prof. dr hab. Mikołaj Bojańczyk
Instytut Informatyki
Uniwersytet Warszawski

Warszawa, May 19, 2023

Review of PhD Thesis,
Complexity of solving equations in finite algebraic structures
by Piotr Kawałek

The present thesis is about solving equations. It goes without saying that this is an important topic.

The main idea behind the program pursued in this thesis is to study the computational complexity of solving equations in finite algebras. For example, solving equations in the Boolean algebra $(\{0,1\}, \vee, \wedge)$ is the same as satisfiability for Boolean formulas, and this is the canonical NP-complete problem. On the other hand, solving equations in a finite commutative group is computationally easy, with a polynomial time solution using Gaussian elimination. What about other finite algebras?

This kind of question makes a lot of sense. There are many finite algebras, so clearly the problem setting is rich, but they do have a general theory developped by McKenzie and many other in the universal algebra community (including important contributions from the Kraków group). As a source of inspiration for this question, one can consider the success story behind the Feder and Vardi dichotomy conjecture for constraint satisfaction problems (CSP). This conjecture says that every CSP problem is either in P or NP-complete, and intermediate complexities cannot arise. It has been observed early on that the dichotomy conjecture is essentially a problem about classifying finite algebras, to which one can apply deep tools of universal algebra. The dichotomy conjecture was recently proved by Bulatov and Zhuk, and is a success story in theoretical computer science, in particular a success story for universal algebra.

As far as I can tell, the program pursued in this thesis was launched by Idziak and Krzaczkowski in a LICS 2016 paper, where they proposed to study satisfiability for equations over abstract finite algebras. An important prior source of inspiration for this program was a paper of Goldmann and Russell from 2002, which gave sufficient and necessary conditions for polynomial time solvability of equations in the important case of finite algebras that are groups. The conditions lower and upper bounds of Goldmann and Russell did not quite match, and in particular the grey area between them (the "solvable-versus-nilpotent"

gap) is one of the main topics of the present thesis, mainly in the context of general abstract, but also with a number of results that are purely about groups. It is worth mentioning that the format of equations used by Goldmann and Russell (called PolSat) is not quite the same as the format used by Idziak and Krzaczkowski (CSAT), with the latter format seeming to make more sense, especially for abstract algebras.

As mentioned above, the program to classify the computational complexity of solving equations in abstract finite algebras is a very natural idea, and a hihgly ambitious project. There are, however, two limitations on the program that could arguably be seen as making it slightly less natural.

1. Usually when thinking of solving equations, one thinks about solving systems of equations. It turns out that solving systems of equations over a given algebra can be reduced to solving CSP's, and therefore the dichotomy theorem of Bulatov and Zhuk can be used to give a full classification of finite algebras for which solving systems of equations is in polynomial time. Therefore, since that problem is solved, one is left with investigating the arguably less interesting problem of solving individual equations; this is the problem from this thesis.

   A further issue with solving individual equations is that the complexity of the problem is somewhat sensitive to the way that the equations are presented. In particular, the thesis uses at least three variants, namely PolSat and CSAT, and, to a lesser extent, Barrington's non-uniform automata for groups, which are also discussed in the thesis. To the non-expert reader, such as myself, the multitude of versions can be confusing. It does not help that the thesis is simply a collection of conference papers, without any significant attempt at unification.

2. The class of all finite algebras does have a lot of highly non-trivial theory, but apparently that theory is not sufficient for the purposes of the present project. For this reason, as far as I can tell, all results in this thesis are about finite algebras that generate congruence modular varieties. This is a certain technical assumption which is satisfied by many finite algebras, in particular groups or rings, but it is not satisfied by some important finite algebas, such as many semigroups. Under this assumption, one can apply a sophisticated theory of commutators developped by Frieze and McKenzie. That being said, requiring this assumption does somewhat diminish the appeal of the program.

Results in the thesis

The thesis contains a wealth of material, distributed across seven papers, which I describe below.

