

On the van der Waerden numbers $w(2; 3, t)$

Tanbir Ahmed^a, Oliver Kullmann^b, Hunter Snevily^{c,1}

^a*Department of Computer Science and Software Engineering, Concordia University, Montréal, Canada.
ta_ahmed@cs.concordia.ca*

^b*Computer Science Department, Swansea University, Swansea, UK.
O.Kullmann@Swansea.ac.uk*

^c*Department of Mathematics, University of Idaho - Moscow, Idaho, USA.*

Abstract

In this paper we present results and conjectures on the ordinary van der Waerden numbers $w(2; 3, t)$ and on the new *palindromic van der Waerden numbers* $\text{pdw}(2; 3, t)$. We have computed the exact value of the previously unknown number $w(2; 3, 19) = 349$, and we provide new lower bounds for $20 \leq t \leq 39$, where for $20 \leq t \leq 30$ we conjecture these bounds to be exact. The lower bounds for $w(2; 3, t)$ with $24 \leq t \leq 30$ refute the conjecture that $w(2; 3, t) \leq t^2$ as suggested in [14]. Based on the known values of $w(2; 3, t)$, we investigate regularities to better understand the lower bounds of $w(2; 3, t)$. Motivated by such regularities, we introduce palindromic van der Waerden numbers $\text{pdw}(k; t_0, \dots, t_{k-1})$, which are defined as the ordinary numbers $w(k; t_0, \dots, t_{k-1})$, but where only palindromic solutions are considered, reading the same from both ends. Different from the situation for ordinary van der Waerden numbers, these “numbers” need actually to be pairs of numbers. We compute $\text{pdw}(2; 3, t)$ for $3 \leq t \leq 27$, and we provide bounds for $t \leq 39$, which we believe to be exact for $t \leq 35$. All computations are based on SAT solving, and we discuss the various relations between SAT solving and Ramsey theory. Especially we introduce a novel (open-source) SAT solver, the `tawSolver`, which performs best on the SAT instances studied here, and which is actually the original DLL-solver ([18]), but with an efficient implementation and a modern heuristic typical for look-ahead solvers (applying the theory developed in [48]).

Contents

1	Introduction	2
1.1	Using SAT solvers	4
1.1.1	Informed versus uninformed SAT solving	5
1.1.2	Parallel/distributed SAT solving	5
1.1.3	Synergies between Ramsey theory and SAT	6
1.2	The results of this paper	7
2	The <code>tawSolver</code>	8
2.1	The basic structure	8
2.2	Look-ahead solvers	9
2.3	From <code>tawSolver-1.0</code> to <code>tawSolver-2.6</code>	9
2.4	The implementation	11
2.5	The optimal projection: the τ -function	12

¹Hunter Snevily passed away on November 11, 2013 after his long struggle with Parkinsons disease. He was an inspiring mathematician. We have lost a great friend and colleague. He will be heavily missed and fondly remembered

3	Computational results on $w(2; 3, t)$	12
3.1	$w(2; 3, 19) = 349$	12
3.2	Some new conjectures	13
3.3	A conjecture on the upper bound	14
4	Patterns in the good partitions	14
4.1	Number of 0's and 00's	15
4.2	Number of 1's	15
4.3	How can it help for SAT solving?	16
5	Palindromes	17
5.1	Palindromic vdW-hypergraphs	18
5.2	Precise values	20
5.3	Conjectured values and bounds	21
5.4	Open problems	21
5.5	Remarks on the use of symmetries	22
6	Experiments with SAT solvers	22
6.1	Complete solvers	22
6.1.1	Cube-and-Conquer	23
6.1.2	VdW-problems	24
6.1.3	Palindromic vdW-problems	26
6.2	Incomplete solvers (stochastic local search)	28
7	Conclusion	29
Appendix A	Certificates	32
Appendix A.1	Conjectured precise lower bounds for $w(2; 3, t)$	32
Appendix A.2	Further lower bounds for $w(2; 3, t)$	33
Appendix A.3	Good palindromic partitions	34
Appendix B	Using the OKlibrary	37
Appendix B.1	Numbers and certificates	37
Appendix B.2	Hypergraphs	40
Appendix B.3	SAT instances	41
Appendix B.4	The SAT solvers	42
Appendix B.4.1	tawSolver	42
Appendix B.4.2	satz	44
Appendix B.4.3	march_pl	45
Appendix B.4.4	OKsolver	45
Appendix B.4.5	Ubcsat	46
Appendix B.5	Running experiments	47

1. Introduction

We consider Ramsey theory and its connections to computer science (see [58] for a survey) by exploring a rather recent link, especially to algorithms and formal methods, namely to SAT solving. SAT is the problem of finding a satisfying assignment for a propositional formula. Since Ramsey problems can naturally be formulated as SAT problems, SAT solvers can be used to compute numbers from Ramsey theory. In the present article, we consider van der Waerden numbers ([69]), where SAT had its biggest success in Ramsey theory, namely the determination of $w(2; 6, 6) = 1132$ in [43], the first new diagonal van der Waerden (short “vdW”) number after almost 30 years.

Definition 1.1. We use $\mathbb{N} = \{x \in \mathbb{Z} : x \geq 1\}$, $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$. An *arithmetic progression* of length $t \in \mathbb{N}$ is a subset $p \subset \mathbb{N}$ of length $|p| = t$ and of the form $p = \{a + i \cdot d : i \in \{0, \dots, t-1\}\}$ for some $a, d \in \mathbb{N}$. A *block partition* of length $k \in \mathbb{N}$ of a set X is a tuple (P_0, \dots, P_{k-1}) of length k of subsets of X (possibly empty) which are pairwise disjoint ($P_i \cap P_j = \emptyset$ for $i \neq j$) and with $P_0 \cup \dots \cup P_{k-1} = X$. The *van der Waerden number* $w(k; t_0, t_1, \dots, t_{k-1}) \in \mathbb{N}$ for $k, t_0, \dots, t_{k-1} \in \mathbb{N}$ is the smallest $n \in \mathbb{N}$ such that for any block partition (P_0, \dots, P_{k-1}) of length k of $\{1, \dots, n\}$ there exists a $j \in \{0, \dots, k-1\}$ such that P_j contains an arithmetic progression of length t_j .

That we have $w(k; t_0, t_1, \dots, t_{k-1}) > n$ can be certified by an appropriate block partition of $\{1, \dots, n\}$; such partitions are the solutions of the SAT problems to be constructed, and we call them “good partitions”:

Definition 1.2. A *good partition* of $\{1, \dots, n\}$ (where $n \in \mathbb{N}_0$) w.r.t. parameters t_0, t_1, \dots, t_{k-1} is a block partition (P_0, \dots, P_{k-1}) of $\{1, \dots, n\}$ containing no block P_j with an arithmetic progression of length t_j (for any j).

So there exists a good partition of $\{1, \dots, n\}$ if and only if $n < w(k; t_0, t_1, \dots, t_{k-1})$. For every $k, t_0, \dots, t_{k-1} \in \mathbb{N}$ the only block partition of $\{1, \dots, 0\} = \emptyset$ is $(\emptyset, \dots, \emptyset)$, and this is a good partition. In this paper, we are interested in the specific van der Waerden numbers $w(2; 3, t)$, $t \geq 3$. Specialising the general definition we obtain:

$w(2; 3, t)$ is the smallest $n \in \mathbb{N}$, such that
for all $P_0, P_1 \subseteq \{1, \dots, n\}$ with $P_0 \cap P_1 = \emptyset$ and $P_0 \cup P_1 = \{1, \dots, n\}$
either P_0 has an arithmetic progression of size 3 or P_1 has an arithmetic progression of size t , or both.

The known exact values of $w(2; 3, t)$ are shown in Table 1 (with our contribution in bold).

t	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
$w(2; 3, t)$	9	18	22	32	46	58	77	97	114	135	160	186	218	238	279	312	349

Table 1: Known values for $w(2; 3, t)$

As references and for relevant information on the above numbers, see Chvátal [15], Brown [13], Beeler and O’Neil [7], Kouril [43], Landman, Robertson and Culver [53], and Ahmed [2, 3, 4, 5].² Recently, Kullmann [49]³ reported the following lower bounds

$$w(2; 3, 19) \geq 349, w(2; 3, 20) \geq 389, w(2; 3, 21) \geq 416.$$

We confirm the exact value of $w(2; 3, 19) = 349$, and we extend the list of lower bounds up to $t = 39$. Brown, Landman, and Robertson [14], showed the lower bound $w(2; 3, t) > t^{2-1/\log \log t}$ for $t \geq 4 \cdot 10^{316}$, and observed that $w(2; 3, t) \leq t^2$ for $5 \leq t \leq 16$, suggesting that this might hold for all t . Our lower bounds in Subsection 3.2 however prove that there are t with $w(2; 3, t) > t^2$. We provide an improved upper bound $1.675t^2$ in Subsection 3.3 (satisfying all known values and lower bounds of $w(2; 3, t)$).

We also present a new type of van-der-Waerden-like numbers, namely *palindromic number-pairs*, obtained by the constraint on good partitions that they must be symmetric under reflection at the mid-point of the interval $\{1, \dots, n\}$. Perceived originally only as a heuristic tool for studying ordinary van der Waerden numbers, it turned out that these numbers are interesting objects on their own. An interesting phenomenon is that we no longer have the standard behaviour of the SAT instances with increasing n , where

- first all instances are satisfiable (for $n < w(k; t_0, \dots, t_{k-1})$), and from a certain point on (the van der Waerden number) all instances are unsatisfiable (for $n \geq w(k; t_0, \dots, t_{k-1})$),
- but now first again all instances are satisfiable (for $n \leq p$), then we have a region with strict alternation between unsatisfiability and satisfiability, and only from a second point on all instances are unsatisfiable (for $n \geq q$).

²This sequence is <http://oeis.org/A007783> in the “On-Line Encyclopedia of Integer Sequences”.

³the conference article [50] contains only material related to Green-Tao numbers and SAT

These two turning points constitute the palindromic “number” $\text{pdw}(2; 3, t) = (p, q)$ as pairs of natural numbers. We were able to compute $\text{pdw}(2; 3, t)$ for $t \leq 27$. We also provide (conjectured) values for $t \leq 39$.⁴ The full definition is in Section 5, while the special case experimentally studied in this paper is defined as follows:

In $\text{pdw}(2; 3, t) = (p, q)$,
the number q is the smallest number such that for all $n \geq q$ and
for all $P_0, P_1 \subseteq \{1, \dots, n\}$ with $P_0 \cap P_1 = \emptyset$ and $P_0 \cup P_1 = \{1, \dots, n\}$ with the property,
that for all $v \in \{1, \dots, n\}$ we have $v \in P_0 \Leftrightarrow n + 1 - v \in P_0$ and $v \in P_1 \Leftrightarrow n + 1 - v \in P_1$,
either P_0 has an arithmetic progression of size 3 or P_1 has an arithmetic progression of size t , or both.
While p is the largest number such that for all $n \leq p$ and for all such (P_0, P_1)
neither P_0 has an arithmetic progression of size 3 nor P_1 has an arithmetic progression of size t .

In the ordinary case of plain partitions (without the additional symmetry condition) we have $p + 1 = q$, and thus one uses just one number (instead of a pair), however here we can have a “palindromic span”, that is, $p + 1 < q$ can happen for the palindromic case. The reason is that from a good partition of $\{1, \dots, n\}$ we obtain a good partition of $\{1, \dots, n - 1\}$ by simply removing n , however for “good palindromic partitions” besides removing n we also need to remove the corresponding vertex 1 (due to the palindromicity condition).

Apparently the most advanced special algorithm (and implementation) for computing (mixed) van der Waerden numbers is the algorithm/implementation developed in [63]. For computing $w(2; 3, 17) = 279$, with this special algorithm a run-time of 552 days is reported (page 113); the machine used should be at most 30% slower than the machine used in our experiments, and so this should translate into at least 400 days on our machine. As we can see in Table 9, the `tauSolver-2.6` used is 85-times faster, while Table 10 shows, that `Cube & Conquer` is around 40-times faster. These algorithms know nothing about the specific problem, and are just given the generic SAT formulation of the underlying hypergraph colouring problem. So it seems that SAT solving does a good job here.⁵

1.1. Using SAT solvers

As explored in Dransfield et al. [19], Herwig et al. [27], Kouril [43, 42], Ahmed [2, 3], and Kullmann [49, 50], we can generate an instance $F(t_0, \dots, t_{k-1}; n)$ of the satisfiability problem (for definition, see any of the above references) corresponding to $w(k; t_0, t_1, \dots, t_{k-1})$ and integer n , such that $F(t_0, \dots, t_{k-1}; n)$ is satisfiable if and only if $n < w(k; t_0, t_1, \dots, t_{k-1})$. In particular, the instance $F(3, t; n)$ corresponding to $w(2; 3, t)$ with n variables consists of the following clauses:

- (a) $\{x_a, x_{a+d}, x_{a+2d}\}$ with $a \geq 1, d \geq 1, a + 2d \leq n$, and
- (b) $\{\bar{x}_a, \bar{x}_{a+d}, \dots, \bar{x}_{a+d(t-1)}\}$ with $a \geq 1, d \geq 1, a + d(t-1) \leq n$,

where an assignment $x_i = \varepsilon$ encodes $i \in P_\varepsilon$ for $\varepsilon \in \{0, 1\}$ (if x_i is not assigned but the formula is satisfied, then i can be arbitrarily placed in either of the blocks of the partition). The (“positive”) clauses (a) (consisting only of variables), constructed from all arithmetic progressions of length 3 in $\{1, \dots, n\}$, prohibit the existence of an arithmetic progression of length 3 in P_0 . And the (“negative”) clauses (b) (consisting only of negated variables), constructed from all arithmetic progressions of length t in $\{1, \dots, n\}$, prohibit the existence of an arithmetic progression of length t in P_1 . To check the satisfiability of the generated instance, we need to use a “SAT solver”. A complete SAT solver finds a satisfying assignment if one exists, and otherwise correctly says that no satisfying assignment exists and the formula is unsatisfiable. One of the earliest complete algorithms is the DLL algorithm ([18]), and our algorithm for computing $w(2; 3, 19) \leq 349$, discussed in Section 2, actually implements this very basic scheme, using modern heuristics.

SAT solving has progressed much beyond this simple algorithm, and the handbook [12] gives an overview (where [70] discusses some applications of SAT to combinatorics). There in [17] we find a general overview on complete SAT algorithms, while [40] gives an overview on incomplete algorithms. For complete algorithms especially the

⁴The sequence $\text{pdw}(2; 3, t)$ is <http://oeis.org/A198684>, <http://oeis.org/A198685> in the “On-Line Encyclopedia of Integer Sequences” (the first and the second components).

⁵As discussed in Subsection 2.1, for enumerating all solutions for $n = w(2; 3, 17) - 1 = 278$ with `tauSolver-2.6` we need at most the time needed for determining unsatisfiability; in this special case we have actually precisely one solution.

algorithms derived from the DLL algorithm are of importance, and there are two families, namely the (earlier) “look-ahead solvers” outlined in [31], and the (later) “conflict-driven solvers” (or “CDCL” like “conflict-driven clause-learning”) outlined in [55]. In Section 6 we will discuss how general SAT solvers perform on the problems from this article. The motivation for our choice of the most basic DLL algorithm for tackling the unsatisfiability of the instance $F(3, 19; 349)$, already employed in [3] and discussed in Subsection 3.1, is, that on these special problems classes this basic algorithm together with a modern heuristic is very competitive — best on ordinary problem instances, and beaten on palindromic instances only by the the Cube & Conquer method.⁶ And then it is also instructive to use such an algorithm, which due to its simplicity might enable greater insight. Another advantage of its simplicity is, that it can also count and enumerate the solutions, but in this article we focus mostly on mere SAT solving; see [22] for an overview on counting solutions.

Local-search based incomplete algorithms (see UbcSAT-suite [67]) are generally faster than a DLL-like algorithm in finding a satisfying assignment (on such combinatorial problems), and this is also the case for the instances of this article. However they may fail to deliver a satisfying assignment when there exists one, and they can not prove unsatisfiability. If they succeed on our instances, then they deliver a good partition, and thus a lower bound for a certain van der Waerden number. So such incomplete algorithms are used for obtaining good partitions and improving lower bounds of van der Waerden numbers. When they fail to improve the lower bound any further, we need to turn to a complete algorithm.

1.1.1. Informed versus uninformed SAT solving

We use general SAT solvers, and the new solvers developed by us are also general SAT solvers, which can run without modification on any SAT problem; these solvers just run on the naked and natural SAT formulation of the problem, without giving them further information. More specifically, to show unsatisfiability we have developed the `tawSolver` (Section 2) and the Cube & Conquer-method (Subsection 6.1.1), while to find satisfying assignments we have selected local-search algorithms (Subsection 6.2).

On the other end of the spectrum is [43, 42], which uses a highly specialised method, which involves a variety of specialised SAT solvers on specialised hardware, in combination with some special insights into the problem domain. For finding satisfying assignments we have the methods developed in [27, 30, 29]. For more examples on informed search to compute van der Waerden numbers, see also Section 2 of [5].

Our “uninformed approach” has stronger bearings on general SAT solving, while the informed approach can be more efficient for producing numerical results (however it seems to need a lot of effort to beat general SAT solvers (by specialised SAT solvers); as we have already reported, our general methods are at least on the instances of this paper much faster than the dedicated (non-SAT-based) method in [63]).

1.1.2. Parallel/distributed SAT solving

The problems we consider are computationally hard, and for the hardest of them in this paper, computation of $w(2; 3, 19) = 349$, a single processor, even when run for a long time, is not enough. Hence some form of parallelisation or distribution of the work is needed. Four levels of parallelisation have been considered for general-purpose SAT solving (in a variety of schemes):

- (i) Processor-level parallelisation: This helps only for very special algorithms, and can only achieve some relatively small speed-up; see [33] for an example which exploits parallel bit-operations. It seems to play no role for the problems we are considering.
- (ii) Computer-level parallelisation: Here it is exploited that currently a single (standard) computer can contain up to, say, 16 relatively independent processing units, working on shared memory. So threads (or processes) can run in parallel, using one (or more) of the following general forms of collaboration:
 - (a) Partitioning the work via partitioning the instance (see below); [71, 39] are “classical” examples.
 - (b) Using the same algorithm running in various nodes on the same problem, exploiting randomisation and/or sharing of learned results; see [36, 24] for recent examples.

⁶The Cube & Conquer method, developed originally on the instances of this article, combines a look-ahead solver with a conflict-driven solver, and is faster by a factor of two on palindromic instances.

(c) Using some portfolio approach, running different algorithms on the same problem, exploiting that various algorithms can behave very differently and unpredictably; see [23] for the first example.

Often these approaches are combined in various ways; see [62, 21, 37, 38] for recent examples. Approaches (b) and (c) do not seem to be of much use for the well-specified problem domain of hard instances from Ramsey theory. Only (a) is relevant, but in a more extreme form (see below). In the context of (ii), still only relatively “easy” problems (compared to the hard problems from Ramsey theory) are tackled.

- (iii) Parallelisation on a cluster of computers: Here up to, say, 100 computers are considered, with restricted communication (though typically still non-trivial). In this case, the approach (ii)(a) becomes more dominant, but other considerations of (ii) are still relevant. For hard problems this form of computation is a common approach.
- (iv) Internet computation, with completely independent computers, and only very basic communication between the centre and the machines: In principle, the number of computers is unbounded. Since progress must be guaranteed, and the instances for which Internet computation is applied would be very hard, at the global level only (ii)(a) is applicable (while at a local level all the other schemes can in principle be applied). Yet there is no real example for a SAT computation at this level.

We remark that the classical area of “high performance computing” seems to be of no relevance for SAT solving, since the basic SAT algorithms like unit-clause propagation are, different from typical forms of numerical computation, inherently sequential (compare also our remarks to (i)). However using dedicated hardware with specialised algorithms has been utilised in [43, 42], yielding the currently most efficient machinery for computing van der Waerden numbers.

A major advantage of the DLL solver architecture (which has been further developed into so-called “look-ahead” SAT solvers) is that the computation is easily parallelisable and distributable: Just compute the tree only up to a certain depth d , and solve the (up to) 2^d sub-problems at level d . Only minimal interaction is required: The sub-problems are solved independently, and in case one sub-problem has been found satisfiable, then the whole search can be aborted (for the purpose of mere SAT-solving; for counting all solutions of course the search needs to be completed). And the sub-problems are accessible via the partial assignment constituting the path from the root to the corresponding leaf, and thus also require only small storage space. This is the core of method (ii)(a) from above, and will be further considered in Subsection 3.1 (for our special example class).

In the subsequent subsection we will discuss the general merits of applying SAT solving to (hard) Ramsey problems. One spin-off of this combination lies in pushing the frontier of large computations. As a first example we have developed in [32, 68], motivated by the considerations of the present article, an improved method for (ii)(a) called “Cube & Conquer”, which is also relevant for industrial problems (typically from the verification area). One aspect exploited here is that for extremely hard problems, splitting into millions of sub-instances is needed. In the literature until now (see above for examples) only splitting as required, by at most hundreds of processors, has been performed, while it turned out that the above “extreme splitting”, when combined with “modern” (CDCL) SAT solvers, is even beneficial when considered as a (hybrid) solver on a single-processor, and this for a large range of problem instances.

1.1.3. Synergies between Ramsey theory and SAT

For Ramsey-numbers (see [57] for an overview on exact results), relatively precise asymptotic bounds exist, and due to the inherent symmetry, relatively specialised methods for solving concrete instances have an advantage. Van-der-Waerden-like numbers seem harder to tackle, both asymptotically and exactly, and perhaps the only way ever to know the precise values is by computation (and perhaps this is also true for Ramsey-numbers, only more structures are to be exploited). SAT solvers are especially suited for the task, since the computational problems are hypergraph-colouring problems, which, at least for two colours, have canonical translations into SAT problems (as only considered in this paper). For more colours, see the approach started in [50], while for a general theory of multi-valued SAT close to hypergraph-colouring, see [51, 52].

Through applying and improving SAT solvers (as in the present article), Ramsey theory itself acquires an applied side. Perhaps unknown to many mathematicians is the fact, that whenever for example a recent microchip is employed, this likely involves SAT solving, playing an important (though typically hidden) role in its development, by providing the underlying “engines” for its verification; see the recent handbook [12] to get some impression of this astounding development. Now we believe that problem instances from Ramsey theory are good benchmarks, serving to improve SAT solvers on hard instances:

- Unlike with random instances (see [1] for an overview), instances from Ramsey theory are “structured” in various ways. One special structure which one finds in all these instances is that they are layered by the number of vertices (the same structural pattern is repeated again and again, on growing scales).
- A major advantage of random instances is their scalability, that is, we can create relatively easily instances of the same “structure” and different sizes. With instances from Ramsey theory we can also vary the parameters, however due to the possibly large and unknown growth of Ramsey-like numbers, controlling satisfiability and hardness is more complicated here. This possible disadvantage can be overcome through computational studies like in this paper, which serve to calibrate the scale via precise numerical data, so that the field of SAT instances from Ramsey-theory becomes accessible (one knows for initial parameter values the satisfiability status and (apparent) solving complexity, and gets a feeling what happens beyond that).
- In this paper, we consider two instance classes: instances related to ordinary van der Waerden numbers $w(2; 3, t)$ and instances related to the palindromic forms $pdw(2; 3, t)$. Now already with these two classes, the two main types of complete SAT solvers, “look-ahead” (see [31]) and “conflict-driven” (see [55]), are covered in the sense that they dominate on one class each (and are (relatively) efficient); see Section 6 for further details. On the other hand, for random instances only look-ahead solvers are efficient (for complete solvers).
- Especially for local-search methods (see [40] for an overview), these problems are hard, but not overwhelmingly so (for the ranges considered), and thus all the given lower bounds can trigger further progress (and insight) into the solution process in a relatively simple engineering-like manner (by studying which algorithms work best where).
- On the other hand, for upper bounds we need to show unsatisfiability, which is much harder (we can only solve much smaller instances). All applications of SAT solving in hardware verification are “unsatisfiability-driven” (see [9, 44] for introductions). So future progress in solving hard Ramsey instances might trigger a breakthrough in tackling unsatisfiability, and should then also improve these industrial applications.

