Habilitationsschrift

Proofs and Constraints

Moritz Müller
Kurt Gödel Research Center
Universität Wien

Preface

This is a cumulative habilitation thesis collecting the following seven articles:

1. *Lower bounds for DNF-refutations of a relativized weak pigeonhole principle.*

   With Albert Atserias and Sergi Oliva.

   The Journal of Symbolic Logic 80 (2): 450-476, 2015.

2. *Partially definable forcing and bounded arithmetic.*

   With Albert Atserias.

   Archive for Mathematical Logic 54 (1): 1-33, 2015.

3. *Hard instances of algorithms and proof systems.*

   With Yijia Chen and Jörg Flum.

   ACM Transactions on Computation Theory 6 (2): Article No. 7, 2014.

4. *Consistency, optimality and incompleteness.*

   With Yijia Chen and Jörg Flum.

   Annals of Pure and Applied Logic 164 (12): 1224-1235, 2013.

5. *Topological dynamics of unordered Ramsey structures.*

   With Andraś Pongrácz.

   Fundamenta Mathematicae 230 (1): 77-98, 2015.

6. *An algebraic preservation theorem for $\aleph_0$-categorical quantified constraint satisfaction.*

   With Hubie Chen.

   Logical Methods in Computer Science 9 (1:15), 2013.

7. *The fine classification of conjunctive queries and parameterized logarithmic space.*

   With Hubie Chen.

   ACM Transactions on Computation Theory 7 (2): Article No. 7, 2015.

These articles are concerned with questions of mathematical logic in computational complexity theory. They treat a rather broad range of topics, roughly organized in two categories. The first four articles belong to proof complexity, the last three to constraint satisfaction.

The selection represents the development of my research activities since my doctoral dissertation. It is formally required that the selection should

(A) witness that this activity significantly transcends the topics treated in my doctoral dissertation, and

(B) contain works that stand in a certain thematic coherence.

(A) The first requirement is obviously met. I use this preface to describe the personal, professional circumstances leading to the seven articles collected in the present thesis.

My dissertation was titled *Parameterized Randomization* and treated topics in parameterized complexity theory, mainly the theory of randomized computations and kernelizations. The seven articles selected are not directly concerned with parameterized complexity theory except for article No. 7 which treats the theory of parameterized logarithmic space.

Already with my first postdoc employment in Sy-David Friedman's *Infinity Project* at the Centre de Recerca Matemàtica in Bellaterra (near Bercelona) I had to switch topics. This project aimed to find cross connections between different branches of mathematical logic. My part within this project was to study applications of methods from set-theoretic forcing in proof complexity. My work on this ultimately lead to article No. 2 in the present thesis.

This article together with article No. 1 concern proof complexity. Proof complexity asks to establish lower bounds on the size of refutations of certain contradictions in certain propositional refutation systems, or, seen from a different perspective, it asks to establish independence of certain simple sentences from certain weak fragments of arithmetic. This work is done in collaboration with Albert Atserias from the Universitat Polytècnica de Catalunya in Barcelona and Sergi Oliva who was his PhD student at the time.

Simultaneously I continued to work together with Jörg Flum and Yijia Chen, a collaboration rooted in my time as a PhD student in Freiburg. Articles No. 3 and 4 treat proof complexity from a more abstract perspective of computational complexity theory, more precisely, we studied the question of the existence of optimal proof systems and algorithms.

The second, main scientific contact from my time in Barcelona is Hubie Chen at the time at the Universitat Pompeo Fabra. We started a fruitful and lasting collaboration on constraint satisfaction leading to articles No. 6 and 7.

Article No. 5 concerns Ramsey properties of amalgamation classes, a topic showing up in some recent developments in constraint satisfaction theory. I came in contact with my co-author András Pongrácz following an invitation of Manuel Bodirsky at the École Polytechnique in Paris.

(B) The second requirement, that the articles collected stand in a certain thematic coherence, is not obviously met. The following introduction is meant to show up such a coherence.

Article No. 1 belongs to proof complexity, it proves a lower bound on the length of certain propositional refutations of a particular sequence of contradictions. Article No. 7 treats the parameterized space complexity of homomorphism problems. These topics are trivially related via certain collapses of complexity classes, e.g. NP and P, and one could easily and quite cheaply say both are concerned with and motivated by the P versus NP problem. It would be easy to argue some sort of coherence via such implications. The following introduction intends to say something more informative.

It explains the contents of each of the seven articles in order and tries to explain the motivation for each of them not only with respect to some open questions in current research but to questions of, as I would like to argue, genuine interest to mathematical logic. This gives a thematic coherence at least in the sense of a common narrative that leads smoothly from the first to the last article. To make this work, the introduction contains proofs of a handful of propositions. These are original to this introduction in the sense that I am not aware of a reference, but I consider each of them folklore in the sense that I expect them to be known to the experts.

The introduction is written in a language understandable by any logician. It intends to provide to the non-expert a glimpse into a selection of subdisciplines of mathematical logic and aims to explain the interest in the respective questions.

In the course, we display one theorem from each of the seven articles and roughly describe its context in current research. A full description of content and context is to be found in the introductions opening each of the articles, and this information is only partially repeated here.

Acknowledgments

I am grateful to the people mentioned in this preface. The present thesis rests on their support.

Introduction

Pigeonhole principles

Proof complexity seeks to show that certain propositional contradictions do not admit short refutations in certain propositional refutation systems. Here, short means polynomial in the size of the contradiction refuted.

We start with some examples. The *drosophila* of proof complexity is the family of pigeonhole principles. For natural numbers $m > n > 0$ the principle PHP^m_n expresses that, if m pigeons fly to n holes, then some hole is doubly occupied. More precisely, PHP^m_n is a contradictory CNF written in propositional variables p_{ij} for $i < m$ and $j < n$ whose truth value is meant to be that of the statement "pigeon i flies to hole j". It consists in the following clauses:

$$\begin{array}{rl} \bigvee_{j<n} p_{ij} & \text{for all } i < m; \\ \neg p_{ij} \vee \neg p_{kj} & \text{for all } i,k < m \text{ with } i \neq k \text{ and all } j < n. \end{array}$$

Maybe the most cited result in proof complexity is Haken's [26] stating that Resolution refutations of PHP^{n+1}_n require size $2^{\Omega(n)}$. Lower bounds for *weak* pigeonhole principles, i.e. PHP^m_n with $m >> n$, are considerably harder to show and achieved by Razborov [47] who gave a $2^{\Omega(n/\log^2 m)}$ lower bound.

Resolution is an important refutation system due to its connections to SAT solvers. We refer to [43] for a recent survey. But it is a rather weak refutation system in that it operates with clauses only. A straightforward extension is $R(k)$ that operates with k-DNFs, that is, disjunctions of conjunctions of up to k many literals. More formally, one can define the systems $R(k)$ as obtained from the usual propositional Gentzen sequent calculus (see e.g. [56]) by restricting the cut rule to conjunctions of at most k literals. Allowing conjunctions of arbitrarily many literals gives the system $R(\infty)$, a system operating with arbitrary DNFs. The powerful AC_0-*Frege* systems cut on formulas of some fixed finite $\wedge, \vee$-alternation rank. The full system without any restriction on the cut rule is called *Frege*.

