

Backtracking (the) Algorithms on the Hamiltonian Cycle Problem

Joeri Slegers

Mice & Man Software and A.I. Development
Amsterdam
The Netherlands
jslegers@hotmail.com

Daan van den Berg

Yamasan Science & Education
Amsterdam
The Netherlands
daan@yamasan.nl

Abstract—Even though the Hamiltonian cycle problem is NP-complete, many of its problem instances are not. In fact, almost all the hard instances reside in one area: near the Koplós-Szemerédi bound, where randomly generated graphs have an approximate 50% chance of being Hamiltonian. If the number of edges is either much higher or much lower, the problem is not hard – most backtracking algorithms decide such instances in (near) polynomial time. Recently however, targeted search efforts have identified very hard Hamiltonian cycle problem instances *very far away* from the Koplós-Szemerédi bound. In that study, the used backtracking algorithm was Vandegriend-Culberson’s, which was supposedly the most efficient of all Hamiltonian backtracking algorithms. In this paper, we make a unified large scale quantitative comparison for the best known backtracking algorithms described between 1877 and 2016. We confirm the suspicion that the Koplós-Szemerédi bound is a hard area for all backtracking algorithms, but also that Vandegriend-Culberson is indeed the most efficient algorithm, when expressed in consumed computing time. When measured in recursive effectiveness however, the algorithm by Frank Rubin, almost half a century old, performs best. In a more general algorithmic assessment, we conjecture that edge pruning and non-Hamiltonicity checks might be largely responsible for these recursive savings. When expressed in system time however, denser problem instances require much more time per recursion. This is most likely due to the costliness of the extra search pruning procedures, which are relatively elaborate. We supply large amounts of experimental data, and a unified single-program implementation for all six algorithms. All data and algorithmic source code is made public for further use by our colleagues.

Keywords—Hamiltonian Cycle; exact algorithm; exhaustive algorithm; heuristic; phase transition; order parameter; data analytics; instance hardness; replication.

I. PREAMBLE

Traversing a crack in the fabric of the scientific spacetime-continuum, this paper finds itself in the unusual position that its designated conclusions have already been overthrown. Following a replication study [1], these extended results should have been published earlier, but as history unfolded, they simply were not. In any case, the study by Cheeseman, Kanefsky & Taylor (henceforth: ‘Cetal’) was first [2]. Sharpening the resolution of the $P \stackrel{?}{=} NP$ problem, they showed that for various NP-complete problems, Hamiltonian cycle, graph colouring, satisfiability and the asymmetric traveling sales-

man problem¹, instances vary greatly in their computational hardness. Sporting over 1400 citations, the paper became a landmark in the field.

It took nearly 30 years for Cetal’s results on traveling salesman to be overthrown, appearing to have been flawed by an overlooked roundoff error [3] (also see the accompanying videos: [4] [5]). Their results on the Hamiltonian cycle problem however, were successfully replicated and published at IARIA’s Data Analytics 2018 conference, and brought the data, sourcecode and figures alive in online interactively publicly accessible resources [1] [6]. The results of that extended study showed that for Cetal’s algorithm, and two others, the hardest instances of the Hamiltonian cycle were located along an area known as the “Koplós-Szemerédi bound”. So far so good.

But the lingering question was how large the influence of the solvers was. Were the found hard instances hard specifically for the solving algorithms used? For the pruning methods? For the branching heuristic? And to make matters worse, a recent follow-up study showed that for the allegedly most efficient Hamiltonian cycle backtracker, the Vandegriend-Culberson (henceforth: ‘Vacul’)-algorithm, (which still requires exponential time in the worst case), the hardest instances were located *very far away* from the Koplós-Szemerédi bound [7] [8]. Only findable by sophisticated evolutionary algorithms and a whole lot of computing power, these very hard instances never showed up in earlier studies.

So is there an insurmountable contradiction here? Probably not. The authors conjecture three reasons for their counterintuitive results: “A first explanation for these surprising results is that these results are specific for the backtracking algorithm we used. This is unlikely however, as the algorithm minimizes most other backtracking algorithms found in literature (yet unpublished results)”. Put differently: chances are very high that these graphs are *also* hard for other complete backtracking algorithms but evidence pending, it remained unconfirmed as yet. This study will at the very least post a very serious sidenote to that hypothesis. A second explanation given by these authors might be that in most studies on Hamiltonian cycle backtracking algorithms, runs are cutoff after a preset

¹Even though the authors themselves dubbed traveling salesman as “NP-complete”, they solved the NP-hard version of the problem.

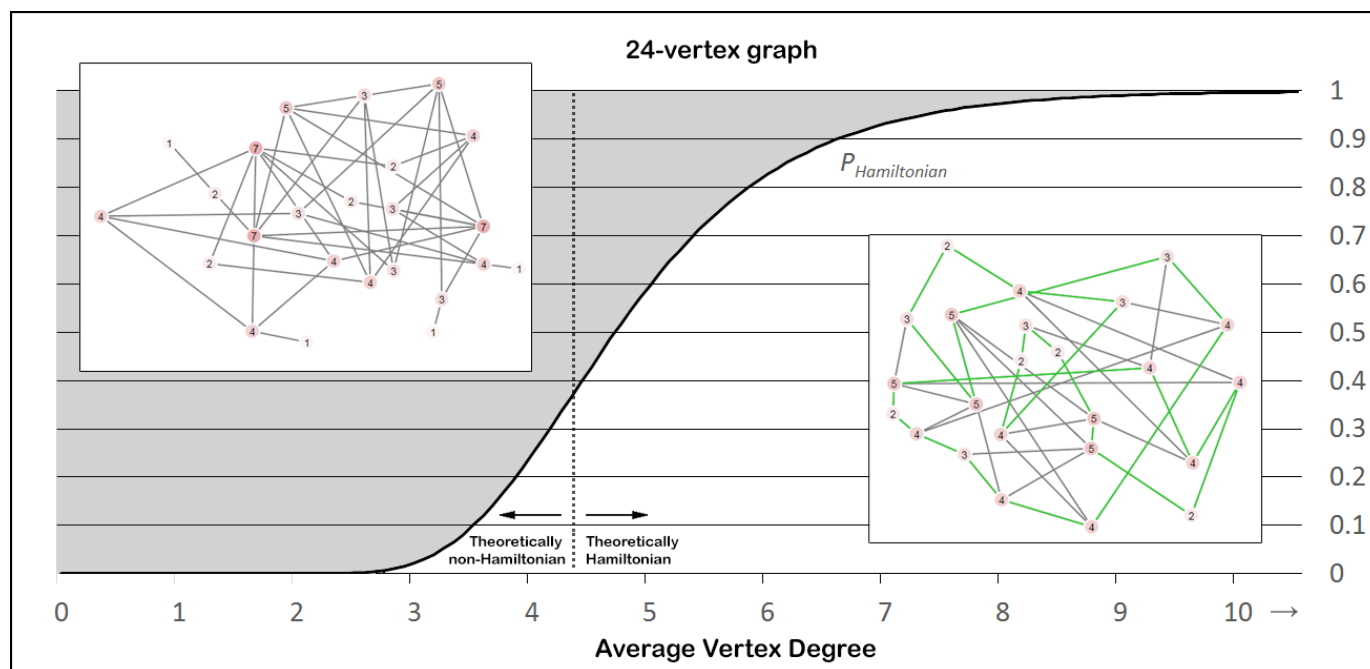


Fig. 1. The probability of a randomly generated graph being Hamiltonian depends on the average vertex degree, and is sigmoidally shaped around the ‘threshold point’ of $\ln(V) + \ln(\ln(V))$. Top-left inset is a non-Hamiltonian random graph, bottom-right inset is a Hamiltonian graph with the Hamiltonian cycle itself being highlighted.

number of recursions, as even small graphs can take up significant decision time. Although theoretically feasible, in practice these cutoff points are usually situated near the Komlós-Szemerédi bound, and not in edge-dense regions far away where the extremely hard instances were located.

The third and most compelling explanation might therefore come from their observation that on first glance, these very hard instances might have low Kolmogorov complexity – they are in some sense *structured* graphs. And as unstructured objects in randomly generated ensemble vastly outnumber the structured objects, the chances of being created by a stochastic process (which is the case in most large-scale comparative studies) are extremely small. Put differently: one would simply not find the very hard instances unless knowing exactly where to look. Or as the authors more poetically phrased: “These graphs are an isolated island of structured hardness in a ocean of unstructured easiness. Whether more such islands exist, and what they look like, awaits further exploration.”² [7].

So to understand their results, a couple of things need to happen. First, Vacul’s algorithm must *evidently* minimize the other backtracking algorithms. Then, a global evolutionary search algorithm should look for the hardest instances for *all* these backtracking algorithms. Third, a hardness hierarchy should be made for all six algorithms, another large quantitative study. Fourth, an aggregated theory should explain *why* some instances of the Hamiltonian cycle problem are harder than others, and what their relation to the Komlós-Szemerédi

²The literate reader might remind Aldous Huxley’s famous quote “Consider the horse. They considered it.” – *A Brave New World*, 1932. [9]

bound is (or is not!). It is the authors’ belief that such an aggregated theory exists. Beyond that, as a fifth point, might lie implications for computability, and the $P \stackrel{?}{=} NP$ problem itself, but developments in this area will depend on foregoing results, and are as yet too close to call.