We believe that for better SAT solving, established hard problem instances are needed in a great variety, and we believe that Ramsey theory offers this potential. To begin the process of applying Ramsey theory in this direction, problem instances from this paper (as well as related to [50]) have been used in the SAT 2011 competition (<http://www.satcompetition.org/2011/>). As already mentioned in the previous subsection, the first fruits of the collaboration between SAT and Ramsey theory appeared in [32, 68], yielding a method for tackling hard problems with strong scalability.

Finally, the interaction between Ramsey theory and SAT should yield new insights for Ramsey theory itself:

1. The numerical data can yield conjectures on growth rates; see Subsection 3.3.
2. The good partitions found can yield conjectures on patterns; see Section 4.
3. New forms of Ramsey problems can be found through algorithmic considerations; see Section 5.
4. The SAT solving process, considered *in detail*, acts like a microscope, enabling insights into the structure of the problem instances which are out of sight for Ramsey theory yet. For approaches towards structures in SAT instances, which we hope to study in the future, see [61, 41].

1.2. The results of this paper

In Section 2, we present the new SAT solver, `tawSolver-2.6`, with superior performance on the instances considered in this paper (only for palindromic instances the new hybrid method `Cube & Conquer` is superior). Section 3 contains our results on the numbers $w(2; 3, t)$. We discuss the computation of the one new van der Waerden number, and present further conjectures regarding precise values⁷ and the growth rate. In Section 4, we investigate some patterns we found in the good partitions (establishing the lower bounds). In Section 5, we introduce palindromic problems and the corresponding palindromic number-pairs. Finally in Section 6, we discuss the observations on the use of the various SAT solvers involved.

⁷to establish these conjectures will require major advances in SAT solving

In this paper, we represent partitions of $w(2; 3, t)$ as bitstrings. For example, the partition $P_0 = \{1, 4, 5, 8\}$ and $P_1 = \{2, 3, 6, 7\}$, which is an example of a good partition of $\{1, 2, \dots, 8\}$, where $8 = w(2; 3, 3) - 1$, is represented as 01100110, or more compactly as $01^20^21^20$, using exponentiation to denote repetition of bits.

2. The `tawSolver`

We now discuss the `tawSolver`, an open-source SAT solver, created by the first author with a special focus on van der Waerden problems (version 1.0), and improved by the second author through an improved branching heuristic (version 2.6).⁸ Algorithm 1 shows that the basic algorithm of the `tawSolver` is the simplest possible (reasonable) DLL-scheme, just branching on a variable plus unit-clause propagation. As we can see in Section 6, it is the strongest SAT solver on the instances considered in this paper, only beaten on palindromic problems by the new hybrid scheme `Cube & Conquer`, which came out as a result on research on the instances of this paper.

2.1. The basic structure

Algorithm 1 specifies the `tawSolver`, which for input F (a formula or “clause-set”) decides satisfiability:

1. Lines 3-5 is “unit-clause propagation” (UCP), denoted by the function r_1 , which sets literals x in the current F to true while there are unit-clauses $\{x\} \in F$.
 - (a) Setting a literal x to true in a clause-set F is performed by first removing all clauses from F containing x , and removing the element \bar{x} from the remaining clauses.
 - (b) r_1 finds a contradiction (Line 4) by finding two unit-clauses $\{v\}$ and $\{\bar{v}\}$ (i.e., $v \wedge \neg v$).
 - (c) While r_1 finds a satisfying assignment (Line 5) if all clauses vanished (have been satisfied).
2. Lines 6-7 give the branching heuristic, which yields the branching literal x , first set to true, then to false, in the recursive call of the `tawSolver`.
 - (a) $p(a, b) \in \mathbb{R}_{>0}$ for $a, b \in \mathbb{R}_{>0}$ in Line 6 is the “projection”, and we consider three choices p_+ , p_* , p_τ .
 - (b) $w_F(x)$ for literal x is a heuristical value, measuring in a sense the “progress achieved” when setting x to FALSE (“progress” in the sense of the instance becoming more constrained, so that more unit-clause propagations are to be expected).
 - (c) The details are specified in Subsections 2.3, 2.5.
3. The implementation is discussed in Subsection 2.4.
4. The tree of recursive calls made by the solver is called the *DLL-tree* of F .

Besides the choice of the heuristic, this is the basic SAT solver as published in [18]. The implementation is optimised for the needs of the branching heuristic, which requires to know from each (original) clause in the input F whether it has been satisfied meanwhile, and if not, what is its current length.

Algorithm 1 `tawSolver`

- 1: Global variable F , initialised by the input.
 - 2: **function** `DLL()` : returns SAT or UNSAT for the current F
 - 3: Update F to $r_1(F)$
 - 4: If contradiction found via r_1 , then goto 12
 - 5: If satisfying assignment found via r_1 , then return SAT
 - 6: Choose variable v with maximal $p(w_F(v), w_F(\bar{v}))$
 - 7: If $w_F(v) \geq w_F(\bar{v})$, then $x := v$, else $x := \bar{v}$
 - 8: Set x to TRUE in F ; if `DLL()` = SAT, then return SAT
 - 9: Undo assignment of x
 - 10: Set \bar{x} to TRUE in F ; if `DLL()` = SAT, then return SAT
 - 11: Undo assignment of \bar{x}
 - 12: Undo assignments made by r_1
 - 13: Return UNSAT
 - 14: **end function**
-

⁸<http://sourceforge.net/projects/tawsolver/>, and in the `OKlibrary`:
<https://github.com/OKullmann/oklibrary/blob/master/Satisfiability/Solvers/TawSolver/tawSolver.cpp>

With a small modification, namely just continuing when a satisfying assignment was found, the `tawSolver` can also count all satisfying assignments, or output them; this is available as a compile-time option for the solver. In Section 4, we will discuss some patterns which we found in satisfying assignments for $F(3, t; n)$ with $n < w(2; 3, t)$. We do not report run-times for determining (or counting) all solutions in Section 6, but for $n = w(2; 3, t) - 1$ (empirically) the run-time is at most the run-time needed to determine unsatisfiability for $n = w(2; 3, t)$; for numerical values of solution-counts see [42].

2.2. Look-ahead solvers

It is useful for the general picture to consider the general r_k -operations, as introduced in [45] and further studied in [25, 26]. These operations transform a clause-set F into a satisfiability-equivalent clause-set via application of some forced assignments (i.e., where the opposite assignments would yield an unsatisfiable clause-set). Let \perp be the empty clause, which stands for a trivial contradiction. r_0 just maps F to $\{\perp\}$ in case of $\perp \in F$, while otherwise F is left unchanged. Now we can recognise r_1 as an operation which is applied recursively to the result of F with literal x set to true if setting \bar{x} to true yields $\{\perp\}$ via r_0 . This scheme yields also the general r_k for $k \in \mathbb{N}$: as long as there is a literal x such that F with \bar{x} set to true yields $\{\perp\}$ via r_{k-1} , set x to true and iterate. The final result, denoted by $r_k(F)$, is uniquely determined. Besides the ubiquitous unit-clause propagation r_1 also r_2 , called “failed literal elimination”, is popular for SAT solving, and even r_3 , typically called “double look-ahead”, is used in some solvers (always partially, testing the reductions only for selected variables).

The general scheme for a look-ahead solver (as stipulated in [48]) now generalises the DLL-procedure from Algorithm 1, by replacing the reduction $F \rightsquigarrow r_1(F)$ in Line 3 by the general $F \rightsquigarrow r_k(F)$ for some $k \geq 1$. Furthermore, for the inspection of a branching variable and the computation of the heuristical values $w(v)$ and $w(\bar{v})$, now the effects of setting v resp. \bar{v} to true and performing r_{k-1} reduction are considered. This explains also the notion of “look-ahead”: the r_k -reduction can be partially achieved at the time when running through all variables v , setting v resp. \bar{v} to true and applying r_{k-1} — if this yields $\{\perp\}$, then performing the opposite assignment is justified. Since r_1 is the standard for reduction of a branch, (partial) r_2 is the default for the reduction at a node.⁹

We see that `tawSolver` uses $k = 1$ (so the “look-ahead” uses $k = 0$, and in this sense `tawSolver` is a “look-ahead solver with zero look-ahead”). The prototypical solver for using $k = 2$ is the `OKsolver` ([46]). In a rather precise sense the `tawSolver` can be considered at the level-1-version of the `OKsolver` (or the latter as the level-2-version of the `tawSolver`). Also for the branching heuristic, which is discussed in the following subsection, `tawSolver` uses the same scheme as the `OKsolver`, appropriately simplified to the lower level. Both `tawSolver` and `OKsolver` are solvers with a “mathematical meaning”, precisely implementing an algorithm to full extent, with the only magical numbers the clause-weights used in the branching heuristic.

The general scheme for the branching heuristic of a look-ahead solver, as developed in [48] (Subsection 7.7.2), is as follows: For a clause-set F and its direct successor F' on a branch (applying the branching assignment and further reductions), a “distance measure” $d(F, F') \in \mathbb{R}_{>0}$ is chosen, with the meaning the bigger this distance, the larger the decrease in complexity. The branching heuristic considers for each variable v its two successor F', F'' and computes the distances $d(F, F'), d(F, F'')$. Then via a “projection” $p : \mathbb{R}_{>0}^2 \rightarrow \mathbb{R}_{>0}$ one heuristical value $h_v := p(d(F, F'), d(F, F''))$ is obtained. Finally some v with maximal h_v is chosen. Choosing which of v or \bar{v} to be processed first (important for satisfiable instances) is done via a second heuristic, estimating the satisfiability-probabilities of F', F'' in some way.

2.3. From `tawSolver-1.0` to `tawSolver-2.6`

We are now turning to the discussion of the branching heuristic in `tawSolver-2.6` (lines 6, 7 in Algorithm 1), the version developed for this article. For `tawSolver-1.0` (used in [2, 3]) the “Two-sided Jeroslaw-Wang” (2sJW) rule by Hooker and Vinay [34] was used, which chooses v such that the weighted sum of the number of clauses of F containing v is maximal, where the weight of a clause of length k is 2^{-k} .¹⁰ As discussed in [48], the ideas from [34] are

⁹The look-ahead solvers `satz` and `march_p1` run through the variables once (actually also only considering “interesting” variables by some criterion), and so they do not compute r_2 , but only an approximation. The only solver to completely compute r_2 is the `OKsolver` (while `satz` and `march_p1` search also for some r_3 reductions on selected variables).

¹⁰We do not care much here about the order of branching, since the algorithm is only effective on unsatisfiable problems, where the order does not matter (while on satisfiable problems local search is much faster).

actually rather misleading, and this is demonstrated here again by obtaining a large speed-up through the replacement of the branching heuristic, as can be seen by the data in Section 6 (comparing `tawSolver-1.0` with `tawSolver-2.6`).

For a literal x , a clause-set F and $k \in \mathbb{N}$ let $\text{ld}_F^k(x) := |\{C \in F : x \in C \wedge |C| = k\}|$ be the “literal degree” of x in the k -clauses of F . The 2sJW-rule consists of three components:

1. The weight $w_F(x)$ of literal x is set as $w_F(x) := \sum_k 2^{-k} \cdot \text{ld}_F^k(x)$.
2. A variable v with maximal $p_+(w_F(v), w_F(\bar{v}))$ for $p_+(a, b) := a + b$ is chosen.
3. The literal $x \in \{v, \bar{v}\}$ to be set first to true is given by the condition $w_F(x) \geq w_F(\bar{x})$.

This approach has the following fundamental flaws:

1. The choice of the first branch (v or \bar{v}) is mixed up with the choice of v itself, but very different heuristics are needed:
 - (a) For the choice of the first branch, some form of approximated *satisfiability*-probability must be maximised,
 - (b) while the branching-variable must minimise some approximated tree-size for the worst case, the *unsatisfiable* case.

In 2sJW the weights 2^{-k} are only motivated by satisfiability-probabilities, but are used for the choice of v itself.

2. Once total weights $w_F(v), w_F(\bar{v})$ have been determined, one number (the projection) must be computed from this (to be maximised). 2sJW uses the sum, which, as demonstrated in [48], corresponds to minimising a *lower bound* on the DLL-tree-size — much better is the product $p_*(a, b) := a \cdot b$, which corresponds to minimising an *upper bound* on the tree-size.

So the improved heuristic (which nowadays, when extended appropriately to take the look-ahead into account, is the basis for all look-ahead solvers) chooses clause-weights $w_2, w_3, \dots \in \mathbb{R}_{>0}$, from which the total weight

$$w_F(x) := \sum_k w_k \cdot \text{ld}_F^k(x)$$

is determined, and chooses a variable v with maximal

$$p_*(w_F(v), w_F(\bar{v})) = w_F(v) \cdot w_F(\bar{v}).$$

The meaning of these weights is completely different from the argumentation in [34]: as mentioned, satisfiability-probabilities have no place here. The underlying distance measure is $\sum_k w_k \cdot \nu^k(F')$, where F' is the resulting clause-set after performing the branch-assignment and the subsequent r_k -reduction, while $\nu^k(F')$ is the number of *new* k -clauses in F' . When setting literal x to true, then $\text{ld}_F^k(\bar{x})$ is an “approximation” of the number of new clauses of length $k - 1$ (since in the clauses containing \bar{x} this literal is removed).

The weights w_k^{OK} for the `OKsolver` have been experimentally determined as roughly 5^{-k} . Since the value of the first weight is arbitrary, the weights are rescaled to $w_2^{\text{OK}} = 1$, obtaining then each new weight by multiplication with $1/5$. Now w_2 for the `tawSolver` is a stand-in for the number of new 1-clauses, which are handled in the `OKsolver` by the look-ahead; accordingly it seems plausible that now w_2 needs a relatively higher weight. We rescale here the weights to $w_3 = 1$ (note that for the `tawSolver` the weight w_k concerns new clauses of length $k - 1$). Empirically we determined $w_2 = 4.85$, $w_4 = 0.354$, $w_5 = 0.11$, $w_6 = 0.0694$, and thereafter a factor of $\frac{1}{1.46}$; thus starting with w_2 the next weights are obtained by multiplying with (rounded) $1/4.85, 1/2.82, 1/3.22, 1/1.59, 1/1.46, \dots$

For the choice of the first branch there are two main schemes, as discussed in [48] (Subsection 7.9). Roughly, the target now is to get rid off (satisfy) as many short clauses as possible (since shorter clauses are bigger obstructions for satisfiability).¹¹ Both schemes amount to choose literal $x \in \{v, \bar{v}\}$ with $w'_F(x) \geq w'_F(\bar{x})$ for some weights w'_k . For the Franco-estimator we have $w'_k = -\log(1 - 2^{-k})$, while for the Johnson-estimator we have $w'_k = 2^{-k}$. In the `OKsolver` the Franco-estimator is used. But for the `tawSolver` with its emphasis on unsatisfiable instances, while the computation of the heuristic is very time-sensitive (much more so than for the `OKsolver`), actually just the same weights $w'_k = w_k$ are used.

¹¹While for a good branching variable we want to *create* as many short clauses as possible (via setting literals to false)!

As one can see from the data in Section 6, on ordinary van der Waerden problems the new heuristic yields a reduction in the size of the DLL-tree by a factor increasing from 2 to 5 for $t = 12, \dots, 16$ (comparing `tawSolver-2.6` with `tawSolver-1.0`), and for palindromic problems by a factor increasing from 5 to 20 for $t = 17, \dots, 23$.¹² We do not present the data, but most of the reduction in node-count is due to the replacement of the sum as projection by the product (the optimised clause-weights only further improve the node reduction by at most 50% for the biggest instances, compared with a simple but reasonable scheme like 2^{-k}).

2.4. The implementation

The `tawSolver` is written in modern C++ (C++11, to be precise), with around 1000 lines of code, with complete input- and output-facilities, error handling and various compile-time options for implementations. The code is highly optimised for run-time speed, but at the same time expressing the concepts via appropriate abstractions, relying on the expressiveness of C++ both at the abstraction- and the implementation-level, so that the compiler can do a good job producing efficient code.

Look-ahead solvers are often “eager”, that is, they represent the clause-set at each node of the DLL-tree in such a way, that the current (“residual”) clause-set is visible to the solver, and precisely the current clauses can be accessed. On the other hand, conflict-driven solvers are all “lazy”, that is, the initial clause-set is not updated, and the state of the current clause-set has to be inferred via the current assignment to the variables. The representation of the input clause-set F by the `tawSolver` now is “mostly lazy”:

1. Assignments to variables are entered into a global array,
2. Via the usual occurrence lists, for each literal x one obtains access to all the clauses $C \in F$ with $x \in C$.
3. This representation of F is static (is not updated), and in this sense we have a lazy datastructure.
4. But the status of clauses, which is either inactive (when satisfied) or active, and their length (in the active case) is handled eagerly, by storing status and length for each clause and updating this information appropriately. So at each node, when running through the occurrence lists (still as in the input), for each clause we can see directly whether the clause is active and in this case its current length.
5. When doing an assignment, then the clause-lengths are updated: if a literal is falsified in a clause, the length is decreased by one, and if a literal is satisfied, the status of the clause is set to inactive.
6. For each active clause containing a variable which is assigned, there is exactly one change (either decrease in length or going from active to inactive). This change is entered into a change-list.
7. When backtracking, then the assignment is simply undone by going through the change-list in reverse order, and undoing the changes to the clauses.

No counters are maintained for the literal degrees $\text{ld}^k(x)$. Instead, the heuristic is computed by running through all literal occurrences in the original input for the unassigned literals, and adding the contributions of the clauses which are still active (this is the use of maintaining the length of a clause).

When doing unit-clause propagation, the basic choice is whether performing a BFS search, by using a first-in-first-out strategy for the processing of derived unit-clauses, of a DFS search, using a last-in-first-out strategy. BFS is slightly easier to implement, but on the palindromic vdW-instances needs roughly 10% more unit-clauses to propagate¹³, while on ordinary vdW-instances it uses less propagations, though the difference is less than 2%, and thus DFS is the default. This can also be motivated by the consideration that newly derived unit-clauses can be considered to be “more expensive”, and thus should be treated as soon as possible.

Look-ahead solvers in general rely on the distance for branch-evaluation to be positive, while a zero value should indicate that a special reduction can be performed. And indeed, when counting new clauses, then the weighted sum being zero means that an autarky has been found, a partial assignment not creating new clauses, which means that all touched clauses are satisfied; see [41].¹⁴ Thus starting with the `OKsolver`, look-ahead solvers looked out for such autarkies, and applied them when found ([31, 48]). Now for a zero-look-ahead solver like the `tawSolver`, these autarkies are just pure literals (only occurring in one sign, not in the other). Their elimination causes a slight run-time increase, without changing much anything else, and so by default they are not eliminated but not chosen for branching (if there are still non-pure literals).

¹²`tawSolver-2.6` additionally has the implementation improved, so that nodes are processed now twice as fast as with `tawSolver-1.0`.

¹³the final result is uniquely determined, but in general there are many ways to get there

¹⁴The point about autarkies is that they can be applied satisfiability-equivalently.

2.5. The optimal projection: the τ -function

In [48] it is shown that the τ -function is the best generic projection in the following sense:

- The τ -function is defined for arbitrary tuples $a \in \mathbb{R}^n$, $n \in \mathbb{N}$, namely $\tau(a) \in \mathbb{R}_{>0}$ is the unique $x \geq 1$ such that $\sum_{i=1}^n x^{-a_i} = 1$.
- This projection induces a linear order on the set of all such “branching tuples” a (of arbitrary length) by defining $a \leq b$ if $\tau(a) \leq \tau(b)$; here “ $a \leq b$ ” means that a is better than b .
- Theorem 7.5.3 in [48] shows that when imposing some general consistency-constraints on the comparison of branching tuples (where it is of importance that branching tuples can have arbitrary length), then there is precisely one such linear order on the set of branching tuples, namely the one induced by τ .

Now specific solvers might have a special built-in bias, and, more importantly, the theorem is not applicable when considering only branching tuples of length 2 (as it is the case for ordinary boolean SAT solving). But nevertheless, considering the τ -function as projection (more precisely, since we maximised projection values, $1/\tau$ is used) is an interesting option, and leads to the `tauSolver-2.6` (with “ τ ” in place of “ t ”):

$$p_\tau(w_F(v), w_F(\bar{v})) := 1/\tau(w_F(v), w_F(\bar{v})).$$

In this context it makes sense to definitely forbid distance-values 0, and thus pure literals are now eliminated.¹⁵ In Section 6 we see that `tauSolver-2.6` is faster than `tauSolver-2.6` on large palindromic problems due to a much reduced node-count, but on ordinary problems the node-count stays basically the same, and then the overhead for computing p_τ makes the `tauSolver-2.6` slower.

The weights for `tauSolver-2.6` have been empirically determined as $w_2 = 7$, $w_4 = 0.31$, $w_5 = 0.19$, and then a factor of $\frac{1}{1.7}$; so starting with w_2 the next weights are obtained by multiplying with $1/7$, $1/3.22$, $1/1.63$, $1/1.7$, \dots ¹⁶

3. Computational results on $w(2; 3, t)$

This section is concerned with the numbers $w(2; 3, t)$. The discussion of the computation of $w(2; 3, 19)$ is the subject of Subsection 3.1. Conjectures on the values of $w(2; 3, t)$ for $20 \leq t \leq 30$ are presented in Subsection 3.2, and also further lower bounds for $31 \leq t \leq 39$ are given there. Finally in Subsection 3.3, we update the conjecture on the (quadratic) growth of $w(2; 3, t)$.

3.1. $w(2; 3, 19) = 349$

The lower bound $w(2; 3, 19) \geq 349$ was obtained by Kullmann [49] using local search algorithms and it could not be improved any further using these incomplete algorithms (because, as we now know, the bound is tight). An example of a good partition of the set $\{1, 2, \dots, 348\}$ is as follows:

$$\begin{aligned} &1^4 01^6 01^{18} 01^3 01^4 01^5 01^4 01^{11} 01^9 01^3 01^6 01^7 01^5 01^{14} 01^{16} 01 01^2 01^2 01^{15} 01^4 01^{12} 0 \\ &1^{15} 01^2 01^5 01^7 01^{10} 01^{13} 01^2 01^{15} 01^{12} 01^4 01^{15} 01^2 01^2 01 01^9 01^6 01^{14} 01^5 01^{14} 01^2. \end{aligned}$$

To finish the search, i.e., to decide that a current lower bound of a certain van der Waerden number is exact, one might require many years of CPU-time. Discovering a new van der Waerden number has always been a challenge, as it requires to explore the search space completely, which has a size exponential in the number of variables in the corresponding satisfiability instance. To prove that an instance with n variables is unsatisfiable, the DLL algorithm has to implicitly enumerate all the 2^n cases. So the algorithm systematically explores all possible cases, however without actually explicitly evaluating all of them — herein lies the strength (and the challenge) for SAT solving.

¹⁵That is, only eliminating those literals (by setting \bar{x} to true) with $w_F(x) = 0$; these eliminations might create further pure literals, which will be eliminated when in the child node the branching variable is computed, and so on.

¹⁶We consider the values for the weights as reasonable all-round values. A deeper understanding, based on the theory developed in [48], is left for future investigations.

In Subsection 1.1.2, we gave an overview on the area of distributing hard SAT problems from a general SAT perspective, and we are concerned here with method (ii)(a), applied to `tawSolver`. We find the simplest division of the computation of the search into parts, that have no inter-process communication among themselves, together with the observation of some patterns, very successful. Namely a level (depth) $L \in \mathbb{N}_0$ of the DLL-tree is chosen, where the level considers only the decisions (ignoring the variables inferred via unit-clause propagation), and the 2^L subtrees rooted at that level are distributed among the processors.

To show the unsatisfiability of $F(3, 19; 349)$, we have used `tawSolver-1.0` and 2.2 GHz AMD Opteron 64-bit processors (200 of them) from the `cirrus` cluster at Concordia University for running the distributed branches of the DLL-tree. The value $L = 8$ was chosen, splitting the search space into $2^8 = 256$ independent parts (subtrees) P_0, \dots, P_{255} . The total CPU-time of all processor together was roughly 196 years (the first part P_0 alone has taken roughly 60 years of CPU-time).¹⁷ For the prediction of run-times for the sub-tasks, the following observation made in Ahmed [3] was used. Recall that for `tawSolver-1.0` (Algorithm 1) the branching rule was to select a variable with maximal $w_F(v) + w_F(\bar{v}) = \sum_k (\text{ld}_F(v) + \text{ld}_F(\bar{v})) \cdot 2^{-k}$, where for the first branch $x \in \{v, \bar{v}\}$ with $\sum_k \text{ld}_F(v) \cdot 2^{-k} \geq \sum_k \text{ld}_F(\bar{v})$ is chosen. Now the observation is that the parts (sub-trees of the DLL-tree) $P_0, P_1, P_2, P_4, P_8, P_{16}, P_{32}, P_{64}, P_{128}$ are bigger than the others parts, and P_0 is the biggest.