Lower bound proofs against $R(k)$ typically employ the method of random restrictions: one randomly assignes truth values to a subset of the variables with the effect that a k-DNF probably simplifies in the sense of getting small "width" (e.g. a shallow decision tree); unless a given refutation is long, all its formulas are simplified simultaneously by a single restriction (probabilistic argument);

finally, one shows that refutations of small width cannot exist (e.g. directly by a game theoretic argument [46]). The probabilistic analysis often requires non-trivial combinatorics based on expander graphs and switching lemmas. An example is the result from Segerlind, Buss and Impagliazzo [53] stating that $R(k)$ refutations of PHP^{2n}_n require size $2^{n^{\Omega(1)}}$ for k as large as $\sqrt{\log n / \log\log n}$. Razborov [48] improved on this allowing k to be as large as $\epsilon \log n/\log\log n$ for some $\epsilon > 0$.

Concerning upper bounds, Maciel, Pitassi and Woods [40] found quasipolynomial size $2^{\log^{O(1)} n}$ refutations of PHP^{2n}_n in $R(k)$ for $k \leq \log^{O(1)} n$. The main open question here is whether PHP^{2n}_n has polynomial size proofs in AC_0-Frege. In fact this is a long standing open question, it goes back to Paris, Wilkie and Woods [44]. We come back to this later.

In Article No. 1 we give a lower bound for $R(\infty)$-refutations of the *relativized* weak pigeonhole principle $\mathrm{PHP}^{n^2,2n}_n$ expressing that, if at least $2n$ out of n^2 pigeons fly into n holes, then some hole must be doubly occupied. Technically, the relativization allows the construction of a restriction akin to those in the seminal paper [23] on circuit lower bounds. We show the following:

Theorem 1 *For every $c < 1.5$ and every sufficiently large n, every $R(\infty)$-refutation of $\mathrm{PHP}^{n^2,2n}_n$ has size at least $2^{\log^c n}$.*

This is almost tight: based on the upper bounds of [40] and standard propositional simulations of relativized bounded arithmetic $T^2_2(\alpha)$ yield a quasipolynomial upper bound even in $R(k)$ for $k \leq \log^{O(1)} n$.

It is worth pointing out that in the above theorem the precise choice of the inference rules defining $R(\infty)$ does not matter much. The lower bound applies to *semantic* DNF-refutations, that is, sequences of DNFs ending in the empty clause such that each DNF is either a clause from the CNF refuted or "strongly" implied by two earlier DNFs; strong implication is a property shared by most common inference rules. We refer to Article No. 1 for the definition. I am not aware of any other superpolynomial lower bounds for semantic DNF-refutations. For $R(\infty)$ the only other known lower bounds come from the lower bounds for the much stronger AC_0-Frege (not semantic).

Ajtai's theorem

Where does the mentioned long standing open question come from? Paris, Wilkie and Woods [44] studied the question whether $I\Delta_0$ can prove the infinitude of primes. The theory $I\Delta_0$ is the fragment of arithmetic based on induction for Δ_0-formulas. They arrived at the question whether this theory can refute

$$\begin{aligned} wphp_\varphi(x,\vec{x}) \quad := \quad & x > 0 \ \wedge\ \forall u < 2x \exists v < x \varphi(u,v,\bar{x}) \\ & \wedge\ \forall u u' v(\varphi(u,v,\bar{x}) \wedge \varphi(u',v,\bar{x}) \to u = u') \end{aligned}$$

for each Δ_0-formula $\varphi(u,v,\vec{x})$. An easier question would be to refute the formula $php_\varphi(x,\vec{x})$ which is similarly defined but replacing the bound $2x$ by $x+1$. One might even hope that $I\Delta_0$ can refute all these formulas by one and the same refutation, a refutation where the formula φ is treated as a black box. More precisely, one might hope to refute

$$\mathit{php}_R(x) := x > 0 \ \wedge\ \forall u < x+1\ \exists v < x\ Ruv \ \wedge\ \forall uu'v(Ruv \wedge Ru'v \to u = u')$$

for a new binary relation symbol R. The theory now is allowed to use induction for $\Delta_0(R)$-formulas, bounded formulas possibly mentioning the new symbol R.

As already mentioned that Ajtai [1] proved that refutations of PHP^{n+1}_n in AC_0-Frege require superpolynomial size $n^{\omega(1)}$. This implies that $I\Delta_0(R)$ does not prove $\neg \mathit{php}_R(x)$ due to the *propositional simulation* of $I\Delta_0(R)$ by AC_0-Frege. We sketch what this means.

Given a $\Delta_0(R)$-formula $\varphi(x)$ and $n \in \mathbb{N}$ it is straightforward to express the truth of $\varphi(n)$ (in the standard model) as a propositional combination $\langle\varphi\rangle_n$ of atomic sentences $Rk\ell$ for parameters $k, \ell \in \mathbb{N}$. In fact, the parameters and the size of the propositional formula $\langle\varphi\rangle_n$ are polynomially bounded in n. If $I\Delta_0(R)$ refutes $\varphi(x)$, a cut elimination argument yields a refutation consisting entirely of $\Delta_0(R)$-formulas. One can propositionally translate all formulas appearing in the refutation and fill in some propositional proof steps to arrive at a propositional refutation of $\langle\varphi\rangle_n$. The depth of the propositional formulas appearing in this refutation depends only on φ, not on n. This way one obtains AC_0-Frege refutations of polynomial size (see [32] for details).

Ajtai showed something stronger. He showed how to expand an arbitrary nonstandard model of true arithmetic M to a model (M, R^M) such that for some (nonstandard) $b, n \in M$

(a) R^M is a bijection between (the initial segment of M up to) $n+1$ and n;

(b) (M, R^M) satisfies the least number principle up to b for $\Delta^b_0(R)$-formulas;

(c) $n < b^{o(1)}$.

Here, $\Delta^b_0(R)$ denotes the formulas all of whose quantifiers are bounded by b; (b) means that if a set defined in (M, R^M) by such a formula (possibly with parameters from M) is satisfied by an element $< b$ then it contains a minimal element; the $o(1)$ in (c) denotes an infinitesimal.

From Ajtai's model one can infer not only that $I\Delta_0(R) \not\vdash \neg\mathit{php}_R(x)$ but also directly that AC_0-Frege refutations of $\langle \mathit{php}_R\rangle_m$ or PHP^{m+1}_m require superpolynomial size. We refer to to [35] for a recent survey of nonstandard model constructions in complexity theory.

Ajtai refers to his construction as "done according to the general ideas of Cohen's method of forcing" [1]. The purpose of Article No. 2 is to explain this. Let L^* be a countable language extending a language L. We give a general framework to construct *generic* L^*-expansions of a countable L-structure M by forcing, namely, for a suitable notion of genericity these expansions are determined by generic filters in a given forcing frame $\mathbb{P}$ and a so-called *conservative* notion of forcing $\Vdash$. We refer to Article No. 2 for the definitions.

For applications in weak arithmetics the relevant forcing frames are typically not definable in M. Generally, the forcing frame is just a second structure.

We verify the usual forcing lemmas familiar from set theory [38] except the definability lemma. If M interprets a linear order $<^M$ we ask for expansions satisfying the least number principle for as many L^*-formulas as possible.

Writing $\bar{a} <^M b$ for $\bar{a}$ a tuple from M and $b \in M$ means that each component of $\bar{a}$ is $<^M b$. We show:

Theorem 2 *Let $b \in M$. Assume Φ is a class of L^*-formulas such that for every $\varphi(\vec{x})$ and every $p \in \mathbb{P}$ the set $\{\bar{a} <^M b \mid p \Vdash \varphi(\bar{a})\}$ is definable in M. Then every generic expansion of M satisfies the least number principle for Φ up to b.*

In set theory the analogous statement is that generic extensions of models of ZF again model ZF. In set theory forcing is so useful because it reduces the task to prove independence from ZF to a combinatorial task of designing forcing frames. The above theorem similarly reduces the task to prove independence from weak arithmetics to the task of designing a suitable forcing frame. A more combinatorial formulation of the definability condition above is obtained by what we call the *antichain method*. This reduces the task to finite combinatorics but unfortunately quite difficult ones – usually referred to as *switching lemmas*.