In this paper, which is an extension of the IARIA’18 paper mentioned in the second paragraph [1], we will address step one: show that in large random ensembles, the hardest instances of the Hamiltonian cycle problem lie around the Komlós-Szemerédi bound for *all* well-known general backtracking algorithms found in literature. We will make exact calculations on large ensembles, and perform a rigorous comparative analysis. While reading the paper, the reader should keep two things in mind: first, other backtracking algorithms than those in this study are well existable and second: it is possible that much harder instances exist for any backtracker in this study or elsewhere. As shown earlier, extremely hard instances are likely extremely rare, extremely hard to find, but also extremely important for the $P \stackrel{?}{=} NP$ problem, as it is these instances that etch upper bounds on these algorithms’ runtimes. Finding these, possibly with targeted evolutionary algorithms, is the second step and will hopefully be done in the near future. For now, we will work with completely unbiased random ensembles of instances to map out the gross features of the complexity landscape.

II. INTRODUCTION

The “Great Divide” between P and NP has haunted computer science and related disciplines for over half a century. Problems in P are problems for which the runtime of the best

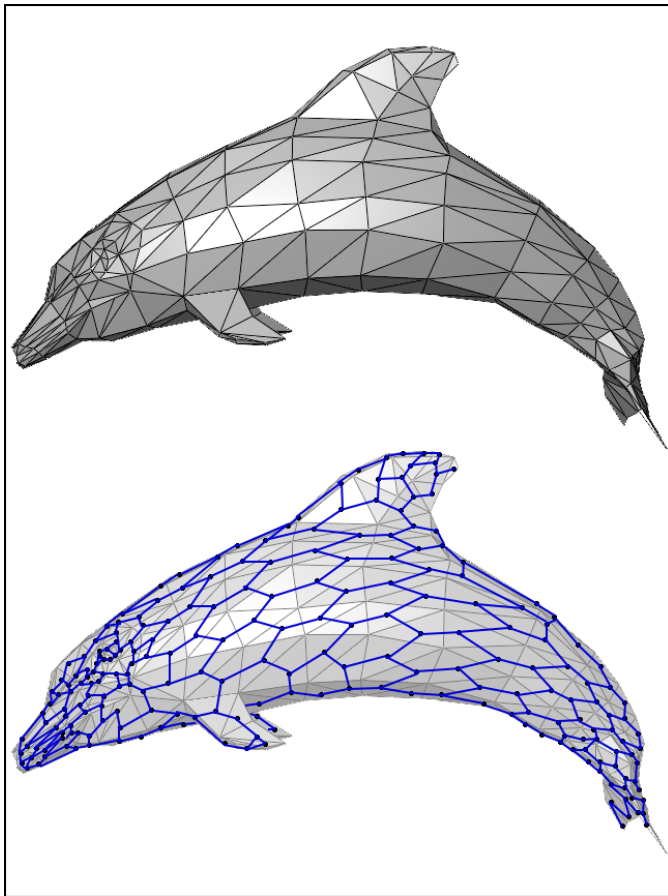


Fig. 2. Fast rendering of triangle mesh 3D images critically depends on finding Hamiltonian cycles through the corresponding 'cubic' graphs, in which every vertex has a maximum degree of three.

known algorithm increases polynomially with the problem size, for example, calculating the average of an array of numbers. If the array doubles in size, so does the runtime of the best known algorithm - a polynomial increase. A problem in NP however, has no such polynomial-time algorithm and it is an open question whether one will ever be found. An example hereof is 'satisfiability' (sometimes abbreviated to SAT), in which an algorithm assigns values 'true' or 'false' to variables in Boolean formulas like $(a \vee \neg b \vee d) \wedge (b \vee c \vee \neg d)$. The task is to choose the variable assignment so that the formula as a whole is satisfied (becomes 'true'), and returning that assignment, or making sure that no such assignment exists. Algorithms that do this, algorithms that are guaranteed to give a solution whenever it exists and return 'no' otherwise, are called *exact algorithms*³

Being exact is a great virtue for an algorithm, but it comes at a hefty price. Often, these algorithms operate by *brute-force* procedures: exhaustively trying all combinations for all variables until a solution is found, which usually takes vast amounts of time. Depth-first search is the keystone example

³Technically speaking, exact algorithms for decision problems such as the Hamiltonian cycle problem should be called *complete*. But since 'exact' seems more in schwing recently, we will stick to that.

of this study; it is exhaustively exact and indeed consumes harrowing amounts of computational power (see Figures 4 and 5). Smarter algorithms exist too; clever search pruning can speed things up by excluding large sections of state space at the cost of some extra computational instructions, an investment that usually pays off. Heuristic algorithms are also fast but not necessarily exact - so it is not guaranteed a solution is found if one exists. After decades of research, the runtimes of even the most efficient complete SAT-algorithm known today still increases exponentially with the number of variables - much worse than polynomial, even for low exponents. Therefore, SAT is in NP, a class of 'Notorious Problems' that rapidly become unsolvable as their size increases. In practice, this means that satisfiability problems (and other problems in NP) with only a few hundred variables are practically unsolvable, whereas industries such as chip manufacture or program verification in software engineering could typically employ millions ~ [10]⁷ [11].

The problem class NP might be considered "the class of dashed hopes and idle dreams" [12], but nonetheless scientists managed to pry loose a few bricks in the great wall that separates P from NP. Most notably, the seminal work "Where the *Really Hard Problems Are*" by Cheeseman, Kanefsky and Taylor (henceforth abbreviated to 'Cetal'), showed that although runtime increases non-polynomially for problems in NP, some *instances* of these hard problems might actually be easy to solve ~ [2]. Not every formula in SAT is hard - easily satisfiable formulas exist too, even with many variables, but the hard ones keep the problem as a whole in NP. But Cetal's great contribution was not only to expose the huge differences in instance hardness within a single NP-problem, they also showed *where* those really hard instances are - and how to get there. Their findings were followed up numerous times and truly exposed some of the intricate inner anatomy of instance hardness, and problem class hardness as a whole.

So where *are* these hard problem instances then? According to Cetal, they are hiding in the phase transition. For the problems in their study, instances suddenly jump from 'having many solutions' to 'having no solutions' when their constrainedness changes. For an example in satisfiability, most randomly generated SAT-formulas of two clauses and four variables such as our formula $(a \vee \neg b \vee d) \wedge (b \vee c \vee \neg d)$ are easily satisfiable; they have many assignments that make them true. But as soon as the order parameter, the ratio of clauses versus variables α , passes ≈ 4.26 , (almost) no satisfiable formulas exist [13]~ [14]. So if we randomly generate a formula with 20 or more clauses on these same four variables, it is almost certainly unsatisfiable and those rare formulas that *are* satisfiable beyond the phase transition have very few solutions - which counterintuitively enough makes them easy again. So, for most exact algorithms, both extremes are quickly decided: for the highly satisfiable formulas in $\alpha \ll 4.26$, a solution is quickly found, and unsatisfiable formulas in $\alpha \gg 4.26$ are quickly proven as such. But in between, just around $\alpha = 4.26$, where the transition from satisfiable to unsatisfiable takes