Meanwhile our result $w(2; 3, 19) = 349$ has been reproduced in [42], via an alternative SAT solving approach (see Subsection 1.1.1). At least at this time there seems to be no competitive alternative to SAT solving. See Section 6 for further remarks on SAT solving for these instances in general. It would be highly desirable to be able to substantially compress the resolution proofs obtained from the solver runs, so that a proof object would be obtained which could be verified by certified software (and hardware); see [16] for some recent literature.

3.2. Some new conjectures

In Table 2, we provide conjectured values of $w(2; 3, t)$ for $t = 20, 21, \dots, 30$. We have used the `UbcSAT` suite [67] of local-search based satisfiability algorithms for generating good partitions, which provide a proof of these lower bounds; see Appendix A.1 for the certificates. In Subsection 6.2 we provide details of the algorithms used to find the good partitions. The characteristics of the searches were such that we believe these values to be optimal, namely with the right settings, these bounds can be found rather quickly, and in the past, all such conjectures turned out to be true (though, as discussed below, the situation gets weaker for $t = 29, 30$). However, since local search based algorithms are incomplete (they may fail to deliver a satisfying assignment, and hence a good partition when there exists one), it remains to prove exactness of the numbers using a complete satisfiability solver or some complete colouring algorithm.

t	20	21	22	23	24	25	26	27	28	29	30
$w(2; 3, t) \geq$	389	416	464	516	593	656	727	770	827	868	903

Table 2: Conjectured precise lower bounds for $w(2; 3, t)$

We observe that for $t = 24, 25, \dots, 30$ we have $w(2; 3, t) > t^2$, which refutes the possibility that $\forall t : w(2; 3, t) \leq t^2$, as suggested in [14], based on the exact values for $5 \leq t \leq 16$ known by then. Further (strict) lower bounds we found are in Table 3 (where now we think it is likely that these bounds can be improved; see Appendix A.2 for the certificates).

t	31	32	33	34	35	36	37	38	39
$w(2; 3, t) >$	930	1006	1063	1143	1204	1257	1338	1378	1418

Table 3: Further lower bounds for $w(2; 3, t)$

¹⁷Comparing `tawSolver-1.0` with `tawSolver-2.6`, as we can see in Table 9, the series of quotients $q_i = \text{old-time} / \text{new-time}$, for $t = 12, \dots, 16$ is (rounded) 4.3, 5.6, 6.8, 9.4, 12.8. This can be approximated well by the law $q_{i+1} = 1.3 \cdot q_i$, which would yield for $t = 19$ the factor $12.8 \cdot 1.3^3 \approx 28.1$. So we would expect with `tawSolver-2.6` at least a speed-up by a factor 20, which would reduce the 200 years to 10 years. Another approximation is obtained by considering Table 9: we see that for each step from t to $t + 1$ the run-time always increases by less than a factor of 10, while for $t = 17$ we use less than five days, which would yield at most 500 days for $t = 19$ with `tawSolver-2.6`.

That we conjecture the data of Table 2 to be true, that is, that the used local-search algorithm is strong enough, while for the data of Table 3 that algorithm seems too weak to reach the solution, has the following background in the data: As we report in Subsection 6.2, in the range $24 \leq t \leq 33$ the local-search algorithm RoTS from the UbcSAT suite was found best-performing. This algorithm is used in an incremental fashion, initialising the search by known solutions for smaller n . This approach for $t = 28$, with a cut-off $5 \cdot 10^6$ rounds, found a solution for $n = 826$, and in 1000 independent runs (non-incremental) two solutions were found. But with cut-off 10^7 in 1000 runs and cut-off $2 \cdot 10^7$ in 500 runs no solutions were found. From our experience this seems “pretty safe” for a conjecture. We are entering now a transition period. For $t = 29$ the iterative approach with cut-off $5 \cdot 10^6$ found the solution for $n = 867$, while cut-off 10^7 found no solution for $n = 868$ in 1000 runs. For $t = 30$ the iterative approach managed to find a solution for $n = 897$; restarting it with cut-off 10^8 found a solution for $n = 902$, while for $n = 903$ no solution with that cut-off was found in 300 runs. So we see that already $t = 30$ is stretching it. However for $t = 31$ the iterative approach with cut-off 10^8 only reached $n = 919$ (despite restarts), while we happen to have a palindromic solution for $n = 930$ (these are much easier to find; see Subsection 5.3). So here now we believe we definitely over-stretched the abilities of the algorithm.

3.3. A conjecture on the upper bound

An important theoretical question is the growth-rate of $t \mapsto w(2; 3, t)$. Although the precise relation “ $w(2; r, t) \leq t^2$ ” has been invalidated by our results, quadratic growth still seems appropriate (see [49] for a more general conjecture on polynomial growth for van der Waerden numbers in certain directions of the parameter space; indeed in some directions linear growth is proven there):

Conjecture 3.1. There exists a constant $c > 1$ such that $w(2; 3, t) \leq ct^2$.

See Conjecture 4.4 for a strengthening. To determine the current best guess for c , and to give some heuristic justification for Conjecture 3.1, we observe the known exact values and lower bounds, and we arrive at the following possible recursion:

$$w(2; 3, t) \leq w(2; 3, t-1) + d(t-1),$$

for $4 \leq t \leq 39$ and some $d > 0$, with $w(2; 3, 3) = 9$. So we make the Ansatz $w(2; 3, t) \leq w_t := 9 + \sum_{i=3}^{t-1} d \cdot i$, for $t \geq 3$, where $d := \max_{t=4}^{39} \frac{w(2;3,t)-w(2;3,t-1)}{t-1}$; in case $w(2; 3, t)$ is not known, we use the lower bounds from Tables 2, 3. From our data we obtain $d = \frac{593-516}{23} = \frac{77}{23}$ (see Appendix B.1). We have (geometric sum) $w_t = \frac{d}{2}t^2 - \frac{3}{2}dt + 9 - 2d < \frac{d}{2}t^2$, and so we obtain

$$w(2; 3, t) \leq \frac{d}{2}t^2 = \frac{77}{46}t^2 < 1.675t^2,$$

which satisfies all data regarding $w(2; 3, t)$ presented so far.

4. Patterns in the good partitions

In this section, we investigate the set of all good partitions corresponding to certain van der Waerden numbers $w(2; 3, t)$ for patterns. As described in Section 1.1.1, the motivation behind this section is to obtain more problem-specific information on the solution-patterns, which may help to design heuristics to reduce search-space while computing specific van der Waerden numbers.

Let $S(t)$ denote the set of all binary strings each of which represents a good partition of the set $\{1, 2, \dots, w(2; 3, t) - 1\}$. Generating $S(t)$ involves traversing the respective search space completely. Let $n_0(B)$, $n_1(B)$, and $n_{00}(B)$ be the number of zeros, ones, and double-zeros, respectively, in a bitstring B (note that three consecutive zeros are not possible in any $B \in S(t)$). Let $\text{EP1S}(B)$ denote the sequence of powers of 1 in a bitstring B . Let $n_p(B)$ and $n_v(B)$ denote the number of peaks (local maxima) and valleys (local minima), respectively, in $\text{EP1S}(B)$ (not necessarily strict). For example, for the compact bitstring $1^8 001^6 01^3 01^1 01^3 001^5 01^8 01^5 001^3 01^1 01^3 01^6 001^8$ (with $n_0 = 16$, $n_1 = 60$ and $n_{00} = 4$), we have the following EP1S, with p and v , marking peaks and valleys, respectively, corresponding to changes in magnitudes.

$$\begin{array}{cccccccccccccccc} 8 & 6 & 3 & 1 & 3 & 5 & 8 & 5 & 3 & 1 & 3 & 6 & 8 \\ p & & v & & p & & v & & p & & & & & \end{array}$$

And for $B = 1^1 0 1^1 0 1^2 0 1^2 0 1^3 0 1^3$ we have $n_0(B) = 5$, $n_1(B) = 12$, $n_{00}(B) = 0$, while there is one valley followed by one peak, and thus $n_v(B) = n_p(B) = 1$.

4.1. Number of 0's and 00's

In this section, we determine the number $\min\{n_0(B) : B \in S(t)\}$, $\max\{n_0(B) : B \in S(t)\}$, and $\max\{n_{00}(B) : B \in S(t)\}$ for $3 \leq t \leq 14$. Observations in Table 4 lead us to Conjectures 4.1 and 4.2.

Table 4: Zeros in good partitions of $\{1, 2, \dots, w(2; 3, t) - 1\}$

$w(2; 3, t)$	$(\min\{n_0(B) : B \in S(t)\}, \max\{n_0(B) : B \in S(t)\})$	$\max\{n_{00}(B) : B \in S(t)\}$
w(2; 3, 3)	(4, 4)	2
w(2; 3, 4)	(6, 6)	2
w(2; 3, 5)	(7, 9)	2
w(2; 3, 6)	(8, 10)	4
w(2; 3, 7)	(11, 12)	3
w(2; 3, 8)	(14, 14)	1
w(2; 3, 9)	(16, 16)	4
w(2; 3, 10)	(19, 21)	5
w(2; 3, 11)	(19, 22)	5
w(2; 3, 12)	(22, 22)	1
w(2; 3, 13)	(25, 29)	5
w(2; 3, 14)	(29, 29)	4

It seems that there is little variation concerning the total number of zeros:

Conjecture 4.1. There exists a constant $c > 0$ such that $|n_0(B) - n_0(B')| \leq ct$, $\forall B, B' \in S(t)$ with $t \geq 3$.

And there seem to be very few consecutive zeros:

Conjecture 4.2. There exists a constant $c > 0$ such that $n_{00}(B) < ct$, $\forall B \in S(t)$ with $t \geq 3$.

4.2. Number of 1's

In this section, we determine $T = \min\{n_p(\text{EPIS}(B)) + n_v(\text{EPIS}(B)) : B \in S(t)\}$, as well as minimum and maximum values of $n_1(B)$ over all $B \in S(t)$. The observations in Table 5 lead us to Conjectures 4.3, 4.4, and Questions 4.1 and 4.2.

Table 5: Selected good-partitions of $\{1, 2, \dots, w(2; 3, t) - 1\}$

$w(2; 3, t)$	A good partition B corresponding to T	T	$(\min\{n_1(B) : B \in S(t)\}, \max\{n_1(B) : B \in S(t)\})$
w(2; 3, 3) = 9	$1^2 0 0 1^2 0 0$ (2 2)	1	(4, 4)
w(2; 3, 4) = 18	$1^3 0 0 1^1 0 1^3 0 0 1^1 0 1^3$ (3 1 3 1 3)	5	(11, 11)
w(2; 3, 5) = 22	$0 0 1^3 0 0 1^1 0 1^4 0 0 1^4 0 1^1$ (3 1 4 4 1)	4	(12, 14)
w(2; 3, 6) = 32	$0 1^5 0 0 1^5 0 1^3 0 0 1^5 0 0 1^5$ (5 5 3 5 5)	3	(21, 23)
w(2; 3, 7) = 46	$1^1 0 1^1 0 1^4 0 1^2 0 1^5 0 1^4 0 1^1 0 0 1^3 0 1^5 0 1^2 0 1^5 0$ (1 1 4 2 5 4 1 3 5 2 5)	8	(33, 34)
w(2; 3, 8) = 58	$1^4 0 1^2 0 1^4 0 1^1 0 1^4 0 1^3 0 1^5 0 0 1^5 0 1^3 0 1^4 0 1^1 0 1^4 0 1^2 0 1^1$ (4 2 4 1 4 3 5 5 3 4 1 4 2 1)	12	(43, 43)

Continued on Next Page...

Table 5: Selected good-partitions of $\{1, 2, \dots, w(2; 3, t) - 1\}$

$w(2; 3, t)$	A good partition B corresponding to T	T	$\min\{n_1(B) : B \in S(t)\},$ $\max\{n_1(B) : B \in S(t)\}$
$w(2; 3, 9) = 77$	$1^8 001^6 01^3 01^1 01^3 001^5 01^8 01^5 001^3 01^1 01^3 01^6 001^8$ (8 6 3 1 3 5 8 5 3 1 3 6 8)	5	(60, 60)
$w(2; 3, 10) = 97$	$1^7 01^4 01^2 01^5 001^2 001^7 01^4 01^8 01^1 01^8 01^4 001^6 001^2 001^8 01^9$ (7 4 2 5 2 7 4 8 1 8 4 6 2 8 9)	13	(75, 77)
$w(2; 3, 11) = 114$	$01^{10} 01^4 001^6 01^{10} 01^2 001^9 01^6 01^1 01^9 001^1 001^{10} 01^6 001^{10} 01^{10}$ (10 4 6 10 2 9 6 1 9 1 10 6 10 10)	11	(91, 94)
$w(2; 3, 12) = 135$	$1^9 01^8 01^9 01^2 01^3 01^1 01^7 01^2 01 01^3 01^{11} 0^2$ $1^{11} 01^3 01 01^2 01^7 01^1 01^3 01^2 01^9 01^8 01^9$ (9 8 9 2 3 1 7 2 1 3 11 11 3 1 2 7 1 3 2 9 8 9)	17	(112, 112)
$w(2; 3, 13) = 160$	$1^1 01^6 01^{12} 01^4 001^{11} 001^6 01^{10} 01^2 01^4 01^{11} 01^1 0$ $1^6 01^9 01^2 01^3 01^7 01^{10} 01^1 001^5 01^{12} 01^5 01^4 01^2$ (1 6 12 4 11 6 10 2 4 11 1 6 9 2 3 7 10 1 5 12 5 4 2)	15	(130, 134)

Again, there seems little variation concerning the total number of ones:

Conjecture 4.3. There exists a constant $c > 0$ such that $|n_1(B) - n_1(B')| \leq ct, \forall B, B' \in S(t)$ with $t \geq 3$.

Stronger than Conjecture 4.3, the number of ones seems very close to the vdW-number for the previous t :

Conjecture 4.4. There exists a constant $c > 0$ such that $|w(2; 3, t - 1) - n_1(B)| < ct, \forall B \in S(t)$.

This conjecture also implies the earlier conjecture on the quadratic growth of $w(2; 3, t)$:

Lemma 4.1. Conjecture 4.4 implies Conjecture 4.3 and Conjecture 3.1.

Proof. Conjecture 4.3 follows by the triangle inequality. Conjecture 3.1 follows, if for t large enough we can show $n_0(B) \leq n_1(B)$ for all $B \in S(T)$, and this is a special case of Szemerédi’s Theorem ([66]), which for arithmetic progressions of size 3 was already proven in [59]¹⁸, namely that the relative size of maximum independent subsets of the hypergraph of arithmetic progressions of size 3 in the numbers $1, \dots, t$ goes to 0 with $t \rightarrow \infty$. \square

We turn to the growth of the number of peaks and valleys:

Question 4.1. For each positive constant c does there exist a t' such that for all $t \geq t', n_p(\text{EPIS}(B)) + n_v(\text{EPIS}(B)) \geq ct, (t \geq 3) \forall B \in S(t)$? (We conjecture yes).

We conclude with the observation, that for $t > 3$ there do not seem to be long plateaus for the numbers of ones:

Question 4.2. Is there a good partition $B \in S(t), (t \geq 4)$ with 3 consecutive numbers equal in $\text{EPIS}(B)$? (Note that, for $t = 3$, the partition $1^1 01^1 001^1 01^1$ has four consecutive exponents, which are the same.)

4.3. How can it help for SAT solving?

If one of the above conjectures (or some other conjecture) turns out to be true, and if moreover the numerical constants have good estimates, then they can be used to restrict the search space. When using a general purpose SAT solver, this can be achieved by adding further constraints. It seems however that these constraints do not help with the search, even if we assume that they are true, since they are too difficult to handle for the solver. It seems the problem is that these constraints do not mix well with the original problem formulation, and a deeper integration is needed. Such an integration was achieved in the case of the palindromic constraint, which is the subject of the following section — here an organic new problem formulation could be established, where the additional restriction doesn’t appear as an “add-on”, but establishes a natural new problem class.

¹⁸see <http://rothstheorem.wikidot.com/on-certain-sets-of-integers>

5. Palindromes

Recall Definitions 1.1, 1.2:

1. for given $k \in \mathbb{N}$ (the number of “colours”),
2. t_0, \dots, t_{k-1} (the lengths of arithmetic progressions),
3. and $n \in \mathbb{N}$ (the number of vertices)

we consider block partitions (P_0, \dots, P_{k-1}) of $\{1, \dots, n\}$ such that no P_i contains an arithmetic progression of length t_i — these are the “good partitions”, and $w(k; t_0, \dots, t_{k-1}) \in \mathbb{N}$ is the smallest n such that no good partition exists. If (P_0, \dots, P_{k-1}) is a good partition of $\{1, \dots, n\}$ w.r.t. t_0, \dots, t_{k-1} , then for $1 \leq n' \leq n$ we obtain a good partition of $\{1, \dots, n'\}$ w.r.t. t_0, \dots, t_{k-1} by just removing vertices $n'+1, \dots, n$ from their blocks. Thus $w(k; t_0, \dots, t_{k-1})$ completely determines for which $n \in \mathbb{N}$ good partitions exist, namely exactly for $n < w(k; t_0, \dots, t_{k-1})$.

Definition 5.1. For $n \in \mathbb{N}$ let $m_n : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ (with “m” like “mirror”) defined by $m_n(v) := n + 1 - v$. This map is extended to $S \subseteq \{1, \dots, n\}$ as usual: $m_n(S) := \{m_n(v) : v \in S\}$.

Now if (P_0, \dots, P_{k-1}) is a good partition w.r.t. n , then also $(m_n(P_0), \dots, m_n(P_{k-1}))$ is a good partition w.r.t. n . So it is of interest to consider self-symmetric partitions (with $m_n(P_i) = P_i$ for all i):

Definition 5.2. A *good palindromic partition* of $\{1, \dots, n\}$ w.r.t. parameters t_0, \dots, t_{k-1} , where $n, t_0, \dots, t_{k-1} \in \mathbb{N}$, is a good partition of $\{1, \dots, n\}$ w.r.t. t_0, \dots, t_{k-1} such that for all $j \in \{0, \dots, k-1\}$ holds $m_n(P_j) = P_j$.

We call these special good partitions “palindromic”, since a block partition can be represented as a string of numbers over $\{0, \dots, k-1\}$, and then the block partition is palindromic iff the string is a palindrome (reads the same forwards and backwards). For example, the string 01^2001^20 represents a good palindromic partition for $k = 2$, $t_0 = t_1 = 3$ and $n = 8$, namely $(\{1, 4, 5, 8\}, \{2, 3, 6, 7\})$, and so does $(\{1, 3, 6, 8\}, \{2, 4, 5, 7\})$, represented by 0101^2010 , while $(\{1, 2, 5, 6\}, \{3, 4, 7, 8\})$, represented by 001^2001^2 , is a good partition which is not palindromic.

For given k and t_0, \dots, t_{k-1} again we want to completely determine (in theory) for which n do good palindromic partitions exist and for which not. The key is the following observation (which follows also from Lemmas 5.2, 5.3).

Lemma 5.1. Consider fixed k, t_0, \dots, t_{k-1} , and $n \geq 3$. From a good palindromic partition (P_0, \dots, P_{k-1}) of $\{1, \dots, n\}$ we obtain a good palindromic partition (P'_0, \dots, P'_{k-1}) of $\{1, \dots, n-2\}$ by removing vertices $1, n$ and replacing the remaining vertices v by $v-1$, that is, $P'_i := \{v-1 : v \in P_i \setminus \{1, n\}\}$.

Proof. The notion of a good partition of $\{1, \dots, n\}$ w.r.t. $w(k; t_0, \dots, t_{k-1})$, as defined in Definition 1.2, can be generalised to good partitions of arbitrary $T \subseteq \mathbb{Z}$ by demanding that for every block partition (P_0, \dots, P_{k-1}) of T into k parts no part P_j contains an arithmetic progression of size t_j . In the remainder of the proof we omit the “w.r.t. t_0, \dots, t_{k-1} ”.

If T has a good partition, then also every subset has a good partition, by restricting the blocks accordingly, and for every $d \in \mathbb{Z}$ also $d + T = \{d + x : x \in T\}$ has a good partition, by shifting the blocks as well.

We can also generalise the notion of a good palindromic partition to intervals $T = \{a, a+1, \dots, b\} \subset \mathbb{Z}$ for $a < b$, defining now the mirror-map $m_{a,b} : T \rightarrow T$ via $v \in T \mapsto b + a - v$ (m_n in Definition 5.1 is the special case $m_n = m_{1,n}$).

Again, if T has a good palindromic partition, then $d + T$ for $d \in \mathbb{Z}$ has as well. But for subsets of T we can only consider sub-intervals $T' = \{a', \dots, b'\}$, where from both sides we have taken away equal amounts. That is, for $a \leq a' < b' \leq b$ with $a' - a = b - b'$ we have, that from a good palindromic partition for T we can obtain a good palindromic partition for T' (by just restricting the blocks).

So from a good palindromic partition of $\{1, \dots, n\}$ we obtain a good palindromic partition of $\{1, \dots, n-2\}$ by first restricting to $\{2, \dots, n-1\}$ and then shifting by -1 . \square

Corollary 5.1.1. If there is no good palindromic partition of $\{1, \dots, n\}$, then there is no good palindromic partition of $\{1, \dots, n + 2 \cdot i\}$ for all $i \in \mathbb{N}_0$.

Proof. If there would be a good palindromic partition of $\{1, \dots, n + 2 \cdot i\}$, then by repeated applications of Lemma 5.1 we would obtain a good palindromic partition of $\{1, \dots, n\}$. \square

Since by van der Waerden's theorem we know there always exists some n such that for all $n' \geq n$ no good palindromic partition exists, we get that the existence of good palindromic partitions w.r.t. fixed t_0, \dots, t_{k-1} is determined by two numbers, the endpoint p of “always exists” resp. q of “never exists”, with alternating behaviour in the interval in-between:

Corollary 5.1.2. Consider the maximal $p \in \mathbb{N}_0$ such that for all $n \leq p$ good palindromic partitions exist, and the minimal $q \in \mathbb{N}$ such that for all $n \geq q$ no good palindromic partitions exist. Then $q - p$ is an odd natural number, where no good palindromic partition exists for $p + 1$, but $p + 2$ again has a good palindromic partition, and so on alternately, until from q on no good palindromic partition exists anymore.

Proof. By Corollary 5.1.1 there is no good palindromic partition for $p + 1 + 2i$ and all $i \in \mathbb{N}_0$. Now for the first $i \in \mathbb{N}_0$, such that $p + 2 + 2i$ has no good palindromic partition, we let $q' := (p + 2 + 2i) - 1$. We have a good palindromic partition for $q - 1$ by definition of i (as the smallest such i) resp. in case of $i = 0$ by definition of p . We have $q' + 2j = (p + 2 + 2i) - 1 + 2j = p + 1 + 2(i + j)$ for $j \in \mathbb{N}_0$, and thus there is no good palindromic partition for $q' + 2j$. And if there would be a good palindromic partition for $q' + 1 + 2j = p + 2 + 2i + 2j$, then by Corollary 5.1.1 there would be a good palindromic partition for $p + 2 + 2i$. So we have $q' = q$. \square

Definition 5.3. The *palindromic van-der-Waerden number* $\text{pdw}(k; t_0, \dots, t_{k-1}) \in \mathbb{N}_0^2$ is defined as the pair (p, q) such that p is the largest $p \in \mathbb{N}_0$ with the property, that for all $1 \leq n \leq p$ there exists a good palindromic partition of $\{1, \dots, n\}$, while q is the smallest $q \in \mathbb{N}$ such that for no $n \geq q$ there exists a good palindromic partition of $\{1, \dots, n\}$. We use $\text{pdw}(k; t_0, \dots, t_{k-1})_1 = p$ and $\text{pdw}(k; t_0, \dots, t_{k-1})_2 = q$. So $0 \leq \text{pdw}(k; t_0, \dots, t_{k-1})_1 < \text{pdw}(k; t_0, \dots, t_{k-1})_2 \leq w(k; t_0, \dots, t_{k-1})$.

