n(j)}(h_{m(j)})$ by $Z_1, \dots Z_{w(j)}$. From Lemma 7.4 on page 88(g) it follows that $f_j = C_{n(j)}(h_{m(j)})(z, X)$ for a suitable choice of $z \in k^{w(j)}$. Thus f_j is a projection of $C_{n(j)}(h_{m(j)})$, and we have proved the VP^h-completeness of $(C_n(h_m))$.

Let now (q_j) be a p-definable family relative to h, say

$$q_j(X_1, \dots, X_{v(j)}) = \sum_{e \in \{0,1\}^{u(j)-v(j)}} f_j(X_1, \dots, X_{v(j)}, e_{v(j)+1}, \dots, e_{u(j)}) \quad, \tag{49}$$

where the family (f_j) is p-computable relative to h. From before we know that for each j

$$f_j(X) = C_{n(j)}(h_{m(j)})(z, X) \tag{50}$$

for p-bounded $m(j), n(j)$ and some $z \in k^{w(j)}$. Therefore, we see that q_j can be obtained from $D_{n(j)}(h_{m(j)})$ by the substitution $d_{v(j)} \mapsto 1$, $d_v \mapsto 0$ for $d \neq v(j)$, and $Z_i \mapsto z_i$ for all i. This shows that (q_j) is a p-projection of $(D_n(h_m))$.

It remains to prove that $(D_n(h_m))$ is p-definable relative to h. Let us abbreviate the a,b,c-variables occuring in $C_n(h_m)$ by $Z_1, \dots, Z_w$ and define

$$g_{m,n} := \sum_{v=0}^{n} d_v E_1 \cdots E_v C_n(h_m)(Z, X_1, \dots, X_v, E_{v+1}, \dots, E_n) \quad, \tag{51}$$

where $E_1, \dots, E_n$ are new indeterminates. As $(C_n(h_m))$ is p-computable relative to h, so is $(g_{m,n})$. The following equality

$$D_n(h_m) = \sum_{e \in \{0,1\}^n} g_{m,n}(Z, X_1, \dots, X_n, e_1, \dots, e_n) \tag{52}$$

proves that indeed $(D_n(h_m)) \in \text{VNP}^h$. □

We call a p-family (h_n) *monotone*, iff h_n is a projection of h_{n+1} for all n.

Corollary 7.6 *(a) If h is monotone and h_0 contains the multiplication, then $(C_n(h_n))$ and $(D_n(h_n))$ are VP^h-complete, resp. VNP^h-complete with respect to p-projection.*

(b) For any p-family h there exist VP^h-complete and VNP^h-complete families with respect to p-projection.

Proof. (a) This is an immediate consequence of Thm. 7.5 on the page before and the monotonicity of $(C_n(h))$ and $(D_n(h))$ expressed in Lemma 7.4 on page 88(i).

(b) Let $h = (h_m)$ be any p-family and U, V, W be new indeterminates. The polynomials $\tilde{h}_m := Uh_m + (1 - U)VW$ contain the multiplication and the p-family $\tilde{h} = (\tilde{h}_m)$ satisfies $\mathrm{VP}^{\tilde{h}} = \mathrm{VP}^h$. Now apply Thm. 7.5 on the page before. □

While there are many natural examples of VNP-complete families (generating functions of NP-complete graph problems, cf. [7]), we don't know of any interesting example of a VP-complete family. The family of determinants is a possible candidate (see [9, Problem 21.3]).

The next result is inspired by Baker, Gill, and Solovay [1]. Its proof is based on Thm. 7.5 on the preceding page combined with some diagonalization argument.

Theorem 7.7 *There exists a p-family h such that $\mathrm{VP}^h = \mathrm{VNP}^h$.*

Proof. First note the following: By Lemma 7.4 on page 88(e) the degree as well as the number of variables of $D_n(h)$ are bounded from above by $p(n) := 2n^4 + 5n + 2$, provided the number of variables and the degree of the polynomial h are bounded by n.

By induction, we are going to construct a monotone sequence of polynomials $h = (h_n)$ such that the number of variables as well as the degree of h_n are bounded from above by n for $n \geq 2$. We define $m_i := i$ and $h_i := X_1 X_2$ for $i \leq 2$. Assume we have already constructed $h_0, \ldots, h_{m_t}$ ($t \geq 2$). We set $m_{t+1} := p(m_t)$ and define $h_j := h_{m_t}$ for $m_t < j < m_{t+1}$ and put $h_{m_{t+1}} := D_{m_t}(h_{m_t})$. By Lemma 7.4 on page 88(h) we have $h_{m_t} \leq h_{m_{t+1}}$ which guarantees the monotonicity. Moreover, the degree and the number of variables of $h_{m_{t+1}}$ are bounded by $m_{t+1} = p(m_t)$.

We claim that $(D_n(h_n))$ is a p-projection of h. In fact, let $n \geq 2$ be given, say $m_t \leq n < m_{t+1}$. The monotonicity of $D_n(h)$ expressed in Lemma 7.4 on page 88(i) implies that

$$D_n(h_n) \leq D_{m_{t+1}}(h_{m_{t+1}}) = h_{m_{t+2}} \quad . \tag{53}$$

But $m_{t+2} = p(p(m_t)) \leq p(p(n))$ and the composition of p with itself is clearly p-bounded. This shows the claim.