Expressive Power, Satisfiability and Equivalence of Circuits over Nilpotent Algebras. The first paper in the thesis was published at MFCS 2018,

and its authors are Idziak, Kawałek, and Krzaczkowski. (Krzaczkowski is a co-author on all papers in the thesis, Idziak comes in most of the time, and the two papers about groups also have Weiss.) The paper is about satisfiability and equivalence problems, where the input uses circuit representation (the CSAT variant mentioned above); i.e. the size of an expression is the number of distinct sub-expressions. This is a natural size measure; in particular it guarantees that the complexity of the decision problems is invariant under the choice of basic operations. Such invariance does not hold for the other size measure PolSAT, in which the size of an expression is the number of symbols used in it (the distinction between the two size measures is the same as circuit size vs formula size in computational complexity). For example, in the PolSAT measure, the complexity might change if we allow the commutator operation $xyx^{-1}x-1$; such a cituation does not arise under the CSAT measure. In this sense, the circuit measure seams to be the "right" one to use, especially for abstract algebras which need not have a chosen set of operations enshrined by centuries of tradition.

The algebras in question are finite algebras with a certain extra assumption: the variety that they generate is congruence modular. This extra assumption is an important property in universal algebra; in particular it enables the use of results of Frieze and McKenzie from their book "Commutator theory for congruence modular varieties". Many results from this thesis are based on a characterisation of Frieze and McKenzie about nilpotent algebras with the congruence modular property, which allow a certain wreath product decomposition of the algebras in question; the essence of the MFCS 2019 paper is a (non-trivial) analysis of the consequences of the characterisation for the decision problems about equations. The wreath product decomposition describes the algebras in question using one module acting on the other. In the case mainly studied by the MFCS 2019 papers, are two modules, with one acting on the other. Theorems 4.2 and 5.2 show polynomial time algorithms for circuit equivalence and circuit satisfiability, under the assumption that the two modules have a special kind – they are 1-dimensional vector spaces over prime fields with different chartacteristics. The algorithms depend on the operations in the algebra being given by their multiplication tables. In a different – and arguably less natural setup – where the operations are given by Turing machine oracles that run in polynomial time, the problems can get higher complexities, such as NP-complete or coNP-complete.

Even Faster Algorithms for CSAT over Supernilpotent Algebras. The second paper in the thesis was published at MFCS 2020, and its authors are Kawałek and Krzaczkowski. This paper shows how to solve the CSAT problem for supernilpotent algebras from congruence modular varieties. Supernilpotent are a class of algebras that were previously known to have equations solvable in deterministic polynomial time. The contribution of the paper is that: (a) the degree of the polynomial in the running time of the algorithm can be improved; and (b) if a randomized Monte Carlo algorithm is allowed, then linear time is possible. I presume that the title is a bit tounge-in-cheek with regards to the

deterministic algorithm, since the degree of the polynomial in the running time is still rather high, but who am I to judge, given that most algorithms in own research are at least exponential.

The randomized algorithm in (b) arises from a reduction of solving equations in the algebra to solving equations of polynomials in a field, such as:

$$\underbrace{3x^2 + 2xy^5z + x^6yz^3 = 2.}_{\text{we want to know if this equation has at least one solution in some fixed finite field, such as } \mathbb{F}_7}$$

The algorithms in (a) and the reduction used in (b) are based on results of Freese, McKenzie and Kearnes, which show how supernilpotent algebras can be decomposed into wreath products of modules. Thanks to those results, the contribution of the paper is devoted to an analysis of solving equations in wreath products of modules. The linear time Monte Carlo algorithm then arises from the observation that if a constant degree polynomial in a finite field achieves a given value for at least one input, then it achieves that value for a constant fraction of inputs. I would be curious to know if this observation (Lemma 1.4, given without proof, like many other results in the thesis) is not some folklore result.

