Electronic Colloquium on Computational Complexity, Report No. 63 (2008)

Valiant-Vazirani Lemmata for Various Logics

Moritz Müller [*]
Albert-Ludwigs-Universität Freiburg

Abstract

We show analogues of a theorem due to Valiant and Vazirani [16] for intractable parameterized complexity classes such as W[P], W[SAT] and the classes of the W-hierarchy as well as those of the A-hierarchy. We do so by proving a general "logical" version of it which may be of independent interest.

1 Introduction

1.1 The Valiant-Vazirani Lemma

In classical complexity theory the famous Valiant-Vazirani Lemma states that "the problems of distinguishing between instances of SAT having zero or one solution, or finding solutions to instances of SAT having unique solutions, are as hard as SAT itself." ([16, Abstract]) Here hardness refers to randomized reductions with one-sided error. This result can be shaped as a probabilistic statement about Boolean logic:

Theorem 1 (Valiant, Vazirani 1985) *There is a polynomial time algorithm computing for any Boolean formula $\alpha(\bar{x})$ a Boolean formula $\beta(\bar{x}\bar{y})$ such that, if α is satisfiable, then*

$$\Pr_{\bar{b}\in\{0,1\}^{|\bar{y}|}}\left[\models \exists^{=1}\bar{x}\,(\alpha\wedge\beta)\frac{\bar{b}}{\bar{y}}\right]\geq\frac{1}{8|\bar{x}|}.$$

It is easy to see that we can also allow α to come from Boolean logic extended by various quantifiers. In this work we intend to show such results for others than Boolean logics.

Our main result states that given a structure $\mathcal{A}$ and a formula ϕ of least fixed-point logic LFP we can do the following in polynomial time:

First we *enlarge* $\mathcal{A}$ by increasing its universe and declare some additional (mainly arithmetical) relations on it; second we compute a formula ψ of roughly the same logical complexity as ϕ such that if you randomly assign values to some distinguished variables of ψ, then with "good" probability ψ in the new

[*] Email: moritz.mueller@math.uni-freiburg.de

ISSN 1433-8092

structure singles out exactly one of the solutions of ϕ in $\mathcal{A}$ provided there are any; furthermore we have one-sided error in the sense that ψ has no solution in case ϕ has none.

Theorem 2 *Let τ be a vocabulary. There is a relational vocabulary τ^* and a polynomial time algorithm which, given a τ-structure $\mathcal{A}$ and a formula $\phi = \phi(\bar{x}) \in \mathrm{LFP}[\tau]$, computes a τ^*-enlargement $\mathcal{A}^*$ of $\mathcal{A}$ and a* quantifier free *formula $\rho = \rho(\bar{x}\bar{y}\bar{z}) \in \mathrm{FO}[\tau^*]$ with parameters in A^* such that, if $\phi(\mathcal{A}) \neq \emptyset$, then*

$$\Pr_{\bar{b}\in(A^*)^{|\bar{y}|}}\left[\mathcal{A}^* \models \exists^{=1}\bar{x}\bar{z}\left(\phi^{\dot{A}}(\bar{x}) \wedge \rho(\bar{x}\bar{y}\bar{z})\right)\frac{\bar{b}}{\bar{y}}\right] \geq \frac{1}{|A^*|^2}.$$

(Here $\phi^{\dot{A}}$ is the relativization of ϕ to the unary predicate $\dot{A} \in \tau^* \setminus \tau$ which is interpreted in $\mathcal{A}^*$ by the universe A of $\mathcal{A}$.)

This allows us to move in polynomial time to a "probably unique" formula without essentially increasing the logical complexity of the input formula.

The logic LFP does not play an essential role here. We state the result for LFP just because this suffices for the applications we have in mind. Our proof applies to any logic in the sense of [8] tractably closed under relativizations and under conjunction with quantifier free first order formulas.

1.2 Applications to complexity theory

Theorem 1 has proven useful in many respects. For example it is a main step in the original proof of Todas theorem [15] stating that $\mathrm{PH} \subseteq \mathrm{BP}\oplus\mathrm{P} \subseteq \mathrm{P}^{\#\mathrm{P}[1]}$.

An immediate corollary of Theorem 1 concerns the problem UNIQUE-SAT: does a given Boolean formula have exactly one satisfying assignment?

Theorem 3 (Valiant, Vazirani 1985) UNIQUE-SAT *is* NP*-hard under randomized polynomial time reductions with one-sided error.*[1]

Downey, Fellows and Regan [7] ask for analogues of these results in the parameterized setting. The first difficulty is how to state such analogues.

In a parameterized problem instances x come along with a parameter $\kappa(x) \in \mathbb{N}$ which is expected to be small in typical applications. Intuitively the parameter encodes some knowledge we have about the inputs which we want to exploit algorithmically. To allow for full exploitation the notion of tractability is adjusted accordingly: parameterized problems decidable in in time $f(\kappa(x)) \cdot |x|^{O(1)}$ for an arbitrary computable $f : \mathbb{N} \to \mathbb{N}$ are *fixed-parameter tractable.* Asking for fpt algorithms means trying to confine the exponential running time needed to solve many natural problems to some 'small' parameter.