The *palindromic gap* is

$$\text{pdg}(k; t_0, \dots, t_{k-1}) := w(k; t_0, \dots, t_{k-1}) - \text{pdw}(k; t_0, \dots, t_{k-1})_2 \in \mathbb{N}_0,$$

while the *palindromic span* is defined as

$$\text{pds}(k; t_0, \dots, t_{k-1}) := \text{pdw}(k; t_0, \dots, t_{k-1})_2 - \text{pdw}(k; t_0, \dots, t_{k-1})_1 \in \mathbb{N}.$$

To certify that $w(k; t_0, \dots, t_{k-1}) = n$ holds means to show that there exists a good partition of $\{1, \dots, n - 1\}$ and that there is no good partition of n . For palindromic number-pairs we need to double the effort:

Theorem 5.1. To certify that $\text{pdw}(k; t_0, \dots, t_{k-1}) = (p, q)$ holds, exactly the following needs to be shown for (arbitrary) $p \in \mathbb{N}_0, q \in \mathbb{N}$ with $p < q$:

- (i) there are good palindromic partitions of $\{1, \dots, p - 1\}$ and $\{1, \dots, q - 1\}$ w.r.t. t_0, \dots, t_{k-1} ;
- (ii) there are no good palindromic partitions of $\{1, \dots, p + 1\}$ and $\{1, \dots, q + 1\}$ w.r.t. t_0, \dots, t_{k-1} .

Proof. The given conditions are necessary for $\text{pdw}(k; t_0, \dots, t_{k-1}) = (p, q)$ by the defining properties of p and q . We show that they are sufficient to establish $\text{pdw}(k; t_0, \dots, t_{k-1}) = (p, q)$. First we have by Corollary 5.1.1 that $q - p$ is odd, since otherwise $p + 1$ having no good palindromic partitions would yield that $q - 1$ would have no good palindromic partition. Then, again by Corollary 5.1.1, all $n \geq q + 1$ have no good palindromic partition, while all $n \leq p - 1$ have good palindromic partitions. By Corollary 5.1.2 we must now have $\text{pdw}(k; t_0, \dots, t_{k-1}) = (p, q)$. \square

5.1. Palindromic vdW-hypergraphs

Recall that a finite hypergraph G is a pair $G = (V, E)$, where V is a finite set (of “vertices”) and E is a set of subsets of V (the “hyperedges”); one writes $V(G) := V$ and $E(G) := E$. The essence of the (finite) van der Waerden problem (which we will now often abbreviate as “vdW-problem”) is given by the hypergraphs $\text{ap}(t, n)$ of arithmetic progressions with progression length $t \in \mathbb{N}$ and the number $n \in \mathbb{N}_0$ of vertices:

- $V(\text{ap}(t, n)) := \{1, \dots, n\}$
- $E(\text{ap}(t, n)) := \{p \subseteq \{1, \dots, n\} : p \text{ arithmetic progression of length } t\}$.

For example $\text{ap}(3, 5) = (\{1, 2, 3, 4, 5\}, \{\{1, 2, 3\}, \{2, 3, 4\}, \{1, 3, 5\}, \{3, 4, 5\}\})$. Considering hypergraphs, the reader might wonder how determination of vdW-numbers fits with hypergraph colouring. While the determination of diagonal vdW-numbers is an ordinary hypergraph colouring problem, for general vdW-numbers a more general concept of hypergraph colouring is to be used, involving the simultaneous colouring of several hypergraphs in the following sense: The diagonal vdW-number $w(k; t, \dots, t)$ for $k, t \in \mathbb{N}$ is the smallest $n \in \mathbb{N}$ such that the hypergraph $\text{ap}(t, n)$ is not k -colourable, where in general a k -colouring of a hypergraph G is a map $f : V(G) \rightarrow \{1, \dots, k\}$ such that no hyperedge is “monochromatic”, that is, every hyperedge gets at least two different values by f . For the general vdW-number $w(k; t_0, \dots, t_{k-1})$ we now consider for each colour $i \in \{0, \dots, k-1\}$ the hypergraph $\text{ap}(t_i, n)$, and we forbid (to formulate “good partition”) for each i that there is a hyperedge in $\text{ap}(t_i, n)$ monocoloured with colour i (while we do not care about the other colours here). Accordingly the SAT-encoding of “ $w(2; 3, t) > n$?”, as discussed in Subsection 1.1, exactly consists of the two hypergraphs $\text{ap}(3, n)$ and $\text{ap}(t, n)$ represented by positive resp. negative clauses.

The task now is to define the palindromic version $\text{pdap}(t, n)$ of the hypergraph of arithmetic progressions, so that for diagonal palindromic vdW-numbers $\text{pdw}(k; t, \dots, t) = (p, q)$ we have, that q is minimal for the condition that for all $n \geq q$ the hypergraph $\text{pdap}(t, n)$ is not k -colourable, while p is maximal for the condition that for all $n \leq p$ the hypergraph is k -colourable. Furthermore we should have that for two-coloured problems (i.e., $k = 2$) the SAT-encoding of “ $\text{pdw}(2; t_0, t_1) > n$?” (satisfiable iff the answer is yes) consists exactly of the two hypergraphs $\text{pdap}(t_0, n)$, $\text{pdap}(t_1, n)$ represented by positive resp. negative clauses (while for more than two colours generalised clause-sets can be used; see [50]).

Consider fixed $t \in \mathbb{N}$ and $n \in \mathbb{N}_0$. Obviously $\text{pdap}(t, 0) := \text{ap}(t, 0) = (\{\}, \{\})$, and so assume $n \geq 1$. Recall the permutation $m = m_n$ of $\{1, \dots, n\}$ from Definition 5.1. As every permutation, m induces an equivalence relation \sim on $\{1, \dots, n\}$ by considering the cycles, which here, since m is an involution (self-inverse), just has the equivalence classes $\{1, \dots, n\} / \sim = \{\{v, f(v)\}\}_{v \in \{1, \dots, n\}}$ of size 1 or 2 comprising the elements and their images. Note that m has a fixed point (an equivalence class of size 1) iff n is odd, in which case the unique fixed point is $\frac{n+1}{2}$. The idea now is to define $m' : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$, which chooses from each equivalence class one representative (so $m'(v) \in \{v, m(v)\}$ and $v \sim w \Leftrightarrow m'(v) = m'(w)$), and to let $\text{pdap}(t, n)$ be the image of $\text{ap}(t, n)$ under m' , that is, $(m'(V(\text{ap}(t, n))), \{m'(H)\}_{H \in E(\text{ap}(t, n))})$. Naturally we choose $m'(v)$ to be the smaller of v and $m(v)$. Now it occurs that images of arithmetic progressions under m' can subsume each other, i.e., for $H_1, H_2 \in E(\text{ap}(t, n))$ with $H_1 \neq H_2$ we can have $m'(H_1) \subset m'(H_2)$, and so we define $\text{pdap}(t, n)$ as the image of $\text{ap}(t, n)$ under m' , where also all subsumed hyperedges are removed (so we only keep the minimal hyperedges under the subset-relation).

Definition 5.4. For $t \in \mathbb{N}$ and $n \in \mathbb{N}_0$ the hypergraph $\text{pdap}(t, n)$ is defined as follows:

- $V(\text{pdap}(t, n)) := \{1, \dots, \lceil \frac{n}{2} \rceil\}$
- $E(\text{pdap}(t, n))$ is the set of minimal elements w.r.t. \subseteq of the set of $m'_n(H)$ for $H \in E(\text{ap}(t, n))$, where $m'_n : \{1, \dots, n\} \rightarrow V(\text{pdap}(t, n))$ is defined by $m'_n(v) := v$ for $v \leq \lceil \frac{n}{2} \rceil$ and $m'_n(v) := n + 1 - v$ for $v > \lceil \frac{n}{2} \rceil$.

Using $\text{ap}(3, 5) = (\{1, 2, 3, 4, 5\}, \{\{1, 2, 3\}, \{2, 3, 4\}, \{1, 3, 5\}, \{3, 4, 5\}\})$ as above, we have $m'(\{1, 2, 3\}) = \{1, 2, 3\}$, $m'(\{2, 3, 4\}) = \{2, 3\}$, $m'(\{1, 3, 5\}) = \{1, 3\}$ and $m'(\{3, 4, 5\}) = \{1, 2, 3\}$, whence $\text{pdap}(3, 5) = (\{1, 2, 3\}, \{\{1, 3\}, \{2, 3\}\})$.

Lemma 5.2. Consider $t \in \mathbb{N}$ and $n \in \mathbb{N}_0$. The hypergraph $\text{pdap}(t, n)$ is embedded into the hypergraph $\text{pdap}(t, n + 2)$ via the map $e : V(\text{pdap}(t, n)) \rightarrow V(\text{pdap}(t, n + 2))$ given by $v \mapsto v + 1$.

Proof. First we note that $|V(\text{pdap}(t, n + 2))| = |V(\text{pdap}(t, n))| + 1$, and so the range of e is $V(\text{pdap}(t, n + 2)) \setminus \{1\}$. Let G be the hypergraph with vertex set $V(\text{pdap}(t, n + 2)) \setminus \{1\}$, whose hyperedges are all those hyperedges $H \in E(\text{pdap}(t, n + 2))$ with $1 \notin H$. We show that e is an (hypergraph-)isomorphism from $\text{pdap}(t, n)$ to G , which proves the assertion.

Now obviously the underlying hypergraph $\text{ap}(t, n)$ is embedded into the underlying $\text{ap}(t, n + 2)$ via the underlying map $v \in V(\text{ap}(t, n)) \mapsto v + 1 \in V(\text{ap}(t, n + 2))$, where the image of this embedding is given by the hypergraph with vertex set $V(\text{ap}(t, n + 2)) \setminus \{1, n + 2\}$, and where the hyperedges are those $H \in E(\text{ap}(t, n + 2))$ with $1, n + 2 \notin H$. Since $m'_{n+2}(n + 2) = 1$ and $m'_n(v) = m'_{n+2}(v + 1) - 1$ for $v \in \{1, \dots, n\}$, the assertion follows from the fact that there are no hyperedges $H, H' \in E(\text{ap}(t, n + 2))$ with $H \cap \{1, n + 2\} \neq \emptyset$, $H' \cap \{1, n + 2\} = \emptyset$ and $m'_{n+2}(H) \subset m'_{n+2}(H')$ (thus $m'_{n+2}(H')$ can only be removed from $\text{pdap}(t, n + 2)$ by subsumptions already at work in $\text{pdap}(t, n)$), and this is trivial since $1 \in m'_{n+2}(H)$ but $1 \notin m'_{n+2}(H')$. \square

The SAT-translation of “Is there a good palindromic partition of $\{1, \dots, n\}$ w.r.t. t_0, t_1 ?” is accomplished similar to the translation of “ $w(2; t_0, t_1) > n$?”, now using $\text{pdap}(t_0, n)$, $\text{pdap}(t_1, n)$ instead of $\text{ap}(t_0, n)$, $\text{ap}(t_1, n)$:

Lemma 5.3. Consider $t_0, t_1 \in \mathbb{N}$, $t_0 \leq t_1$, and $n \in \mathbb{N}_0$. Let the boolean clause-set $F^{\text{pd}}(t_0, t_1, n)$ be defined as follows:

- the variable-set is $\{1, \dots, \lceil \frac{n}{2} \rceil\}$ ($= V(\text{pdap}(t_0, n)) = V(\text{pdap}(t_1, n))$);
- the hyperedges of $\text{pdap}(t_0, n)$ are directly used as positive clauses;
- the hyperedges H of $\text{pdap}(t_1, n)$ yield negative clauses $\{\bar{v}\}_{v \in H}$.

Then there exists a good palindromic partition if and only if $F^{\text{pd}}(t_0, t_1, n)$ is satisfiable, where the satisfying assignments are in one-to-one correspondence to the good palindromic partitions of $\{1, \dots, n\}$ w.r.t. (t_0, t_1) . \square

For more than two colours, Lemma 5.3 can be generalised by using generalised clause-sets, as in [50], and there one also finds the “generic translation”, a general scheme to translate generalised clause-sets (with non-boolean variables) into boolean clause-sets (see also [51, 52]).

5.2. Precise values

See Subsection 6.1 for details of the computation.

Table 6: Palindromic vdW-numbers $\text{pdw}(2; 3, t)$

t	$\text{pdw}(2; 3, t)$	$\text{pds}(2; 3, t)$	$\text{pdg}(2; 3, t)$
3	(6, 9)	3	0
4	(15, 16)	1	2
5	(16, 21)	5	1
6	(30, 31)	1	1
7	(41, 44)	3	2
8	(52, 57)	5	1
9	(62, 77)	15	0
10	(93, 94)	1	3
11	(110, 113)	3	1
12	(126, 135)	9	0
13	(142, 155)	13	5
14	(174, 183)	9	3
15	(200, 205)	5	13
16	(232, 237)	5	1
17	(256, 279)	23	0
18	(299, 312)	13	0
19	(338, 347)	9	2
Continued on Next Page...			

Table 6: Palindromic vdW-numbers $\text{pdw}(2; 3, t)$

t	$\text{pdw}(2; 3, t)$	$\text{pds}(2; 3, t)$	$\text{pdg}(2; 3, t)$
20	(380, 389)	9	≥ 0
21	(400, 405)	5	≥ 11
22	(444, 463)	19	≥ 1
23	(506, 507)	1	≥ 9
24	(568, 593)	25	≥ 0
25	(586, 607)	21	≥ 49
26	(634, 643)	9	≥ 84
27	(664, 699)	35	≥ 71

5.3. Conjectured values and bounds

For $28 \leq t \leq 39$ we have reasonable values on $\text{pdw}(2; 3, t)$, which are given in Table 7, and which we believe to be exact for $t \leq 35$. These values have been computed by local-search methods (see Subsection 6.2), and thus for sure we can only say that they present lower bounds. We obtain conjectured values for the palindromic span (which might however be too large or too small) and conjectured values for the palindromic gap (which additionally depend on the conjectured values from Subsection 3.2, while for $t \geq 31$ we only have the lower bounds from Subsection 3.2).

Table 7: Conjectured palindromic vdW-numbers $\text{pdw}(2; 3, t)$

t	$\text{pdw}(2; 3, t) \geq$	$\text{pds}(2; 3, t) \sim$	$\text{pdg}(2; 3, t) \sim$
28	(728, 743)	15	84
29	(810, 821)	11	47
30	(844, 855)	11	48
31	(916, 931)	15	0
32	(958, 963)	5	44
33	(996, 1005)	9	59
34	(1054, 1081)	27	63
35	(1114, 1155)	41	50
36	(1186, 1213)	27	45
37	(1272, 1295)	23	44
38	(1336, 1369)	33	10
39	(1406, 1411)	5	8

For the certificates for these lower bounds see Appendix A.3.

5.4. Open problems

The relation between ordinary and palindromic vdW-numbers are of special interest:

- It seems the palindromic span can become arbitrarily large — also in relative terms? Perhaps the span shows a periodic behaviour, oscillating between small and large?
- Similar questions are to be asked for the gap. Does it attain value 0 infinitely often?

Do the hypergraphs $\text{pdap}(t, n)$ have interesting properties (more basic than their chromatic numbers)? A basic exercise would be to estimate the number of hyperedges and their sizes. In the subsequent Subsection 6.1 we find data that SAT solvers behave rather different on palindromic vdW-problems (compared to ordinary problems). It seems that palindromic problems are more “structured” than ordinary problems — can this be made more precise? Perhaps the hypergraphs $\text{pdap}(t, n)$ show characteristic differences to the hypergraphs $\text{ap}(t, n)$, which could explain the behaviour of SAT solvers?

5.5. Remarks on the use of symmetries

The heuristic use of symmetries for finding good partitions has been studied in [27, 30, 29] (while for symmetries in the context of general SAT solving see [60]). Especially we find there an emphasis on “internal symmetries”, which are not found in the problem, but are imposed on the solutions.

The good palindromic partitions introduced in this section are more restricted in the sense, that they are based on the symmetries m from Subsection 5.1 of the clause-sets F expressing “ $w(k; t_0, t_1, \dots, t_{k-1}) > n$?” (i.e., we have $m(F) = F$; recall Subsection 1.1), which then is imposed as an internal symmetry on the potential solution by demanding that the solutions be self-symmetric. In [27] “reflection symmetric” certificates are mentioned, which for even n are the same as good palindromic partitions, however for odd n they ignore vertex 1, not the mid-point $\lceil \frac{n}{2} \rceil$ as we do. This definition in [27] serves to maintain monotonicity (i.e., a solution for $n + 1$ yields a solution for n , while we obtain one only for $n - 1$ (Lemma 5.1)). We believe that palindromicity is a more natural notion, but further studies are needed here to compare these two notions.

Other internal symmetries used in [27, 30, 29] are obtained by modular additions and multiplications (these are central to the approaches there), based on the method from [56] for constructing lower bounds for diagonal vdW-numbers. No generalisations are known for the mixed problems we are considering.

Finally we wish to emphasise that we do not consider palindromicity as a mere heuristic for finding lower bounds, but we get an interesting variation of the vdW-problem in its own right, which hopefully will help to develop a better understanding of the vdW-problem itself in the future.

6. Experiments with SAT solvers

We conclude by summarising the experimental results and insights gained by running SAT solvers on the instances considered in this paper. All the solvers (plus build environments), generators and the data are available in the `OKLibrary` ([47]); see Appendix B for more information.

For determining unsatisfiability we consider complete SAT solvers in Subsection 6.1. In general, for (ordinary) vdW-problems look-ahead solvers seem to perform better than conflict-driven solvers, while for palindromic problems it seems to be the opposite. However `tauSolver-2.6` is the best (single) solver for both classes.

The hybrid approach, `Cube & Conquer`, was developed precisely on the instances of this paper, as discussed in [32] (further developments one finds in [68]). This approach is third-best on vdW-problems (after `tauSolver-2.6` and `rawSolver-2.6`), and best on palindromic vdW-problems (before `rawSolver-2.6` and `tauSolver-2.6`).

We conclude this section in Subsection 6.2 by remarks on incomplete SAT solvers, used to obtain lower bounds (determine satisfiability).

For the experiments we used a 64-bit workstation with 32 GB RAM and Intel i5-2320 CPUs (6144 KB cache) running with 3 GHz, where we only employed a single CPU.

6.1. Complete solvers

Complete SAT solvers exist in mainly two forms, “look-ahead solvers” and “conflict-driven solvers”; see [55, 31] for general overviews on these solver paradigms. Besides the `tauSolver` (see Section 2), for our experimentation we use the following (publicly available) complete solvers, which give a good coverage of state of the art SAT solving and of the winners of recent SAT competitions and SAT races¹⁹:

¹⁹The (parent) SAT competition homepage is at <http://www.satcompetition.org> with links to each individual competition.

- Look-ahead solvers:

- `OKsolver` ([46]), a solver with well-defined behaviour, no ad-hoc heuristics, and which applies complete r_2 (at every node). This solver won gold at the SAT 2002 competition.
- `satz` ([54]), a solver which applies partial r_2 and r_3 . In the `OKlibrary` we maintain version 215, with improved/corrected in/output and coding standard.
- `march_pl` ([28]), a solver applying partial r_2 , r_3 , and resolution- and equivalence-preprocessing. `march_pl` contains the same underlying technology as its sibling solvers `march_{rw,hi,ks,dl,eq}`, which won gold, silver and bronze at the 2004 to 2011 SAT competitions and SAT races. We use the `pl` (partial lookahead) version.

- Conflict-driven solvers:

- `MiniSat` family:
 - * `MiniSat` ([20]), version 2.0 and 2.2, the latest version of this well-established solver, used as starting point for many new conflict-driven solvers. Previous versions won gold at the SAT Race 2006 and 2008, as well as numerous bronze and silver awards at the SAT competition 2007.
 - * `CryptoMiniSat` ([64]), a `MiniSat` derivative designed specifically to tackle hard cryptographic problems. This solver won gold at SAT Race 2010 and gold and silver at the SAT competition 2011. We use version 2.9.6.
 - * `Glucose` ([6]), a `MiniSat` derivative utilising a new clause scoring scheme and aggressive learnt-clause deletion. This solver won gold in both SAT 2011 competition and SAT Challenge 2012. We use versions 2.0 and 2.2.
- `Lingeling` family:
 - * `PicoSAT` ([8]), a conflict-driven solver using an aggressive restart strategy, compact data-structures, and offering proof-trace options to allow for unsatisfiability checking. This solver won gold and silver at the SAT competition 2007. We use the latest version 913.
 - * `PrecoSAT` ([10]), integrates the `SATeLite` preprocessor into `PicoSAT`, applying various reductions including partial r_2 at certain nodes in the search tree. This solver won gold and silver at the SAT 2009 competition. We use the latest version 570.
 - * `Lingeling`([11]), based on `PrecoSAT`, focuses further on integrating preprocessing and search, introducing new algorithms and data-structures to speed up these techniques and reduce memory footprint. As with `PrecoSAT`, this solver applies partial r_2 at specially chosen nodes in the search tree. This solver won bronze at the SAT 2011 competition and silver at the SAT Race 2010. We use the latest version `ala-b02aa1a-121013`.

6.1.1. *Cube-and-Conquer*

The `Cube & Conquer` method uses a look-ahead solver as the “cube-solver”, splitting the instance into subinstances small enough such that the “conquer-solver”, a conflict-driven solver, can solve almost all sub-instances in at most a few seconds. We use the `OKsolver` as the cube-solver and `MiniSat` as the conquer-solver. The main (and single) parameter is $D \in \mathbb{N}_0$, the cut-off depth for the `OKsolver`: the DLL-tree created by the `OKsolver` is cut off when the number of assignments reaches D , where it is important that this includes *all assignments* on the path, not just the decisions, but also the forced assignments found by r_1 and r_2 — only in this way a relatively balanced load is guaranteed. The data reported in Tables 10, 14 shows first data on the cube-phase, namely

- D (cut-off depth),
- the number of nodes in the (truncated) DLL-tree of the `OKsolver`,
- the time needed (this includes writing the partial assignments representing the sub-instances to files),
- and the number N of sub-instances.

For the conquer-phase we have:

- the median and maximum time for solving the sub-instances by MiniSat,
- the sum of conflicts over all sub-instances,
- and the total time used by MiniSat.

Finally the overall total time is reported, which does not include the time used by the processing-script, which applies the partial assignments to the original instance and produces so the sub-instances: this adds an overhead of nearly 20% for the smallest problem, but this proportion becomes smaller for larger problems, and is less than 1% for the largest problems.

6.1.2. VdW-problems

We consider the (unsatisfiable) instances to determine the upper bounds for $w(2; 3, t)$ with $12 \leq t \leq 17$; in Table 8 we give basic data for these instances (plus $t = 18, 19$).

t	n	c	c_3	c_t	ℓ
12	135	5,251	4,489	762	22,611
13	160	7,308	6,320	988	31,804
14	186	9,795	8,556	1,239	43,014
15	218	13,362	11,772	1,590	59,166
16	238	15,812	14,042	1,770	70,446
17	279	21,616	19,321	2,295	96,978
18	312	26,889	24,180	2,709	121,302
19	349	33,487	30,276	3,211	151,837

Table 8: Instance data for $F(3, t, n)$, where n is the number of vertices as well as the number of variables, $c = c_3 + c_t$ is the number of clauses, c_i the number of clauses of length i , and $\ell = 3c_3 + tc_t$ is the number of literal occurrences.

In Table 9, we see the running times and number of nodes/conflicts for the SAT solvers. We see that in general look-ahead solvers here have the upper hand over conflict-driven solvers, with the `tawSolver-2.6` with a large margin the fastest solver. Regarding conflict-driven solvers, we see that version 2.2 for MiniSat is superior over version 2.0, while for Glucose it is the opposite. The low node-count for `march.pl` seems due to the preprocessing phase, which adds a large number of resolvents to the original instance: this reduces the node-count, but increases the run-time. Compared to the other look-ahead solvers, the strength of `tawSolver-2.6` is that the number of nodes is just larger by a factor of most 3, while processing of each node happens much faster. Compared with the strongest conflict-driven solver, MiniSat-2.2, we see that the node-count of `tawSolver-2.6` is considerably less than the number of conflicts used by MiniSat, and that one node is processed somewhat faster than one conflict.