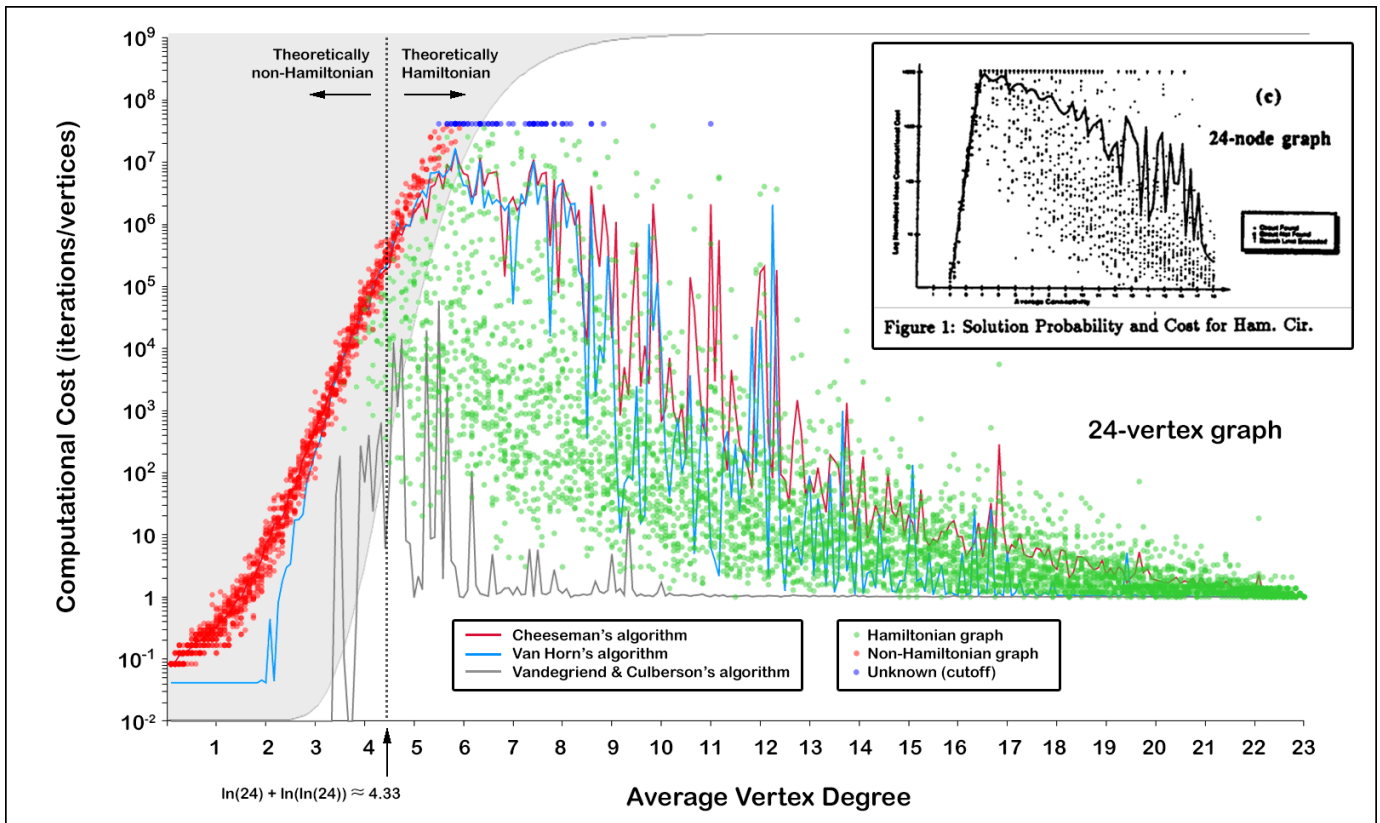


Fig. 3. Results from an earlier replication of Cetal's seminal work on Hamiltonian cycle hardness, extended with algorithms by Van Horn and Vacul. The top-right inset is Cetal's original figure, and it covers no data points. Note how the computational cost is highest along the Komlós-Szemerédi bound.

place, are formulas that take the longest to decide upon. This is where the really hard problem instances are: hiding in the phase transition that separates the solvable from the unsolvable region.

Cetal identify this order parameter not only for SAT; the Hamiltonian cycle problem has one too, and so does k -colorability. For the Hamiltonian cycle problem though, one should understand that the phase transition is somewhat inverse to SAT: as one adds edges to a graph, it becomes *solvable* rather than unsolvable. But the general principle still holds: both extremes are easy, and the phase transition is where the really hard problem instances are. And although Cetal's seminal results are relatively coarse, they have been followed up in more detail, and they are solid ~ [14]–[17]. To put it in a quote by Ian Gent and Toby Walsh: "[Indeed, we have yet to find an NP-complete problem that *lacks* a phase transition]" ~ [18].

The ubiquity of phase transitions throughout the class is not a complete surprise. Satisfiability, k -colorability and the Hamiltonian cycle problem are *NP-complete* problems; a subset of problems in NP that with more or less effort can be transformed into each other ~ [19]. This means a lot. This means that if someone finds a polynomial exact algorithm for just one of these problems, all of them become easy and the whole hardness class will simply evaporate. That person would also be an instant millionaire thanks to the Clay

Mathematics Institute that listed the $P \stackrel{?}{=} NP$ -question as one of their 'Millennium Problems' ~ [20]. But the intricate relations inside NP-completeness might also stretch into the properties of phase transitions and instance hardness. Or, to pour it into another fluid expression by Ian Gent and Toby Walsh "[Although any NP-complete problem can be transformed into any other NP-complete problem, this mapping does *not* map the problem space uniformly]" ~ [18]. So, a phase transition in say, satisfiability, does *not* guarantee the existence of a phase transition in Hamiltonian Cycle or in Vertex Coloring. The fact is though, that Cetal do find them for all three, and their results are solid.

III. THE HAMILTONIAN PHASE TRANSITION

The Hamiltonian cycle problem comes in many different shapes and forms, but in its most elementary formulation involves finding a path (a sequence of distinct edges) in an undirected and unweighted graph that visits every vertex exactly once, and forms a closed loop. The probability of a random graph being Hamiltonian (i.e., having a Hamiltonian Cycle), has been thoroughly studied ~ [21]–[23]. In the limit, it is a smooth function of vertex degree and therefore the probability for a random graph of v Vertices and e edges being

Hamiltonian can be calculated analytically⁴:

$$P_{Hamiltonian}(v, e) = e^{-e^{-2c}} \quad (1)$$

in which

$$e = \frac{1}{2}v \cdot \ln(v) + \frac{1}{2}v \cdot \ln(\ln(v)) + c \cdot v \quad (2)$$

Like the phase transition around α in SAT, the Hamiltonian phase transition is also sigmoidally shaped⁵ across a ‘threshold point’, the average degree of $\ln(v) + \ln(\ln(v))$ for a graph of v vertices (Figure 1). The phase transition gets ever steeper for larger graphs and becomes instantaneous at the threshold point as v goes to infinity. For this (theoretical) reason, the probability of being Hamiltonian at the threshold point is somewhat below 0.5 at $e^{-1} \approx 0.368$.

The probability of being Hamiltonian is one thing, deciding whether a given graph actually *has* a Hamiltonian cycle is quite another. A great number of exact algorithms have been developed through the years, the earliest being exhaustive methods that could run in $O(v!)$ time [24]. A dynamic programming approach, quite advanced for the time, running in $O(n^{22^n})$ was built by Michael Held & Richard Karp, and by Richard Bellman independently [25] [26]. Some early edge pruning efforts and check routines can be found in the work of Silvano Martello and Frank Rubin whose algorithms still run in $O(v!)$ but are in practice much faster (as we will show in Figures 4 and 5) [27] [28]. Many of their techniques eventually ended up in the algorithm by Vandegriend & Culberson [29]. Algorithms by Bollobás and Björklund run faster than Bellman–Held–Karp, but are technically speaking not exact for finite graphs [30] [31]. The 2007 algorithm by Iwama & Nakashima [32] runs in $O(2^{1.251v})$ time on cubic graphs, thereby improving Eppstein’s 2003 algorithm that runs in $O(2^{1.260v})$. While these kind of marginal improvements on specific instances are typical for the progress in the field, these two actually deserve some extra attention.

The cubic graph, in which every vertex has a maximum degree of three, is of special importance in the generation of 3D computer images. Many such images are built up from triangle meshes, and as specialized hardware render and shade triangles at low latencies, the performance bottleneck is actually in feeding the triangular structure into the hardware. A significant speedup can be achieved by not feeding every triangle by itself, but by combining them into triangle strips. An adjacent triangle can be defined by only one new point from the previously fed triangle, and therefore adjacent triangles combined in a single strip can speedup the feeding procedure by a maximum factor three for each 3D object. Finding a single strip that incorporates all triangles in the mesh is equivalent to finding a Hamiltonian cycle through the corresponding cubic graph in which every triangle is a

vertex, which makes both Eppstein’s and Iwama&Nakashima’s result of crucial importance for the 3D imagery business (see Figure 2).

So concludingly, none of the exact algorithms for the Hamiltonian cycle problem runs faster than exponential on all instances (with vertices of any degree), and the Bellman–Held–Karp “is still the strongest known”, as Andreas Björklund states on the first page of his 2010-paper [31]. Summarizingly, all backtracking algorithms still run in $O(v!)$ time. In the next sections, we will have a closer look at the algorithmic details of all *exact* algorithms from previous paragraphs. We will first make a structural and historical comparison, and after that bang out a lot of data to make a rigorous quantitative comparison as well. Finally, we will draw some conclusions, discuss the impact of the work, and project a trajectory for future research.

IV. ALGORITHMICS

For this extended investigation, we generated large numbers of problem random instances for the Hamiltonian cycle problem, varying in numbers of vertices and edges. We then solved these using almost every authoritative exact algorithm we could find: **Depth-first search** invented far before any modern day computer in 1882, **Rubin’s algorithm** from 1974, **Martello’s algorithm ’595’** published in 1983, **Cetal’s algorithm**, used in their 1991 seminal work on instance hardness, Vandegriend & Culberson’s, abbreviated to **Vacul’s algorithm**, an elaborate machinery published in 1998 and finally **Van Horn’s algorithm**, directly derived from Cetal’s, published in 2018 [33] [28] [27] [2] [29] [1].

There are at least several more, some of which are very old and without Google Scholar index or available pdfs. One prominent candidate missing in our list is the dynamic programming implementation built by Michael Held & Richard Karp, and by Richard Bellman independently [25] [26]. While the algorithm is complete, and with a time complexity of $O(v^22^v)$ has a better worst case performance than our six algorithms which all have $O(v!)$, it has significant memory requirements. Furthermore, it is not a backtracker and therefore does not perform ‘recursions’ as such, making a direct comparison to the other six slightly more difficult. It might be a viable candidate for a future comparison but for now, we will stick to backtracking algorithms, and canalize towards a generalized approach.

As it turns out, these six algorithms have several similarities and differences. But even though the respective papers have significant numbers of citations, some of the (especially earlier) authors of these particular algorithms seemed to be unaware of each others’ progress. The main exception is Vacul, who include a reference to both Martello and Cetal, but missed Frank Rubin’s work, which might simply be due to geographical dispersity and the lack of internet. Or conversely: the relatively recent development of computational resources, and proliferation of global communication enables us *now* to make a large structural comparison between them relatively easily. In

⁴An unfortunate convention: please note that the e in the right hand side of Eq. 1 is the base of the natural logarithm, whereas in the left hand side, e is the number of edges

⁵Remember this transition is *inverted* with respect to SAT: it goes ‘from no to yes’ when getting denser, where SAT goes ‘from yes to no’.

any case, there appear to be some globally emerging algorithmic design patterns for this problem, which are dispersedly found across algorithms. We will structurally compare all, and assess their effectiveness regarding the Hamiltonian cycle problem.