There is no class of intractable parameterized problems playing a role as predominant as NP in the classical setting. Instead we face several hierarchies of intractable classes, most prominently the W-hierarchy. Other important

[1] In fact UNIQUE-SAT is complete for D^{P} under these reductions [16, Corollary 5]. See [2] for a discussion.

classes are W[SAT] and W[P] on the top of the W-hierarchy and those of the A-hierarchy.

$$\begin{array}{ccccccccccc}
 & & \mathrm{A}[1] & \subseteq & \mathrm{A}[2] & \subseteq & \ldots & & & & \\
 & & \| & & \rotatebox{-90}{$\supseteq$} & & & & & & \\
\mathrm{FPT} & \subseteq & \mathrm{W}[1] & \subseteq & \mathrm{W}[2] & \subseteq & \ldots & \subseteq & \mathrm{W}[\mathrm{SAT}] & \subseteq & \mathrm{W}[\mathrm{P}]
\end{array}$$

We may ask for parameterized analogues of the Valiant-Vazirani Lemma for each of these intractable classes. Furthermore, the notion of parameterized randomized reduction can be given various interesting renderings [7, 13, 14]. E.g. we can restrict the random complexity available on input x to $f(\kappa(x)) \cdot \log |x|$ many random bits. We call such a reduction W[P]-*randomized.* If we impose no bound on the random complexity we speak of paraNP-*randomization.*

Now, what is known? Downey et al. [7] showed

Theorem 4 (Downey, Fellows, Regan 1998) *Let* $t \geq 1$. *There are* paraNP-*randomized reductions with one-sided error from* W[t] *to* UniqueW[t].

The classes UniqueW[t] for $t \geq 1$ have been introduced in [7]. The reductions there use $\kappa(x) \cdot |x| \cdot \log |x|$ random bits. With a view to the goal to find some parameterized analogue of Todas theorem Downey et al. asked how to derandomize this to, in the present terminology, W[P]-randomized reductions. We answer this question here. We give reductions using less than $\kappa(x)^3 \cdot \log |x|$ random bits:

Theorem 5 *Let* $t \geq 1$. *There are* W[P]*-randomized reductions with one-sided error*

1. *from* A[t] *to* UniqueA[t]*,*
2. *from* W[t] *to* UniqueW[t]*,*
3. *from* W[P] *to* UniqueW[P] *and*
4. *from* W[SAT] *to* UniqueW[SAT].

In [7] it is asked for the complexity of the parameterized problem p-UNIQUE-CLIQUE: given a graph and a parameter k, does the graph contain exactly one clique of size k? As a corollary we get

Theorem 6 p-UNIQUE-CLIQUE *is* W[1]*-hard under* W[P]*-randomized reductions with one-sided error.*

2 Logical preliminaries

First-order logic A vocabulary τ is a *finite* set of relation symbols, function symbols and constants. Relation or function symbols have an associated *arity.* τ*-atoms* are of the form $t_1 = t_2$ or $Rt_1 \cdots t_r$ for a relation symbol $R \in \tau$ of arity

r and τ-terms $t_1,\dots,t_r$. *τ-terms* a build from *individual variables* $x_1, x_2, \dots$, constants and function symbols. (First order) *τ-formulas* $\mathrm{FO}[\tau]$ are build from atoms using Boolean connectives and existential and universal quantification.

A τ-structure $\mathcal{A}$ consists in an *universe* A, i.e. a *finite* non-empty set, and for each relation symbol R (function symbol f) from τ of arity r a relation $R^{\mathcal{A}} \subseteq A^r$ (function $f^{\mathcal{A}} : A^r \to A$), and for any constant $c \in \tau$ an element $c^{\mathcal{A}} \in A$. A vocabulary is *relational* if and only if it contains no function symbols (but may contain constants). Accordingly we speak of relational formulas and relational structures.

Let $t, u \in \mathbb{N}$. Π_t is the class of *relational* first order formulas of the form $\forall \bar{y}_1 \exists \bar{y}_2 \cdots Q\bar{y}_t\phi$, where ϕ is quantifier-free and Q is $\exists$ for even t and $\forall$ for odd t. If additionally we have $|\bar{y}_i| \leq u$ for all $i \in [t]$ we say that the formula belongs to the class $\Pi_{t,u}$.

Parameters and solutions A *τ-formula with parameters in $\mathcal{A}$* is one containing besides symbols from τ also *parameters in $\mathcal{A}$*, i.e.constants a with $a = a^{\mathcal{A}} \in A$. Such a formula can be interpreted only in structures containing its parameters. We write $\phi\frac{\bar{a}}{\bar{x}}$ for the formula obtained from ϕ by substituting in ϕ the parameters $\bar{a}$ in $\mathcal{A}$ for the free occurences of the variables $\bar{x}$.

Definition 7 For $\phi = \phi(\bar{x})$ the set $\phi(\mathcal{A})$ of *solutions* of ϕ in $\mathcal{A}$ is the set of all $\bar{a} \in A^{|\bar{x}|}$ such that $\mathcal{A} \models \phi\frac{\bar{a}}{\bar{x}}$. For a subtupel $\bar{x}_1$ of $\bar{x}$, say for simplicity $\bar{x} = \bar{x}_1\bar{x}_2$, call $(\bar{x}_1, \bar{a}_1)$ a *partial solution* of $(\mathcal{A}, \phi)$ if and only if $\bar{a}_1\bar{a}_2 \in \phi(\mathcal{A})$ for some $\bar{a}_2$; in this case $\bar{a}_1\bar{a}_2$ is a *satisfying extension* of $(\bar{x}_1, \bar{a}_1)$.

Definition 8 A *τ^*-enlargement* of a τ-structure $\mathcal{A}$ is a $(\tau \cup \tau^*)$-structure $\mathcal{A}^*$ which is an expansion of an extension[2] of $\mathcal{A}$ such that there is some unary relation symbol $\dot{A} \in \tau^* \setminus \tau$ with $\dot{A}^{\mathcal{A}^*} = A$.

Least fixed-point logic Least-fixed point logic LFP is FO extended by formulas of the form $[\mathrm{lfp}_{\bar{x},X}\phi]\bar{t}$. For simplicity we recall syntax and semantics only for the case where $\bar{t} = \bar{z}$ for a tuple of variables $\bar{z}$.