One aspect important here for the superiority of look-ahead solver is the “tightness” of the problem formulation. Consider for example $t = 12$, not with $n = 135$ as in Tables 8, 9, but with $n = 1000$; this yields $c = 294,455$, $c_3 = 249,500$, $c_{12} = 44,955$, and $\ell = 1,287,960$, which is now a highly redundant problem instance. For `tawSolver-2.6` we obtain 1,311,511 nodes and 2,868 sec, and for `tauSolver-2.6` we get 935,475 nodes and 2,452 sec, while for MiniSat-2.2 we get 1,140,616 conflicts and 159 sec. We see that MiniSat-2.2 was able to utilise the additional clauses to determine unsatisfiability with fewer conflicts, and with a run-time not much affected by the large increase in problem size, while for `tawSolver-2.6` the run-time (naturally) explodes, and the number of nodes stayed the same.²⁰ If we consider a typical branching-heuristics for look-ahead solvers (as discussed in Subsection 2.3), then we see that locality of the search process is not taken into consideration, and thus for non-tight problem formulations the solver can “switch attention” again and again. This is very different from heuristics for conflict-driven solvers, which

²⁰The `OKsolver` yields a more extreme example: the run was aborted after 657,648 sec and 603,177 nodes, where yet only %49.2 of the search space was visited (so that the solver was still working on completing the first branch at the root of the tree, making very slow progress towards %50). This shows the big overhead caused by the r_2 -reduction, and the danger of a heuristic which (numerically) sees opportunities “all over the place”, and thus can not focus on one relevant part of the input.

via “clause-activity” have a strong focus on locality of reasoning. Furthermore, look-ahead solvers consider much more of the whole input, for example the `tawSolver` considers always all remaining variables and their occurrences for the branching heuristic, while conflict-driven solvers do not use such global heuristics.

$t =$	12	13	14	15	16	17
<code>tawSolver-2.6</code>	11 961,949	83 5,638,667	673 35,085,795	5,010 194,035,915	42,356 1,462,429,351	401,940 10,258,378,909
<code>rawSolver-2.6</code>	19 953,179	143 5,869,055	1,068 35,668,687	7,607 200,208,507	59,585 1,479,620,647	
<code>tawSolver-1.0</code>	47 1,790,733	463 13,722,975	4,577 102,268,511	47,006 774,872,707	532,416 8,120,609,615	
<code>satz</code>	77 262,304	711 1,698,185	6,233 10,822,316	54,913 66,595,028	562,161 599,520,428	
<code>march_pl</code>	185 47,963	1,849 279,061	17,018 1,975,338	175,614 11,959,263		
<code>OKsolver</code>	216 281,381	3,806 2,970,723	47,598 22,470,241			
<code>MiniSat-2.2</code>	107 5,963,349	1,716 63,901,998	16,836 463,984,635	190,211 3,205,639,994		
<code>MiniSat-2.0</code>	273 1,454,696	3,022 9,298,288	33,391 60,091,581	274,457 314,678,660		
<code>PrecoSAT</code>	211 2,425,722	2,777 16,978,254	47,624 140,816,236			
<code>PicoSAT</code>	259 9,643,671	4,258 82,811,468	48,372 576,692,221			
<code>Glucose-2.0</code>	58 1,263,087	781 8,377,487	84,334 163,500,051			
<code>Lingeling</code>	519 1,659,607	7,651 24,124,525	107,243 176,909,499			
<code>CryptoMiniSat</code>	212 2,109,106	4,630 18,137,202	141,636 205,583,043			
<code>Glucose-2.2</code>	94 1,444,017	1,412 10,447,051	>940,040 aborted			

Table 9: Complete solvers on unsatisfiable instances $F(3, t; n)$ for computing $w(2; 3, t)$ (with $t = 12, \dots, 16$ and $n = 135, 160, 186, 218, 238$). The first line is run-time in seconds, the second line is the number of nodes for look-ahead solvers resp. number of conflicts for conflict-driven solvers.

Finally we consider `Cube & Conquer`, with the `OKsolver` as `Cube-solver` and `MiniSat-2.2` as `Conquer-solver`, in Table 10. We see that the combination is vastly superior to each of the two solvers involved, and approaches in performance the best solver, the `tawSolver-2.6` (but still slower by a factor of two).

$t =$	13	14	15	16	17
D	20	30	35	40	50
nds	3,197	27,053	64,663	209,593	1,399,505
t	10	146	821	3,248	23,546
N	1599	13,527	32,331	104,797	699,751
t : med, max	0.06, 0.49	0.06, 0.68	0.16, 3.9	0.46, 29.6	0.8, 199
Σ cfs	8,479,987	59,402,586	361,511,501	3,723,995,162	35,931,491,146
Σ t	120	961	6,888	80,056	1,006,718
total t	130	1,107	7,709	83,304	1,030,264
factor	13.2	15.2	24.7	NA	NA

Table 10: Cube & Conquer, via the OKsolver as the cube-solver, and MiniSat-2.2 as the conquer-solver. Times are in seconds. “factor” is run-time of MiniSat-2.2, divided by total time of Cube & Conquer. The run-times of the OKsolver includes writing all data-files (the partial assignments), the run-times of MiniSat include reading the files. 10^6 seconds are roughly 11.6 days.

6.1.3. Palindromic vdW-problems

The data for the palindromic problems we considered is shown in Table 11. Recall that for palindromic problems, that is, the determination of $\text{pdw}(2; 3, t) = (n_1, n_2)$, we have to determine two numbers: the n_1 such that all $F^{\text{pd}}(3, t, n)$ with $n \leq n_1$ are satisfiable, while $F^{\text{pd}}(3, t, n_1 + 1)$ is unsatisfiable, and n_2 for which $F^{\text{pd}}(3, t, n)$ is unsatisfiable for all $n \geq n_2$, while $F^{\text{pd}}(3, t, n_2 - 1)$ is satisfiable. In order to do so, as shown in Theorem 5.1, the main unsatisfiable instances are for n_1 and $n_2 + 1$. To reduce the amount of data, we don’t show the data for these two critical points, but for n_2 , which is easier than $n_2 + 1$ (in our range by a factor of around five; possible due to the fact that except of one case n_2 happens to be odd here, as discussed in the next paragraph), and harder than n_1 .

For $F^{\text{pd}}(3, t, n)$ with odd n we can determine that the middle vertex $\frac{n+1}{2}$ can not be element of the first block of the partition (belonging to progression-size 3), since then no other vertex could be in the first block (due to the palindromic property and the symmetric position of the middle vertex), and then we would have an arithmetic progression of size t in the second block. Due to this (and there might be other reasons), palindromic problems for odd n are easier (running times can go up by a factor of 10 for even n).²¹

t	n	v	c	ℓ	c_2	c_3	$c_{\lceil t/2 \rceil}$	$c_{\lceil t/2 \rceil + 1}$	c_t
17	279	140	10,536	45,139	185	9,357	25	0	969
18	312	156	13,277	58,763	52	11,954	9	0	1,262
19	347	174	16,208	70,414	230	14,586	28	0	1,364
20	389	195	20,327	88,944	258	18,393	10	19	1,647
21	405	203	21,950	96,305	269	19,958	29	0	1,694
22	463	232	28,650	126,560	308	26,171	11	21	2,139
23	507	254	34,289	152,236	337	31,448	34	0	2,470
24	593	297	46,881	209,792	394	43,156	12	24	3,295
25	607	304	48,979	219,525	404	45,237	37	0	3,301
26	643	322	54,843	246,503	428	50,813	12	24	3,566
27	699	350	64,719	292,102	465	60,133	38	0	4,083

Table 11: Instance data for $F^{\text{pd}}(3, t, n)$, where v is the number of variables, $c = c_2 + c_3 + c_{\lceil t/2 \rceil} + c_{\lceil t/2 \rceil + 1} + c_t$ is the number of clauses, c_i the number of clauses of length i , and ℓ is the number of literal occurrences.

First we consider the look-ahead solvers in Table 12. Comparing `tawSolver` with the other solvers, we see a similar behaviour as with (ordinary) vdW-problems, but more extreme so. The node-count of `tawSolver-2.6` and `rawSolver-2.6` is not much worse than the “real” look-ahead solvers, with exception of `march_pl` (where again a large number of inferred clauses is added by the solver). The weak performance of the OKsolver is (likely) explained

²¹We remark that while for example `PrecoSAT` determines this forced variable right at the beginning, this is not the case for the `MiniSat` versions, which infer that fact rather late, and they are helped by adding the corresponding unit-clause to the instance.

by the instances not having many r_2 -reductions (recall that OKsolver is *completely* eliminating failed literals, as the only solver), and so the overhead is prohibitive (the savings in node-count don't pay off). `satz` only investigates %10 of the most promising variables for r_2 -reductions, and additionally looks for some r_3 -reductions. This strategy here works far better than OKsolver's "strategy" (but the OKsolver deliberately doesn't employ a "strategy" here, since the aim is to have a stable and "mathematical meaningful" solver); nevertheless still the overhead is too large.

An interesting aspect is that for larger t the more complex heuristic (i.e., projection) of `rawSolver-2.6` compared to `tawSolver-2.6` pays off. This is different from ordinary vdW-problems. And as the comparison with `tawSolver-1.0` shows, the heuristic (mostly the projection) is of great importance here (more pronounced than for ordinary vdW-problems).

t	<code>rawSolver-2.6</code>	<code>tawSolver-2.6</code>	<code>satz</code>	<code>tawSolver-1.0</code>	<code>march_pl</code>	OKsolver
17	1 32,855	0.8 32,697	12 16,466	7 143,319	35 1,448	18 5,023
18	11 276,249	8 279,309	182 208,873	60 1,063,979	269 12,289	335 100,803
19	13 283,229	10 285,037	143 123,199	134 2,009,635	500 12,423	322 62,009
20	48 894,777	39 897,529	701 459,899	738 9,076,261	1,980 39,681	1,419 206,617
21	115 2,144,743	101 2,239,371	2,592 1,567,736	2,541 30,470,349	5,053 99,493	3,536 490,841
22	564 8,427,503	525 8,683,035	9,418 4,393,139	18,306 170,414,771	25,841 376,285	47,593 3,197,173
23	1,547 19,858,971	1,695 21,565,129	35,633 12,587,868	86,869 573,190,251	77,763 876,315	132,150 7,461,907
24	8,558 79,790,419	26,724 198,685,857				
25	22,841 219,575,127					

Table 12: Look-ahead solvers on unsatisfiable instances $F^{Pd}(3, t; n)$ for computing $w(2; 3, t)$ (with $t = 17, \dots, 25$ and $n = 279, \dots, 607$). The first line is run-time in seconds, the second line is the number of nodes.

The conflict-driven solvers are shown in Table 13. We see that they are not competitive with `tawSolver-2.6` or `rawSolver-2.6`, however now most of them are better than the "real" look-ahead solvers. Here MiniSat-2.2 is better than MiniSat-2.0, and Glucose-2.2 is better than Glucose-2.0, so we show only data for the newest versions. With Glucose we see a pattern which we observed also at other (hard) instance classes: for smaller instances Glucose is better than MiniSat, but from a certain point on the performance of Glucose becomes very bad. This is likely due to the more aggressive restart strategy, which pays off for smaller instances, but from a certain point on the solver becomes essentially incomplete.

t	MiniSat	Glucose	PrecoSAT	Lingeling	CryptoMiniSat
17	0.8	0.8	1.2	3.7	3.6
	34,426	34,826	41,961	57,306	59,443
18	19	14	25	59	78
	607,908	340,568	506,793	919,123	871,916
19	19	15	24	61	72
	568,924	336,861	485,357	915,107	765,301
20	118	66	131	355	384
	2,852,150	1,132,012	1,799,145	3,633,502	3,071,462
21	423	228	445	1,060	1,418
	9,179,642	2,903,573	4,687,589	8,672,073	8,458,496
22	3,151	1,631	2,825	8,428	14,321
	51,582,064	13,397,451	22,283,651	41,696,062	49,716,762
23	8,191	6,817	9,280	28,543	55,544
	108,028,217	36,314,064	54,951,563	104,007,799	141,249,316
24	54,678	> 992,540	82,750	152,076	
	476,716,936	> 1,100,664,795 aborted	261,084,988	285,546,948	

Table 13: Conflict-driven solvers on unsatisfiable instances $F^{\text{pd}}(3, t; n)$ for computing $w(2; 3, t)$ (with $t = 17, \dots, 24$ and $n = 279, \dots, 593$). The first line is run-time in seconds, the second line is the number of conflicts.

Finally we consider Cube & Conquer in Table 14. We see that this is now the fastest solver overall. Glucose-2.2 is %10 faster, but since this is only a small amount, for consistency we stick with MiniSat-2.2.

$t =$	23	24	25	26	27
D	25	35	45	55	65
nds	1,717	5,559	17,633	77,161	220,069
t	106	500	1,752	7,889	25,478
N	859	2,780	8,817	38,581	110,032
t : med, max	0.95, 17.6	1.2, 27	0.81, 47	0.95, 58	0.82, 125
Σ cfs	27,308,572	93,831,664	258,829,555	1,231,383,588	3,423,841,749
Σ t	1,095	4,466	11,822	55,306	172,033
total t	1,201	4,966	13,574	63,195	197,511
factor	1.3	1.7	1.7	NA	NA

Table 14: Cube & Conquer, via the OKsolver as the cube-solver, and MiniSat-2.2 as the conquer-solver. Times are in seconds. “factor” is run-time of best solver, i.e., $\tau_{\text{awSolver-2.6}}$, divided by total time of Cube & Conquer. 10^5 seconds are roughly 1.2 days.

6.2. Incomplete solvers (stochastic local search)

In the OKlibrary we use the UbcSAT suite (see [67]) of local-search algorithms in version 1-2-0. The considered algorithms are GSAT, GWSAT, GSAT-TABU, HSAT, HWSAT, WALKSAT, WALKSAT-TABU, WALKSAT-TABU-NoNull, Novelty, Novelty+, Novelty++, Novelty+p, Adaptive Novelty+, RNovelty, RNovelty+, SAPS, RSAPS, SAPS/NR, PAWS, DDFW, G2WSAT, Adaptive G2WSat, VW1, VW2, RoTS, IRoTS, SAMD. The performance of local-search algorithms is very much instance-dependent, and so a good choice of algorithms is essential. Our experiments yield the following selection criteria:

- For standard problems (Section 3) the best advice seems to use GSAT-TABU for $t \leq 23$, to use RoTS for $t > 23$, and to use Adaptive G2WSat for $t > 33$ (also trying DDFW then).
- For the palindromic problems (Section 5) GSAT-TABU is the best algorithm.

For a given t in principle we let these algorithm run for $n = t + 1, t + 2, \dots$, until the search seems unable to find a solution. But running these algorithms from scratch on these vdW-problems is much less effective than using an incremental approach, based on a solution found for $n - 1$, respectively for palindromic vdW-problems on a solution found for $n - 2$ (according to Lemma 5.1), as initial guess, and repeating this process for the next n : this helps to go much quicker through the easier part of the search space (of possible n), and also seems to help for the harder problems. Finally, we recall that in Subsection 3.2 we explained how we made the distinction between lower bounds we conjecture to be exact and sheer lower bounds.

7. Conclusion

This article presented the following contributions to the fields of Ramsey theory and SAT solving:

- Study of $w(2; 3, t)$:
 1. determination of $w(2; 3, 19) = 349$;
 2. lower bounds for $w(2; 3, t)$ with $20 \leq t \leq 30$, conjectured to be exact;
 3. further lower bounds for $31 \leq t \leq 39$;
 4. improved conjecture on the growth rate of $w(2; 3, t)$;
 5. various observations on structural properties of good partitions.
- Introduction and study of $\text{pdw}(2; 3, t)$:
 1. basic definitions and properties;
 2. determination of $\text{pdw}(2; 3, t)$ for $t \leq 27$;
 3. lower bounds for $\text{pdw}(2; 3, t)$ with $28 \leq t \leq 35$, conjectured to be exact;
 4. further lower bounds for $36 \leq t \leq 39$.
- SAT solving:
 1. introduced the new SAT-solver `tawSolver`, with the basic implementation given by `tawSolver-1.0`, and the versions with improved heuristic by `tawSolver-2.6` and `tauSolver-2.6`;
 2. experimental comparison with current look-ahead and conflict-driven solvers;
 3. comparison and data for the new `Cube & Conquer` method;
 4. experimental determination of good local-search algorithms for lower bounds.

We hope that these investigations contribute to a better understanding of the connections between Ramsey theory and SAT solving. The following seem relevant research directions for future investigations:

- Showing $w(2; 3, 20) = 389$ (recall Subsection 3.2) should be in reach with `tawSolver-2.6`, while showing $w(2; 3, 21) = 416$ seems to require new (algorithmic) insight (when using similar computational resources).
- Conjecture 3.1 states that the lower bound from [14] for $w(2; 3, t)$ is tight up to a small factor.
- In Section 4 four conjectures on patterns in good partitions are presented (one implying Conjecture 3.1).
- In Subsection 5.4 various open problems on palindromic van der Waerden numbers are stated.
- Considering SAT solving:
 1. Understand the differences between ordinary and palindromic problems:
 - Why is the projection relatively more important for the palindromic problems? (So that the difference between `tawSolver-2.6` and `tawSolver-1.0` is more pronounced, and `tauSolver-2.6` becomes faster than `tawSolver-2.6` on bigger instances.)
 - Why do we have different behaviour of look-ahead versus conflict-driven solvers?
 2. Can the branching heuristic of `tawSolver` for the instances of this paper be much further improved? Especially can we gain some understanding of the weights?
 3. How to understand the success of `Cube & Conquer`? Does its success indicate that there are important dag-like structures in good resolution refutations of the instances of these classes, which are dispersed locally, so that ordinary conflict-driven solvers have problems exploiting them?

Acknowledgements

The authors would like to thank Donald Knuth, the Editor and the anonymous referees for their valuable suggestions and helpful comments.

References

- [1] Dimitris Achlioptas. Random satisfiability. In Biere et al. [12], chapter 8, pages 245–270.
- [2] Tanbir Ahmed. Some new van der Waerden numbers and some van der Waerden-type numbers. *Integers*, 9:65–76, 2009. #A6.
- [3] Tanbir Ahmed. Two new van der Waerden numbers: $w(2;3,17)$ and $w(2;3,18)$. *Integers*, 10:369–377, 2010. #A32.
- [4] Tanbir Ahmed. On computation of exact van der Waerden numbers. *Integers*, 12(3):417–425, 2012. #A71.
- [5] Tanbir Ahmed. Some more van der Waerden numbers. *Journal of Integer Sequences*, 16(4), 2013. #13.4.4.
- [6] Gilles Audemard and Laurent Simon. Predicting learnt clauses quality in modern SAT solvers. In *IJCAI'09 Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 399–404. AAAI, 2009.
- [7] Michael D. Beeler and Patrick E. O’Neil. Some new van der Waerden numbers. *Discrete Mathematics*, 28(2):135–146, 1979.
- [8] Armin Biere. Picosat essentials. *Journal on Satisfiability, Boolean Modeling and Computation*, 4:75–97, 2008.
- [9] Armin Biere. Bounded model checking. In Biere et al. [12], chapter 14, pages 455–481.
- [10] Armin Biere. P{re,i}coSAT@SC’09. <http://fmv.jku.at/precosat/preicosat-sc09.pdf>, 2009.
- [11] Armin Biere. Lingeling and friends entering the SAT Challenge 2012. In Adrian Balint, Anton Belov, Daniel Diepold, Simon Gerber, Matti Järvisalo, and Carsten Sinz, editors, *Proceedings of SAT Challenge 2012: Solver and Benchmark Descriptions*, volume B-2012-2 of *Department of Computer Science Series of Publications B*, pages 33–34. University of Helsinki, 2012. https://helda.helsinki.fi/bitstream/handle/10138/34218/sc2012_proceedings.pdf.
- [12] Armin Biere, Marijn J.H. Heule, Hans van Maaren, and Toby Walsh, editors. *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, February 2009.
- [13] T. C. Brown. Some new van der Waerden numbers (preliminary report). *Notices of the American Mathematical Society*, 21:432, 1974.
- [14] Tom Brown, Bruce M. Landman, and Aaron Robertson. Bounds on some van der Waerden numbers. *Journal of Combinatorial Theory, Series A*, 115:1304–1309, 2008.
- [15] V. Chvátal. Some unknown van der Waerden numbers. In R.K. Guy, editor, *Combinatorial Structures and their Applications*, pages 31–33. Gordon and Breach, New York, 1970.
- [16] Scott Cotton. Two techniques for minimizing resolution proofs. In Strichman and Szeider [65], pages 306–312. ISBN-13 978-3-642-14185-0.
- [17] Adnan Darwiche and Knot Pipatsrisawat. Complete algorithms. In Biere et al. [12], chapter 3, pages 99–130.
- [18] Martin Davis, George Logemann, and Donald Loveland. A machine program for theorem-proving. *Communication of the ACM*, 5:394–397, 1962.
- [19] Michael R. Dransfield, Lengning Liu, Victor W. Marek, and Mirosław Truszczyński. Satisfiability and computing van der Waerden numbers. *The Electronic Journal of Combinatorics*, 11(#R41), 2004.
- [20] Niklas Eén and Niklas Sörensson. An extensible SAT-solver. In Enrico Giunchiglia and Armando Tacchella, editors, *Theory and Applications of Satisfiability Testing 2003*, volume 2919 of *Lecture Notes in Computer Science*, pages 502–518, Berlin, 2004. Springer. ISBN 3-540-20851-8.
- [21] Luís Gil, Paulo Flores, and Luís Miguel Silveira. PMSat: a parallel version of MiniSAT. *Journal on Satisfiability, Boolean Modeling and Computation*, 6:71–98, 2009.
- [22] Carla P. Gomes, Ashish Sabharwal, and Bart Selman. Model counting. In Biere et al. [12], chapter 20, pages 633–654.
- [23] Jun Gu. The multi-SAT algorithm. *Discrete Applied Mathematics*, 96-97:111–126, 1999.
- [24] Long Guo, Youssef Hamadi, Said Jabbour, and Lakhdar Sais. Diversification and intensification in parallel SAT solving. In *CP’10 Proceedings of the 16th international conference on Principles and practice of constraint programming*, volume 6308 of *Lecture Notes in Computer Science*, pages 252–265. Springer-Verlag, 2010.
- [25] Matthew Gwynne and Oliver Kullmann. Generalising and unifying SLUR and unit-refutation completeness. In Peter van Emde Boas, Frans C. A. Groen, Giuseppe F. Italiano, Jerzy Nawrocki, and Harald Sack, editors, *SOFSEM 2013: Theory and Practice of Computer Science*, volume 7741 of *Lecture Notes in Computer Science (LNCS)*, pages 220–232. Springer, 2013.
- [26] Matthew Gwynne and Oliver Kullmann. Generalising unit-refutation completeness and SLUR via nested input resolution. *Journal of Automated Reasoning*, 52(1):31–65, January 2014.
- [27] P.R. Herwig, M.J.H. Heule, P.M. van Lambalgen, and H. van Maaren. A new method to construct lower bounds for van der Waerden numbers. *The Electronic Journal of Combinatorics*, 14(#R6), 2007.
- [28] Marijn Heule, Mark Dufour, Joris van Zwieten, and Hans van Maaren. March_eq: Implementing additional reasoning into an efficient look-ahead SAT solver. In Hoos and Mitchell [35], pages 345–359. ISBN 3-540-27829-X.
- [29] Marijn Heule and Toby Walsh. Internal symmetry. In Pierre Flener and Justin Pearson, editors, *The 10th International Workshop on Symmetry in Constraint Satisfaction Problems (SymCon’10)*, pages 19–33, 2010.
- [30] Marijn Heule and Toby Walsh. Symmetry within solutions. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)*, pages 77–82, 2010.
- [31] Marijn J. H. Heule and Hans van Maaren. Look-ahead based SAT solvers. In Biere et al. [12], chapter 5, pages 155–184.
- [32] Marijn J.H. Heule, Oliver Kullmann, Siert Wieringa, and Armin Biere. Cube and conquer: Guiding CDCL SAT solvers by lookaheads. In Kerstin Eder, João Lourenço, and Onn Shehory, editors, *Hardware and Software: Verification and Testing (HVC 2011)*, volume 7261 of *Lecture Notes in Computer Science (LNCS)*, pages 50–65. Springer, 2012. <http://cs.swan.ac.uk/~csoliver/papers.html#CuCo2011>.
- [33] Marijn J.H. Heule and Hans van Maaren. Parallel SAT solving using bit-level operations. *Journal on Satisfiability, Boolean Modeling and Computation*, 4:99–116, 2008.