Using a known switching lemma [37] the method yields a model (M, R^M) as in Ajtai's theorem but with b in (c) as large as $2^{n^{o(1)}}$. This implies independence of $\neg php_R(x)$ from full bounded arithmetic $S_2(R)$, and, on the propositional side, an exponential lower bound $2^{m^{\Omega(1)}}$ on the size of AC_0-refutations of PHP^{m+1}_m[3].

Independence in mathematical logic

Ajtai's result remains somewhat isolated. For example, it is not known whether one one ca find such a model where in (a) we have R^M violating $wphp_R(x)$ instead $php_R(x)$. This would imply superpolynomial lower bounds on AC_0-refutations of PHP^{2n}_n, and this answer the open question mentioned before.

More generally, one would like to violate other combinatorial principles in (a), but only a handful are known to be amenable to Ajtai's proof (see [32, Ch. 12]). Second, even slightly enlarging the class $\Delta^b_0(R)$ in (b) to allow formulas with parity quantifiers is one of the long standing open problems in proof complexity. Given the failure to repeat Ajtai's argument in different contexts, one can say that Ajtai's proof has not been properly understood as providing a *method* to prove independence from weak arithmetics. Doing a step in this direction is the main motivation behind Article No. 2.

This motivation becomes clear from a wider perspective. In Pudlák's words "There has been impressive success in proving independence results for set theory. [...] This may give the impression that logic is doing very well in studying the independence phenomenon which is only partly true" [45, Section 1]. He says "we are not satisfied with the current methods of proving independence results. The main reason is that except for Gödel's theorem which gives only some special formulas, no general method is known for proving independence of Π_1 sentences." [45, Section 3]

Article No. 2 intends to show that a large part of Ajtai's argument can be mimicked by general forcing arguments. The combinatorial core of the argument, the switching lemma, is interpreted as a statement about the existence of certain antichains. This formulates the combinatorics in a language familiar to set theorists, and, in fact, the hope expressed in Article No. 2 is that some of their vast forcing experience might be transferable to proof complexity.

Another recent attempt to turn Ajtai's argument into a method is given by Krajíček in his book on forcing with random variables [34]. In this setting the role of the switching lemma is to provide a certain form of bounded quantifier elimination.

Cook's program and optimality

We now take a more abstract view. Define a *proof system* for a set of binary strings Q to be a polynomial time computable function F from binary strings to binary strings whose image is Q. An *F-proof* of a string $x \in Q$ is a string w such that $F(w) = x$. Note $Q \in$ NP, if F is *polynomially bounded* in the sense that there is a polynomial p such that every $x \in Q$ has an F-proof w of size $|w| \leq p(|x|)$. Conversely, given an NP-machine $\mathbb{A}$ for a nonempty Q, one gets a polynomially bounded proof system $F_{\mathbb{A}}$ for Q: map any w coding an accepting run of $\mathbb{A}$ to the input it accepts, and any other w to some fixed $x_0 \in Q$.

A *refutation system* is a proof system for Unsat, the set of contradictory CNFs, and proofs are called *refutations*. E.g. Resolution "is" the refutation system that maps (F, π) to F if π codes a Resolution refutation of the CNF F; strings not encoding such pairs are mapped to some fixed contradictory CNF.

As observed in Cook and Reckhow's founding paper of proof complexity [13], NP equals coNP if and only if there exists a polynomially bounded refutation system. Proof complexities task, or *Cook's program*, to establish lower bounds on refutations in stronger and stronger systems is often dubbed an approach to ultimately show NP $\neq$ coNP. It depends on the reader's optimism whether this can count as a motivation. In any case, it would be interesting to have an optimal refutation system R such that superpolynomial lower bounds to R imply such bounds for any other system.

A proof system F for Q is *optimal* if for every other proof system F' for Q there exists a polynomial p such that for all w' there is w with $|w| \leq p(|w'|)$ and $F(w) = F'(w')$. Trivially, polynomially bounded proof systems are optimal. The hypothesis that there is no optimal refutation systems thus stengthens the hypothesis that NP $\neq$ coNP. In fact, in one of the most important papers of proof complexity [36], Krajíček and Pudlák showed that it implies NE $\neq$ coNE, and subsequent work [4, 31] showed it implies seemingly even stronger separations of complexity classes. Here, NE denotes nondeterministic time $2^{O(n)}$.

That a refutation system R is not polynomially bounded is witnessed by a set X of contradictory CNFs such that minimal refutation length of R is not polynomially bounded on X. If we find such a *hard set* even in NP, it witnesses that R is not optimal. Krajíček and Pudlák [36] showed a partial converse: there are no optimal refutation systems if *and only if* every refutation system has a hard set in NP. Indeed, then for each R one can feasibly produce contradictory CNFs which are hard to refute for R, i.e. R has a hard sequence (see below).

Following [41], Article No. 3 asks to what extent the non-optimality of proof systems and, for a similar notion going back to Levin [39], the non-optimality of deterministic algorithms for a problem Q can be feasibly witnessed as above, that is, by hard sets and/or hard sequences. We display a positive result from Article No. 3. A sequence $(x_n)_{n\geq 1}$ is *(quasi-)polynomially hard* for a proof system F of Q if x_n can be computed from n in time polynomial in n, each x_n is in Q, and $\min\{|w| \mid F(w) = x_n\}$ is not (quasi-)polynomially bounded in n.

For $t \geq 1$ let Π^p_t denote the "universal" t-th level of the polynomial hierarchy.

Theorem 3 *Let $t \geq 1$ and Q be Π^p_t-complete. There is no optimal proof system for Q if and only if every proof system for Q has a polynomially hard sequence.*

The proof proceeds by considering the complexity of the problem to show that a given nondeterministic algorithm is sound for Q. As this is undecidable, we consider a bounded version. This can be considered an abstract version of reflection principles of refutation systems (see e.g. [32, Ch. 14]):

SOUND$_Q$
Instance: a nondeterministic algorithm $\mathbb{A}$ and $n \in \mathbb{N}$ in unary.
Problem: does Q contain all strings of length at most n which $\mathbb{A}$ accepts within at most n steps?

The statements in the above theorem are equivalent to the statement that there does not exists a nondeterministic algorithm accepting SOUND$_Q$ that accepts each "yes" instance $(\mathbb{A}, n)$ in time $n^{f(|\mathbb{A}|)}$ for some arbitrary $f : \mathbb{N} \to \mathbb{N}$.

As shown in Article No. 3, a similar equivalences hold true for deterministic algorithms instead of proof systems. The equivalences might fail for arbitrary Q if one assumes the so-called measure hypothesis.

Note the complexity analysis of SOUND$_Q$ transcends classical computational complexity theory in that the time complexity is measured not by a unary function in the instance length but by a binary function taking into account a "parameter" $|\mathbb{A}|$. This is the frame-work of parameterized complexity theory going back to Downey and Fellows ([18] is a recent monograph). We note in passing that this frame-work revolutionarized computational complexity theory, especially the algorithmics of NP-hard problems. It naturally appears in the present, somewhat different context. We shall see more examples.