$[\mathrm{lfp}_{\bar{x},X}\phi]\bar{z}$ is a τ-formula of LFP, $\phi \in \mathrm{LFP}[\tau]$, whenever $\phi = \phi(\bar{x}\bar{y})$ is a $\tau \cup \{X\}$-formula of LFP, X is a $|\bar{x}|$-ary relation symbol occuring positively in ϕ and $|\bar{x}| = |\bar{z}|$. It has free variables $\bar{y}\bar{z}$. It is satisfied by $\bar{b}\bar{c}$ in a τ-structure $\mathcal{A}$ if and only if $\bar{c}$ is in the least fixed-point reached when, starting with $B = \emptyset$, iterating the operation

$$B \mapsto \phi\frac{\bar{b}}{\bar{y}}((\mathcal{A}, B)).$$

Here $(\mathcal{A}, B)$ is the $\tau \cup \{X\}$-structure obtained from $\mathcal{A}$ by interpreting X by B.

[2] This means $A \subseteq A^*$ and $s^{\mathcal{A}^*} = s^{\mathcal{A}}$ for a relation smybol or a constant $s \in \tau$ and for a function symbol $f \in \tau$ the restriction of $f^{\mathcal{A}^*}$ to A is $f^{\mathcal{A}}$.

3 Proof of the main result

In this section we prove Theorem 2. Let a formula $\phi(\bar{x}) \in \mathrm{LFP}[\tau]$ and a τ-structure $\mathcal{A}$ be given. Let $k := |\bar{x}|$. If $|A| = 1$, there is nothing to do – we take $\rho := y = y$ and as $\mathcal{A}^*$ any τ^*-enlargement of $\mathcal{A}$ with $A^* = A$. The same works in case $k = 0$, since then $\phi(\mathcal{A})$ either contains nothing or exactly the empty tuple.

So we assume $k \geq 1, |A| \geq 2$ and $A = \{0, \ldots, |A| - 1\}$. We can further assume that $k \geq 3$ and

$$\bar{0} := \underbrace{0 \cdots 0}_{k \text{ times}} \notin \phi(\mathcal{A}).$$

To see this, note that $\tilde{\phi}(\bar{x}x_1x_2) := \phi(\bar{x}) \wedge x_1 = 1 \wedge x_2 = 1$ (where 1 is a parameter) satisfies this assumption and $|\tilde{\phi}(\mathcal{A})| = |\phi(\mathcal{A})|$; if $\mathcal{A}^*$ and $\tilde{\rho}(\bar{x}x_1x_2\bar{y}\bar{z})$ satisfy our claim for $\tilde{\phi}$, then $\mathcal{A}^*$ and $\rho(\bar{x}\bar{y}\bar{z}') := x_1 = 1 \wedge x_2 = 1 \wedge \tilde{\rho}$ with $\bar{z}' := \bar{z}x_1x_2$ satisfies our claim for $\phi(\bar{x})$.

3.1 Construction

Let p be the smallest prime bigger than $|A|$ and $k + 2$. Since by Bertrands Postulate there is a prime between n and $2n$ for every $n > 1$ we know that p can be computed in polynomial time. The structure $\mathcal{A}^*$ has as universe

$$A^* := \{0, \ldots, k \cdot p - 1\}.$$

We set $\tau^* := \{R_+, R_\times, R_{\mathrm{mod}}, \dot{A}, \leq\}$ for ternary $R_+, R_\times, R_{\mathrm{mod}}$, binary $\leq$ and unary $\dot{A}$. We let $\mathcal{A}^*$ be a τ^*-enlargement of $\mathcal{A}$ with

$$\begin{array}{rcl} \dot{A}^{\mathcal{A}^*} & := & A \\ R_+^{\mathcal{A}^*} & := & \{(a,b,c) \in (A^*)^3 \mid a + b \equiv c \bmod p\} \\ R_\times^{\mathcal{A}^*} & := & \{(a,b,c) \in (A^*)^3 \mid a \cdot b \equiv c \bmod p\} \\ R_{\mathrm{mod}}^{\mathcal{A}^*} & := & \{(a,b,c) \in (A^*)^3 \mid a \equiv b \bmod c\} \\ \leq^{\mathcal{A}^*} & := & \text{the natural order} \end{array}$$

In $\mathcal{A}^*$ the set $\{0, \ldots, p-1\}$ carries the structure of $\mathbb{F}_p$, the field with p elements. The set of solutions $\phi(\mathcal{A}) \subseteq A^k \subseteq \mathbb{F}_p^k$ can now be viewed as living in the k-dimensional vectorspace over $\mathbb{F}_p$. We identify vectors with their representations as k-tuples over $\mathbb{F}_p$ with relation to the standard basis. For a vector $\bar{a} \in \mathbb{F}_p^k \setminus \{\bar{0}\}$ let $\langle \bar{a} \rangle^\perp$ denote the hyperplane of vectors orthogonal to $\bar{a}$. If we translate $\langle \bar{a} \rangle^\perp$ by $\bar{a}$ we get $H_{\bar{a}}$, that is

$$H_{\bar{a}} := \langle \bar{a} \rangle^\perp + \bar{a} = \left\{ \bar{b} + \bar{a} \mid \bar{b} \in \langle \bar{a} \rangle^\perp \right\}.$$

Since this is a hyperplane in $\mathbb{F}_p^k$ we know

$$|H_{\bar{a}}| = p^{k-1}. \tag{1}$$

Let $\bar{a}, \bar{b} \in \mathbb{F}_p^k \setminus \{\bar{0}\}$ be different. If $\bar{a}$ and $\bar{b}$ are linearly independent, the hyperplanes $H_{\bar{a}}$ and $H_{\bar{b}}$ are not parallel and so the affine subspace $H_{\bar{a}} \cap H_{\bar{b}}$ has dimension $k-2$. If $\bar{a}$ and $\bar{b}$ are linearly dependent, then $H_{\bar{a}} \cap H_{\bar{b}} = \emptyset$. It is herefore that we use the translate $H_{\bar{a}}$ instead of $\langle \bar{a} \rangle^{\perp}$. Especially we get

$$|H_{\bar{a}} \cap H_{\bar{b}}| \leq p^{k-2}. \tag{2}$$

The proof strategy typically persued to get the Valiant-Vazirani theorem is by subsequently deviding the space where the solutions live (there a cartesian power of $\{0,1\}$) randomly in half for a random number of times. Then the probability that exactly one solution remains is bounded from below.