- [34] John N. Hooker and V. Vinay. Branching rules for satisfiability. *Journal of Automated Reasoning*, 15:359–383, 1995.
- [35] Holger H. Hoos and David G. Mitchell, editors. *Theory and Applications of Satisfiability Testing 2004*, volume 3542 of *Lecture Notes in Computer Science*, Berlin, 2005. Springer. ISBN 3-540-27829-X.
- [36] Antti E. Hyvärinen, Tommi Junttila, and Ilkka Niemelä. Incorporating clause learning in grid-based randomized SAT solving. *Journal on Satisfiability, Boolean Modeling and Computation*, 6:223–244, 2009.
- [37] Antti E. Hyvärinen, Tommi Junttila, and Ilkka Niemelä. Partitioning search spaces of a randomized search. In *AI*IA 2009: Proceedings of the XIth International Conference of the Italian Association for Artificial Intelligence Reggio Emilia on Emergent Perspectives in Artificial Intelligence*, volume 5883 of *Lecture Notes in Computer Science*, pages 243–252. Springer-Verlag, 2009.
- [38] Antti E. Hyvärinen, Tommi Junttila, and Ilkka Niemelä. Partitioning SAT instances for distributed solving. In *LPAR’10 Proceedings of the 17th international conference on Logic for programming, artificial intelligence, and reasoning*, volume 6397 of *Lecture Notes in Computer Science*, pages 372–386. Springer-Verlag, 2010.
- [39] Bernard Jurkowiak, Chu Min Li, and Gil Utard. A parallelization scheme based on work stealing for a class of SAT solvers. *Journal of Automated Reasoning*, 34(1):73–101, January 2005.
- [40] Henry A. Kautz, Ashish Sabharwal, and Bart Selman. Incomplete algorithms. In Biere et al. [12], chapter 6, pages 185–203.
- [41] Hans Kleine Büning and Oliver Kullmann. Minimal unsatisfiability and autarkies. In Biere et al. [12], chapter 11, pages 339–401.
- [42] Michal Kouril. Computing the van der Waerden number $w(3, 4) = 293$. *INTEGERS: Electronic Journal of Combinatorial Number Theory*, 12(A46):1–13, 2012.
- [43] Michal Kouril and Jerome L. Paul. The van der Waerden number $W(2, 6)$ is 1132. *Experimental Mathematics*, 17(1):53–61, 2008.
- [44] Daniel Kroening. Software verification. In Biere et al. [12], chapter 16, pages 505–532.
- [45] Oliver Kullmann. Investigating a general hierarchy of polynomially decidable classes of CNF’s based on short tree-like resolution proofs. Technical Report TR99-041, Electronic Colloquium on Computational Complexity (ECCC), October 1999.
- [46] Oliver Kullmann. Investigating the behaviour of a SAT solver on random formulas. Technical Report CSR 23-2002, Swansea University, Computer Science Report Series (available from <http://www-compsci.swan.ac.uk/reports/2002.html>), October 2002. 119 pages.
- [47] Oliver Kullmann. The OKLibrary: Introducing a “holistic” research platform for (generalised) SAT solving. *Studies in Logic*, 2(1):20–53, 2009.
- [48] Oliver Kullmann. Fundamentals of branching heuristics. In Biere et al. [12], chapter 7, pages 205–244.
- [49] Oliver Kullmann. Exact Ramsey theory: Green-Tao numbers and SAT. Technical Report arXiv:1004.0653v2 [cs.DM], arXiv, April 2010.
- [50] Oliver Kullmann. Green-Tao numbers and SAT. In Strichman and Szeider [65], pages 352–362. ISBN-13 978-3-642-14185-0.
- [51] Oliver Kullmann. Constraint satisfaction problems in clausal form I: Autarkies and deficiency. *Fundamenta Informaticae*, 109(1):27–81, 2011.
- [52] Oliver Kullmann. Constraint satisfaction problems in clausal form II: Minimal unsatisfiability and conflict structure. *Fundamenta Informaticae*, 109(1):83–119, 2011.
- [53] Bruce Landman, Aaron Robertson, and Clay Culver. Some new exact van der Waerden numbers. *INTEGERS: Electronic Journal of Combinatorial Number Theory*, 5(2):1–11, 2005. #A10.
- [54] Chu Min Li and Anbulagan. Heuristics based on unit propagation for satisfiability problems. In *Proceedings of 15th International Joint Conference on Artificial Intelligence (IJCAI’97)*, pages 366–371. Morgan Kaufmann Publishers, 1997.
- [55] Joao P. Marques-Silva, Ines Lynce, and Sharad Malik. Conflict-driven clause learning SAT solvers. In Biere et al. [12], chapter 4, pages 131–153.
- [56] John R. Rabung. Some progression-free partitions constructed using Folkman’s method. *Canadian Mathematical Bulletin*, 22(1):87–91, 1979.
- [57] Stanisław P. Radziszowski. Small Ramsey numbers. *The Electronic Journal of Combinatorics*, August 2009. Dynamic Surveys DS1, Revision 12; see <http://www.combinatorics.org/Surveys>.
- [58] Vera Rosta. Ramsey theory applications. *The Electronic Journal of Combinatorics*, December 2004. Dynamic Surveys DS13, Revision 1; see <http://www.combinatorics.org/Surveys>.
- [59] K.F. Roth. On certain sets of integers. *Journal of the London Mathematical Society*, 28:245–252, 1953.
- [60] Kareem A. Sakallah. Symmetry and satisfiability. In Biere et al. [12], chapter 10, pages 289–338.
- [61] Marko Samer and Stefan Szeider. Fixed-parameter tractability. In Biere et al. [12], chapter 13, pages 425–454.
- [62] Tobias Schubert, Matthew Lewis, and Bernd Becker. PaMiraXT: Parallel SAT solving with threads and message passing. *Journal on Satisfiability, Boolean Modeling and Computation*, 6:203–222, 2009.
- [63] Pascal Schweitzer. *Problems of Unknown Complexity: Graph isomorphism and Ramsey theoretic numbers*. PhD thesis, Universität des Saarlandes, Saarbrücken, 2009. Revised version, April 2012; http://www.mpi-inf.mpg.de/~pascal/docs/thesis_pascal_schweitzer.pdf.
- [64] Mate Soos, Karsten Nohl, and Claude Castelluccia. Extending SAT solvers to cryptographic problems. In Oliver Kullmann, editor, *Theory and Applications of Satisfiability Testing - SAT 2009*, volume 5584 of *Lecture Notes in Computer Science*, pages 244–257. Springer, 2009. http://planete.inrialpes.fr/~soos/publications/Extending_SAT_2009.pdf.
- [65] Ofer Strichman and Stefan Szeider, editors. *Theory and Applications of Satisfiability Testing - SAT 2010*, volume LNCS 6175 of *Lecture Notes in Computer Science*. Springer, 2010. ISBN-13 978-3-642-14185-0.
- [66] E. Szemerédi. On sets of integers containing no k elements in arithmetic progression. *Acta Arithmetica*, 27:299–345, 1975.
- [67] Dave A.D. Tompkins and Holger H. Hoos. UBCSAT: An implementation and experimentation environment for SLS algorithms for SAT and MAX-SAT. In Hoos and Mitchell [35], pages 306–320. ISBN 3-540-27829-X.
- [68] Peter van der Tak, Marijn J. H. Heule, and Armin Biere. Concurrent Cube-and-Conquer. In Alessandro Cimatti and Roberto Sebastiani, editors, *Theory and Applications of Satisfiability Testing - SAT 2012*, volume 7317 of *Lecture Notes in Computer Science (LNCS)*, pages 475–476. Springer, 2012.
- [69] B.L. van der Waerden. Beweis einer Baudetschen Vermutung. *Nieuw Archief voor Wiskunde*, 15:212–216, 1927.
- [70] Hantao Zhang. Combinatorial designs by SAT solvers. In Biere et al. [12], chapter 17, pages 533–568.

[71] Hantao Zhang, Maria Paola Bonacina, and Jie Hsiang. PSATO: a distributed propositional prover and its application to quasigroup problems. *Journal of Symbolic Computation*, 11:1–18, 1996.

Appendix A. Certificates

Appendix A.1. Conjectured precise lower bounds for $w(2; 3, t)$

The certificate for $t = 20$ is also palindromic (while for $t = 24$ a palindromic certificate is given in Subsection Appendix A.3).

$w(2; 3, 20) \geq 389$.

$$1^{19}01^{11}01^401^70101^401^{13}01^901^40101^{14}0^{21}30^2101^901^{18}01^901^40101^501^301^{10}01^{16}01^801^{16}01^{10}01^301^50101^401^9$$

$$01^{18}01^9010^21^30^21^{14}0101^401^901^{13}01^40101^701^401^{11}01^{19}$$

$w(2; 3, 21) \geq 416$.

$$1^801^{17}0101^60101^901^{12}0^21^{19}01^{18}01^30101^201^{12}0^21^50^2101^{10}01^{18}01^3010^21^801^{12}01^{20}01^{18}01^30101^701^801^{12}0^2$$

$$101^{17}01^{18}01^3010^21^801^501^60^21^{12}01^601^{15}0101^40101^{11}01^{20}$$

$w(2; 3, 22) \geq 464$.

$$1^20^21^{17}0^21^901^{12}0101^{12}01^{15}01^201^{10}01^401^701^501^{12}01^90101^301^801^30101^901^{12}01^501^701^401^{10}01^201^{15}01^{12}$$

$$0101^{12}01^90^21^{17}0^21^901^{12}0101^{12}01^{15}01^201^{10}01^401^701^501^{12}01^90101^301^801^501^901^{12}01^501^701^{11}$$

$w(2; 3, 23) \geq 516$.

$$1^40^21^201^{17}01^40101^{15}01^{16}01^401^501^{20}0101^201^80^2101^401^{15}01^201^401^{16}01^901^{10}0101^901^{10}01^701^{17}01^60101^{19}01^{16}0$$

$$1^{21}0^21^{19}01^601^201^{12}010^21^40101^{20}01^{13}0^21^{11}01^90^21^601^401^{13}0101^301^801^901^{20}01^501^{18}01^301$$

$w(2; 3, 24) \geq 593$.

$$1^{21}01^{18}01^{16}01^401^701^60101^{14}01^30^21^801^701^301^201^{20}0^21^301^701^{15}01^701^30^21^{20}01^201^301^701^901^{18}0101^601^{21}01^70$$

$$1^{10}01^701^{21}01^60101^{18}01^901^701^301^201^{20}0^21^301^701^{15}01^701^30^21^{20}01^201^301^701^80^21^301^{14}0101^601^701^401^{16}01^{18}01^{21}$$

$w(2; 3, 25) \geq 656$.

$$1^{16}01^201^{19}01^801^701^{19}01^801^40101^501^{17}01^{10}01^{21}0^21^20101^{10}01^701^{12}01^401^301^601^701^{11}01^301^{13}01^401^{17}0101^3$$

$$01^60^21^601^{17}01^801^701^{13}01^{14}01^201^401^{13}0^21^801^701^{19}01^801^40101^701^{15}01^{10}01^901^{11}0^21^20101^{10}01^501^{14}0$$

$$1^401^301^601^701^{11}01^301^{13}01^401^{17}0101^301^60^21^{24}01^801^9$$

$w(2; 3, 26) \geq 727$.

$$1^{10}01^{23}01^{10}0101^{20}01^401^{11}01^601^{11}0^21^201^501^601^501^{23}01^40101^701^{17}01^{16}0101^{11}0^21^20101^301^401^201^{18}$$

$$01^301^501^{14}01^{12}01^{16}01^401^{19}01^8010^21^401^{13}01^{14}0101^{20}01^401^{18}01^{11}0^21^201^501^601^501^{23}01^40101^701^{15}0101^{16}$$

$$0101^{11}0^21^20101^301^401^201^{18}01^901^{14}01^{12}01^{16}01^401^{19}01^801^201^{18}01^301^{25}$$

$w(2; 3, 27) \geq 770$.

$$1^{24}01^301^501^{18}01^{17}0101^201^{21}01^301^701^201^{20}0^21^{12}01^{15}01^{10}01^{11}01^301^901^601^{13}01^{22}01^30^2$$

$$1^801^501^{20}01^60101^{16}01^701^301^5010^2101^{21}01^20^21^{14}01^901^{17}0101^201^301^{17}01^301^50101^201^{20}0^21^{12}01^{15}01^{22}$$

$$01^301^901^601^{13}01^{22}01^30^21^801^501^{20}01^801^{16}01^701^301^5010^2101^{24}01^601^801^{20}01^{15}01^{16}01^5$$

$w(2; 3, 28) \geq 827$.

$$1^{27}01^{10}01^{22}0101^{16}01^{13}01^{16}01^{20}01^401^{16}0101^{11}0^21^20101^301^60^21^{18}01^301^501^{14}01^{12}01^{21}01^{14}0$$

$$1^{13}010^21^401^{23}01^40101^{25}01^{18}01^{11}01^30101^301^401^701^{17}01^{10}0101^701^{17}01^{11}01^401^{13}0^21^20101^301^7$$

$$01^{18}01^301^501^{14}01^{12}01^{16}01^401^{14}01^{13}010^21^{18}01^901^40101^{25}01^{18}01^{11}01^901^40101^{23}01^{10}01^901^{12}01^{16}01^{10}$$

$w(2; 3, 29) \geq 868$.

$$1^{15}01^{21}01^{18}01^{17}01^{18}01^{21}01^201^701^{25}0101^{12}01^{17}01^601^70^21^201^{17}01^50^2101^{10}01^601^{13}0101^201^{17}01^{16}01^701^4$$

$$01^{20}0101^201^501^{11}01^{25}01^{10}01^{25}01^{20}2^1^{16}01^301^{10}01^401^{20}0101^201^601^{20}01^{14}0^21^201^{17}01^50^2101^{10}01^601^{13}0101^2$$

$$01^901^701^901^601^701^401^{20}0101^201^{17}01^{25}01^201^701^{25}01^301^{16}01^601^{12}01^{25}01^601^{27}01^8$$

$w(2; 3, 30) \geq 903$.

$$1^{22}01^{16}01^{22}01^{26}01^70101^80101^{22}01^{12}01^{16}01^701^601^901^{11}01^60^21^{15}0^21^{13}01^701^80101^{23}01^601^{28}01^701^301^401^{22}$$

$$01^80101^20^21^{11}0101^{22}01^{11}01^{14}01^301^{19}01^401^{16}0^21^{11}0101^501^{28}01^60^2101^{10}01^201^{14}01^701^{10}01^{23}01^601^{28}$$

$$01^701^501^20^21^{18}01^201^80101^301^{11}0101^{16}01^{25}0101^401^{23}01^{16}01^40^21^{13}01^{12}01^301^{27}01^{10}01^{14}$$

Appendix A.2. Further lower bounds for $w(2; 3, t)$

The certificate for $t = 31$ is also palindromic.

$w(2; 3, 31) > 930$.

$$1^{12}01^{19}01^{16}01^{10}01^901^301^401^{14}01^{12}01^{14}01^{18}0^21^{10}01^401^{13}01^801^50101^{10}01^{17}01^{15}01^{16}01^801^501^80101^401^801^{19}$$

$$01^{16}01^501^{18}01^401^{25}01^{14}01^{16}01^301^901^601^{13}01^{11}01^201^{11}01^{13}01^601^901^301^{16}01^{14}01^{25}01^401^{18}01^501^{16}01^{19}01^8$$

$$01^40101^801^501^801^{16}01^{15}01^{17}01^{10}0101^501^801^{13}01^401^{10}0^21^{18}01^{14}01^{12}01^{14}01^401^301^901^{10}01^{16}01^{19}01^{12}$$

$w(2; 3, 32) > 1006$.

$$1^{15}01^{26}01^501^{11}01^401^{14}01^401^{16}01^{28}01^{22}01^201^501^701^{19}01^20101^{16}01^{13}01^501^301^701^401^{13}01^901^{14}01^{28}01^701^{15}01^{10}$$

$$010^21^{18}0^21^201^{13}01^601^{21}01^{27}01^20101^801^701^{14}01^{13}01^201^{24}01^{10}01^50101^{17}01^{10}0^21^401^{28}01^901^{11}0^21^801^601^70$$

$$1^{11}01^301^{10}010^21^{28}01^{14}01^{17}01^301^{12}01^{13}01^501^{16}01^{13}01^301^{10}01^{22}01^201^{12}01^701^{28}01^{20}01^201^501^401^{22}01^801^{28}$$

$w(2; 3, 33) > 1063$.

$$1^{29}01^601^{14}01^{11}01^{21}01^{14}01^201^{18}0101^{15}01^{12}01^{25}01^{16}0101^{11}0^21^401^{10}01^501^{13}01^{16}01^{20}01^201^{13}01^{22}01^501^40101^{19}01^90$$

$$1^401^901^{15}0101^{18}01^{11}0^21^401^301^401^201^{28}01^601^701^{29}01^401^{19}01^{10}0^21^{18}01^{14}0101^501^{14}01^{16}01^{28}01^601^{20}01^801^401^7$$

$$01^{14}01^401^{18}01^{11}01^901^401^201^401^{17}0101^301^{14}01^{29}01^{18}0^21^201^{10}01^201^{19}01^201^{10}01^{27}01^{18}01^{12}01^401^{20}01^901^{24}01^{13}$$

$w(2; 3, 34) > 1143.$

$1^{32}01^70101^801^{29}01^{23}01^30^21^{10}01^301^{17}01^201^901^501^{16}01^{15}01^{20}0^21^201^80^2101^{31}01^{18}01^601^{12}01^601^{17}01^301^70^21^{26}01^601^20101^{24}0^21^701^201^301^{21}01^{24}01^{10}01^{11}01^{20}01^{10}01^{18}01^60^21^{28}01^701^301^{14}01^{27}010^21^{20}0101^401^701^201^{20}0^21^{28}01^701^{18}01^901^{10}01^801^{28}01^5010^2101^301^{27}01^201^501^{32}01^{24}01^601^{25}01^{10}01^201^{23}01^20101^901^{13}01^{10}01^{11}01^{20}01^{12}01^{16}01^701^301^{33}$

$w(2; 3, 35) > 1204.$

$1^{34}01^{24}01^801^{22}01^30^21^{10}01^{29}01^701^401^{14}01^{15}01^701^{21}01^601^501^{13}01^{22}0101^{12}01^201^{12}01^701^{11}0^21^401^301^701^{18}01^501^301^401^201^601^{34}01^{19}01^80101^501^{17}01^{10}01^201^{18}01^301^{18}0101^{29}01^{18}01^301^{32}0^21^201^501^{27}01^{16}01^{24}01^301^{12}0101^{17}01^80101^{20}01^{13}01^{22}01^801^601^701^{21}01^{12}01^{19}01^{11}01^{22}0101^80101^501^{13}01^{14}01^901^{15}0101^{10}01^701^{15}01^{13}01^401^{13}01^{24}01^{12}01^{21}01^{23}01^60^21^{24}01^{34}$

$w(2; 3, 36) > 1257.$

$1^{10}01^{33}01^{16}01^{12}01^601^{11}01^{25}0101^401^501^201^{12}01^601^{29}0101^401^{17}01^{26}01^{21}0101^{12}01^30101^20^21^{13}01^501^{16}01^{32}01^201^{19}0^2101^401^301^{25}01^{10}01^{32}01^501^201^{14}01^701^{10}01^{17}01^{16}01^401^{20}01^{16}01^{13}01^501^{12}01^301^{11}0101^{34}01^901^{10}01^{30}01^301^501^{26}01^901^{11}01^601^201^{33}01^201^{10}01^201^{20}01^601^301^{24}01^{16}01^{19}0^2101^401^301^{17}01^701^{16}01^{21}01^{26}0101^501^{16}01^{26}01^301^40^21^{13}01^401^201^80101^{22}01^{16}01^{20}01^{28}0101^{33}$

$w(2; 3, 37) > 1338.$

$1^501^{33}01^{22}01^201^{12}0^21^{18}01^{14}01^901^{12}01^{15}01^{24}01^{19}01^{17}01^70^21^70101^{12}01^201^{10}01^{22}01^801^{32}01^301^601^{19}01^{14}01^{21}0^21^201^{12}01^701^{26}01^{22}01^20101^301^{32}01^301^{14}01^{11}01^{17}01^{18}0^21^20101^301^701^{22}01^{32}01^{22}0101^{28}0^21^401^{31}01^601^{17}01^301^{14}01^{21}01^401^701^{14}01^{35}01^501^{16}01^{19}01^301^40101^{23}01^{10}01^{31}01^{12}0101^{11}01^301^{12}01^501^{14}01^201^{10}01^{27}01^{30}0^21^20101^801^201^401^{27}010^21^701^{29}01^401^{14}01^{16}01^{18}01^{14}01^{21}0^21^{11}01^5$

$w(2; 3, 38) > 1378.$

$1^{34}01^{14}01^701^{13}01^{22}0101^{12}01^{23}01^{12}01^{13}010^21^{28}01^401^201^{24}01^{10}01^701^{29}01^401^{19}01^801^201^{18}0^21^{21}01^{14}01^{13}01^701^{17}0^21^401^{26}01^30101^{10}01^{20}01^{16}01^{18}01^901^{12}0101^{21}01^601^{37}01^601^{15}01^{13}01^601^{12}01^801^{32}01^401^501^{19}01^301^{10}01^{21}0^21^401^{18}01^{28}01^{37}01^601^{11}01^301^{10}01^701^{13}0^21^401^{34}01^501^401^{25}01^50^21^301^{27}01^{10}01^501^{23}01^401^{22}01^801^{12}01^{16}01^{13}0^21^401^{20}01^{10}01^{33}01^70101^{14}01^501^{27}01^70101^60^21^{19}01^{16}01^{30}$

$w(2; 3, 39) > 1418.$

$1^201^401^{13}01^{34}0101^{20}01^{21}01^{19}01^701^201^{36}01^{20}01^20^21^801^{24}01^{12}01^501^{27}01^{18}01^{11}01^701^601^{30}01^501^{16}01^401^{26}01^40^21^{28}01^201^401^201^601^{18}01^{26}01^{17}01^301^701^{16}01^{11}01^501^201^{22}01^301^601^401^{36}01^{24}01^801^{16}01^601^{13}0101^{34}01^{26}01^{36}01^601^201^501^{16}01^{15}01^301^901^701^{11}01^201^{24}01^{15}01^{27}0101^201^301^901^{17}010^21^{36}01^{22}0^21^{11}01^{24}01^{28}01^601^{29}0^21^50^21^301^{32}01^20^21^501^201^{15}01^401^701^401^{29}01^801^{22}01^401^{32}01^{12}01^{30}$

Appendix A.3. Good palindromic partitions

Table A.15 gives good palindromic partitions for $n_1 - 1, n_2 - 1$ with $(n_1, n_2) = \text{pdw}(2; 3, t)$, $3 \leq t \leq 39$, according to the values in Tables 6, 7. Due to the palindromic property, for the corresponding n we only show the partition of $\{1, \dots, \lceil \frac{n}{2} \rceil\}$, so that for example the good palindromic partition 01^2001^20 for $t = 3$ and $n = 8$ is compressed to 01^20 . Note that such compressed partitions correspond exactly to the solution of $F^{\text{pd}}(3, t, n)$ as defined in Lemma 5.3.