On the other hand, we know from Cor. 7.6 on the preceding page(a) that $(D_n(h_n))$ is VNP^h-complete. As $(D_n(h_n))$ is also contained in VP^h, it follows that $\mathrm{VP}^h = \mathrm{VNP}^h$. □

Up to now we have not succeeded in establishing a p-family h such that $\mathrm{VP}^h \neq \mathrm{VNP}^h$. A promising approach for this is as follows (compare Bennett and Gill [4]). For each n choose independently $h_n \in k[X_1, \ldots, X_n]$ of degree most n *at random* according to some probability distribution. Since the classes VP^h and VNP^h are invariant under finite variation of h, the event $\mathcal{E} = \{h \mid \mathrm{VP}^h \neq \mathrm{VNP}^h\}$ is a so-called tail event. Kolmogorov's zero-one law (cf. Feller [14, Chap. 4]) implies therefore that $\mathrm{Prob}(\mathcal{E}) \in \{0, 1\}$. We conjecture that this probability is one if the h_n are chosen with independent $0, 1$-coefficients.

Acknowledgements

References

[1] T. Baker, J. Gill, and R. Solovay, Relativizations of the $P =?NP$ question, *SIAM J. Comp.*, 4, 431–442, 1975.

[2] J. L. Balcázar, J. Díaz, and J. Gabarró, *Structural Complexity I*, Springer Verlag, 1988.

[3] S. Ben-David, K. Meer, and C. Michaux, A note on non-complete problems in NP_R, Preprint, 1996.

[4] C.H. Bennett and J. Gill, Relative to a random oracle A, $P^A \neq NP^A \neq co - NP^A$ with probability 1, *SIAM J. Comp.*, 10, 96–113, 1981.

[5] L. Blum, F. Cucker, M. Shub, and S. Smale, Algebraic Settings for the Problem "$P \neq NP$?", In *The mathematics of numerical analysis*, number 32 in Lectures in Applied Mathematics, Amer. Math. Soc., 125–144, 1996.

[6] L. Blum, M. Shub, and S. Smale, On a theory of computation and complexity over the real numbers, *Bull. Amer. Math. Soc.*, 21, 1–46, 1989.

[7] P. Bürgisser, *Completeness and Reduction in Algebraic Complexity Theory*, Habilitationsschrift Universität Zürich, To appear in Springer, 1998.

[8] P. Bürgisser, Cook's versus Valiant's hypothesis, *Theoret. Comp. Sci.*, To appear.

[9] P. Bürgisser, M. Clausen, and M.A. Shokrollahi, *Algebraic Complexity Theory*, Number 315 in Grundlehren der mathematischen Wissenschaften, Springer Verlag, 1997.

[10] J. Cai and L.A. Hemachandra, On the power of parity polynomial time, In *Proc. STACS'89*, number 349 in LNCS, Springer Verlag, 229–239, 1989.

[11] O. Chapuis and P. Koiran, Saturation and Stability in the Theory of Computation over the Reals, *Annals of Pure and Applied Logic*, To appear.

[12] S.A. Cook, The complexity of theorem proving procedures, In *Proc. 3rd ACM STOC*, 151–158, 1971.

[13] T. Emerson, Relativizations of the P=?NP question over the reals (and other ordered rings), *Theoret. Comp. Sci.*, 133, 15–22, 1994.

[14] W. Feller, *An introduction to probability theory and its applications*, volume 2, John Wiley & Sons, 1971.

[15] J. von zur Gathen, Feasible arithmetic computations: Valiant's hypothesis, *J. Symb. Comp.*, 4, 137–172, 1987.

[16] J. Heintz and J. Morgenstern, On the intrinsic complexity of elimination theory, *Journal of Complexity*, 9, 471–498, 1993.

[17] R.M. Karp and R.J. Lipton, Turing machines that take advice, In *Logic and Algorithmic: An international Symposium held in honor of Ernst Specker*, Monogr. No. 30 de l'Enseign. Math., 255–273, 1982.

[18] R.E. Ladner, On the structure of polynomial time reducibility, *J. ACM*, 22, 155–171, 1975.

[19] Landweber, Lipton, and Robertson, On the structure of sets in NP and other complexity classes, *Theoret. Comp. Sci.*, 15, 181–200, 1981.

[20] G. Malajovich and K. Meer, On the structure of NP$_C$, *SIAM J. Comp.*, To appear.

[21] C.H. Papadimitriou, *Computational Complexity*, Addison-Wesley, 1994.

[22] C.H. Papadimitriou and S. Zachos, Two remarks on the power of counting, In *Proc. 6th GI conference in Theoretical Computer Science*, number 145 in LNCS, Springer Verlag, 269–276, 1983.

[23] U. Schöning, A uniform approach to obtain diagonal sets in complexity classes, *Theoret. Comp. Sci.*, 18, 95–103, 1982.

[24] U. Schöning, Minimal pairs for *P*, *Theoret. Comp. Sci.*, 31, 41–48, 1984.

[25] S. Smale, Complexity theory and numerical analysis, *Acta Numerica*, 6, 523–551, 1997.

[26] L.G. Valiant, Completeness classes in algebra, In *Proc. 11th ACM STOC*, 249–261, 1979.

[27] L.G. Valiant, The complexity of computing the permanent, *Theoret. Comp. Sci.*, 8, 189–201, 1979.

[28] L.G. Valiant, Reducibility by algebraic projections, In *Logic and Algorithmic: an International Symposium held in honor of Ernst Specker*, volume 30, Monogr. No. 30 de l'Enseign. Math., 365–380, 1982.

[29] L.G. Valiant and V.V. Vazirani, NP is as easy as detecting unique solutions, *Theoret. Comp. Sci.*, 47, 85–93, 1986.