Here we argue similarly. Our solutions live in $\mathbb{F}_p^k$. An idea now is to cut down this space to a $1/p$-fraction using random hyperplanes and to do so subsequently for a random number of times. We do not select random hyperplanes but a random number of random vectors from $\mathbb{F}_p^k$ and look at the event that all these are contained in a hyperplane $H_{\bar{a}}$ associated with a solution $\bar{a} \in \phi(\mathcal{A})$. Arguing with (1) and (2) similarly to the argument known from the classical setting we can bound from below the probability that the above event holds for exactly one solution. It is here where we need that $\bar{0} \notin \phi(\mathcal{A})$. Details follow.

Intuitively the following formula wants $\bar{x}$ from A such that the first $u+2$ of $\bar{y}_1, \ldots, \bar{y}_{k+2} \in \mathbb{F}_p^k$ are in $H_{\bar{x}}$:

$$\begin{aligned} \rho' \quad := \quad & \bigwedge_{j \in [k]} \dot{A} x_j \wedge \bigwedge_{i \in [k+2]} \text{“}\bar{y}'_i = \bar{y}_i \bmod p\text{”} \wedge R_{\text{mod}} u u' (k+1) \\ & \wedge \bigwedge_{i \in [k+2]} \big(\max\{i-2, 0\} \leq u' \to \text{“}\bar{y}'_i \in H_{\bar{x}}\text{”} \big). \end{aligned}$$

Here $\max\{i-2,0\}, (k+1)$ and p are parameters. “$\bar{y}'_i = \bar{y}_i \bmod p$” abbreviates the formula $\bigwedge_{j \in [k]} R_{\text{mod}} y_{ij} y'_{ij} p$, where we write e.g. $\bar{y}_1 = y_{11} \cdots y_{1k}$.

We now explain for what formula “$\bar{y} \in H_{\bar{x}}$” stands for. Observe that – loosely written – $\bar{y} \in H_{\bar{x}}$ if and only if $\sum_j (y_j - x_j) \cdot x_j = 0$. We introduce auxiliary variables for all intermediate results obtained when computing this sum, namely we intend $u_j = y_j - x_j$, the v_j's to denote the products summed and the w_j's to denote the partial sums obtained when adding up the v_j's one after the other.

We let “$\bar{y} \in H_{\bar{x}}$” stand for the following formula with parameter 0:

$$\bigwedge_{j \in [k]} R_+ u_j x_j y_j \wedge \bigwedge_{j \in [k]} R_\times u_j x_j v_j \wedge \bigwedge_{j \in [k-1]} R_+ w_j v_j w_{j+1} \wedge w_1 = v_1 \wedge w_k = 0.$$

So “$\bar{y} \in H_{\bar{x}}$” is a formula in the variables $\bar{x}\bar{y}\bar{u}\bar{v}\bar{w}$. In ρ' we use for each $i \in [k+2]$ different auxiliary variables $\bar{u}_i \bar{v}_i \bar{w}_i$ in “$\bar{y}'_i \in H_{\bar{x}}$”, that is “$\bar{y}'_i \in H_{\bar{x}}$” $=$ “$\bar{y}'_i \in H_{\bar{x}}$”$(\bar{x}\bar{y}'_i\bar{u}_i\bar{v}_i\bar{w}_i)$.

We aim at a formula ρ such that with good probability for a random assignment to the $\bar{y}_i$ and u we have exactly one solution of $\phi^{\dot{A}} \wedge \rho$. Sofar this cannot work since we can assign whatever we want to the auxiliary variables $\bar{u}_i\bar{v}_i\bar{w}_i$ for i larger that the value of u' plus 2. We set

$$\rho := \rho' \wedge \bigwedge_{i \in [k+2]} (\neg \max\{i-2, 0\} \leq u' \to \text{“}\bar{u}_i\bar{v}_i\bar{w}_i = \bar{0}\text{”}).$$

In the notation of our claim we let $\bar{y}$ comprise all the $\bar{y}_i$'s and u and we let $\bar{z}$ comprise all primed variables as well as all the auxiliary variables $\bar{u}_i\bar{v}_i\bar{w}_i$.

This completes the construction of $\mathcal{A}^*$ and ρ. It is clear that $\mathcal{A}^*$ and ρ can be computed in polynomial time given $\mathcal{A}$ and $\bar{x}$.[3]

3.2 Probability analysis

Assume $\phi(\mathcal{A}) \neq \emptyset$. Let $b \in A^*$ and $\bar{b}_1, \ldots, \bar{b}_{k+2} \in (A^*)^k$ be arbitrary. Abbreviate

$$\psi := (\phi^{\dot{A}} \wedge \rho)\frac{b}{u}\frac{\bar{b}_1 \cdots \bar{b}_{k+2}}{\bar{y}_1 \cdots \bar{y}_{k+2}}.$$

Then $(\bar{x}, \bar{a})$ is a partial solution of $(\mathcal{A}^*, \psi)$ if and only if

$$\bar{a} \in \phi(\mathcal{A}) \text{ and } (\bar{b}_i \bmod p) \in H_{\bar{a}} \text{ for all } i \in [(b \bmod (k+1)) + 2].$$

Furthermore, if $(\bar{x}, \bar{a})$ is a partial solution of $(\mathcal{A}^*, \psi)$, then it has exactly one satisfying extension. Especially

$$|\psi(\mathcal{A}^*)| = \left|\left\{\bar{a} \in \phi(\mathcal{A}) \mid \bar{b}_i \bmod p \in H_{\bar{a}} \text{ for all } i \in [(b \bmod (k+1)) + 2]\right\}\right|. \quad (3)$$

We define a function f by

$$f\left(b\bar{b}_1 \cdots \bar{b}_{k+2}\right) := \left|(\phi^{\dot{A}} \wedge \rho)\frac{b}{u}\frac{\bar{b}_1 \cdots \bar{b}_{k+2}}{\bar{y}_1 \cdots \bar{y}_{k+2}}(\mathcal{A}^*)\right|.$$

Think of the set of the $b\bar{b}_1 \cdots \bar{b}_{k+2}$'s as carrying a probability space with the uniform probability measure. Declare two points of this space $b\bar{b}_1 \cdots \bar{b}_{k+2}$ and $b'\bar{b}'_1 \cdots \bar{b}'_{k+2}$ to be *equivalent* if and only if $b \equiv b' \bmod (k+1)$ and componentwise $\bar{b}_i \equiv \bar{b}'_i \bmod p$ for all $i \in [k+2]$.