The spectrum problem and a conjecture of Riis

Ajtai's result states that $(\mathrm{PHP}^{n+1}_n)_{n\geq 1}$ is a polynomially hard sequence for AC$_0$-Frege. We vaguely asked for other "combinatorial principles". hard for AC$_0$-Frege or other refutation systems. To make this precise consider the problem

SPEC
Instance: a first-order sentence φ and $n \in \mathbb{N}$.
Problem: does φ have a model with universe of size n?

Recall, the *spectrum of* φ is the set $\mathrm{spec}(\varphi) := \{n \in \mathbb{N} \mid (\varphi, n) \in \textsc{Spec}\}$. If numbers are encoded in unary, then each spectrum is in NP. As propositional translation we take a polynomial time function mapping (φ, n), with n given in unary, to a CNF $\langle\varphi\rangle_n$ which is satisfiable if and only if $n \in \mathrm{spec}(\varphi)$. If $\mathrm{spec}(\varphi)$ is empty, we get a sequence of contradictions $(\langle\varphi\rangle_n)_{n\geq 1}$. For example, if φ states that a binary function symbol is injective, then a natural choice of the translation would yield a CNF $\langle\varphi\rangle_n$ which, maybe up to renaming variables, equals $\mathrm{PHP}_n^{n^2}$ plus "functionality clauses" $\neg p_{ij} \vee \neg p_{ik}$ for all $i < n^2$ and $j, k < n, i \neq k$.

Asking for the length of refutations of $\langle\varphi\rangle_n$ is tantamount for asking for the nondeterministic time complexity of the complement of $\mathrm{spec}(\varphi)$. This is the spectrum problem: Scholz asked 1952 for "a necessary and sufficient condition" [52] for a subset of $\mathbb{N}$ to be a spectrum, and Asser asked 1955 whether spectra are closed under complementation [2]. Jones and Selman [28] showed that, under some natural identification of natural numbers with binary strings, the set of spectra equals NE. So Asser really asked whether NE equals coNE. We refer to [19] for a comprehensive, historical survey of the spectrum problem.

Call a first-order sentence φ *(quasi-)polynomially hard for* a refutation system R if for all $c \in \mathbb{N}$ there is $n \notin \mathrm{spec}(\varphi)$ such that there is no R-refutation of $\langle\varphi\rangle_n$ of size n^c (of size $2^{\log^c n}$). Let NEXP denote nondeterministic time $2^{n^{O(1)}}$.

Proposition 1 *The following are equivalent.*

(a) *There is a first-order sentence φ that is quasi-polynomially hard for every refutation system R.*

(b) *For every refutation system R there is a first-order sentence φ that is quasi-polynomially hard for R.*

(c) NEXP $\neq$ coNEXP.

Proof: The implication form (a) to (b) is trivial. To see that (b) implies (c), assume NEXP = coNEXP. This implies that there is an NEXP machine $\mathbb{A}$ that accepts, given φ and n *in binary*, precisely if $n \notin \mathrm{spec}(\varphi)$. Say, $\mathbb{A}$ runs in time $2^{(|\varphi|\cdot\log n)^{O(1)}}$. Let R_0 be an arbitrary refutation system. Define R to map a string $0w$ to $R_0(w)$, and a string $1w$ to $\langle\varphi\rangle_n$ or $R_0(w)$ depending on whether w encodes an accepting run of $\mathbb{A}$ on (φ, n). Then R is a refutation system that has quasi-polynomial size refutations of each contradiction of the form $\langle\varphi\rangle_n$. Thus R witnesses that (b) fails.

To show (c) implies (a), assume (a) fails, that is, for every first-oder φ there is a refutation system R_φ and $c_\varphi \in \mathbb{N}$ such that for every $n \notin \mathrm{spec}(\varphi)$ there is a R_φ-refutation of $\langle\varphi\rangle_n$ of size $2^{\log^{c_\varphi} n}$. Let $Q \in$ NEXP. Choose $d \in \mathbb{N}$ such that $\{x01^{|x|^d} \mid x \in Q\}$ is in NE. For a string y let $\mathrm{num}(y)$ be the natural with binary representation $1y$. By Jones and Selman's theorem cited earlier, there exists a first-order φ such that $\mathrm{spec}(\varphi) = \{\mathrm{num}(x01^{|x|^d}) \mid x \in Q\}$. The machine that on input x guesses w and checks $R_\varphi(w) = \langle\varphi\rangle_{\mathrm{num}(x01^{|x|^d})}$ accepts the complement of Q. Each $x \notin Q$ is accepted in time $2^{|x|^{O(1)}}$, so $Q \in$ coNEXP. $\square$

This proposition improves on [50] where it is shown that (c) implies (a) with quasi-polynomial hardness weakened to polynomial hardness. A similar argument shows that (b) with quasi-polynomial hardness weakened to polynomial hardness is equivalent to NE $\neq$ coNE.

Riis [49] conjectured that a stronger version of (b) holds, namely, one additionally requires $\mathrm{spec}(\varphi) = \emptyset$. Then all $\langle\varphi\rangle_n$ are contradictory. In other words, Riis conjectured that for every refutation system there exists a first-order sentence φ such that $(\langle\varphi\rangle_n)_{n\geq 1}$ is a polynomially hard sequence for R. He noted that his conjecture implies NP $\neq$ coNP. It seems to be much stronger:

Proposition 2 *Riis' conjecture holds if and only if there is no optimal refutation system.*

Proof: (Sketch) The forward direction is easy to see. Conversely, assume Riis' conjecture fails, that is, there exists a refutation system R such that for every first-order φ with empty spectrum there exists $c_\varphi \in \mathbb{N}$ such that R has a refutation of $\langle\varphi\rangle_n$ of size n^{c_φ}.

Let $Q := \textsc{Unsat}$. By the comment after Theorem 3 it suffices to find a nondeterministic $\mathbb{V}$ accepting $\textsc{Sound}_Q$ that accepts every "yes" instance $(\mathbb{A}, n)$ in time $n^{f(|\mathbb{A}|)}$ for some function f. It is not hard to see that it suffices to find $\mathbb{V}$ accepting $\textsc{Sound}_Q$ such that for every input $(\mathbb{A}, n)$ *such that* $\mathbb{A}$ *accepts a subset of* Q there exists $d_{\mathbb{A}} \in \mathbb{N}$ such that $\mathbb{V}$ accepts $(\mathbb{A}, n)$ in time $n^{d_{\mathbb{A}}}$.

For every nondeterministic $\mathbb{A}$ the set $\{n \in \mathbb{N} \mid (\mathbb{A}, n) \notin \textsc{Sound}_Q\}$ is in NE (input n coded in binary). Standard means allow to compute from $\mathbb{A}$ a first-order $\varphi_{\mathbb{A}}$ such that this set equals the spectrum of $\varphi_{\mathbb{A}}$. In particular, if $\mathbb{A}$ accepts a subset of Q, then this spectrum is empty.

The desired algorithm $\mathbb{V}$ on input $(\mathbb{A}, n)$ with $n \geq 1$ simulates an arbitrary nondeterministic algorithm accepting $\textsc{Sound}_Q$ and in parallel does the following: compute $\varphi_{\mathbb{A}}$, guess a string w and check $R(w) = \langle\varphi_{\mathbb{A}}\rangle_n$. □

The foundational crisis

Central questions of modern mathematical logic at its birth during the so-called foundational crisis of mathematics were whether formalized mathematics can be proven to be consistent (Hilbert's program), and whether it is algorithmically decidable whether a given sentence is valid (Hilbert's Escheidungsproblem). In a lecture at the European Congress of Mathematics 2004, Krajíček pointed out that problems of proof complexity can be seen as "quantitative versions" [33] of these questions.