In this study we generalize the approach, unifying similar procedural subroutines to the same piece of source code. The only feature that was removed is the random restart option from Vacul's algorithm. Surely a good way of shortening the average runtime on many NP-complete problems [34], it also makes one-off comparisons a lot harder on large randomized instance ensembles such as ours. As most algorithms predate the internet, none came with readily implementable source code. We like to emphasize our belief that text such as you are reading now is an inferior medium for communicating algorithmics. As the smallest of details can make the largest differences when it comes to runtimes. Therefore, one should *always supply publicly accessible source code* when writing about algorithms. We will try to set the example by supplying ours [35]. It is quite possible that algorithmic details as we chose them are different from the original authors' implementations. It is also thinkable that even where we are precise, further improvements are possible. In any case, let us go forward with public source code along our publications on algorithms. It is better than text.

In the next section, we will describe the six backtrack algorithms used, but also the subroutines, many of which are common to multiple algorithms. The subjects are somewhat 'interwoven' between the algorithmic explanation, simply because it seems to make the most sense storywise. Operationally speaking, it makes more sense to compartmentalize these subroutines, so that is how the reader will find them in the source code [35]. For a general overview, please refer to Table I.

V. THREE BASIC ALGORITHMS

Depth-first search can in every way be considered as the basis for all algorithms in this study. It is surely the oldest, gaining widespread popularity from Tarjan's paper [36], but was designed roughly a century earlier by Frenchman Charles Pierre Trémaux, and mentioned in a publication from 1882 as a strategy for solving mazes [37]. As (planar) mazes are in many ways equivalent to (planar) graphs, its success as an algorithm for graph traversal is hardly surprising.

In modern-day computers, depth-first search is readily implemented either by a recursive function or via a stack data structure, the latter of which is usually a constant factor slower. Depth-first search is an exact algorithm in optimization problems, meaning it will always return the best possible answer (e.g. the shortest route, or the optimal timetable). For decision problems like the Hamiltonian cycle problem or satisfiability, it will always return a solution if the problem instance has one, or ensure it does not exist. For finding a Hamiltonian cycle (or any other particular vertex order) in a graph, its ominous runtime complexity is $O(v!)$ in the number

of vertices v , which makes it practically unusable for any number of v over two digits. The algorithm can just as well be deployed for solving SAT-formulas, running in slightly less daunting exponential complexity of $O(2^n)$ in n , the number of Boolean variables.

Depth-first search is a constructive algorithm that starts at the first vertex in the data structure that holds the randomly generated graph, thereby not having any specified preference for degree. From there, its recursive step is to add the first adjacent vertex (again, in the order of the data structure) that is not already in the path. This step is repeated either until no adjacent vertices are available. This can mean two things: either all vertices are in the path, and a final check for closure confirms the existence of a Hamiltonian Cycle after which the algorithm halts, or the algorithm backtracks, removing the last vertex from the path and adding the next adjacent vertex in its place. If none such vertex exists it backtracks again. If this happens at the root level then all possibilities are exhausted, and the algorithm halts and returns "no Hamiltonian Cycle". As a typical property of (uninformed) exhaustive exact algorithms, it thereby tries all permutations of vertices if necessary and has a time complexity, or worst-case runtime, of $O(v!)$.

A big *practical* difference however, materializes by not just adding the next vertex from the data structure, but preferring either vertices of high degree or low degree. This can be done runtime, or by sorting the data structure upon reading the graph ahead of recursing. In our implementation, we always chose the former option, resorting during the run. As this next-vertex-preference is a somewhat rule-of-the-thumb, we will call it the **branching heuristic** and can be instantiated with categorical⁶ parameter values $\{none, high, low\}$ (see Table I). In its simplest implementation, it involves only changing a '>' to a '<' in the source code of the algorithm. Nonetheless the impact on the algorithm's performance for large ensembles of instances such as ours can be enormous [1].

Cetal's algorithm follows the exact same paradigmatic lines as plain depth-first search, but prefers higher degree vertices when branching. Ties are unaddressed, and whether a more sophisticated order of preference has any significant impact on the algorithm's runtime remains an open question, especially for larger graphs. Cetal's algorithm still runs in $O(v!)$, but its somewhat more fine-grained time complexity might be $O(v \cdot \log(v))!$, with the additional term accounting for sorting the vertices to descending degree numbers. Cetal's paper was published with four experiments, among which the Hamiltonian cycle problem, and like this study, contained experimental results on an ensemble of random graphs in regular degree intervals.

Van Horn's algorithm is in many ways the opposite of Cetal's algorithm, starting at the vertex with the lowest degree and preferring lower degree vertices over higher degree vertices when recursing. It therefore also runs in $O(v!)$ time,

⁶Sometimes referred to as 'symbolic parameter' or 'qualitative parameter'.

but its fine-grained time complexity is $v \cdot \log(v) + v!$. Its only algorithmic difference is that it traverses the data structure backwardly, or alternatively, that its data structure is sorted in inverse order. Van Horn’s algorithm was introduced during an extended replication of Cetal’s Hamiltonian study, and thereby contained experimental results on random graphs of various average degree. The ensemble was likely larger than Cetal’s⁷, and made publicly available, along with the source code of the algorithms used⁸.

VI. SEARCH PRUNING

A. Search Pruning: Edge Pruning

One of the most prolific enhancements of depth-first search on NP-complete problems is **search pruning**: cutting off branches from the search tree that cannot hold a solution. If a contradictive pair of clauses is found from some variable assignment in CNF3SAT, the search backtracks immediately, departing from further assignments in the subtree. If a perfect rectangle packing problem contains an unfillable gap, it cannot be solved regardless of the rest of the configuration, so depth-first will halt and backtrack [38]. As such, search pruning can be a valuable way of saving recursions, but the (sometimes huge) drawback is that the pruning procedure *itself* takes computational resources too. Even though most pruning procedures are subexponential routines, they can theoretically be invoked in every recursion. The art therefore, is to balance an investment in pruning so that it pays enough time dividend. In many cases, this is well possible, or, abiding by Steven Skiena’s famous words: “Clever pruning can make short work of surprisingly hard problems” [39].

For the Hamiltonian cycle problem, search pruning comes in two forms. The first, **edge pruning**, involves removing edges from the graph, which can be done either in the preprocessing stage or during a recursive step. By removing edges that *cannot possibly be in any Hamiltonian cycles*, the nodes in the recursive search tree get a lower degree, leading to fewer recursions. Interestingly enough, the prunative removal of edges is always related to the existence of *required edges* that are adjacent to a vertex with degree two and therefore *must* be in any Hamiltonian cycle that might exist in the graph. It can

⁷Cetal’s ensemble size is unknown, but an estimate is given by Van Horn et al.

⁸We provide no reference to their source code to avoid confusion. Our implementation is in many ways a generalization of theirs. Of course, they do own the credit for open sourcing their implementation

be slightly confusing to keep search pruning and edge pruning apart, but it can be remembered like this: edge pruning means cutting away edges, search pruning involves speeding up the search process. Edge pruning therefore, partially instantiates search pruning.

From literature, we find exactly three edge pruning methods which are all implemented in this study. The first, **neighbour pruning**, finds a vertex with two required edges, and removes all others. Second, the **path pruning** method looks for paths of required edges, which might eventually become a Hamiltonian cycle, and removes edges that would prematurely close it. Third, **solution pruning** removes all edges from the second-last vertex of the partial Hamiltonian path so far. Since path pruning and neighbour pruning might reciprocally facilitate each other’s operation, both are repeated until no further pruning occurs. Finally, edges that are removed during preprocessing before recursing are definitely gone, but edges pruned in the recursive step *have* to be put back when the algorithm backtracks.

B. Search Pruning: Non-Hamiltonicity Checks

The second category of search pruning does not involve the removal of edges, but checking whether a Hamiltonian cycle is achievable at all. There are four such *non-Hamiltonicity checks* and it should be noted that indeed all of them are negative: they can only decide graphs to be either ‘surely non-Hamiltonian’ or ‘undecided’. Although checks that give ‘surely Hamiltonian’ or ‘undecided’ are certainly imaginable, they are not described in the literature regarding the six general backtrack algorithms used in this study.

The first is a plain and simple **degree check** which verifies whether any vertex in the graph has edge degree one or zero. If this is the case, the graph cannot be Hamiltonian, and the algorithm needs to backtrack. The second graph configuration that cannot contain a Hamilton Cycle, is a graph where required edges form a closed loop smaller than v . This graph configuration can be checked by iterating over all required edges and thus required $O(v)$ computational effort. We call this check the **premature closure** check. The third check, the **disconnectedness check** determines whether a graph is broken up into two or more pieces. The subroutine involves picking a vertex and adding it to an empty list of found vertices. From that list, it picks the next item and adds all its adjacent vertices to end of the list, unless they are already added. When the subroutine reaches the end of the list, it counts the number

Algorithm	Pruning	Heuristic	Non-Hamiltonicity Checks
Depth-First	None	None	None
Cetal’s	None	High	None
Van Horn’s	None	Low	None
Martello’s	Solution, Path	Low	Degree
Rubin’s	Solution, Path, Neighbour	None	Degree, One-Connectedness, Disconnectedness, Premature Closure
Vacul’s	Solution, Path, Neighbour	Low	Degree, One-Connectedness, Disconnectedness

TABLE I. Overview of all the techniques used by the algorithms that are examined in this research. For the pruning methods, it is displayed if an algorithm uses that method. For the branching heuristic it is presented if an algorithm uses one and if so which one.