Then the event $\{f = 1\}$ is a union of such equivalence classes. All equivalence classes have the same size since both p and k devide $|A^*|$. It is herefore that we chose A^* as we did. Thus the probability that f is 1 on a random argument $b\bar{b}_1 \cdots \bar{b}_{k+2}$ is the same as the probability that f is 1 on an argument chosen uniformly at random from those $b\bar{b}_1 \cdots \bar{b}_{k+2}$ with $b < k+1$ and $\bar{b}_i \in \mathbb{F}_p^k$ for all $i \in [k+2]$.

Let $B, B_1, \ldots, B_{k+2}$ be independent random variables such that B is uniformly distributed in $\{0, \ldots, k\}$ and each B_i is uniformly distributed in $\mathbb{F}_p^k$. Say, these random variables are defined on a probability space with measure Pr. We aim to bound $\Pr[f(BB_1 \cdots B_{k+2}) = 1]$ from below.

[3] Note that ρ depends only weakly on ϕ.

Call m *good* if and only if

$$p^m \leq |\phi(\mathcal{A})| \leq p^{m+1}.$$

Since $1 \leq |\phi(\mathcal{A})| \leq |A|^k \leq p^k$ for any good m we have $0 \leq m \leq k$. Hence $\Pr[B \text{ is good}]$ is at least $1/(k+1)$. Note that, if we find some t such that $\Pr[f(mB_1 \cdots B_{k+2}) = 1] \geq t$ for all good m, then we know that for at least an $1/(k+1)$ fraction of b's we find at least a t fraction of $\bar{b}_1 \cdots \bar{b}_{k+2}$'s such that f is 1 and hence $\Pr[f(BB_1 \cdots B_{k+2}) = 1] \geq t/(k+1)$.

Thus to prove our theorem it suffices to establish the following Claim I. It is only for the sake of some commodity here why we assumed $k \geq 3$ since this implies $1/((k+1)2p^2) \geq 1/|A^*|^2$.

Claim I: If m is good, then $\Pr\left[f(mB_1 \cdots B_{k+2}) = 1\right] \geq 1/(2p^2)$.

Proof of Claim I: Let m be good. By equation (3)

$$f(mB_1 \cdots B_{k+2}) = |\{\bar{a} \in \phi(\mathcal{A}) \mid B_1, \ldots, B_{m+2} \in H_{\bar{a}}\}|.$$

Hence $f(m\bar{b}_1 \cdots \bar{b}_{k+2}) = 1$ if and only if there is a solution $\bar{a} \in \phi(\mathcal{A})$ such that $H_{\bar{a}}$ contains all $\bar{b}_1, \ldots, \bar{b}_{m+2}$ but there is no other solution with this property. Define for $\bar{a} \in \phi(\mathcal{A})$ the event

$$A(\bar{a}) := \{B_1, \ldots, B_{m+2} \in H_{\bar{a}}\} \ \cap \bigcap_{\bar{a}' \in \phi(\mathcal{A}) \setminus \{\bar{a}\}} \ \bigcup_{i \in [m+2]} \{B_i \notin H_{\bar{a}'}\}.$$

Then $A(\bar{a}) \cap A(\bar{a}') = \emptyset$ for different $\bar{a}, \bar{a}' \in \phi(\mathcal{A})$ and

$$\{f(mB_1 \cdots B_{k+2}) = 1\} = \dot{\bigcup_{\bar{a} \in \phi(\mathcal{A})}} A(\bar{a}). \tag{4}$$

Using the following Claim II we get what we want:

$$\begin{aligned} \Pr\left[f(mB_1 \cdots B_{k+2}) = 1\right] &\overset{(4)}{=} \sum_{\bar{a} \in \phi(\mathcal{A})} \Pr[A(\bar{a})] \overset{\text{Claim II}}{>} \sum_{\bar{a} \in \phi(\mathcal{A})} \frac{p-1}{p^{m+3}} \\ &\overset{m \text{ good}}{\geq} \frac{p^m(p-1)}{p^{m+3}} \geq \frac{1}{2p^2}. \end{aligned}$$

Claim II: If m is good, then $\Pr[A(\bar{a})] > (p-1)/p^{m+3}$ *for all* $\bar{a} \in \phi(\mathcal{A})$.

Proof of Claim II: Let m be good and $\bar{a} \in \phi(\mathcal{A})$. Write $\bar{B} := (B_1, \ldots, B_{m+2})$. By (1)

$$\Pr\left[\bar{B} \in H_{\bar{a}}^{m+2}\right] = \prod_{i \in [m+2]} \Pr\left[B_i \in H_{\bar{a}}\right] = 1/p^{m+2}. \tag{5}$$

Let $\bar{a}, \bar{a}' \in \phi(\mathcal{A}), \bar{a}' \neq \bar{a}$. Then by (2)

$$\Pr\left[\bar{B} \in (H_{\bar{a}} \cap H_{\bar{a}'})^{m+2}\right] \leq 1/p^{2(m+2)}. \tag{6}$$

By (5) and (6)

$$\Pr\left[\bar{B} \in H_{\bar{a}'}^{m+2} \mid \bar{B} \in H_{\bar{a}}^{m+2}\right] \leq \frac{1/p^{2(m+2)}}{1/p^{m+2}} = 1/p^{m+2}. \tag{7}$$