Table A.15: Good palindromic partitions for $\text{pdw}(2; 3, t) - 1$ according to Theorem 5.1, Part (i)

t	$\text{pdw}(2; 3, t) - 1$	Good palindromic partitions
3	(5, 8)	$1^20, 10^21$
4	(14, 15)	$0101^30, 1^2010^21^2$
5	(15, 20)	$1^40^21^2, 1^20101^40$
6	(29, 30)	$1^201^50^21^301, 1^50^21^50^21$
7	(40, 43)	$01^40101^30^21^501, 0^21^401^201^501^401$
8	(51, 56)	$1^701^201^30^2101^401^3, 1^201^201^40101^401^301^50$
9	(61, 76)	$1^401^30^21^501^201^501^401, 1^80^21^601^30101^30^21^501^4$
10	(92, 93)	$1^901^4010^21^90^210^21^901^4, 1^901^80^21^20^21^60^21^401^801$
11	(109, 112)	$1^301^301^9010^21^80^21^30101^801^401^4, 1^20101^30101^401^70^21^501^301^60101^1001$
12	(125, 134)	$1^{11}0101^8010^2101^{10}01^801^501^40101^2, 1^901^801^901^201^30101^701^20101^301^{11}0$
13	(141, 154)	$101^201^{11}01^{10}010^21^401^{10}01^{11}010^21^601^2, 1^{10}01^401^{11}01^401^{10}010^21^30^2101^{10}01^401^201^4$
14	(173, 182)	$1^201^701^601^301^701^601^{10}0101^{10}01^50101^601^20^21^7, 1^40^21^201^80101^701^801^501^801^70101^801^20^21^{13}01^2$
15	199 204	$1^501^{10}01^70^21^401^201^501^701^801^{10}01^301^{12}01^501^201^401$ $1^{11}01^601^901^201^{11}01^301^501^20^21^201^{11}0101^{10}01^{13}01^2$
16	231 236	$1^{14}01^901^601^501^20^21^401^{12}01^301^701^20^21^{12}01^801^501^{10}01$ $1^{15}01^901^20^21^401^9010^21^901^{14}01^90^21^{11}01^501^201^901^3$
17	255 278	$1^{14}01^{16}0101^201^{16}0101^501^601^301^901^201^{11}0101^701^601^30^21^8$ $1^401^601^501^301^{10}01^601^501^80^21^{10}0101^601^70101^801^{12}0101^{16}01^90^21$
18	298 311	$1^201^201^{16}0101^{16}0^21^401^{17}01^201^50101^{13}01^201^601^201^501^{12}01^{10}01^7$ $1^601^801^301^{17}0^21^30^21^{12}01^{16}01^701^90^210^21^401^901^701^401^{14}01^{10}01^6$
19	337 346	$1^{11}01^20101^{15}0^21^301^201^{16}01^{17}0^2101^8010^21^{17}01^{16}01^201^30^2101^{15}01^201^{10}01^3$ $1^{11}01^{18}01^50101^601^70101^301^201^{18}01^90101^20^21^401^301^20101^{12}01^30101^{12}01^{13}01^5$
20	379 388	$1^{16}0^2101^{10}01^{18}01^501^60^21^201^501^{15}01^{10}0101^{16}01^30101^{16}01^201^90^21^{19}0101^201^9$ $1^{19}01^{11}01^401^70101^401^{13}01^901^40101^{14}0^21^30^2101^901^{18}01^40101^501^301^{10}01^{16}01^4$
21	399 404	$1^60101^60^21^{11}01^301^401^601^{10}01^{19}01^701^201^301^{10}01^60^21^401^601^201^{12}01^{16}0^21^{19}0101^501^801^7$ $1^{18}0101^50^21^301^{13}01^90^21^{19}0101^80101^{16}01^70101^{14}01^80101^{20}0101^301^{19}0101^{17}01$
22	443 462	$1^301^801^701^801^{19}01^801^70101^{12}01^70101^801^{16}0^21^{16}01^501^20101^{20}0101^{16}01^{19}01^301^20101^6$ $1^{18}01^{17}01^601^401^7010^21^501^{15}01^{10}01^{14}0^2101^60^21^201^801^{14}0^21^701^401^{10}01^{21}01^{10}01^{11}01^40^21^9$
23	505 506	$1^{11}01^{22}01^{20}0^21^901^401^80^21^601^{22}0^21^30^21^601^{10}01^{11}01^201^{11}01^{10}01^60^21^30^21^{22}01^601^901^401^{10}01^{10}$ $1^{22}0^21^{18}01^201^{19}01^301^201^80^21^{18}01^901^{12}01^20^21^901^{19}01^20101^601^{10}01^401^{10}01^{19}010^21^601^{14}01^{10}$
24	567 592	$01^{18}01^80^21^{15}01^801^{21}0^21^801^{16}010^21^601^201^{14}01^{10}01^{20}01^{10}01^{14}01^201^60^2101^{16}01^80^21^{21}01^80^2$ $1^{14}0^21^201^3$ $1^{21}01^{18}01^{16}01^401^701^60101^{14}01^30^21^801^701^301^201^{20}0^21^301^701^{15}01^701^30^21^{20}01^201^301^701^901^{18}$ $0101^601^{21}01^701^5$
25	585 606	$1^{19}01^901^{18}0^21^{21}01^{15}0^2101^{10}01^{14}01^{13}01^501^9010^21^801^601^80^2101^901^501^{13}01^{14}01^{10}010^21^{15}01^{21}$ $0^21^401^901^2$ $1^{24}01^401^{10}01^{17}01^801^401^{23}0101^{12}01^501^{12}01^90^21^301^801^601^{21}01^601^801^30^21^901^{12}01^501^{12}0101^{23}$ $01^401^70^21^301^{10}01$

Continued on Next Page...

Table A.15: Good palindromic partitions for $\text{pdw}(2; 3, t) - 1$ according to Theorem 5.1, Part (i)

t	$\text{pdw}(2; 3, t) - 1$	Good palindromic partitions
26	≥ 633	$1^3 01^{13} 01^4 01^{12} 01^{20} 01^{21} 01^6 01^{18} 01^7 01^6 01^7 0^2 1^{12} 01^{14} 0^2 1^4 01^2 01^{14} 01^{25} 01^2 01^4 0^2 1^{14} 01^{12} 0^2$ $1^{14} 01^4 01^2 01^{13} 01^4 01^9 01^{18} 01$
	≥ 642	$1^{15} 01^{23} 01^6 01^{13} 01^4 01^{14} 01^2 01^{13} 01^{11} 01^{15} 0^2 1^6 01^{12} 01^5 01^7 01^{14} 0^2 1^{11} 0^2 101^{13} 01^4 01^{17} 01^4 01^{13} 01^6$ $01^7 01^8 01^4 01^6 01^{13} 01^{17} 01^5$
27	≥ 663	$1^{18} 01^{18} 01^{22} 01^5 01^{13} 0^2 1^2 01^4 01^9 0^2 1^2 01^8 01^2 01^3 01^8 0^2 1^{12} 01^{18} 01^9 01^{14} 01^5 01^9 01^{18} 01^{12} 0^2$ $1^{12} 01^{14} 0^2 1^9 01^2 01^{25} 01^4 01^4$
	≥ 698	$1^{22} 01^6 01^{25} 01^{10} 01^{25} 01^{11} 0^2 1^9 01 01^5 01^3 01^7 0^2 1^9 01^{14} 01^{23} 01^3 0^2 1^9 01^{17} 0^2 1^3 01^5 01^{12} 01^{11} 01^7 0$ $1^9 0^2 1^7 01^{21} 01^{12} 01^{13} 01^6 01$
28	≥ 727	$1^{26} 01^5 01^{13} 01^8 01^{19} 01^{13} 01^{10} 01^{19} 01^{12} 01^{15} 01^7 0^2 1^{12} 01 01^{12} 0^2 1^4 01^9 01^2 01^{15} 01^{22} 01^{15} 01^2 01^9 0$ $1^5 01^{12} 01 01^{13} 01^{14} 01^8 01^{12} 01^4 01^{10} 01^2$
	≥ 742	$1^{21} 01^{22} 01 01^4 01^{25} 01^{10} 01 01^{23} 01^{11} 0^2 1^9 01^7 01^3 01^7 0^2 1^9 01^{14} 01^{27} 0^2 1^9 01^{17} 0^2 1^3 01^5 01^{12} 01^{11} 01^7$ $01^9 0^2 1^7 01^{21} 01^{12} 01^{13} 01^6 01$
29	≥ 809	$1^{22} 01^{20} 01^{28} 01^{13} 01^3 01^6 01^4 01^{16} 01^{22} 01^{16} 01^{11} 01^{12} 0^2 1^2 01^7 01^5 01^{10} 01^{11} 01^3 01^5 01^{22} 01 01^3 01^{24} 0$ $1^{11} 01^3 01^6 01^9 01^{11} 01^{12} 01^{11} 01^2 01^{24} 0^2 1^3 01^{12}$
	≥ 820	$1^{13} 01 0^2 1^{16} 01^{15} 01^{14} 01 01^4 01^2 01^{27} 01^4 0^2 1^{14} 01^{12} 01^{16} 01^6 01^7 01^{10} 01^7 01^6 01^{16} 01^{12} 01^{20} 01^9 01^4 01^{21} 0$ $1^{15} 01^7 01^6 01^{13} 01^6 0^2 1^{11} 01^{20} 01^{21} 01 01^5 01^6 01^4$
30	≥ 843	$1^{29} 01^{25} 0^2 1^{17} 01^9 01^5 0^2 1^{22} 01^7 01^{10} 01^9 01^4 01^9 0^2 1^{17} 01^{13} 01^8 01^{28} 01^5 01^{10} 01 01^{10} 01^5 01^8 01^{19} 0$ $1^{12} 01^{24} 01^{13} 01^4 01^{10} 01^2 01^{15} 01^9 01 0^2 1^{15} 01^{11}$
	≥ 854	$1^{29} 01 01^2 01 01^7 01^4 01^{16} 01^{18} 01^{14} 01^{21} 0^2 1^7 01^3 01^2 01^{13} 01^{14} 01^4 01^{16} 01^3 01^{11} 01^{24} 01^4 01^{10} 01^7 01^{25} 01^{10}$ $0^2 1^4 01^{18} 0^2 1^{15} 01^{11} 01^5 01 01^4 01^3 01^{21} 01^{18} 01^9 01^{13}$
31	≥ 915	$1^{12} 01^3 01^4 01^7 01^{11} 01^8 01^{16} 01^{10} 01^4 01^3 01^4 01^{19} 01^{29} 01^6 01^{26} 01^3 01 01^{24} 01^8 01^{16} 01 0^2 1^{25} 01^3 01^8 0$ $1^2 01^{16} 01^2 01 01^3 01^{24} 01^7 01^3 01^{24} 01^2 01 01^{28} 01^{20} 01^{11} 01 01^{16} 01^5$
	≥ 930	$1^{12} 01^{19} 01^{16} 01^{10} 01^9 01^3 01^4 01^{14} 01^{12} 01^{14} 01^{18} 0^2 1^{10} 01^4 01^{13} 01^8 01^5 01 01^{10} 01^{17} 01^{15} 01^{16} 01^8 01^5 01^8$ $01 01^4 01^8 01^{19} 01^{16} 01^5 01^{18} 01^4 01^{25} 01^{14} 01^{16} 01^3 01^9 01^6 01^{13} 01^{11} 01$
32	≥ 957	$1^{24} 01^3 01^7 01^{19} 01^6 01^4 01^{17} 01^{30} 0^2 1^6 01^5 01^{24} 0^2 1^{11} 01^{18} 01^4 01^{17} 01^{19} 01^{10} 01^{11} 01^7 01^{17} 0^2 1^{13} 0$ $1^{23} 01^{12} 01^{11} 01^7 0^2 1^3 01^{25} 01^4 01^5 01^7 01^{22} 01^7 01^{12} 01^{30} 01^2 01^6 01$
	≥ 962	$1^{11} 0^2 1^{18} 01^{21} 01^5 01^2 01^{17} 01^{19} 01^{16} 01^{13} 01^3 01^{24} 01^5 01^2 01^{23} 01^{12} 01^{23} 0^2 1^5 01^{24} 01 01^5 0^2 101^9 01^{17}$ $01^{15} 01^6 01^{19} 01^9 01^{24} 01^5 01^2 01^5 0^2 1^{23} 01^{27} 01^8 01^{19} 01^2$
33	≥ 995	$1^{31} 01 01^{20} 01^7 0^2 1^{12} 01^4 01^{30} 01^7 01^9 01^8 01^{11} 01^3 01^{13} 01^{18} 01^{12} 01^{11} 01^3 01 01^{10} 0^2 1^4 01^7 01^{21} 01^{14} 01^{28}$ $01 01^{11} 01^3 0^2 1^{17} 01^{30} 01^6 01 01^{28} 01^{32} 01^2 01^8 01^5 01^3 01^8 01^3 01^{13}$
	≥ 1004	$1^2 01^6 01^{22} 01^4 01^{32} 01^2 01^9 0^2 101^7 01^5 01^{23} 01^{18} 01^{30} 01 0^2 1^{16} 01^{13} 01^9 01^6 01^9 01^{19} 01^{11} 01^4 01^{17} 01^{10} 0$ $1^{26} 0^2 1^{18} 01^{10} 01^{23} 01^{12} 01^{15} 01^9 01^{12} 01^6 0^2 1^8 0^2 1^3 01^3 01^{12}$
34	≥ 1053	$1^9 01^{14} 01^{21} 01 01^4 01^{27} 01^8 01^{24} 01^{32} 01^3 01^4 01^{16} 01^{13} 01 01^2 01^{33} 01^3 01^4 01^3 01^{16} 01 01^{13} 01^3 01^{16} 01^6$ $01^{33} 0^2 1^9 01^4 01^{31} 01^2 01 01^{16} 01^4 01^{31} 01^3 01^2 01^7 0^2 1^{22} 01 01^8 01^{30} 01^3$
	≥ 1080	$1^{23} 0^2 1^9 01^{22} 01^{13} 01^{25} 0^2 1^{21} 01^8 01^{15} 01^{11} 01^{18} 0^2 1^4 01^{20} 01 01^2 01^5 01^{21} 01^8 01^4 01^{23} 01^4 01^8 01^{21}$ $01^{10} 01^{20} 01^4 0^2 1^2 01^{27} 01^{10} 01^4 01^2 01^{26} 01 01^{25} 01^8 01^9 01^{17} 01^4 0^2 1^{12} 01^{21} 0^2 1^3 01^3$
35	≥ 1113	$1^7 01^{12} 01^{21} 01^{28} 01^3 01^{16} 01 01^{23} 01^4 01 01^{10} 01^8 01^{14} 01^{28} 01^9 01^{12} 01^{21} 01^{23} 0^2 101^5 01^3 01^{21} 01^2 01^{23} 0$ $1^5 01^{14} 01^{15} 01^3 01 0^2 101^5 01^{17} 01^{10} 01^{23} 01 01^3 01^2 01^9 01^{21} 01^{14} 01^{30} 01^{25} 01^{10} 01^7$
	≥ 1154	$1^{28} 01^8 01^{18} 01^6 01^9 01^{29} 01^6 01^9 01^6 01^{12} 01^5 01 01^{10} 01^{25} 01^{28} 01^{13} 01^8 01^9 01^3 01^{25} 01^{20} 01^3 01^{21} 01^5 0$ $1^{10} 01^2 01^{34} 01 01^3 01^2 01^{15} 01^{17} 01^2 01^{27} 01^{19} 01^8 0^2 1^6 01 01^2 0^2 1^{19} 01^{30} 01^{21} 01^7$
36	≥ 1185	$1^6 01^{28} 01^{10} 01^{16} 01^{28} 01^{35} 01^4 01^7 01^{23} 01 01^2 01^8 01^2 01^3 01^5 01^{15} 01^2 01 01^{19} 01^{30} 01^{25} 01^{12} 01^3 01^{14} 0$ $1^8 01^{11} 01^7 01^4 01^{18} 01 01^7 01^3 01^2 01^4 01^6 01^{21} 01^3 01^{28} 01^{24} 01^{23} 01^3 0^2 101^{32} 01^5 01^{11} 01^{10} 01^{15}$
	≥ 1212	$1^3 01^{25} 01^{13} 01^{24} 01^5 01^4 01^{17} 01^{13} 01^{10} 01^8 01^{28} 01 0^2 1^{32} 01^{10} 01^{13} 01^{20} 01 01^4 01^{12} 01 01^{13} 01^{20} 01^8 01^4 01^{32}$ $01 01^{28} 01^4 01^{19} 01^{32} 01^6 01^{18} 01^2 01^6 0^2 1^3 01^{24} 01^6 01^{20} 01^{13} 01^{14} 01^7 01^{23} 01^4 01^8$
37	≥ 1271	$1^{30} 01^6 01^{15} 01^{13} 01^{23} 01 0^2 1^{15} 0^2 1^{28} 01^7 01^{28} 01^2 01^3 01^{13} 01^{22} 0^2 1^{12} 01^3 01 0^2 1^{36} 01^{28} 01^{16} 0^2 101^{16}$ $01^{23} 01^{13} 01^8 01^3 01^9 01^{12} 01^7 01^{29} 01^{16} 01^5 01 01^{15} 01 0^2 1^{33} 01^{30} 0^2 1^6 01^{12} 01^{17} 01^5 01^9 01^2 01^4 01^6$

Continued on Next Page...

Table A.15: Good palindromic partitions for $\text{pdw}(2; 3, t) - 1$ according to Theorem 5.1, Part (i)

t	$\text{pdw}(2; 3, t) - 1$	Good palindromic partitions
	≥ 1294	$1^7 01^{24} 01^{31} 01^{60} 01^{30} 01^{01} 01^{30} 01^{40} 01^{22} 01^{34} 01^{01} 01^{80} 01^{32} 01^{80} 01^{22} 01^{20} 01^{13} 01^{23} 01^{60} 01^{22} 01^{10} 01^{01} 01^4$ $01^9 01^{19} 01^4 01^{13} 01^7 01^{15} 01^{29} 01^{60} 01^3 01^{01} 01^{10} 01^{22} 01^{28} 01^{12} 01^8 01^{16} 01^8 01^{01} 01^{80} 01^{22} 01^{23} 01^4 01^{12} 01^9$
38	≥ 1335	$1^{10} 01^{35} 01^{12} 01^8 01^{12} 01^7 01^{01} 01^{33} 01^{14} 01^{01} 01^5 01^{13} 01^2 01^7 01^{22} 01^{34} 01^{13} 01^7 01^{26} 01^{01} 01^4 01^{37} 01^{13} 01^{01} 01^4$ $01^{36} 01^{10} 01^{27} 01^{13} 01^{01} 01^7 01^{18} 01^9 01^2 01^{29} 01^{18} 01^2 01^{28} 01^{01} 01^{21} 01^2 01^{19} 01^9 01^{36} 01^6$
	≥ 1368	$1^{11} 01^7 01^{36} 01^{32} 01^3 01^2 01^8 01^{20} 01^{01} 01^{11} 01^{33} 01^4 01^{16} 01^8 01^{28} 01^7 01^4 01^{26} 01^{10} 01^{22} 01^4 01^8 01^4 01^{11} 01^5$ $01^4 01^{11} 01^{13} 01^{24} 01^{17} 01^{13} 01^4 01^{22} 01^7 01^{17} 01^8 01^9 01^2 01^{27} 01^{20} 01^{18} 01^{22} 01^3 01^9 01^2 01^{25} 01^{10} 01^{27} 01$
39	≥ 1405	$101^6 01^5 01^6 01^2 01^6 01^{33} 01^4 01^{10} 01^2 01^{28} 01^6 01^4 01^{29} 01^{18} 01^6 01^{10} 01^{01} 01^{36} 01^{22} 01^{33} 01^{01} 01^{11} 01^2 01^{01} 01^{11} 01^{22}$ $01^{33} 01^4 01^{11} 01^{21} 01^2 01^{13} 01^3 01^{32} 01^{19} 01^{13} 01^8 01^{11} 01^{24} 01^{15} 01^{36} 01^{13} 01^6 01^{26} 01^2 01^3 01^{02} 01^{15} 01^{11} 01^4$
	≥ 1410	$1^3 01^{34} 01^{24} 01^3 01^{32} 01^4 01^{18} 01^2 01^{35} 01^{12} 01^7 01^{16} 01^{11} 01^6 01^9 01^{20} 01^6 01^9 01^{19} 01^5 01^{10} 01^{01} 01^{23} 01^6 01^{31}$ $01^3 01^{12} 01^7 01^9 01^6 01^{28} 01^6 01^{17} 01^{10} 01^2 01^9 01^{18} 01^6 01^2 01^8 01^{20} 01^{17} 01^3 01^{32} 01^5 01^{12} 01^{34} 01^{01} 01^{24} 01^{15} 01^{13} 01^3 01^7$

Appendix B. Using the OKlibrary

The OKlibrary, available at <http://www.ok-sat-library.org> is an open-source research and development platform for SAT-solving and related areas (attacking hard problems); see [47] for some general information. For the purpose of reproduction of all results, one can use the Git ID “4cea9abf851424ca56f2ad0e4b8be2d707b041c2” (package 00147).²² For the purposes of this article the following components are directly relevant:

- The OKlibrary provides an already rather extensive library of functions for the computer algebra system Maxima²³. For example all hypergraph generators discussed in this article, and all vdW- and palindromic vdW-numbers can be computed and investigated at this level.
- For computations which take more time, C++ implementations are available.
- The OKlibrary provides easy access to (original) SAT solvers and related tools (as “external sources”).²⁴
- Finally these components are integrated into tools for running and evaluating experiments.²⁵

In the following sections we demonstrate the use of these tools. Some general technical remarks:

1. The installed OKlibrary lives inside directory OKplatform.
2. Inside this directory the Maxima-installation is called via `oklib --maxima` on the (Linux) command-line.
3. The C++ programs as well as the external sources, here the various SAT solvers, are placed on the path of the (Linux) user, and are thus callable by their name on the command-line (anywhere).

Appendix B.1. Numbers and certificates

All known vdW-numbers and palindromic vdW-numbers and known bounds are available at the computer-algebra level in the OKlibrary (using Maxima). For example the (known) numbers $w(2; 3, t)$ and $\text{pdw}(2; 3, t)$ are printed as follows (where inside the OKlibrary we typically use the letter “k” for the length of an arithmetical progression, not “t” as in this article):

²²Via the Git ID one can identify the versions of programs used in the article. The package provides the sources and a build system. Since building depends on the environment (the operating system to start with), there can not be a guarantee for the build to succeed, but perhaps later (or earlier) packages need to be used.

²³<http://maxima.sourceforge.net/>

²⁴The aim is to serve as a comprehensive collection, also maintaining “historical” versions.

²⁵In general we use the R system for statistical evaluation.

```

OKplatform> oklib --maxima
(%i1) oklib_load_all();
(%i2) output(N) := block([L],
  print(" k   vdw           pdvdw           span           gap"),
  for k : 3 thru N do (L : [3,k],
    printf(true, "~2,d ~12,a ~27,a ~12,a ~12,a~%",
      k, vanderwaarden(L), pdvanderwaarden(L), pd_span(L), pd_gap(L))))$
(%i3) output(40);

```

k	vdw	pdvdw	span	gap
3	9	[6,9]	3	0
4	18	[15,16]	1	2
5	22	[16,21]	5	1
6	32	[30,31]	1	1
7	46	[41,44]	3	2
8	58	[52,57]	5	1
9	77	[62,77]	15	0
10	97	[93,94]	1	3
11	114	[110,113]	3	1
12	135	[126,135]	9	0
13	160	[142,155]	13	5
14	186	[174,183]	9	3
15	218	[200,205]	5	13
16	238	[232,237]	5	1
17	279	[256,279]	23	0
18	312	[299,312]	13	0
19	349	[338,347]	9	2
20	[389,inf-1]	[380,389]	9	[0,inf-390]
21	[416,inf-1]	[400,405]	5	[11,inf-406]
22	[464,inf-1]	[444,463]	19	[1,inf-464]
23	[516,inf-1]	[506,507]	1	[9,inf-508]
24	[593,inf-1]	[568,593]	25	[0,inf-594]
25	[656,inf-1]	[586,607]	21	[49,inf-608]
26	[727,inf-1]	[634,643]	9	[84,inf-644]
27	[770,inf-1]	[664,699]	35	[71,inf-700]
28	[827,inf-1]	[[728,inf-1],[743,inf-1]]	[15,0]	[84,0]
29	[868,inf-1]	[[810,inf-1],[821,inf-1]]	[11,0]	[47,0]
30	[903,inf-1]	[[844,inf-1],[855,inf-1]]	[11,0]	[48,0]
31	[931,inf-1]	[[916,inf-1],[931,inf-1]]	[15,0]	[0,0]
32	[1007,inf-1]	[[958,inf-1],[963,inf-1]]	[5,0]	[44,0]
33	[1064,inf-1]	[[996,inf-1],[1005,inf-1]]	[9,0]	[59,0]
34	[1144,inf-1]	[[1054,inf-1],[1081,inf-1]]	[27,0]	[63,0]
35	[1205,inf-1]	[[1114,inf-1],[1155,inf-1]]	[41,0]	[50,0]
36	[1258,inf-1]	[[1186,inf-1],[1213,inf-1]]	[27,0]	[45,0]
37	[1339,inf-1]	[[1272,inf-1],[1295,inf-1]]	[23,0]	[44,0]
38	[1379,inf-1]	[[1336,inf-1],[1369,inf-1]]	[33,0]	[10,0]
39	[1419,inf-1]	[[1406,inf-1],[1411,inf-1]]	[5,0]	[8,0]
40	unknown	unknown	unknown	unknown

As one can see, if only bounds are known instead of a precise number n resp. number-pair (p, q) , then the numbers $x \in \{n, p, q\}$ are replaced by pairs (a, b) with $a \leq x \leq b$. Here $b = \text{inf}-1$ indicates that the number is finite, but no more

precise upper bounds are known.²⁶ So for example we only know currently that $w(2; 3, 20) \geq 389$, and this is shown by the interval $[389, \text{inf} - 1]$. Span and gap are simply computed according to definition, where in maxima the symbol inf is treated here like an unknown. That implies $\text{inf} - \text{inf} = 0$, and thus for palindromic span and gap the “0” in the second position indicate that the numbers in the first positions could go up or down. For example $w(2; 3, 20) \geq 389$ and $\text{pdw}(2; 3, 20) = (380, 389)$, whence nothing can be said about $\text{pdg}(2; 3, 20) = w(2; 3, 20) - \text{pdw}(2; 3, 20)_2$ except the trivialities that it is at least 0 and less than infinity; the latter becomes $(\text{inf} - 1) - 389 = \text{inf} - 390$.