One way to make this point: recall, Gödel showed 1930 that the set of valid sentences is computably enumerable, Church and Turing showed 1936 it is not decidable. Trakhtenbrot showed 1950 that validity in the finite is not computably enumerable. Asking for the computational complexity of $\textsc{Spec}$ can be regarded as a quantitative version of these questions.

Already Gödel, ahead of his time, asked in a letter to von Neumann whether

BOUNDED ENTSCHEIDUNG
Instance: a first-order sentence φ and $n \in \mathbb{N}$.
Problem: is there a proof of φ of size at most n?

can be decided in polynomial time – see [59] for what Gödel literally asked, and a historical discussion of the letter. The problem is NP-complete even for the version where we demand φ to come from propositional logic [10] (and understand "provable" to refer to a certain Frege system), so we see Gödel adressed the P versus NP problem in his context.

Asking for the nondeterministic time complexity of the complement of BOUNDED ENTSCHEIDUNG can be seen as a quantitative version of Hilbert's program. Chen and Flum showed that this complexity is $n^{f(|\varphi|)}$ for some function f if and only if optimal refutation systems exist [14]. See [14, 15] for more about of the parameterized complexity of BOUNDED ENTSCHEIDUNG and SPEC.

In their already mentioned seminal work [36], Krajíček and Pudlák link the existence of an optimal refutation system to the feasibility of "Hilbert's program in a modified, finitistic sense" [36, p. 1067]. We refer to [36] for details.

Finally, let us mention a more philosophical connection. Hilbert listed 10 problems in his famous speech at the International Congress of Mathematicians in Paris 1900, and 23 problems in the paper of his speech. His notebook contains a 24th problem originally planned to be included in the Paris-list. It asks for criteria for the simplicity of proofs. We refer to [58] for a precise statement and historical discussion of this problem.

Proof complexity addresses it – we cite from the introduction of Cook and Nguyen's monograph [12]: "One purpose of this book is to serve as a basis for a program we call "Bounded Reverse Mathematics".[. . .] From the complexity theory point of view, the idea is to find the smallest complexity class such that the theorem can be proved using concepts in that class."

Independence and optimality

According to the previous citation, proof complexity studies the complexity of proofs in terms of the computational complexity of the concepts involved in the proofs. Stronger theories may allow to reason with more complicated concepts. A problem Q may be decidable by a (deterministic) algorithm whose proof of correctness needs concepts not available in a given theory. It is conceivable that stronger theories may be able to verify faster algorithms.

S. Friedman asked whether $T + Con_T$ can be characterized in this context as a minimal extension of a computably enumerable theory T. Here, Con_T is Gödel's sentence expressing the consistency of T.

Let us fix a decidable problem Q outside P. Say, a theory T *knows* an algorithm $\mathbb{A}$ if it proves "$\mathbb{A}$ decides Q"; we understand this as an arithmetical statement about (the natural number coding) $\mathbb{A}$. Call $\mathbb{A}$ *as fast as* $\mathbb{B}$ if the running time of $\mathbb{A}$ on any input x is bounded by a polynomial in $|x|$ plus the running time of $\mathbb{B}$ on x. Article No. 4 proves the following.

Theorem 4 *For every sufficiently strong, computably enumerable theory T there exists an algorithm $\mathbb{A}_T$ such that:*

(a) $\mathbb{A}_T$ is as fast as any algorithm known by T, and T proves this.

(b) For every theory T^* containing T the following are equivalent:

(i) T^* proves Con_T.

(ii) T^* knows $\mathbb{A}_T$.

(iii) There exists an algorithm which T^* knows and proves to be as fast as any algorithm known by T.

Being sufficiently strong just means to contain a certain finite set of true arithmetical sentences. For suitable choices of Q, one can derive a version of Gödel's second incompleteness theorem. Somewhat curiously, one can then show that ZFC+Con_{ZFC} knows some algorithm which is faster than any algorithm ZFC knows. Here we call $\mathbb{A}$ *faster* than $\mathbb{B}$ if $\mathbb{A}$ is as fast as $\mathbb{B}$ but not vice-versa.

Spectra and descriptive complexity theory

Fagin's work on so-called generalized spectra are foundational for what became known as descriptive complexity theory. For a finite language L, the generalized spectrum $\mathrm{spec}_L(\varphi)$ of a first-order sentence φ is the set of L-structures A that have an expansion A' interpreting the language of φ such that $A' \models \varphi$. In other words $\mathrm{spec}_L(\varphi)$ is the set of models of the existential second-order sentence obtained from φ by existentially quantifying the symbols from the language of φ which are outside L. Fagin [22] showed that the set of these spectra "is" NP.

A binary string "is" an ordered structure interpreting (the order symbol and) a unary relation symbol. Conversely, an ordered finite structure can be canonically coded by a binary string, and thus input to a Turing machine. Thus computational problems "are" classes of finite ordered structures. For a class Q of L-structures (not necessarily ordered), one considers the class $Q^<$ of all expansions of structures in Q by a linear order. Fagin showed that for every $Q^<$ is in NP there is an L-sentence φ of existential second-order logic φ that defines Q, i.e. $A \models \varphi$ if and only if $A \in Q$. Note there is a nondeterministic *model checker* that given an ordered expansion of A and φ accepts exacty if $A \models \varphi$, and runs in time $|A|^{f(|\varphi|)}$ for some function f. This way, φ determines an algorithm witnessing $Q^< \in$ NP. Intuitively, existential second order logic can be seen as a programming language for NP problems.

The central problem of descriptive complexity theory is whether there is a logic *for* P: the logic should have a sentence φ defining Q for each $Q^< \in$ P and have a *deterministic* model checker with time complexity as above. Partial solutions have been obtained when restricting attention to structures from a class $\mathcal{K}$, then asking for a logic for P *on* $\mathcal{K}$: we require φ and the model checker to behave as desired only for inputs $A \in \mathcal{K}$. We refer to [20] for clean definitions.

The Immerman-Vardi theorem states that least fixed-point logic LFP is a logic for P on any class of ordered structures (see [20]). The deepest result in the area is due to Grohe who showed that LFP plus counting is a logic for P on any class $\mathcal{K}$ of graphs with excluded minors. We refer to [25] for a survey.

The KPT correspondence

The following proposition is a simple corollary to a beautiful theorem of Kechris, Pestov and Todorcevic [30]. We first explain this theorem.

Let L be a finite relational language and $\mathcal{K}$ an isomorphism closed class of finite L-strutures. Recall, $\mathcal{K}$ is *Fraïssé* if it is hereditary and has joint embedding and amalgamation (see e.g. [57, Sec. 4.4]). Such classes equal the set of isomorphic copies of finite substructures of their *Fraïssé limit*, a countable ultrahomogenous structure K (each isomorphism between finite substructures can be extended to an automorphism of K).

For $A, B \in \mathcal{K}$ let B^A denote the set of embeddings of A into B. The *Ramsey degree* of $A \in \mathcal{K}$ is the minimal $d \in \mathbb{N}$ such that for all $B \in \mathcal{K}$ and $k \geq 2$ we have $C \to (B)^A_{k,d}$. If no such d exists, the Ramsey degree is ∞. Here, $C \to (B)^A_{k,d}$ means that for every colouring χ of C^A with k colours there exists an embedding $b \in C^B$ such that χ takes at most d colours on the set $\{b \circ a \mid a \in B^A\}$. If $\mathcal{K}$ is Fraïssé with limit K, one can show that A has Ramsey degree $\leq d$ in $\mathcal{K}$ if and only if $K \to (B)^A_{k,d}$ for all $B \in \mathcal{K}$ and $k \geq 2$.