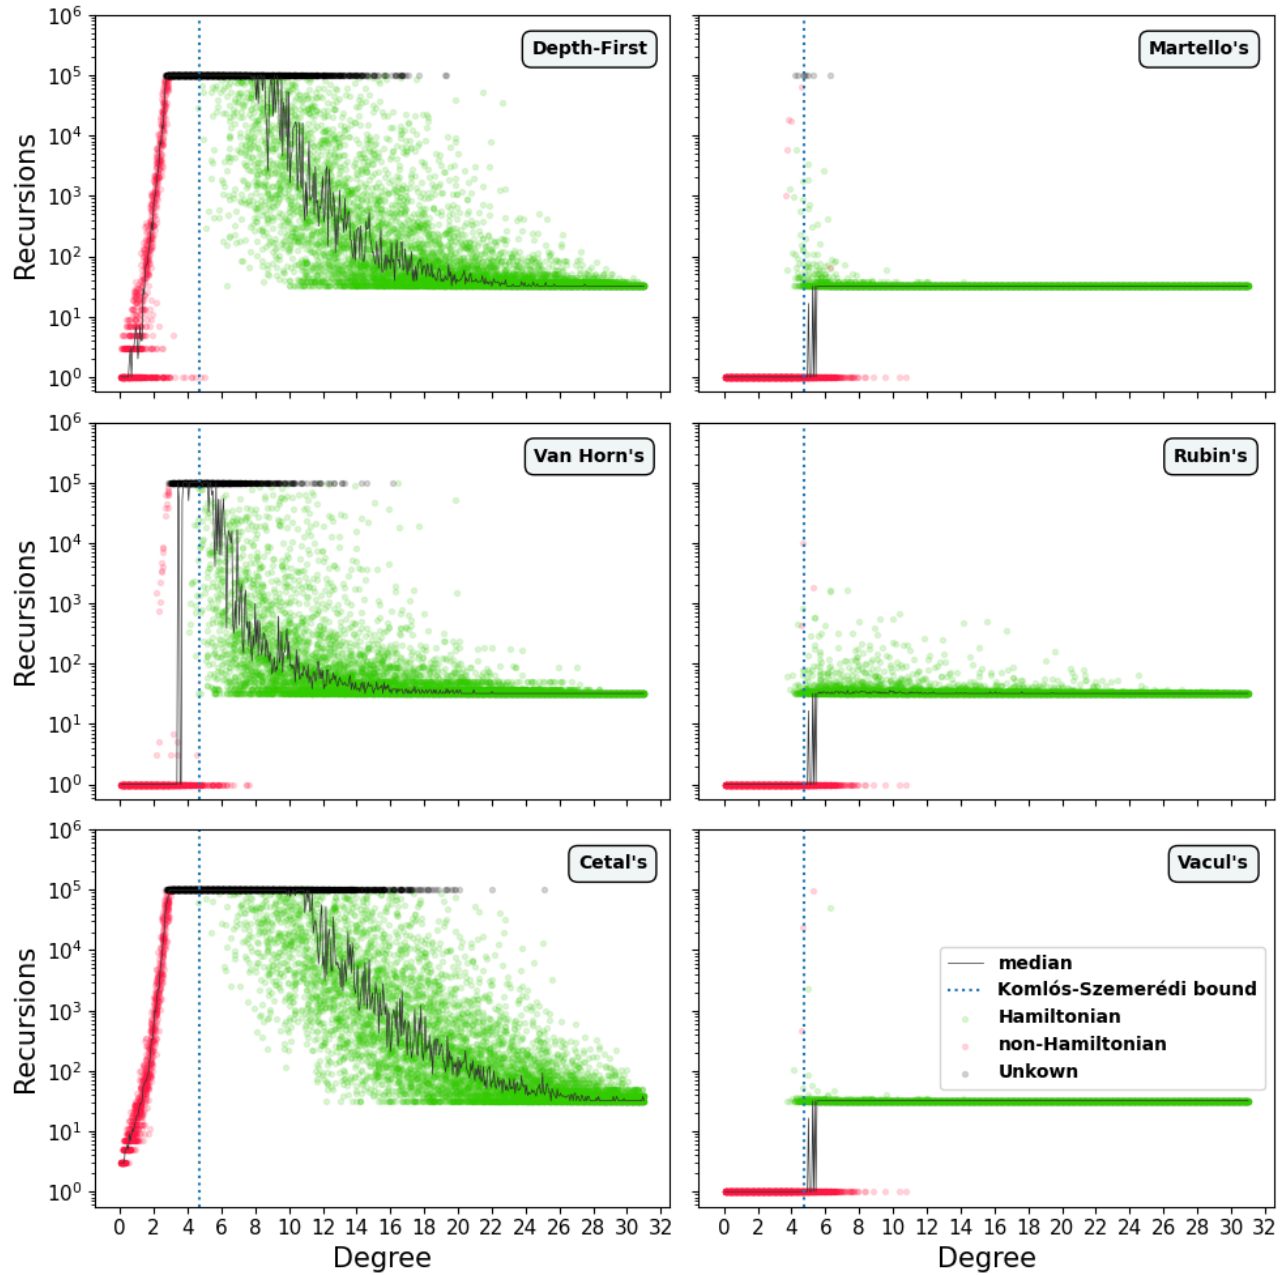


Fig. 4. The number of recursions required for solving the 9920 random graphs of $v = 32$. For all algorithms, the hardest graphs are situated close to the Komlós-Szemerédi bound of $\ln(32) + \ln(\ln(32)) \approx 4.71$ where the probability of being Hamiltonian transitions from zero to one. But while the choice of branching heuristic clearly makes a difference (Depth-first, Van Horn's, Cetal's) the improvement is most dramatic when the algorithm also implements edge pruning and non-Hamiltonicity check procedures (Martello's, Vacul's, Rubin's).

of items in the list. Only if it is equal to v , the graph is connected. Note that in some cases, graphs filtered out by this subroutine can also be filtered out by the degree check; it could be considered a hint that the order in which search pruning subroutines are applied also matters for the computation time. The fourth and final check is the **one-connectedness check**. A graph is one-connected if it contains a vertex that if removed, breaks up the graph into two or more disconnected parts. Such a vertex is commonly dubbed an *articulation point*, and

Tarjan's algorithm finds all articulation points in $O(|v| + |e|)$ time [36]. To construct a Hamilton Cycle, there need to be at least two edge-disjoint paths between any two non-adjacent vertices. Therefore, there cannot be a Hamilton cycle in a 1-connected graph. Rubin's is the only algorithm that deploys this technique, we do not know how, but we assumed it could have been done with Tarjan's algorithm, which was published two years earlier. In any case, that is how we implemented the check, as literature gave us no definitive answer.