Intermediate problems in modular circuits satisfiability. The third paper in the thesis was published at LICS 2020, and its authors are Idziak, Kawałek and Krzaczkowski. This paper, like the MFCS 2018 paper, studies the CSAT problem for algebras that are nilpotent but not supernilpotent. The paper uses a parameter $h \in \{1, 2, \ldots\}$, which quantifies how such an algebra fails to be supernilpotent. An algebra with this parameter can be decomposed as a wreath product of h layers of algebras that use only modulo counting; in each layer the field for modulo counting is distinct from the fields in the adjacent layers. This decomposition is the main property used in the analysis (as was the case in the MFCS 2018 paper). The main contribution of the paper is that the complexity of solving equations in such algebras is the between the quasi-polynomials

$$2^{c \log^{h-1}(\text{size of circuit describing the equation})}$$

and

$$2^{c \log^{h}(\text{size of circuit describing the equation})}.$$

The bound is not exactly tight, because the lower bound uses $h-1$ and the upper bound uses h, but it is close; furthermore it can be made tight if randomized algorithms are allowed. This estimate on the complexity is conditional, with the lower bound using the standard ETH hypothesis, and the upper bound using a hypothesis called SESH that I was unfamiliar with before, and which seems rather tailored to the problem under consideration. The proofs use a highly non-trivial analysis of circuits with modulo gates; in particular a clever way of constructing an and gate using limited resources. (Constructing and gates is the main lower bound technique throughout the thesis.) The bounds established

in this paper can be seen as evidence that there is no clean classification for the computational complexity of the CSAT problem, unlike the CSP problem which is subject to the famous dichotomy proved by Bulatov and Zhuk. Also, the techniques in the paper are applied to give the first example of a solvable group whose PolSat problem is not in P, under ETH. This example uses PolSat, which is the variant of CSAT with the more verbose, non-circuit, format of equations. The example refutes a conjecture of Horvath (and previously Burris and Lawrence), assuming ETH.

Equation satisfiability in solvable groups and **Satisfiability problems for finite groups.** These two papers were published in 2022, in the journal Theory of Computing Systems and in the conference ICALP, respectively. In both cases, the authors are Idziak, Kawałek, Krzaczkowski and Weiss. The journal paper is the only journal paper in the thesis, the remaining papers being refereed conference publications. Both papers are about the classical setting of finite groups, as opposed to the more general setting of finite algebras. (Although some of the lemmas are proved in the more general case of finite algebras, using tame congruence theory.) Both papers provide lower bounds for satisfiability problems in finite groups, conditionally on ETH.

The main contribution of the first ToCS is a strengthening of the counter-example from the previous paper: assuming ETH, PolSat is not in P for every group that has Fitting length at least 3 (Fitting length is referred to as supernilpotent rank in other parts of the thesis). The ICALP paper provides a further series of hardness results, which mainly use yet another variant of satisfiability, namely the ProgramSAT model of Barrington, and its variants. These papers were the hardest for me to asses, all I can say is that they are highly involved technically.

Satisfiability of Circuits and Equations over Finite Malcev Algebras This paper was published at STACS 2022, and its authors are Idziak, Kawałek and Krzaczkowski.

A rather direct corollary of tame congruence theory is that if an algebra is not solvable, then there is a quotient algebra for which satisfiability of equations is NP-complete. This is because a non-solvable algebra must have, by a definition, a pair covering pair of congruences which induces – in the sense of tame congruence theory – a copy of the Boolean algebra $(\{0,1\}, \vee, \wedge)$. The fact that satisfiability for Boolean circuits is NP-complete, and the way in which the copy is encoded in the algebra, leads to a quotient with NP-complete solving of equations. Prior work of Idziak and Krzaczkowski also showed that for Malcev algebras, non-nilpotent algebras will necessarily have a quotient with NP-complete CSAT. The present paper improves this, by showing that the quotient is not necessary, and CSAT is NP-complete for the algebra itself. Further results in this paper use the nilpotent rank, and show that CSAT will not be in P (without necessarily being NP-complete) already when this rank is at least three, assuming ETH. The proofs, as in all papers from the thesis, are hard. In this particular paper,

the techniques are based mainly on the advanced toolbox of universal algebra.