The Ramsey degree of $\mathcal{K}$ is the supremum of the Ramsey degrees of its members. If it is 1, then $\mathcal{K}$ is simply called *Ramsey.* Examples are the class of finite linear orders or the class of finite ordered graphs (see [30]).

Kechris, Pestov and Todorcevic [30] proved that a Fraïssé class $\mathcal{K}$ is Ramsey if and only if the automorphism group $G_{\mathcal{K}}$ of the Fraïssé limit K of $\mathcal{K}$ is *extremely amenable* (to be precise, [30] considers ordered classes and a slightly different definition of the Ramsey property). This means that every continuous action of $G_{\mathcal{K}}$ on a compact Hausdorff space (a $G_{\mathcal{K}}$*-flow*) has a fixed-point. This refers to the usual topology on $G_{\mathcal{K}}$ as a subset of K^K: the set of functions from K to K carrying the Tychonoff topology with K discrete.

Proposition 3 *If $\mathcal{K}$ is Fraïssé and Ramsey, then* LFP *is a logic for* P *on $\mathcal{K}$.*

Proof: Let K be the Fraïssé limit of $\mathcal{K}$. Let LO be the set of linear orders on K with the topology: for finite $A \subseteq K$ a basic open neighborhood of $R \in LO$ is the set of orders that agree with R on A. This makes LO compact and Hausdorff, and in fact a $G_{\mathcal{K}}$-flow with respect to the so-called logic action $(g, R) \mapsto g(R) := \{(g(x), g(y)) \mid (x, y) \in R\}$ of $G_{\mathcal{K}}$ on LO. By the theorem of Kechris, Pestov and Todorcevic, this action has a fixed point $R_0 \in LO$.

This means that $R_0 \subseteq K^2$ is invariant under automorphisms of K. Since K is ultrahomogenous and has a finite relational language, it is straightforward (and well-known how) to find a quantifier-free formula $\varphi(x, y)$ defining R_0 in K.

Every $A \in \mathcal{K}$ embeds into the Fraïssé limit K. Since φ is quantifier-free, it defines an order $\varphi(A)$ on A.

To show LFP is a logic for P on $\mathcal{K}$, let $Q^< \in$ P for Q a class of L-structures. The Immerman-Vardi theorem gives an LFP formula ψ in the language $L \cup \{<\}$

whose model class is $Q^<$. Then ψ is order-invariant in that it holds in either all or none ordered expansions of an L-structure A. Let ψ' be the LFP formula in the language L obtained from ψ by replacing $<$ by φ. Then for all $A \in \mathcal{K}$ we have $A \in Q$ if and only if $(A, \varphi(A)) \models \psi$, if and only if $A \models \psi'$. □

Kechris et al. [30] is the starting point of a series of works linking combinatorial properties of $\mathcal{K}$ to topological properties of $G_{\mathcal{K}}$. A recent, exciting example is [60], and [29] an earlier survey. One of the two main results of Article No. 5 is the following strengthening of the mentioned theorem from [30].

Theorem 5 *Let $d \in \mathbb{N}$ and $\mathcal{K}$ be Fraïssé. Then $\mathcal{K}$ has Ramsey degree at most d if and only if the universal minimal flow of $G_{\mathcal{K}}$ has size at most d.*

We refer to the article for a definition of the universal minimal flow. Here, we mention that being extremely amenable is equivalent to the universal minimal flow having size 1.

The algebraic approach to CSPs

Ramsey theory and the KPT correspondence shows up in some recent developments of constraint satisaftion theory [9]. See [8] for a survey.

For ease of presentation we assume all languages of structures to be relational. Given a structure M the *constraint satisfaction problem* $\textsc{Csp}(M)$ asks to decide the primitive positive theory of M. Recall, the *primitive positive* formulas are built from atoms by conjunction and existential quantification. Many important problems in computer science are naturally phrased as $\textsc{Csp}(M)$ for finite or countable ω-categorical M. See [5] for a survey on the latter.

The initial result of the area is Schäfer's dichotomy [54] stating that $\textsc{Csp}(M)$ for two-element M is in P or NP-complete. The central open problem is Feder and Vardi's [21] conjecture that this holds for all finite M. We talk of an area because this study developed an own body of concepts, methods and results following the so-called algebraic approach early promoted by Jeavons et al. [27].

This approach is based on a preservation theorem: if M is finite or countable ω-categorical, then a relation is primitive positive definable in M if and only if it is preserved by all polymorphisms of M [7]. An *(r-ary) polymorphism* of M is a homomorphism from M^r into M, and it *preserves* a relation R if it is a polymorphism of the expansion (M, R). The set of polymorphisms contains all projections and is closed under composition, so forms a *clone*.

If the polymorphism clone of M contains that of N, then every relation primitive positive definable in M is also so definable in N. This gives a simple reduction from $\textsc{Csp}(M)$ to $\textsc{Csp}(N)$. In this sense the tractablity of $\textsc{Csp}(M)$ depends only on the polymorphism clone of M and the quest for a complexity classification of a given family of CSPs reduces to a classification of clones. See the survey [11] for an elegant proof of Schäfer's theorem along these lines.

Families of interest are those $\textsc{Csp}(M)$ with M ranging over structures *definable* in some fixed structure K (the universe of M equals that of K and each relation of M is first-order definable in K). If K is the limit of a Fraïssé class $\mathcal{K}$

in a finite (relational) language, all these M are ω-categorical. If $\mathcal{K}$ is Ramsey, clone analysis is systematized through the use of so-called canonical functions – see the survey [8]. Here, the KPT correspondence gets useful. These are severe restrictions on K. But even if K is the pure set (i.e. it interprets the empty language) the resulting family of M, so-called *equality templates*, is rich.

Quantified CSPs are considerably harder to understand. Such a problem asks to decide the positive Horn theory of a structure M. Recall, *positive Horn* formulas are built from atoms by conjunctions and existential *and universal* quantification. Article No. 6 gives a preservation theorem:

Theorem 6 *Let M be countable ω-categorical. A relation over M is positive Horn definable over M if and only if it is preserved by all surjective periomorphisms of M.*

The concept of a periomorphism is introduced in Article No. 6. We define the *periodic power* M^{per} as the substructure of $M^{\mathbb{N}}$ consisting of all *periodic* $f \in M^{\mathbb{N}}$, i.e. those f such that for some $n > 0$ we have $f(k) = f(k \bmod n)$ for all k. It is possible to view M^{per} as a direct limit of the finite powers of M. By definition, a *periomorphism* of M is a homomorphism from M^{per} to M.

Based on this preservation theorem Article No. 6 gives a new proof of the complexity classification of quantified CSPs for equality templates [6]: every such problem is in L or NP-complete or coNP-hard.

CSPs seen from the other side

It is straightforward and well-known to associate with a primitive positive sentence φ in a relational language a structure A_φ such that an arbitary structure B satisfies φ if and only if there is a homomorphism from A_φ into B. Therefore, a CSP is sometimes defined as fixing B in the homomorphism problem:

Hom	
Instance:	two relational structures A and B.
Problem:	is there a homomorphism from A into B?