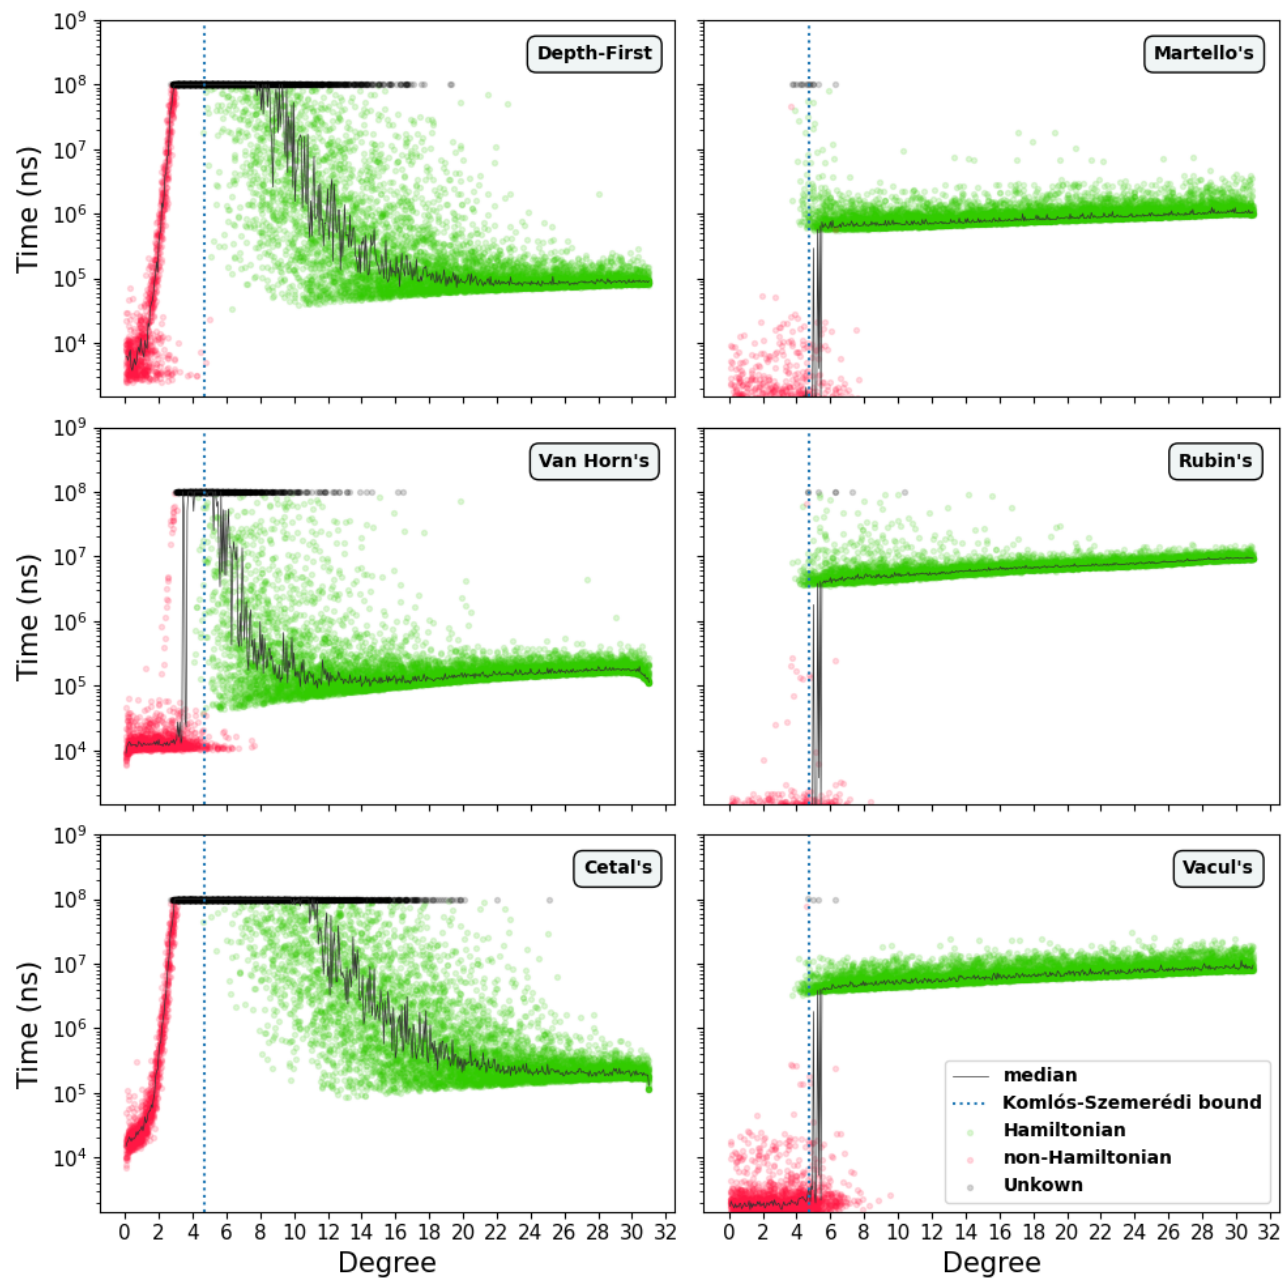


Fig. 5. Procedures for pruning and non-Hamiltonicity checks (Martello's, Vacul's, Rubin's) clearly take 'real' time, but even on graphs as small as these, they pay off – the right hand side triplets have close to zero cutoffs. In these algorithms, the dramatic difference in time consumption between the red and green dots around the Komlós-Szemerédi bound show that these procedures save time especially on *non-Hamiltonian* graphs.

VII. THREE ADVANCED ALGORITHMS

Martello's algorithm '595' is a rare example of an algorithm from the 80's that actually came with written source code in FORTRAN when published. The practice of supplying 'open source' code along scientific experiments is a practice that is much valued today, but was definitely ahead of time in 1983. Martello's algorithm was designed with the purpose of finding *all* Hamiltonian cycles in *directed* graphs. We made the smallest possible changes, adapting the algorithm to undirected graphs, and made it halt after the first solution. Martello's

algorithm has both solution pruning and path pruning, the latter of which is repeatedly applied in each recursion. Furthermore, it uses the *low* degree preference in its branching heuristic, preferring sparselier connected vertices over densier connected vertices.

Like Silvano Martello, Frank Rubin was an influential researcher on early algorithms for NP-complete problems and like Silvano Martello, he designed an algorithm for the *directed* Hamiltonian cycle problem, which we adapted to undirected graphs. **Rubin's algorithm** was quite sophisticated

for the time (1974) and in retrospect would have outperformed most other algorithms. It deploys solution, path and neighbour pruning exhaustively in each recursion, and additionally performs all four checks for non-Hamiltonicity.

The premature closure check in this algorithm is worth giving some thought, as it might be redundant in combination with the check for disconnectness. We think that if a set of required edges forms a closed cycle smaller than v , the graph is automatically disconnected. Furthermore, this algorithm is the only one that performs a check for one-connectedness. In the original paper it is not specified how this check is done, but as stated earlier, it is well possible that this was Tarjan's algorithm and in any case, we implemented it as such. Finally, Rubin does not fully specify his disconnectedness check, but does say it runs in quadratic time. So does ours, and chances are they are nearly identical.

These combined features make it a very advanced algorithm for the time. In fact, no other algorithm in our ensemble deploys so many different subroutines. The only thing it does not employ however, is a branching heuristic. It assumedly just picks the next vertex from the data structure. We think that if Frank Rubin would have just implemented the low-degree branching heuristic too, it would have been the most efficient backtracking algorithm to date. And given that it was contrived nearly half a century ago, these minute details could have profoundly changed the course of history for this problem.

The final algorithm included in this research is the algorithm by Vandegriend and Culberson, abbreviated to **Vacul's algorithm** [29]. It was designed for the undirected Hamiltonian cycle problem and uses solution pruning, path pruning, and neighbor pruning, the last two of which are run exhaustively each recursion. It also uses the degree check, the disconnectness check and the one-connectedness check, 24 after Frank Rubin introduced them in his algorithm. Since Rubin's study is not referenced by VaCul, it is possible they invented these routines themselves. We can only guess the reasons, but the lack of widely accessible internet papers at the time is at least

one likely culprit.

It is sad to see that history missed so many beats, but again, these were the days from before internet and the synchronisation of information; now is the time to make up for it. Important to note is that Vacul's algorithm has an additional feature, a *random restart*, which we left out. Even though stochastically speaking, random restarts on long backtrack runs can save computation time, it makes the algorithm a lot harder to compare to the other algorithms. Vacul's paper comes with a quantitative test on a set of random graphs of variable average degree.

VIII. EXPERIMENT AND RESULTS

Analogous to previous studies by Cetal, Vacul and Van Horn et al., three sets of randomly generated undirected graphs were created: one with $v = 16$ vertices, one with $v = 24$ vertices and one with $v = 32$ vertices (we will only show and discuss $v = 32$, results are comparable for other values of v). We created 20 graphs for every number of edges $e = \{1, 2, 3, \dots, \frac{1}{2}v(v-1)\}$, resulting in 2400 graphs to solve for $v = 16$. For $v = 24$, the procedure resulted in 5520 graphs, and 9920 random graphs were generated for the 32-vertex set. This amounts to a subtotal of 17,840 graphs, each of which has been solved by all six algorithms *twice* - once for recursions and once for system time. This means the whole investigation comprises 214,080 runs and therefore, to keep things a little insightful, we will show and discuss results of the 9920-piece ensemble for $v = 32$ only. It should be clearly understood though, that the results obtained from the different algorithms in both the time subexperiment and the recursion subexperiment, came from the *same* 17,840 source graphs, facilitating a direct comparison.

We first solve (i.e.: decide) all graphs for Hamiltonicity for all six algorithms implemented on a general code base which is publicly accessible [35]. Programming the generalized algorithm required some interpretation, as none of the algorithms was published with directly usable source code,

Algorithm	#cutoff	#finished	avg. recs	stddev. recs	total (x1000)
Depth-First	1959	7961	2992	11661	219,716
Van Horn's	814	9106	844	5890	8,908
Cetal's	2571	7349	4551	13980	29,054
Martello's	8	9912	40	686	1,193
Rubin's	0	9920	31	106	307
Vacul's	0	9920	44	1136	438

TABLE II. Recursions required by all six algorithms on the entire ensemble of random graphs. Note that 'average recursions' and 'stddev recursions' apply to the *finished* problem instances only.

Algorithm	#cutoff	#finished	avg. ns	stddev. ns	total (x 1M)
Depth-First	1916	8004	3054626	11341555	216,049
Van Horn's	813	9107	913606	5524056	89,620
Cetal's	2583	7337	4695537	13464862	292,751
Martello's	14	9906	864532	1502152	9,964
Rubin's	7	9913	6080907	4439895	60,980
Vacul's	4	9916	6125890	3718473	61,144

TABLE III. Runtime required by all six algorithms on the entire ensemble of random graphs. Average and stddev apply to the *finished* problem instances only.

but incorporates the three options for branching heuristic, the three edge pruning routines and the four non-Hamiltonicity checks, all in subroutines that can be toggled individually. As such, there is no difference between the branching heuristic in Depth-first and Rubin's, or the the degree-check in Martello's and Vacul's, even if historically speaking, they might not have been identical. The only clearly different feature is the random restart option, which is removed from Vacul's algorithm. There is a lot to say about this procedure, mostly that it *can* be very useful for decision problem instances that *do* have a solution. A somewhat smaller adaptation might be the disconnectedness check and the one-connectedness check, both from Rubin's algorithm. We suspect our quadratic implementation might be conceptually equal to Rubin's, which is also reported as a quadratic order, but eludes further specification. Finally, the sorting procedure for the branching heuristic can be done in various ways whose theoretical and practical runtimes may differ, though not to extend of becoming superpolynomial. We suspect that this routine too, is equal to all in literature, but we can not be sure. Summarizing, all algorithmic components are uniform, deterministic and adapted to undirected graphs.