Clearly

$$\begin{aligned} \Pr\left[A(\bar{a})\right] &= \Pr\left[\bar{B} \in H_{\bar{a}}^{m+2}\right] \cdot \Pr\Big[\bigcap_{\bar{a}'} \bigcup_{i \in [m+2]} \{B_i \notin H_{\bar{a}'}\} \mid \bar{B} \in H_{\bar{a}}^{m+2}\Big] \\ &\geq \Pr\left[\bar{B} \in H_{\bar{a}}^{m+2}\right] \cdot \Big(1 - \sum_{\bar{a}'} \Pr\left[\bar{B} \in H_{\bar{a}'}^{m+2} \mid \bar{B} \in H_{\bar{a}}^{m+2}\right]\Big). \end{aligned}$$

Here $\bar{a}'$ ranges over $\phi(\mathcal{A}) \setminus \{\bar{a}\}$. Using (5) and (7) we conclude that

$$\Pr\left[A(\bar{a})\right] \geq \frac{1}{p^{m+2}} \cdot \left(1 - \frac{|\phi(\mathcal{A})| - 1}{p^{m+2}}\right) \overset{m \text{ good}}{>} \frac{1}{p^{m+2}} \cdot \left(1 - \frac{p^{m+1}}{p^{m+2}}\right) = \frac{p-1}{p^{m+3}}.$$

□

4 Applications to parameterized complexity theory

In this section we apply Theorem 2 to prove Theorem 5. This is straightforward using the model-checking characterizations of the classes involved.

4.1 Complexity theoretic preliminaries

The parameterized setting We first recall the basic notions from parameterized complexity theory [6, 12]. We lean on [12]. Fix a finite alphabet Σ containing at least two elements. A *parameterized problem* is a pair (P, κ) of a classical problem $P \subseteq \Sigma^*$ and a *parameterization* $\kappa : \Sigma^* \to \mathbb{N}$ computable in polynomial time. (P, κ) belongs to the class FPT of *fixed-parameter tractable* problems if and only if there is an algorithm deciding it in *fpt time*, i.e. on input $x \in \Sigma^*$ it needs time at most $f(\kappa(x)) \cdot |x|^{O(1)}$ where $f : \mathbb{N} \to \mathbb{N}$ is some computable funtion.

An *fpt reduction* from a parameterized problem (P, κ) to another (P', κ') is a (many-one) reduction $r : \Sigma^* \to \Sigma^*$ from P to P' computable in fpt time "which doesn't increase the parameter too much": $\kappa' \circ r \leq g \circ \kappa$ for some computable $g : \mathbb{N} \to \mathbb{N}$. We write $(P, \kappa) \leq^{\text{fpt}} (P', \kappa')$ in case such a reduction exists and

$$[(P, \kappa)]^{\text{fpt}} := \left\{(P', \kappa') \mid (P', \kappa') \leq^{\text{fpt}} (P, \kappa)\right\}.$$

Parameterized randomization A W[P]-*randomized reduction with one-sided error*[4] from (P, κ) to (P', κ') is a probabilistic Turing machine T running in fpt

[4] The choice of this terminology [14] is motivated by the machine characterization [5] of the class W[P].

time such that for some computable $f, g, h : \mathbb{N} \to \mathbb{N}$ and some $c \in \mathbb{N}$ it uses on any $x \in \Sigma^*$ at most $f(\kappa(x)) \cdot \log |x|$ many random bits (coins) and

- if $x \in P$, then $\Pr[T(x) \in P'] \geq \frac{1}{g(\kappa(x))|x|^c}$,
- if $x \notin P$, then $\Pr[T(x) \in P'] = 0$,
- $\Pr[\kappa'(T(x)) \leq h(\kappa(x))] = 1$.

Here, as usual, the probability is taken over the coin tosses of T on x and the random variable $T(x)$ is the output of T on x.

Model-checking problems For a set of formulas Φ of some logic we consider the parameterized model checking problem p-MC(Φ)

> *Input:* a structure $\mathcal{A}$ and a formula $\phi \in \Phi$.
> *Parameter:* $\|\phi\| \in \mathbb{N}$ (the size of ϕ).
> *Question:* Is $\phi(\mathcal{A}) \neq \emptyset$?

This sloppy notation intends to define the parameterized problem consisting of the classical problem $\mathrm{MC}(\Phi) := \{(\mathcal{A}, \phi) \mid \phi \in \Phi \text{ and } \phi(\mathcal{A}) \neq \emptyset\}$ and the parameterization $(\mathcal{A}, \phi) \mapsto \|\phi\|$.

p-var-MC(Φ) is the parameterized problem obtained from MC(Φ) using the parameterization

$$(\mathcal{A}, \phi) \mapsto \text{number of (bound or free) variables of } \phi.$$

Parameterized intractable classes There are well-known characterizations of parameterized intractable classes by model-checking problems. Originally [6] the classes of the W-hierarchy, W[SAT] and W[P] have been defined by weighted satisfiability problems for various classes of Boolean circuits. For the "basic machinery" translating weighted satisfiability problems to model-checking problems and vice-versa see [10]. Let $t \geq 1$.

- $\mathrm{A}[t] = [p\text{-MC}(\Pi_{t-1})]^{\mathrm{fpt}}$,
- $\mathrm{W}[t] = [p\text{-MC}(\Pi_{t-1,1})]^{\mathrm{fpt}}$,
- $\mathrm{W[SAT]} = [p\text{-var-MC}(\Pi_0)]^{\mathrm{fpt}}$

These characterizations can be found in the monograph [12, Chapter 7]. For a characterization of W[P] by model-checking problems let $\Sigma_1\mathrm{LFP}^{[1]}$ denote the class of LFP-formulas of the form $[\mathrm{lfp}_{x,X}\phi]y$, where ($X$ is unary and) ϕ is a first order formula in which at most one variable has bounded occurences. Then [4, Theorem 33(1)]

- $\mathrm{W[P]} = [p\text{-MC}(\Sigma_1\mathrm{LFP}^{[1]})]^{\mathrm{fpt}}$

Clearly all these characterizations although not originally stated this way hold true for formulas with parameters.

