

A Study of Tabu Search for Coloring Random 3-Colorable Graphs Around the Phase Transition

Jean-Philippe Hamiez*, Jin-Kao Hao

Université d'Angers / LERIA, 2 boulevard Lavoisier, 49045 Angers CEDEX 01, France
 {hamiez, hao}@info.univ-angers.fr

Fred W. Glover

OptTek Systems Inc., 2241 17th Street, Boulder, CO 80302, USA
 glover@opttek.com

ABSTRACT

We present an experimental investigation of tabu search (TS) to solve the 3-coloring problem (3-COL). Computational results reveal that a basic TS algorithm is able to find proper 3-colorings for random 3-colorable graphs with up to 11 000 vertices and beyond when instances follow the uniform or equipartite well-known models, and up to 1 500 vertices for the hardest class of flat graphs. This study also validates and reinforces some existing phase transition thresholds for 3-COL.

Keywords: 3-coloring, random graphs, phase transitions, tabu search

1. INTRODUCTION

Given a simple undirected graph $G = (V(G), E(G))$, where $V(G) = \{v_1, v_2, \dots, v_n\}$ is a set of n vertices (n is usually called the “order” of G) and $E(G) \subset V(G) \times V(G)$ a set of m edges, and a set $C = \{c_1, c_2, \dots, c_k\}$ of k colors, a k -coloring of G is any assignment of one of the k available colors from C to every vertex in $V(G)$. More formally, a k -coloring of G is a mapping $c : V(G) \rightarrow C$. The k -coloring problem (k -COL) is to find such a mapping (or prove that none exists) such that adjacent vertices receive different colors (called “proper” k -coloring). More formally, a proper k -coloring of G verifies $\{v_i, v_j\} \in E(G) \rightarrow c(v_i) \neq c(v_j)$. The tightly related optimization version of k -COL is the graph coloring problem (COL): Determine a proper k -coloring of G with k minimum, i.e. the *chromatic number* $\chi(G)$.

k -COL is known to be *NP-complete* when $k \geq 3$ for general graphs (Garey & Johnson, 1979; Karp, 1972). It remains NP-complete even for *particular* classes of graphs, including, for instance, triangle-free graphs with maximum degree 4 (Maffray & Preissmann, 1996). Classes of graphs for which 3-COL can be decided in polynomial time are discussed, for instance, in (Alekseev et al., 2007; Kochol et al., 2003).

Another way to express the difficulty of a combinatorial search problem is to consider the *phase transition* phenomenon which refers to the “easy-hard-easy” transition regions where a problem goes from easy to hard, and conversely (Cheeseman et al., 1991; Dubois et al., 2001; Gent et al., 1996; Hartmann & Weigt, 2005; Hogg et al., 1996; Monasson et al., 1999), see also (Barbosa & Ferreira, 2004; Krzakala et al., 2004; Zdeborová & Krzakala, 2007) for k -COL. Various phase transition

* Corresponding author.

thresholds (noted τ hereafter) have been identified for some classes of random graphs. For 3-COL, τ seems to occur when the edge probability p is such that $2pn/3 \approx 16/3$ according to Petford & Welsh (1989) (referred as τ_w in the rest of the paper), when the mean connection degree $2m/n \approx 5.4$ (τ_c from Cheeseman et al. (1991)), when $7/n \leq p \leq 8/n$ (τ_h from Eiben, van der Hauw, & van Hemert (1998)), when $2m/n \approx 4.6$ (τ_g from Culberson & Gent (2001)), or when $p \approx 3/n + 3(n-3)(1 - 1/6^{2/n})/2n$ (τ_e from Erben (2001)). Note that τ_e and τ_w are similar to the upper bound of $\tau_h(8/n)$. τ_c and τ_g are also similar but τ_c holds only for graphs that are first transformed (before solving) using three “particular reduction operators” (Cheeseman et al., 1991). Additionally, τ_e was characterized just for equipartite graphs and τ_w only for equipartite and uniform graphs (the construction of such graphs is described in Sect. 0). Henceforth, we use the terminology *outside* of τ_h (or τ_c or τ_g , etc.) to indicate parameter values outside of the indicated τ setting.

This paper focuses on an experimental study of finding solution for 3-colorable random graphs around and outside of phase transitions. We are particularly interested in two questions. First, are graphs around phase transitions really difficult to color from a practical solution point of view? Effectively, the different thresholds for phase transition have been established either theoretically or empirically. In both cases, it would be interesting to verify these thresholds by large scale computational experimentation. Notice that, except (Eiben et al., 1998), most experimental studies (see e.g. (Cheeseman et al., 1991; Hogg et al., 1996)) are based only on systematic backtracking search algorithms and small graphs (with no more than 200 vertices). Little is known about the behavior of a (metaheuristic-based) search algorithm on solving large and very large 3-colorable graphs.

Closely related to this first question is another interesting point: Given the phase transition phenomenon, what are the largest sizes of the graphs that can be colored in practice? Actually, the phase transition thresholds distinguish the relative hardness of instances around and outside of the thresholds. They don’t tell much about whether such instances can be solved easily with a practical solution algorithm (such as tabu search) and for which problem sizes a solution is possible.

In this study, we aim to investigate these issues by studying a large range of random graphs generated according to three well-known distributions: Uniform, equipartite, and flat (see next section for more details). For the solution algorithm, we employ a simple tabu search (TS) algorithm (Glover & Laguna, 1997) which can be considered as a baseline reference for the class of metaheuristic (k -) coloring algorithms.