We ran two experiments: one recording the number of recursions, and the other recording system time. Traditionally speaking, number of recursions is the way to go. Measuring recursions is immune to the choice of programming language, compiler used, processor speed, parallelization or incoming resource expenditures such as downloads or updates that might interfere with runtimes. Second, the number of recursions computationally speaking largely outweighs pruning and check procedures, which are usually of lower order complexities - typically polynomial versus factorial. Finally, and following previous two reasons, the order complexity of the algorithm also rest with the number of recursions. As, in a broader view, it is the order complexity of algorithms for NP-complete problems that hinge the $P \stackrel{?}{=} NP$ problem, any results

which might influence the view on this problem must also be expressed in recursions. However, as we wanted to extend beyond the reach of theoretical computer science and simply get an estimate of time it takes *in the real world* to mobilize the alleged 'optimization procedures', and immediately question the 'real' impact. A cutoff for the maximum number of a run was 10^5 recursions, and 10^8 nanoseconds⁹

In the recursion experiment, the cutoff was often reached for the three basic algorithms (Fig. 4), leaving 20% of the instances unsolved for depth-first, 8% for Van Horn's algorithm and a toe-curling 26% for Cetal's algorithm (for exact values, see Table II). Of the more advanced algorithms, Martello's was unable to solve 0.1% of the instances, but Rubin's and Vacul's successfully solved them all. The average number of recursions needed for the solved instances further confirms the hierarchy in the three basic algorithms, and the success of the search pruning efforts. For the advanced algorithms, the ranking slightly changes, with Martello's having the second lowest computational effort within the solved graphs. This changes nothing however for the hierarchy of algorithmic performance when measured in recursions: **1.Rubin's, 2.Vacul's, 3.Martello's, 4.Van Horn's, 5.Depth-first, 6.Cetal's.**

For the time experiment, the cutoff value was again reached often for the three basic algorithms (Fig. 5). These results were largely proportional to the recursions-experiment, leaving 19% of the instances unsolved for depth-first, 8% for Van Horn's algorithm and a painful 26% for Cetal's algorithm (for exact values, see Table III). The picture changes slightly for the more advanced algorithms however, where Martello's, Rubin's and Vacul's algorithms left 14, 7 and 4 graphs unsolved, still well below 0.2% for all three algorithms. Still, it should be noted that *relatively* speaking, this is a huge increase in failure (75%

⁹We could have chosen microseconds, milliseconds or even plain seconds here, but this unit holds the best balance between explanation and visualization such as seen in Fig. 5.

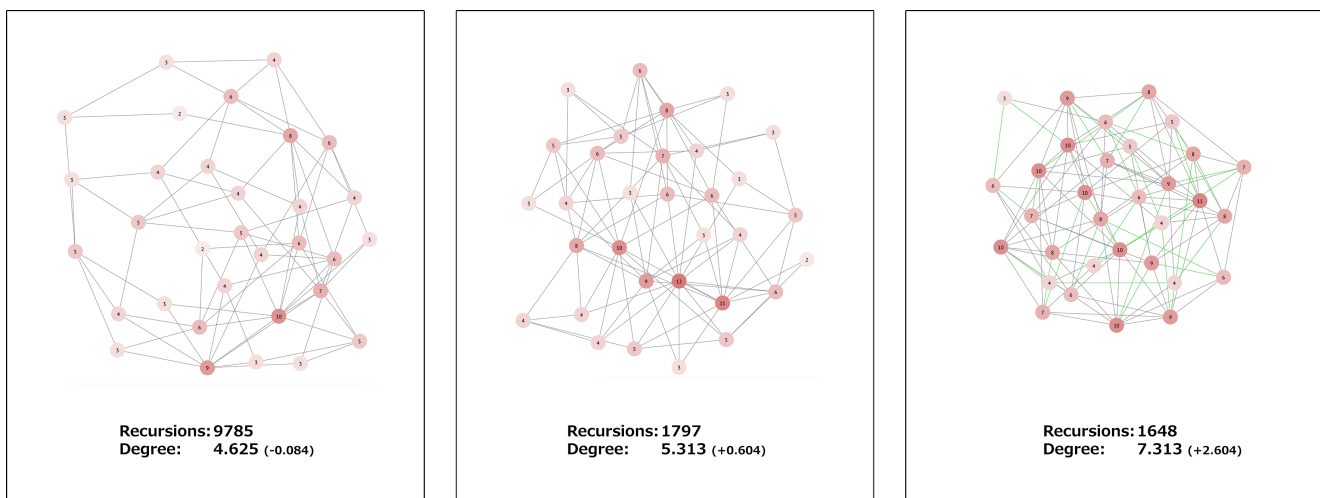


Fig. 6. The three hardest graphs for Rubin's algorithm, the best performing algorithm in this study, support the generalized conclusion that hard instances reside near the Komlós-Szemerédi bound of $\ln(32) + \ln(\ln(32)) \approx 4.71$ (the 'distance' in edge degree is in brackets). The suspicion lingers however, that this is only true for *randomly generated* instances ensembles, and targeted results find harder, denser graphs.

more failures for Martello's alone). It also swaps the top two contestors in the hierarchy of algorithmic performance relative to the recursion-experiment. The best performing algorithms when measured in system time are: **1.Vacul's**, **2.Rubin's**, **3.Martello's**, **4.Van Horn's**, **5.Depth-first**, **6.Cetal's**.

IX. CONCLUSION AND DISCUSSION

It is a fresh awakening that even for graphs of $v = 32$, which is relatively small, and the theoretical computational effort required relatively futile, the optimization procedures still make such a large difference. Generally speaking, these procedures appear to pay off in the number of recursions required to decide a graph, especially further away from the Komlós-Szemerédi bound. When measured in system time however, a different view emerges. The average time to decide a graph hugely increases, especially in denser regions away from the Komlós-Szemerédi bound.

The big question is how these effects scale up, and whether the increase in search pruning effort really pays off on a larger scale, and whether that counts for recursions, for system time, or for both. Furthermore, the distribution of edge pruning techniques, branching heuristics and non-Hamiltonicity checks among the algorithms is largely haphazard. This is because for now, we tried to stick as closely as possible to historical conventions handed down to us from literature, but a more structured approach might be desirable. Furthermore, it is quite surprising to see the Rubin's algorithm, which is the second oldest algorithm, overaging Van Horn's by more than four decades, contained the most sophisticated procedures, but also performed best. On the other side, it is a little desillusioning to find Cetal's algorithm, so often cited, finishing last in this test. Finally, some thought must be given to whether other edge pruning techniques, branching heuristics or checks could be devised, and in what order such procedures should be applied. Last but not least, it *might* still pay off to dynamically rearrange the data structure in some way as to make the degree preference in branching dynamically applicable. It should be noted though that this might be programmatically tough, as backtracking would require un-sorting the dynamically rearranged data structure.

For this study, we made 20 graphs for every possible edge density. Although this procedure is rigorously systematic, it also allows for some skewness in the results: in the end, there is structurally speaking, just one possible graph with one edge. On the other end of the density spectrum, we have 20 'random' graphs with with 496 edges, all being fully connected, and thereby isomorphic. As the number of different graphs for v vertices and e edges is equal to $\binom{v}{e}$, the maximum number of different graphs can be found at 248 edges for 32 vertices, which is exactly halfway the Figures 4 and 5. So if every existible graph of for $v = 32$ was equally likely, most of the graphs would a) be Hamiltonian and b) be easy, at least for the best three of our algorithms. For larger graphs, this effect grows stronger as the Komlós-Szemerédi bound grows (double) logarithmically in v , whereas

the peak existence grows (half) quadratically. A really weird but inescapable conclusion therefore is, that for larger values of v , nearly all instances are Hamiltonian, and many of those might be easy. This does not hold for our results, or that of Cetal, Vacul and Van Horn et al., who all 'columnized' their graphs into different degree categories. But combined, the combinatorial assessment and the columnized results show that the prediction of instance hardness for the Hamiltonian cycle problem critically relies on the a priori availability of information of the graph(s) to be solved, even if one future algorithm turns out to be superior.

And this brings us to the last point, tying the discussion back to the preamble. It has been shown that for at least one algorithm, Vacul's in this case, the hardest instances are non-Hamiltonian, very far away from the Komlós-Szemerédi bound, in a very dense region of the combinatorial space [8] [7] (also see [52]). This observation apparently contradicts nearly everything that was written in the previous paragraph, but can be explained from the randomness in large ensembles of problem instances such as used in this study. The found extremely hard instances were *structured*, and abiding by the teachings of A. N. Kolmogorov, structured objects in large randomized ensembles are rare [40]. So finding these instances paradoxically means knowing where to look, and one way to do that is to use a parameter-unsensitive evolutionary algorithm such as Plant Propagation [41] [42] [43] [44] [45] [50] [51]. For very large graphs, one could better resort to a single-individual search heuristic, such as HillClimbing or simulated annealing [46] [47] [48]. On the other end of the scale, the use of large randomized ensembles for algorithmic performance raises some questions. Although proper benchmarking is becoming a hot topic [49], the issue of randomness herein is still seldomly discussed and needs more attention.

X. FUTURE WORK

After the results of this study, the road ahead took a few sudden turns. What needs to happen next, is to see if one further generalized backtracking algorithm can convincingly dominate all others. Such a study should include combinations of heuristics and search pruning that are not tested as yet, and it is our belief that such an algorithm exists. After that, its hardest instances should be found, possibly by a global evolutionary search algorithm. Considering earlier results, it could resemble the graph found by Slegers & Van den Berg in 2020, as its structure suggests it might be hard for backtracking algorithms in general. Only after that, hardness hierarchies, and possible implications for computability, and the $P \stackrel{?}{=} NP$ problem itself might be found, but progress in this direction will depend on future developments.