Also the certificates (good partitions) are available, in various representations. First the certificate for $w(2; 3, 20) \geq 389$ (see Subsection Appendix A.1), for which we check that it is in fact a palindromic certificate:

```
(%i4) full_certificate_string_vdw_3k(20);
(%o4) ["1^{19}01^{11}01^{4}01^{7}0101^{4}01^{13}01^{9}01^{4}0101^{14}
0^{2}1^{3}0^{2}101^{9}01^{18}01^{9}01^{4}0101^{5}01^{3}01^{10}01^{16}
01^{8}01^{16}01^{10}01^{3}01^{5}0101^{4}01^{9}01^{18}01^{9}010^{2}1^{3}
0^{2}1^{14}0101^{4}01^{9}01^{13}01^{4}0101^{7}01^{4}01^{11}01^{19}"]
(%i5) certificate_pdvdw_p([3,20],388,full_certificate_vdw_3k(20)[1]);
(%o5) true
```

And here certificates for palindromic number-pairs:

```
(%i6) cfull_certificate_string_pdvdw_3k(34);
(%o6) [["1^{9}01^{14}01^{21}0101^{4}01^{27}01^{8}01^{24}01^{32}01^{3}0
1^{4}01^{16}01^{13}0101^{2}01^{33}01^{3}01^{4}01^{3}01^{16}0101^{13}
01^{3}01^{16}01^{6}01^{33}0^{2}1^{9}01^{4}01^{31}01^{2}0101^{16}01^{4}
01^{31}01^{3}01^{2}01^{7}0^{2}1^{22}0101^{8}01^{30}01^{3}"],
["1^{23}0^{2}1^{9}01^{22}01^{13}01^{25}0^{2}1^{21}01^{8}01^{15}0
1^{11}01^{18}0^{2}1^{4}01^{20}0101^{2}01^{5}01^{21}01^{8}01^{4}01^{23}
01^{4}01^{8}01^{21}01^{10}01^{20}01^{4}0^{2}1^{2}01^{27}01^{10}01^{4}0
1^{2}01^{26}0101^{25}01^{8}01^{9}01^{17}01^{4}0^{2}1^{12}01^{21}0^{2}
1^{3}01^{3}"]]]
(%i7) extract_data_certificates_pdvdw_3k(34);
(%o7) [[3,34],1054,1081,
[[10,25,47,49,54,82,91,116,149,153,158,175,189,191,194,228,232,
237,241,258,260,274,278,295,302,336,337,347,352,384,387,389,
406,411,443,447,450,458,459,482,484,493,524]],
[[24,25,35,58,72,98,99,121,130,146,158,177,178,183,204,206,209,
215,237,246,251,275,280,289,311,322,343,348,349,352,380,391,
396,399,426,428,454,463,473,491,496,497,510,532,533,537]]]
```

With the first command we get the representation of the good partitions as used in this paper (where now for the palindromic situation we have two good partition according to Theorem 5.1), while the second command yields a list with five elements: first the parameter tuple, then the two components of the palindromic number-pair, and then two lists with the good partitions available, now represented via the block in the partition for the second colour.

Analysing the patterns according to Section 4, and applying these measurements to the certificates stored in the OKlibrary for $20 \leq t \leq 39$ is done as follows:

```
(%i8) for k : 20 thru 39 do
print(k,firste(vanderwaerden3k(k)),
map(analyse_certificate,full_certificate_vdw_3k(k)));
20 389 [[48,340],[44,45],[4,37],[5,27],[20,1]]]
21 416 [[50,365],[43,44],[7,34],[13,26],[8,1]]]
22 464 [[54,409],[51,52],[3,47],[5,40],[27,1]]]
```

²⁶In principle there exist theoretical upper bounds, but for practical purposes these bounds are completely useless.

```

23 516 [[ [59,456], [53,54], [6,45], [12,36], [17,1] ] ]
24 593 [[ [63,529], [57,58], [6,54], [13,37], [20,1] ] ]
25 656 [[ [74,581], [69,70], [5,64], [11,45], [16,2] ] ]
26 727 [[ [78,648], [72,72], [6,64], [13,42], [21,1] ] ]
27 770 [[ [79,690], [72,73], [7,65], [15,58], [11,2] ] ]
28 827 [[ [79,747], [74,75], [5,64], [11,44], [19,1] ] ]
29 868 [[ [81,786], [76,77], [5,69], [11,57], [27,1] ] ]
30 903 [[ [83,819], [76,77], [7,67], [13,57], [15,1] ] ]
31 931 [[ [82,848], [80,81], [2,77], [5,53], [58,1] ] ]
32 1007 [[ [87,919], [82,83], [5,78], [9,62], [29,1] ] ]
33 1064 [[ [89,974], [85,86], [4,80], [9,58], [25,1] ] ]
34 1144 [[ [96,1047], [87,88], [9,80], [19,63], [23,2] ] ]
35 1205 [[ [95,1109], [91,92], [4,84], [9,67], [41,1] ] ]
36 1258 [[ [101,1156], [97,98], [4,88], [9,72], [42,1] ] ]
37 1339 [[ [105,1233], [97,98], [8,90], [17,65], [30,2] ] ]
38 1379 [[ [104,1274], [96,97], [8,91], [17,73], [26,1] ] ]
39 1419 [[ [105,1313], [98,99], [7,95], [13,72], [46,1] ] ]

```

Per line we print out three items: t , the lower bound on $w(2; 3, t)$ and the list of data for each stored certificate. Now currently we have only stored one certificate for each $20 \leq t \leq 39$, and thus the third item contains just one list, with five pairs for the different statistics.²⁷ These five pairs have the following meaning:

1. First come n_0 and n_1 .
2. Then come the numbers of terms 0^s and 1^s (we don't use "00" here, and so these terms alternate, and thus their numbers differ at most by one).
3. Then from these counts the cases with $s = 1$ are excluded; thus the first element of the pair is n_{00} .
4. Now these exponents s are put in the list, and the sums of the numbers of peaks and valleys are computed; again for block 0 and block 1 of the partition, and thus now the second element of the pair is $n_p + n_v$.
5. Finally for these lists of exponents the maximal size of an interval with constant values is computed; thus if there were a second element of the pair with value 3 or greater, then Question 4.2 would have been answered in the positive.

The value of d from Subsection 3.3 is computed as follows:

```

(%i9) lmax(Delta_1(map(first,create_list(vanderwaarden([3,k]),k,3,39)))
          /create_list(k,k,3,38));
(%o9) 77/23
(%i10) round_fdd(77/23/2,3);
(%o10) 1.674

```

Appendix B.2. Hypergraphs

The hypergraphs are available at Maxima-level, and the computationally expensive palindromic hypergraph also at C++ level:

```

(%i11) arithprog_hg(3,5);
(%o11) [{1,2,3,4,5},{1,2,3},{1,3,5},{2,3,4},{3,4,5}]
(%i12) arithprog_pd_hg(3,5);
(%o12) [{1,2,3},{1,3},{2,3}]

```

> PdVanderWaerden-03-DNDEBUG 3 5

²⁷We found more than one solution in each case, but always very similar to the one stored; there seems to be a clustering of solutions, and perhaps there is always only one (or very few) cluster.

c Palindromised hypergraph with arithmetic-progression length 3
and 5 vertices.

```
p cnf 3 2
1 3 0
2 3 0
```

Appendix B.3. SAT instances

The SAT-instance for considering $w(2; 3, t)$ with n vertices is created by the program call

```
VanderWaerdenCNF-03-DNDEBUG 3 t n,
```

for example for $t = 4$ and $n = 6$

```
> VanderWaerdenCNF-03-DNDEBUG 3 4 6
> cat VanDerWaerden_2-3-4_6.cnf
c Van der Waerden numbers with partitioning into 2 parts;
  SAT generator written by Oliver Kullmann, Swansea, May 2004, October 2010.
c Arithmetical progression sizes  $k_1 = 3$ ,  $k_2 = 4$ .
c Number of elements  $n = 6$ .
c Iterating through the arithmetic progressions in colexicographical order.
p cnf 6 9
1 2 3 0
2 3 4 0
1 3 5 0
3 4 5 0
2 4 6 0
4 5 6 0
-1 -2 -3 -4 0
-2 -3 -4 -5 0
-3 -4 -5 -6 0
```

The SAT-instance for considering $pdw(2; 3, t)$ with n vertices is created by the program call

```
PdVanderWaerdenCNF-03-DNDEBUG 3 t n,
```

for example for $t = 4$ and $n = 9$

```
> PdVanderWaerdenCNF-03-DNDEBUG 3 4 9
> cat VanDerWaerden_pd_2-3-4_9.cnf
c Palindromic van der Waerden problem: 2 parts, arithmetic progressions of
  size 3 and 4, and 9 elements, yielding 5 variables.
p cnf 5 10
1 2 3 0
2 4 0
1 3 4 0
1 5 0
2 5 0
3 5 0
4 5 0
-2 -4 0
-1 -3 -5 0
-3 -4 -5 0
```

Appendix B.4. The SAT solvers

All solvers are installed via the OKlibrary; the OKsolver²⁸ is a solver specific to the OKlibrary, tawSolver-2.6²⁹ was developed in it (starting with version 1.0), and satz³⁰ as well as march_pl are maintained in the OKlibrary. Example output for the column $t = 12$ in Table 9, with the instance produced by

```
> VanderWaerdenCNF-03-DNDEBUG 3 12 135
```

resp.

```
> VanderWaerdenCNF-03-DNDEBUG 3 12 134
```

for the satisfiable case, is provided in the following.

Appendix B.4.1. tawSolver

First tawSolver-2.6 (output with one additional line-break for the url):

```
> tawSolver -v
tawSolver:
authors: Tanbir Ahmed and Oliver Kullmann
url's:
  http://sourceforge.net/projects/tawsolver/
  https://github.com/OKullmann/oklibrary/blob/master/
    Satisfiability/Solvers/TawSolver/tawSolver.cpp
Version: 2.6.6
Last change date: 17.8.2013
Mapping k -> weight, for clause-lengths k specified at compile-time:
  2->4.85 3->1 4->0.354 5->0.11 6->0.0694
Divisor for open weights: 1.46
Option summary = ""
Macro settings:
  LIT_TYPE = std::int32_t (with 31 binary digits)
  UCP_STRATEGY = 1
Compiled without TAU_ITERATION
Compiled without ALL_SOLUTIONS
Compiled without PURE_LITERAL
Compiled with NDEBUG
Compiled with optimisation options
Compilation date: Aug 17 2013 21:38:43
Compiler: g++, version 4.7.3
Provided in the OKlibrary http://www.ok-sat-library.org
Git ID = 237cbfc4d9b772a29e125928959af14cb4495d3e

> tawSolver VanDerWaerden_2-3-12_135.cnf
s UNSATISFIABLE
c number_of_variables           135
c number_of_clauses             5251
c maximal_clause_length        12
c number_of_literal_occurrences 22611
c running_time(sec)            10.58
```

²⁸<https://github.com/OKullmann/oklibrary/tree/master/Satisfiability/Solvers/OKsolver/SAT2002>

²⁹<https://github.com/OKullmann/oklibrary/blob/master/Satisfiability/Solvers/TawSolver/tawSolver.cpp>

³⁰<https://github.com/OKullmann/oklibrary/blob/master/Satisfiability/Solvers/Satz/satz215.2.c>

```

c number_of_nodes          961949
c number_of_binary_nodes   480974
c number_of_1-reductions   11312180
c reading-and-set-up_time(sec) 0.004
c file_name                VanDerWaerden_2-3-12_135.cnf
c options                  ""

```

A “binary node” is one with two children, i.e., where the second branch was not explored since the first branch was found satisfiable. And a “1-reduction” is one assignment of a literal x to true due to a unit-clause $\{x\}$. Calling `rawSolver-2.6` happens via `ttawSolver`, and the counting version is called `ctawSolver`.

```

> ttawSolver VanDerWaerden_2-3-12_135.cnf
s UNSATISFIABLE
c number_of_variables      135
c number_of_clauses       5251
c maximal_clause_length   12
c number_of_literal_occurrences 22611
c running_time(sec)       19.29
c number_of_nodes         953179
c number_of_binary_nodes  476589
c number_of_1-reductions  11285634
c number_of_pure_literals 1317
c reading-and-set-up_time(sec) 0.005
c file_name               VanDerWaerden_2-3-12_135.cnf
c options                 "PT5"

```

```

> ctawSolver VanDerWaerden_2-3-12_135.cnf
s UNSATISFIABLE
c number_of_variables      135
c number_of_clauses       5251
c maximal_clause_length   12
c number_of_literal_occurrences 22611
c running_time(sec)       10.64
c number_of_nodes         961949
c number_of_binary_nodes  480974
c number_of_1-reductions  11312180
c number_of_solutions     0
c reading-and-set-up_time(sec) 0.005
c file_name               VanDerWaerden_2-3-12_135.cnf
c options                 "A19"

```

Options are reported via acronyms: “P” for pure literals, “T” for the tau-heuristics, followed by the number of iterations of the Newton-Raphson method, and “A” for all solutions, followed by the number of decimal digits for counting. Instead of just counting, we can also output all solutions, for example to standard output:

```

> ctawSolver VanDerWaerden_2-3-12_134.cnf -cout
v 1 2 3 4 5 6 7 8 9 -10 11 12 13 14 15 16 17 18 -19 20 21 22 23 24 25 26 27 28
-29 30 31 -32 33 34 35 -36 37 -38 39 40 41 42 43 44 45 -46 47 48 -49 50 -51 52
53 54 -55 56 57 58 59 60 61 62 63 64 65 66 -67 -68 69 70 71 72 73 74 75 76 77
78 79 -80 81 82 83 -84 85 -86 87 88 -89 90 91 92 93 94 95 96 -97 98 -99 100
101 102 -103 104 105 -106 107 108 109 110 111 112 113 114 115 -116 117 118 119
120 121 122 123 124 -125 126 127 128 129 130 131 132 133 134 0
s SATISFIABLE

```

```

c number_of_variables          134
c number_of_clauses           5172
c maximal_clause_length       12
c number_of_literal_occurrences 22266
c running_time(sec)           10.56
c number_of_nodes              968509
c number_of_binary_nodes       484254
c number_of_1-reductions       11308431
c number_of_solutions          1
c reading-and-set-up_time(sec) 0.004
c file_name                    VanDerWaerden_2-3-12_134.cnf
c options                      "A19"

```

The solution is given in the DIMACS format for partial assignments, with positive literals setting the underlying variable to `true`, and negative literals setting them to `false` (so positive literals are the elements of the partition for $t = 12$ here). For all options, use `tawSolver` without arguments, or see the source code. Finally we note that `rawSolver-2.6` and `ctawSolver` are just compilations of the `tawSolver` with specific options set³¹, namely:

```

ttawSolver:
-DPURE_LITERALS -DTAU_ITERATION=5

```

```

ctawSolver:
-DALL_SOLUTIONS

```

```

cttawSolver:
-DTAU_ITERATION=5 -DALL_SOLUTIONS

```

Appendix B.4.2. satz

Now `sat`:

```

> satz215 VanDerWaerden_2-3-12_135.cnf
s UNSATISFIABLE
c sat_status                   0
c number_of_variables          135
c initial_number_of_clauses    5251
c reddiff_number_of_clauses    0
c running_time(sec)            76.73
c number_of_nodes              262304
c number_of_binary_nodes       133373
c number_of_pure_literals       55
c number_of_1-reductions        5482044
c number_of_2-look-ahead        30069498
c number_of_2-reductions        1196400
c number_of_3-look-ahead        563872
c number_of_3-reductions        257097
c file_name                    VanDerWaerden_2-3-12_135.cnf

```

Here “`reddiff`” is the “difference due to reduction” in the number of clauses: clauses can be removed by subsumption (not applicable here), while clauses can be added by resolution (does not happen here).

³¹see <https://github.com/OKullmann/oklibrary/blob/master/Satisfiability/Solvers/TawSolver/makefile> for the `makefile` in the `OKlibrary`

Appendix B.4.3. *march_pl*

Now *march_pl*:

```
> march_pl VanDerWaerden_2-3-12_135.cnf
c main():: nodeCount: 47963
c main():: dead ends in main: 110
c main():: lookAheadCount: 10456897
c main():: unitResolveCount: 274045
c main():: time=184.539993
c main():: necessary_assignments: 5287
c main():: bin_sat: 0, bin_unsat 0
c main():: doublelook: #: 421439, succes #: 321732
c main():: doublelook: overall 4.150 of all possible doublelooks executed
c main():: doublelook: succesrate: 76.341, average DL_trigger: 273.489
s UNSATISFIABLE
```

Appendix B.4.4. *OKsolver*

And to conclude the complete solvers, the *OKsolver*:

```
> OKsolver_2002-03-DNDEBUG VanDerWaerden_2-3-12_135.cnf
s UNSATISFIABLE
c sat_status 0
c initial_maximal_clause_length 12
c initial_number_of_variables 135
c initial_number_of_clauses 5251
c initial_number_of_literal_occurrences 22611
c number_of_initial_unit-eliminations 0
c reddiff_maximal_clause_length 0
c reddiff_number_of_variables 0
c reddiff_number_of_clauses 0
c reddiff_number_of_literal_occurrences 0
c number_of_2-clauses_after_reduction 0
c running_time(sec) 215.8
c number_of_nodes 281381
c number_of_single_nodes 0
c number_of_quasi_single_nodes 0
c number_of_2-reductions 2049274
c number_of_pure_literals 29
c number_of_autarkies 0
c number_of_missed_single_nodes 0
c max_tree_depth 36
c proportion_searched 1.000000e+00
c proportion_single 0.000000e+00
c total_proportion 1
c number_of_table_enlargements 0
c number_of_1-autarkies 490
c number_of_new_2-clauses 0
c maximal_number_of_added_2-clauses 0
c file_name VanDerWaerden_2-3-12_135.cnf
```

Appendix B.4.5. Ubcsat

If we want to run an algorithm from the Ubcsat-suite³² on its own (while running it in the iterative fashion, as discussed in Subsection 6.2, is shown in the following Appendix B.5), for example `gsat-tabu`, then this can be done as follows (using an additional line-break in the command-line, and four additional line-breaks in the first output-line), for a cut-off 10^6 , ten runs, and initial seed 0 (for reproducibility):

```
> ubcsat-okl -alg gsat-tabu -cutoff 1000000 -runs 10 -seed 0
-i VanDerWaerden_2-3-12_134.cnf
# -rclean -r out stdout run,found,best,beststep,steps,seed -r stats stdout
numclauses,numvars,numlits,fps,beststep[mean],steps[mean+max],percentsolve,
best[min+max+mean+median] -runs 10 -cutoff 100000 -rflush
-alg gsat-tabu -cutoff 1000000 -runs 10 -seed 0
-i VanDerWaerden_2-3-12_134.cnf
    sat  min          osteps          msteps          seed
    1  0      1          3588          1000000          0
    2  1      0          543154         543154 1492175541
    3  0      1          5687          1000000 367425000
    4  0      1          3152          1000000 3611176606
    5  0      1          164885         1000000 388711246
    6  0      1          50599          1000000 4160687068
    7  0      1          3533          1000000 533276301
    8  0      1          94759          1000000 1146607069
    9  0      1          2921          1000000 3903233437
   10 0      1          8071          1000000 127100396
```

```
Clauses = 5172
Variables = 134
TotalLiterals = 22266
FlipsPerSecond = 513073
BestStep_Mean = 88034.9
Steps_Mean = 954315.4
Steps_Max = 1000000
PercentSuccess = 10.00
BestSolution_Mean = 0.9
BestSolution_Median = 1
BestSolution_Min = 0
BestSolution_Max = 1
```

Here we use the wrapper-script `ubcsat-okl`³³, which outputs the output for the runs in a style typical for statistical data (easily readable for example by the tool `R`³⁴, as used in the `OKLibrary`):

1. First a comment-line, starting with “#”, showing the parameters passed to the `ubcsat`-program (everything until “-rflush” is the default, coming from `ubcsat-okl`, and after that come the parameters from the command-line (possibly overriding the defaults)).
2. Then a line with the headings for the six output columns (`osteps` is for the number of rounds for reaching the optimum, while `msteps` is for the maximum number of steps).
3. Followed by data for the runs (above, one of the ten runs was successful).
4. Finally summary statistics (this is not readable by tools like `R`, and needed to be removed; however for a quick human-readable overview it is useful).

³²see <http://ubcsat.dtopkins.com/>

³³see link to shell script

³⁴<http://www.r-project.org/>

Appendix B.5. Running experiments

For running UbcSAT-algorithm to determine lower bounds on $w(2; 3, t)$ and $pdw(2; 3, t)$, also providing “conjectures” on the precise values, we have the following tools (using no parameters here serves to print some basic helper-information). First the general tool for $w(2; t_0, t_1)$:

```
> RunVdWk1k2
ERROR[RunVdWk1k2]: Six parameters k1, k2, n0, alg, runs, cutoff
are needed: The progression-lengths k1,k2, the starting number n0 of
vertices, the ubcsat-algorithm, the number of runs, and the cutoff.
An optional seventh parameter is a path for the file containing an
initial assignment for the first ubcsat-run.
```

The special version with $k_1=3$, handling our case $w(2; 3, t)$:³⁵

```
> RunVdW3k
ERROR[RunVdW3k]: Five parameters k, n0, alg, runs, cutoff
are needed: The progression-length k, the starting number n0 of vertices,
the ubcsat-algorithm, the number of runs, and the cutoff.
An optional sixth parameter is a path for the file containing an
initial assignment for the first ubcsat-run.
```

For example

```
> RunVdW3k 27 678 gsat-tabu 1000 10000000
```

starts the investigation of $w(2; 3, 27)$ with $n = 678$ (ad-hoc, no solution given), where the cut-off value (the number of rounds for stochastic local search) is 10^6 , and 1000 runs are executed; from $n = 679$ on the first three runs will use the solution found for $n - 1$, while further runs use a random initial assignment.

Handling palindromic instances is done similarly³⁶:

```
> RunPdVdWk1k2
ERROR[RunPdVdWk1k2]: Five parameters k1, k2, alg, runs, cutoff
are needed: The progression-lengths k1,k2, the ubcsat-algorithm,
the number of runs, and the cutoff.
```

And for running complete solvers on palindromic instances we have³⁷:

```
> CRunPdVdWk1k2
ERROR[CRunPdVdWk1k2]: Three parameters k1, k2, solver, are needed:
The progression-lengths k1, k2 and the SAT solver.
```

³⁵see link to shell script

³⁶see link to shell script

³⁷see link to shell script