4.2 Parameterized Valiant-Vazirani lemmata

For a (classical) model-checking problem MC(Φ) let $\textsc{Unique}$-MC(Φ) be the (classical) problem

> *Input:* a structure $\mathcal{A}$ and a formula $\phi \in \Phi$.
> *Question:* Is $|\phi(\mathcal{A})| = 1$?

Naturally we write p-$\textsc{Unique}$-MC(Φ) for the 'uniqueness-variant' of p-MC(Φ).

For $t \geq 1$ we define the 'uniqueness-variant' of W[t] by setting

$$\mathrm{UniqueW}[t] := [p\text{-}\textsc{Unique}\text{-}\mathrm{MC}(\Pi_{t-1,1})]^{\mathrm{fpt}}.$$

The classes $\mathrm{UniqueA}[t], \mathrm{UniqueW}[\mathrm{P}]$ and $\mathrm{UniqueW}[\mathrm{SAT}]$ are similarly obtained.

Remark 9 In [7] the classes UniqueW[t] were defined via "uniqueness variants" of weighted satisfiability problems for certain circuit classes. That our definition is equivalent follows from the fact that these problems are *parsimoniously* interreducible with the corresponding model-checking problems.[5]

Proof of Theorem 5. Let a model-checking problem MC(Φ) for $\Phi \subseteq$ LFP be given, let $(\mathcal{A}, \phi(\bar{x}))$ be an instance of it and let $k = |\bar{x}|$. Our probabilistic Turing machine first computes deterministically $\mathcal{A}^*$ and $(\phi^{\dot{A}} \wedge \rho)$ from Theorem 2 in polynomial time. It then guesses randomly $\bar{b} \in (A^*)^{|\bar{y}|}$ and outputs $\left(\mathcal{A}^*, (\phi^{\dot{A}} \wedge \rho)\frac{\bar{b}}{\bar{y}}\right)$.

Note that for any Φ coming from any of the model-checking problems complete for one of the classes mentioned in Theorem 5, we have $(\phi^{\dot{A}} \wedge \rho) \in \Phi$, whenever $\phi \in \Phi$ – perhaps modulo a polynomial time transformation. Without loss of generality we can assume that $|A| \geq k + 2$, so $|A^*| < k \cdot 2 \cdot |A|$. In case $\phi(\mathcal{A}) = \emptyset$, clearly $(\phi^{\dot{A}} \wedge \rho)\frac{\bar{b}}{\bar{y}}(\mathcal{A}^*) = \emptyset$. Hence we have one-sided error. Otherwise we have $(\mathcal{A}^*, (\phi^{\dot{A}} \wedge \rho)\frac{\bar{b}}{\bar{y}}) \in \textsc{Unique}$-MC($\Phi$) with probability at least $1/|A^*|^2 > 1/(k^2 \cdot 4 \cdot |A|^2)$.

Note that $|\bar{y}| = (k + 2) \cdot k + 1$ and thus the machine needs only

$$((k + 2) \cdot k + 1) \cdot \lceil \log |A^*| \rceil \text{ random bits.}$$

This number obeys a bound of the form $f(\kappa(x)) \cdot \log |x|$ for κ being the parameterization mapping $x = (\mathcal{A}, \phi)$ to the number k of variables of ϕ and hence also for κ being the parameterization by $\|\phi\|$. □

Proof of Theorem 6. Let $(P, \kappa) \in$ W[1]. Then $(P, \kappa) \leq^{\mathrm{fpt}} p$-MC($\Pi_0$). By Theorem 5 we have a W[P]-randomized reduction from p-MC(Π_0) to its uniqueness variant p-$\textsc{Unique}$-MC(Π_0). Now, there is a *parsimonious* fpt reduction from p-MC(Π_0) to p-$\textsc{Clique}$, i.e. an fpt reduction which on any instance $(\mathcal{A}, \phi)$ of

[5] For $t = 1$ see e.g. [12, Theorems 14.17, 14.12]. For $t > 1$ on the one hand the reductions from [12, Lemmata 7.5, 7.10] can be slightly modified to become parsimonious and on the other hand so it is for the ones from [12, Lemmata 7.23, 7.24].

p-MC(Π_0) produces an instance $(\mathcal{G}, k)$ of p-Clique such that $|\phi(\mathcal{A})|$ equals the number of k-cliques in $\mathcal{G}$. Hence this is an fpt reduction from p-Unique-MC(Π_0) to p-Unique-Clique. □

Thus the uniqueness variant of p-Clique is hard for W[1] under W[P]-randomized reductions with one-sided error.

What is an upper bound on its complexity? It is easy to see that the problem p-Two-Cliques

Input: a graph $\mathcal{G}$ and a natural $k \in \mathbb{N}$.
Question: k.
Question: does $\mathcal{G}$ contain at least two k-Cliques?

is in W[1]. E.g. it is easy to write down a quantifier free first-order formula $\phi(\bar{x}\bar{y})$ in the language of graphs satisfied exactly by those $\bar{a}\bar{b}$ such that $\bar{a}$ and $\bar{b}$ are tuples listing different k-cliques – but that reduces p-Two-Cliques to p-MC(Π_0). But clearly p-Unique-Clique is the intersection of p-Clique and the complement of p-Two-Cliques. Thus, by W[1]-completess of p-Clique, two non-adaptive oracle queries to p-Clique suffice to solve p-Unique-Clique.

Proposition 10 p-Unique-Clique *is non-adaptively fpt Turing-reducible to* p-Clique.

(Here, in fpt Turing reductions the oracle access has to be *balanced*: the parameter of any query has to be effectively bounded in terms of the input parameter.)