XI. ACKNOWLEDGEMENTS

Thanks to Gijs van Horn for maintaining the nice interactive Hamiltonian cycle problem up to date (<https://hamiltoncycle.gijsvanhorn.nl/>).

REFERENCES

- [1] G. van Horn, R. Olij, J. Slegers, and D. van den Berg, "A predictive data analytic for the hardness of hamiltonian cycle problem instances," *DATA ANALYTICS 2018*, p. 101, 2018.
- [2] P. C. Cheeseman, B. Kanefsky, and W. M. Taylor, "Where the really hard problems are," in *IJCAI*, vol. 91, 1991, pp. 331–340.
- [3] J. Slegers, R. Olij, G. van Horn, and D. van den Berg, "Where the really hard problems aren't," *Operations Research Perspectives*, vol. 7, p. 100160, 2020.
- [4] D. van den Berg, "Video showing where cetal's results on atsp are flawed (part1)," <https://www.youtube.com/watch?v=zfqmiLkbCRc>, 2020.
- [5] —, "Video showing where cetal's results on atsp are flawed (part2)," https://www.youtube.com/watch?v=I8W_GVzqDnk, 2020.
- [6] G. van Horn, "Interactively viewable hamiltonian cycle hardness," <https://hamiltoncycle.gijsvanhorn.nl/>, 2020.
- [7] J. Slegers and D. van den Berg, "Plant propagation & hard hamiltonian graphs," *Evo** 2020, p. 10, 2020.
- [8] —, "Looking for the hardest hamiltonian cycle problem instances," in *IJCCI*, 2020, p. 40–48.
- [9] A. Huxley, *Brave new world*. Ernst Klett Sprachen, 2007.
- [10] R. E. Bryant, "On the complexity of vlsi implementations and graph representations of boolean functions with application to integer multiplication," *IEEE transactions on Computers*, vol. 40, no. 2, pp. 205–213, 1991.
- [11] F. Ivančić, Z. Yang, M. K. Ganai, A. Gupta, and P. Ashar, "Efficient sat-based bounded model checking for software verification," *Theoretical Computer Science*, vol. 404, no. 3, pp. 256–274, 2008.
- [12] M. Krötzsch, "Complexity theory, lecture 6: Nondeterministic polynomial time," <https://iccl.inf.tu-dresden.de/w/images/0/08/CT2019-Lecture-06-overlay.pdf>, 2019.
- [13] T. Larrabee and Y. Tsuji, *Evidence for a satisfiability threshold for random 3CNF formulas*. Citeseer, 1992.
- [14] S. Kirkpatrick and B. Selman, "Critical behavior in the satisfiability of random boolean expressions," *Science*, vol. 264, no. 5163, pp. 1297–1301, 1994.
- [15] I. P. Gent and T. Walsh, "Easy problems are sometimes hard," *Artificial Intelligence*, vol. 70, no. 1-2, pp. 335–345, 1994.
- [16] T. Hogg and C. P. Williams, "The hardest constraint problems: A double phase transition," *Artificial Intelligence*, vol. 69, no. 1-2, pp. 359–377, 1994.
- [17] T. Hogg, "Refining the phase transition in combinatorial search," *Artificial Intelligence*, vol. 81, no. 1-2, pp. 127–154, 1996.
- [18] I. P. Gent and T. Walsh, "The tsp phase transition," *Artificial Intelligence*, vol. 88, no. 1-2, pp. 349–358, 1996.
- [19] M. R. Garey and D. S. Johnson, *Computers and intractability*. wh freeman New York, 2002, vol. 29.
- [20] A. M. Jaffe, "The millennium grand challenge in mathematics," *Notices of the AMS*, vol. 53, no. 6, pp. 652–660, 2006.
- [21] P. Erdos and A. Rényi, "On the evolution of random graphs," *Publ. Math. Inst. Hung. Acad. Sci.*, vol. 5, no. 1, pp. 17–60, 1960.
- [22] L. Pósa, "Hamiltonian circuits in random graphs," *Discrete Mathematics*, vol. 14, no. 4, pp. 359–364, 1976.
- [23] J. Komlós and E. Szemerédi, "Limit distribution for the existence of Hamiltonian cycles in a random graph," *Discrete Mathematics*, vol. 43, no. 1, pp. 55–63, 1983.
- [24] S. Roberts and B. Flores, "Systematic generation of Hamiltonian circuits," *Communications of the ACM*, vol. 9, no. 9, pp. 690–694, 1966.
- [25] M. Held and R. M. Karp, "A dynamic programming approach to sequencing problems," *Journal of the Society for Industrial and Applied Mathematics*, vol. 10, no. 1, pp. 196–210, 1962.
- [26] R. Bellman, "Dynamic programming treatment of the travelling salesman problem," *Journal of the ACM (JACM)*, vol. 9, no. 1, pp. 61–63, 1962.
- [27] S. Martello, "Algorithm 595: An enumerative algorithm for finding Hamiltonian circuits in a directed graph," *ACM Transactions on Mathematical Software (TOMS)*, vol. 9, no. 1, pp. 131–138, 1983.
- [28] F. Rubin, "A search procedure for Hamilton paths and circuits," *Journal of the ACM (JACM)*, vol. 21, no. 4, pp. 576–580, 1974.
- [29] B. Vandegriend and J. Culberson, "The gn, m phase transition is not hard for the Hamiltonian cycle problem," *Journal of Artificial Intelligence Research*, vol. 9, pp. 219–245, 1998.
- [30] B. Bollobas, T. I. Fenner, and A. M. Frieze, "An algorithm for finding Hamilton paths and cycles in random graphs," *Combinatorica*, vol. 7, no. 4, pp. 327–341, 1987.
- [31] A. Bjorklund, "Determinant sums for undirected Hamiltonicity," *SIAM Journal on Computing*, vol. 43, no. 1, pp. 280–299, 2014.
- [32] K. Iwama and T. Nakashima, "An improved exact algorithm for cubic graph tsp," in *International Computing and Combinatorics Conference*. Springer, 2007, pp. 108–117.
- [33] S. Even, "Graph algorithms," 1979.
- [34] C. Gomes and B. Selman, "On the fine structure of large search spaces," in *Proceedings 11th International Conference on Tools with Artificial Intelligence*. IEEE, 1999, pp. 197–201.
- [35] J. Slegers, "Source code," <https://github.com/Joeri1324/What-s-Difficult-About-the-Hamilton-Cycle-Poblem->, 2020.
- [36] R. Tarjan, "Depth-first search and linear graph algorithms," *SIAM journal on computing*, vol. 1, no. 2, pp. 146–160, 1972.
- [37] E. Lucas, "Recreations mathematiques, four volumes: Gauthier-villars," Paris, France (1882/1894), pp. 161–197, 1882.
- [38] D. van den Berg, F. Braam, M. Moes, E. Suilen, S. Bhulai et al., "Almost squares in almost squares: solving the final instance," 2016.
- [39] S. S. Skiena, "The algorithm design manual," p. 247, 1998.
- [40] M. Li, P. Vitányi et al., *An introduction to Kolmogorov complexity and its applications*. Springer, 2008, vol. 3.
- [41] A. Salhi and E. S. Fraga, "Nature-inspired optimisation approaches and the new plant propagation algorithm," 2011.
- [42] M. De Jonge and D. van den Berg, "Parameter sensitivity patterns in the plant propagation algorithm," in *IJCCI*, 2020, p. 92–99.
- [43] M. de Jonge and D. van den Berg, "Plant propagation parameterization: Offspring & population size," *Evo** 2020, p. 19, 2020.
- [44] M. Paauw and D. Van den Berg, "Paintings, polygons and plant propagation," in *International Conference on Computational Intelligence in Music, Sound, Art and Design (Part of EvoStar)*. Springer, 2019, pp. 84–97.
- [45] W. Vrieling and D. van den Berg, "Fireworks algorithm versus plant propagation algorithm," 2019.
- [46] R. Geleijn, M. van der Meer, Q. van der Post, and D. van den Berg, "The plant propagation algorithm on timetables: First results," *EVO** 2019, p. 2, 2019.
- [47] R. Dahmani, S. Boogmans, A. Meijs, and D. van den Berg, "Paintings-from-polygons: simulated annealing," in *International Conference on Computational Creativity (ICCC'20)*, 2020.
- [48] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," *science*, vol. 220, no. 4598, pp. 671–680, 1983.
- [49] T. Bartz-Beielstein, C. Doerr, J. Bossek, S. Chandrasekaran, T. Eftimov, A. Fischbach, P. Kerschke, M. Lopez-Ibanez, K. M. Malan, J. H. Moore et al., "Benchmarking in optimization: Best practice and open issues," *arXiv e-prints*, pp. arXiv–2007, 2020.
- [50] W. Vrieling and D. van den Berg, "A Dynamic Parameter for the Plant Propagation Algorithm," *EVO** 2021, p. 5, 2021.
- [51] W. Vrieling and D. van den Berg, "Parameter control for the Plant Propagation Algorithm," *EVO** 2021, p. 1, 2021.
- [52] D. van den Berg and P. Adriaans, "Subset Sum and the Distribution of Information," in *IJCCI*, 2021, p. 134-140