Now we want to classify the complexities of the problems obtained by restricting “the other side” (see the title of [24]), namely we demand A to come from a class of finite structures $\mathcal{A}$. The resulting problem is Hom($\mathcal{A}$).

Via the above reduction, $\mathcal{A}$ determines a subset of the primitive positive sentences. From the perspective of database theory [55], classifying the complexities of the problems Hom($\mathcal{A}$) means classifying the complexities of the problems obtained by putting syntactical restrictions on “conjunctive queries” φ to be evaluated over an arbitrary “database” B. From this perspective typical instances have a small query φ resp. A and a large B. The framework of parameterized complexity allows to take into account this asymmetry of the input, and design fast and practically useful algorithms even when Hom($\mathcal{A}$) is NP-complete. One aims to solve Hom($\mathcal{A}$) in *fpt time*, that is, time $f(|A|) \cdot |B|^{O(1)}$ on instance (A, B) where $f : \mathbb{N} \to \mathbb{N}$ is some (possibly fast growing) computable function.

We restrict attention to decidable classes $\mathcal{A}$ with a finite bound on the arities of relations appearing. Surprisingly, the complexity of Hom($\mathcal{A}$) turns out to be determined by the graph class $\mathcal{G}_{\mathcal{A}}$ of Gaifman graphs of cores of the structures in $\mathcal{A}$. The *Gaifman graph* of a structure A has the same universe as A and an edge between two points if they occur together in some tuple in some relation of A. A finite structure is a *core* if all its endomorphisms are automorphisms. Every finite structure A maps homomorphically to a core obtained from a substructure of A by deleting some tuples from some relations. This core is unique up to isomorphism and called the core *of* A.

Dalmau, Kolaitis and Vardi [16] showed that if $\mathcal{G}_{\mathcal{A}}$ has bounded treewidth, then Hom($\mathcal{A}$) can be solved in fpt (even polynomial) time. By definition, a graph has *treewidth* k if there exists a family of size $\leq k+1$ sets of vertices X_t associated with nodes t of some tree T such that every vertex and every edge of the graph appears in some X_t and for every vertex v the set $\{t \in T \mid v \in X_t\}$ is connected in T. Similarly one defines *pathwidth* using paths instead of trees.

Grohe [24] proved that the result of Dalmau et al. is optimal in the sense that, if $\mathcal{G}_{\mathcal{A}}$ has unbounded treewidth, then Hom($\mathcal{A}$) is not decidable in fpt time unless Hom(Cliques) is too (this would contradict a standard hypothesis in parameterized complexity theory [17]). The proof relies on a deep theorem of Robertson and Seymour [51] stating that a class of graphs of unbounded treewidth contains all grids as minors.

Article No. 7 asks which problems Hom($\mathcal{A}$) can be solved in parameterized logarithmic space, that is, space $f(|A|) + O(\log |B|)$ on instance (A, B) where again $f : \mathbb{N} \to \mathbb{N}$ is some computable function. This is to classical logarithmic space L as fpt time is to P. The corresponding notion of *pl reductions* is straightforwardly defined.

Then the main result of Article No. 7 reads as follows. It uses the notion of treedepth introduced by J. Nešetřil and P. Ossona de Mendez [42]. Roughly said, a graph class has bounded treedepth if it has bounded treewidth and we only need trees T of bounded height to witness this. Bounded treedepth implies bounded pathwidth, implies bounded treewidth.

Theorem 7 *Let $\mathcal{A}$ be a decidable class of structures with bounded arity.*

1. *If $\mathcal{G}_{\mathcal{A}}$ has bounded treewidth and unbounded pathwidth, then* Hom($\mathcal{A}$) *is equivalent to* Hom(coloured trees) *under pl reductions.*
2. *If $\mathcal{G}_{\mathcal{A}}$ has bounded pathwidth and unbounded treedepth, then* Hom($\mathcal{A}$) *is equivalent to* Hom(coloured paths) *under pl reductions.*
3. *If $\mathcal{G}_{\mathcal{A}}$ has bounded treedepth, then* Hom($\mathcal{A}$) *is decidable in parameterized logarithmic space.*

Closing Hom(coloured path) and Hom(coloured trees) under pl reductions defines interesting complexity classes between parameterized logarithmic space and fpt time. They have natural machine definitions and can be seen as parameterized analogues of the classical classes NL and LOGCFL between L and P.

Bibliography

[1] M. Ajtai. The complexity of the pigeonhole principle. Proceedings of the 29th Annual Symposion on the Foundations of Computer Science, pp. 346–355, 1988.

[2] G. Asser. Das Repräsentenproblem in Prädikatenkalkül der ersten Stufe mit Identität. Zeitschrift für mathematische Logik und Grundlagen der Mathematik 1: 252-263, 1955.

[3] P. Beame, J. Krajíček, R. Impagliazzo, T. Pitassi, P. Pudlák and A. Woods. Exponential Lower Bound for the Pigeonhole Principle. Proceedings of the ACM Symposium on Theory of Computing, ACM Press, pp. 200-220, 1992.

[4] S. Ben-David and A. Gringauze. On the existence of optimal propositional proof systems and oracle-relativized propositional logic. Electronic Colloquium on Computational Complexity, Report TR98-021, 1998.

[5] M. Bodirsky. Constraint Satisfaction Problems with Infinite Templates. In N. Creignou, P. G. Kolaitis and H. Vollmer (eds.), Complexity of Constraints - An Overview of Current Research Themes, Lecture Notes in Computer Science 5250, 2008.

[6] M. Bodirsky and H. Chen. Quantified equality constraints. SIAM Journal on Computing 39 (8): 3682-3699, 2010.

[7] M. Bodirsky and J. Nešetřil. Constraint satisfaction with countable homogeneous templates. Journal of Logic and Computation 16 (3): 359-373, 2006.

[8] M. Bodirsky and M. Pinsker. Reducts of Ramsey structures In M. Grohe, J. A. Makowsky, Model Theoretic Methods in Finite Combinatorics, AMS Contemporary Mathematics 558, pp. 489-519, 2011.

[9] M. Bodirsky, M. Pinsker and Tsankov. Decidability of definability. Journal of Symbolic Logic 78 (4): 1036-1054, 2013.

[10] S. R. Buss. On Gödel's theorems on lengths of proofs II: Lower bounds for recognizing k symbol provability. In Feasible Mathematics II, P. Clote and J. Remmel (eds), Birkhauser, pp. 57-90, 1995.

[11] H. Chen. A rendezvous of logic, complexity, and algebra. ACM Computing Surveys 42 (1), Article 2, 2009.

[12] S. Cook and P. Nguyen. Logical Foundations of Proof Complexity. Perspectives in Logic, Cambridge University Press, 2010.

[13] S. A. Cook and R. A. Reckhow. The relative efficiency of propositional proof systems. J. Symb. Logic 44(1):36–50, 1979.

[14] Y. Chen and J. Flum. On slicewise monotone parameterized problems and optimal proof Systems for TAUT. Proceedings of the 19th EACSL Annual Conference on Computer Science Logic, Lecture Notes in Computer Science 6247, pp. 200-214, 2010.

[15] Y. Chen and J. Flum. On the complexity of Gödel's proof predicate. The Journal of Symbolic Logic 75 (1): 239-254, 2010.

[16] V. Dalmau, Ph. G. Kolaitis, and M. Y. Vardi. Constraint satisfaction, bounded treewidth, and finite-variable logics. Proceedings of the 8th International Conference on Principles and Practice of Constraint Programming, Lecture Notes in Computer Science 2470, pp. 310-326, Springer, 2002.