Complexity of Modular Circuits This paper was published at LICS 2022, and its authors are Idziak, Kawałek, and Krzaczkowski. The paper considers a variant of AC^0 circuits, which use modulo gates instead of and/or gates. (This variant appeared already before in the thesis.) The classes of circuits under consideration are denoted by $CC_h[m]$: the height of the circuits is h and the circuits can use gates that count modulo m. The circuits are not allowed to use the classical and/or gates; in fact the lower bound proofs are based on showing how and/or can be simulated, and in a certain sense this is necessarily so (see Theorem 6.1). The main contribution of the paper is that, under ETH, the satisfiability problem (i.e. given a $CC_h[m]$ circuit, decide if there is some satisfying assingment to its input gates) is in PTIME if and only if the height is $h = 1$ or the modulus m is a power of a single prime. For example the modulus $m = 5^7$ is good for any height, but modulus $m = 3 \cdot 5$ is not good unless the trivial height $h = 1$ is used. The lower bound is based on a combination of results from Barrington et al. and on constructions from the MFCS 2018 paper in the present thesis. The upper bound is proved by showing that circuits which use a modulus m that is a power of a single prime will necessarily enjoy a certain "balanced" property (roughly the number of "yes" outputs cannot be very different from the number of "no" outputs), and this property disallows expressing a large and gate, which ensures a simple satisfiability algorithm. The paper also contains some other, rather technical results, including stronger lower bounds that relate to recent work on modulo circuits by Chapman and Williams, as well as a strengthening of the counter-example for Horvath's conjecture that was given in the LICS 2020 paper.

Final comments. What about the success of the project that was purused in this thesis? One can say that this project was influenced by the CSP classification project, and what makes the CSP classification project so attractive is that there is a clean dichotomy between P and NP-complete. After several years of research, with much of it in the present thesis, it seems to be the case that for solving equations there is no such clean dichotomy, or trichotomy, etc. As far as I can tell, it could be the case that the project is converging to a series of every more sophisticated results that tighten the gap between known upper and lower bounds, but this gap might never be closed.

Evaluation

This thesis is highly advanced technically; it could be the most difficult thesis among those that I have reviewed. (Although I would have to say that being difficult is not the ultimate goal in successful mathematics.) It draws one a wide body of techniques, from universal algebra, group theory and computational complexity. All papers in the thesis have co-authors, with Krzaczkowski being a co-author for all papers, and Idziak being a co-author in all papers with one

exception. I do not know about the exact extent to which the author contributed independently to the results, but based on my conversations with him, I can assume that it was a significant one. I would like to congratulate the author on this impressive body of work.

Apart from the intrinsic difficulty of the results, reading the thesis is made harder by its format. It is simply a collection of papers, all of which (with one exception) are conference papers. Conference papers have strict page limits, which enforces a very compressed style, with proofs relegated to appendices. (The author chose not to include the appendices from the conference papers.) Also, different papers use different notation, which makes it hard to spot the redundancies and repetition in the ensemble. For these reasons, I found it very challenging to understand the results, not to mention checking their correctness. Fortunately, I had the opportunity to personally meet the author. During this meeting, we spent several hours discussing the results and proofs. These discussions were very satisfactory and enabled me to more successfully navigate the written document.

As I mentioned before, this is a very advanced thesis with technicaly deep results. The underlying papers were published in selective venues, with two papers at the A* conference LICS, as well as a papers at STACS, ICALP, MFCS(2) and a journal paper in Theory of Computing Systems. This is far above the usual requirements for a PhD thesis. For these reasons, I recommend without hesitation to accept this thesis.

Mikołaj Bojańczyk