5 Further results and questions

For the sake of readability we skipped some notions from parameterized complexity theory and thereby waived some easy further results. The reader familiar with the nondeterministic random access machine model [5] easily sees that the reduction from the proof of Theorem 2 can be implemented by such a machine which is tail-nondeterministic – i.e. it performs nondeterministic (random) steps only in “the end of the computation”. This is called W[1]-randomization in [14].

The argument given for Theorem 5 shows more than stated. It proves analoguous statements also for all classes of the W*-hierarchy, all classes of the W$^{\text{func}}$-hierarchy, all classes of the A-matrix and the classes AW[$*$], AW[SAT] and AW[P]. See [12, Chapter 8] for model-checking characterizations of these classes excepting AW[P]. For AW[P] see [4, Theorem 33(2)].

Our results on the Clique problem do not exhibit a special property of this problem, they examplify more general statements using notions from parameterized counting complexity [12, Chapter 14]. For a parameterized counting problem (F, κ) (i.e. $F : \Sigma^* \to \mathbb{N}$ and κ a parameterization) let Unique-(F, κ) denote the problem of, given $x \in \Sigma^*$, to decide if $F(x) = 1$ (with parameterization κ). Of course we think of a counting problem as coming from a the decision problem D-(F, κ) which asks, given $x \in \Sigma^*$, if $F(x) > 0$. Then

Theorem 11 *Let $t \geq 1$. If a parameterized counting problem (F, κ) is* #W[t]*-complete under parsimonious fpt reductions, then* Unique-(F, κ) *is* W[t]*-hard under* W[1]*-randomized fpt reductions with one-sided error.*

Proposition 12 *Let $t \geq 1$. If a parameterized counting problem (F, κ) is* #W[t]*-complete under parsimonious fpt reductions, then* Unique-(F, κ) *is non-adaptively fpt Turing-reducible to* D-(F, κ).

Let us mention that Theorem 2 has been applied in [3]. Chen and Flum consider for a Π_t-formula $\phi = \phi(X)$ which is negative in a set variable X the parameterized counting problem p-Non-Maximal-WD$_\phi$ to count, given a structure $\mathcal{A}$ and a parameter k, the number of non-maximal solutions of size k. They show that p-Non-Maximal-WD$_\phi$ is not #W[1]-hard under parsimonious reductions unless W[1] is tractable in the sense that each problem in W[1] can be solved by a W[P]-randomized fpt algorithm [3].

Perhaps the most important question is whether in Theorem 2 we can achieve a success probability (let it be one with two sided error) of at least, say, $1/k$? This question comes from the struggle for a parameterized analogue of Todas theorem. Something like this seems to be prerequisite for combinatorics like those persued in [15] for the operators BP and $\oplus$. Possible parameterized analogues [12, Chapter 14] of Todas theorem are statements such as $\mathrm{A}[t] \subseteq \mathrm{FPT}^{\#\mathrm{W}[1]}$ or $\mathrm{W}[t] \subseteq \mathrm{FPT}^{\#\mathrm{W}[1]}$ or $\mathrm{A}[t] \subseteq \mathrm{FPT}^{\#\mathrm{W}[\mathrm{P}]}$.

Acknowledgements I owe much to discussions I had with Juan Andrés Montoya. I thank Jörg Flum for his advice.

References

[1] V.Arvind and V.Raman, *Approximation Algorithms for Some Parameterized Counting Problems.* In I.Bose and P.Morin, ed., Proceedings of the 13th Annual International Symposium on Algorithms and Computation, LNCS 2518, pp. 453-464. Springer, 2002.

[2] R. Chang, J. Kadin, and P. Rohatgi, *On Unique Satisfiability and the Threshold Behavior of Randomized Reductions*, Journal of Computer and System Sciences, 50(3), pp. 359-373, 1995.

[3] Y.Chen, J.Flum, *The Parameterized Complexity of Maximality and Minimality Problems*, Annals of Pure and Applied Logic, 151(1), pp. 22-61, 2008.

[4] Y.Chen, J.Flum and M.Grohe, *Bounded Nondeterminism and Alternation in Parameterized Complexity Theory,* In Proceedings of the 18th IEEE Conference on Computational Complexity, pp.13-29, 2003.

[5] Y.Chen, J.Flum and M.Grohe, *Machine-Based Methods in Parameterized Complexity Theory,* Theoretical Computer Science 339, pp. 167-199, 2005.

[6] R.G.Downey, M.R.Fellows, *Parameterized Comlexity*, Springer, 1999.

[7] R.G.Downey, M.R.Fellows and K.W.Regan, *Parameterized Circuit Complexity and the W Hierachy*, Theoretical Computer Science 191(1-2), pp. 97-115, 1998.

[8] H.D.Ebbinghaus, *Extended Logics: The General Framework*, in: Model-Theoretical Logics, ed. J.Barwise, S.Feferman, Springer-Verlag, pp. 25-76, 1985.

[9] H.D.Ebbinghaus, J.Flum, *Finite Model Theory*, 2nd edition, Springer-Verlag, 1999.

[10] J.Flum, M.Grohe, *Model-Checking Problems as a Basis for Parameterized Intractability*, Logical Methods in Computer Science 1(1), 2005.

[11] J.Flum, M.Grohe, *The Parameterized Complexity of Counting Problems*, SIAM Journal on Computing. 33(4), pp. 892-922, 2004.

[12] J.Flum and M.Grohe, *Parameterized Complexity Theory.* Springer, 2006.

[13] J.A.Montoya, communication.

[14] M.Müller, *Randomized Approximations of Parameterized Counting Problems*, in Proceedings of the 2nd International Workshop on Parameterized and Exact Computation (IWPEC'06), Lecture Notes in Computer Science 4169, pp.50-59, 2006.

[15] S.Toda, *PP is as Hard as the Polynomial Hierarchy*, SIAM Journal on Computing 20(5) , pp. 865-877, 1991.

[16] L.G.Valiant, V.V.Vazirani, *NP is as Easy as Detecting Unique Solutions*, Proceedings of the 17th annual ACM symposium on Theory of Computing, pp. 458-463, 1985.