[17] R. G. Downey and M. R. Fellows. Fixed-parameter tractability and completeness II: On completeness for W[1]. Theoretical Computer Science, 141: 109-131, 1995.

[18] R. G. Downey and M. R. Fellows. Fundamentals of Parameterized Complexity. Springer, 2013.

[19] A. Durand, N. D. Jones, J. A. Makowsky and M. More. Fifty years of the spectrum problem: survey and new results. Bulletin of Symbolic Logic 18 (4): 505-553, 2012.

[20] H.-D. Ebbinghaus, J. Flum. Finite Model Theory, 2nd edition. Perspectives in Mathematical Logic, Springer, 1999.

[21] T. Feder and M. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory. SIAM Journal on Computing 28: 57-104, 1999.

[22] R. Fagin. Generalized first-order spectra and polynomial-time recognizable sets. In R. M. Karp (ed.), Complexity of Computation, SIAM-AMS Proceedings 7, pp.43-73, 1974.

[23] M. L. Furst, J. B. Saxe and M. Sipser. Parity, circuits, and the polynomial-time hierarchy. In Proceedings of 22nd Annual Symposium on Foundations of Computer Science, pp. 260-270, 1981.

[24] M. Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. Journal of the ACM 54 (1), Article No. 1, 2007.

[25] M. Grohe. From polynomial time queries to graph structure theory. Communications of the ACM 54 (6): 104-112, 2011.

[26] A. Haken. The intractability of resolution. Theoretical Computer Science 39: 297-308, 1985.

[27] P. G. Jeavons, D. A. Cohen and M. Gyssens. A Unifying Framework for Tractable Constraints. Proceedings of the 1st International Conference on Principles and Practice of Constraint Programming, Lecture Notes in Computer Science 976, pp. 276-291. 1995.

[28] N. D. Jones and A. L. Selman. Turing machines and the spectra of first-order formulas. The Journal of Symbolic Logic 39: 139-150, 1974.

[29] A. S. Kechris. Dynamics of non-archimedean Polish groups. European Congress of Mathematics, Krakow, 2-7 July, 2012, pp. 375-397, R. Latala et al. (eds.), European Mathematical Society, 2014.

[30] A. S. Kechris, V. Pestov and S. Todorcevic. Fraïssé limits, Ramsey theory, and topological dynamics of automorphism groups. Geometric and Functional Analysis 15: 106-189, 2005.

[31] J. Köbler and J. Messner. Complete problems for promise classes by optimal proof systems for test sets. In Proceedings of the 13th IEEE Conference on Computational Complexity, pp. 132-140, 1998.

[32] J. Krajíček. Bounded Arithmetic, Propositional Logic, and Complexity Theory. Encyclopedia of Mathematics and its Applications, Cambridge University Press, 1995.

[33] J. Krajíček. Proof complexity. In A. Laptev (ed.), European congress of mathematics (ECM), Stockholm, Sweden, June 27–July 2, 2004. European Mathematical Society, pp. 221-231, 2005.

[34] J. Krajíček. Forcing with random variables and proof complexity, vol. 382. London Mathematical Society Lecture Note Series, Cambridge University Press, 2011.

[35] J. Krajíček. Expansions of pseudofinite structures and circuit and proof complexity. arXiv:1505.00118 [math.LO], 2015.

[36] J. Krajíček and P. Pudlák. Propositional proof systems, the consistency of first order theories and the complexity of computations. Journal of Symbolic Logic 54 (3): 1063-1079, 1989.

[37] J. Krajíček, P. Pudlák and A. Woods. An exponential lower bound to the size of bounded depth frege proofs of the pigeonhole principle. Random Structures and Algorithms **7**(1), pp. 15-39. 1995.

[38] K. Kunen. Set Theory. Studies in Logic 34, College Publications, London, revised edition, 2013.

[39] L. Levin. Universal sequential search problems. Problems of Information Transmission 9 (3): 265-266, 1973.

[40] A. Maciel, T. Pitassi, and A. R. Woods. A new proof of the weak pigeonhole principle. Journal of Computer and Systems Sciences 64 (4): 843–872, 2002.

[41] J. Messner. On the Simulation Order of Proof Systems. Ph.D. Dissertation. University of Ulm, 2000.

[42] J. Nešetřil and P. Ossona de Mendez. Tree depth, subgraph coloring, and homomorphism bounds. Euro- pean Journal of Combinatorics 27 (6):1022-1041, 2006

[43] J. Nordström. On the interplay between proof complexity and SAT solving. ACM SIGLOG News 2 (3): 19-44, 2015.

[44] J. B. Paris, A. J. Wilkie and A. R. Woods. Provability of the pigeonhole principle and the existence of infinitely many primes. Journal of Symbolic Logic 53 (4): 1235-1244, 1988.

[45] P. Pudlák. A bottom-up approach to foundations of mathematics. Proceedings Gödel'96, Logical Foundations of Mathematics, Computer Science and Physics – Kurt Gödel's Legacy, Springer Lecture Notes in Logic **6**, pp. 81-97, 1996.

[46] P. Pudlák. Proofs as games. American Mathematical Monthly, pp. 541-550, 2000.

[47] A. A. Razborov. Resolution lower bounds for the weak functional pigeonhole principle. Theoretical Computer Science 303 (1): 233-243, 2003.

[48] A. A. Razborov. Pseudorandom generators hard for k-DNF resolution and polynomial calculus resolution. Annals of Mathematics 181 (2): 415-472, 2015.

[49] S. Riis. A complexity gap for tree-resolution. Computational Complexity 10: 179-209, 2001.

[50] S. Riis and M. Sitharam. Generating hard tautologies using predicate logic and the symmetric group. Logic Journal of the IGPL 8 (6): 787-795, 2000.

[51] N. Robertson and P. D. Seymour. Graph minors V. Excluding a planar graph. Journal of Combinatorial Theory, Series B, 41: 92-114, 1986.

[52] H. Scholz. Ein ungelöstes Problem in der symbolischen Logik. The Journal of Symbolic Logic 17, p. 160, 1952.

[53] N. Segerlind, S. Buss, and R. Impagliazzo. A switching lemma for small restrictions and lower bounds for k-DNF resolution. SIAM Journal on Computing 33 (5): 1171-1200, 2004.

[54] T. J. Schaefer. The complexity of satisfiability problems. Proceedings of the 10th ACM Symposium on Theory of Computing, pp. 216-226, 1978.

[55] N. Schweikardt, T. Schwentick, and L. Segoufin. In Database Theory: Query Languages. In M. J. Atallah and M. Blanton (eds.), Algorithms and Theory of Computation Handbook (2nd ed.), Vol. 2: Special Topics and Techniques. CRC Press, Chapter 19, 2009

[56] A. S. Troelstra and H. Schwichtenberg. Basic Proof Theory. 2nd Edition, Cambridge University Press, 2000.

[57] K. Tent and M. Ziegler. A Course in Model Theory. Cambridge University Press, ASL, Lecture Notes in Logic, 2012.

[58] R. Thiele. Hilbert's twenty-fourth problem. American Mathematical Monthly, January 2003.

[59] A. Urquhardt. Von Neumann, Gödel and complexity theory. The Bulletin of Symbolic Logic 16 (4): 516-530, 2010.

[60] A. Zucker. Topological dynamics of automorphism groups, ultrafilter combinatorics, and the generic point problem. Transactions of the American Mathematical Society, published electronically November 16, 2015.