We report computational results on graphs with up to 11 000 vertices, leading to two main findings. First, the variation of solution difficulty of random graphs around and outside of phase transition thresholds are clearly confirmed throughout the experiments: Graphs around the phase transition thresholds are actually more difficult to color than those outside of the thresholds. Second, for the three classes of graphs (uniform, equipartite and flat), the TS algorithm is able to find solutions for graphs with up to at least 11 000 vertices if the graphs are outside of the phase transitions. For graphs around the phase transitions, the TS algorithm always manages to find solutions for uniform and equipartite graphs with up to at least 11 000 vertices, but for flat graphs, the performance seems limited to graphs of 1 500 vertices.

The next section presents the three classes of 3-colorable random graphs studied in this paper. The TS 3-coloring algorithm is described in Sect. 3. Computational results are given in Sect. 4 before concluding.

2. RANDOM GRAPHS

While many classes of random graphs exist (Bollobás, 2001; Krivelevich & Sudakov, 2006), we focus our study on three well-known classes of 3-colorable graphs: Uniform, equipartite, and flat.

There are several reasons for this choice. These random graphs have been object of a number of theoretical (and sometimes practical) studies and analyses, see e.g. (Bollobás, 2001; Braunstein et al., 2003; Culberson & Gent, 2001; Erben, 2001; Fleurent & Ferland, 1996a; Krzakala et al., 2004; Zdeborová & Krzakala, 2007). There is a publicly available generator from <http://web.cs.ualberta.ca/~joe/Coloring/Generators/generate.html> (newer version). The work reported in (Eiben et al., 1998), the only paper that we are aware of on practical solution of the 3-coloring problem, is based on random graphs generated by the same generator, making it possible to use the results of Eiben et al. (1998) as a reference for reporting the 3-coloring results of our TS algorithm.

Uniform. Vertices are first randomly assigned to one of the 3 colors uniformly and independently. Then, each edge $\{v_i, v_j\}$ verifying $c(v_i) \neq c(v_j)$ appears with probability p . We will refer to these graphs with the $U_{n,p}$ notation (or U , for short). Specify 3 at “K-coloring schemes”, 3 at “partition number”, 0 at “variability”, and 1 at “graph type” prompts when running the generator.

Equipartite. In $E_{n,p}$ graphs, $V(G)$ is first split into 3 subsets $V_{c_i \in C}$ ($C = \{c_1, c_2, c_3\}$ since $k = 3$) such that $|V_{c_i}| = \lfloor n/3 \rfloor$ or $|V_{c_i}| = \lceil n/3 \rceil \forall c_i \in C$ (i.e. all V_{c_i} are nearly equal in size, the smallest subset having one less member than the largest), $v_j \in V_{c_i}$ meaning $c(v_j) = c_i$. Then, edges appear as in U graphs. Specify 2 at “K-coloring schemes”, 3 at “partition number”, and 1 at “graph type” prompts.

Flat. Based on E graphs, the $F_{n,p}$ graphs have an additional property related to the variation of the expected degree of the vertices. Specify 6 at “K-coloring schemes”, 3 at “partition number”, and 0 at “flatness” prompts.

3. TC: A TABU SEARCH ALGORITHM FOR 3-COL

In this section, we describe the components and overall scheme of our tabu search 3-coloring algorithm (called TC) used for our 3-COL experiments. TC is an application to 3-COL of the TS metaheuristic (Glover & Laguna, 1997). Its implementation is based on the TS (k -)coloring algorithms given in (Dorne & Hao, 1998; Fleurent & Ferland, 1996a), which themselves are improved variants of TABUCOL, the first TS algorithm for general (k -)COL introduced in (Hertz & de Werra, 1987)¹.

Starting state. The well known greedy DSATUR algorithm (Brélaz, 1979) is used to build a starting 3-coloring (proper or not) while restricting the number of available colors to 3. Vertices that cannot be assigned any of the 3 colors without generating conflicts are (temporarily) removed from the graph with their incident edges. After running DSATUR, these free vertices are finally randomly assigned one of the 3 authorized colors.

Fitness function. Let \mathbf{C} be the set of all 3-colorings (proper or not) of G and $\tilde{E}(c)$ be the set of conflicting edges (i.e. with endpoints colored the same) of $c \in \mathbf{C} : \tilde{E}(c) = \{\{v_i, v_j\} \in E(G) : c(v_i) = c(v_j)\}$. Any 3-coloring c is evaluated according to the following fitness function to be minimized: $f(c) = |\tilde{E}(c)|$ ($f: \mathbf{C} \rightarrow \{0, 1, \dots, m\}$). Note that c is a proper 3-coloring if $f(c) = 0$.

¹ A C++ source code implementing TABUCOL is available e.g. from www.imada.sdu.dk/~marco/gcp-study.

Move operator. A move m maps a 3-coloring c to another 3-coloring c' (i.e. $m : \mathbf{C} \rightarrow \mathbf{C}$) by changing the color of exactly one vertex v_j to $c'(v_j) \neq c(v_j)$, noted $c' = m_c(v_j, c'(v_j))$. Let $M(c)$ be the set of all potential moves available from c : $M(c) = \{(v_j, c'(v_j)) : c'(v_j) \neq c(v_j)\}$.

Neighborhood. The set of 3-colorings c' reachable from c by applying all potential moves defines the neighborhood $N(c)$ of c . More formally, $N(c) = \{c' = m_c(v_j, c'(v_j)) : (v_j, c'(v_j)) \in M(c)\}$.

Tabu list. When a move m is performed from a 3-coloring c to $c' \in N(c)$, the reverse move $m_c^{-1}(v_j, c(v_j)) = c$ (i.e. assigning to v_j its previous color) is “tabu” (forbidden) for the next $TT = \min\{(k-1)f(c), \alpha|\tilde{V}(c)| + \text{rand}(\gamma)\}$ iterations², where α is a TC parameter, $\text{rand}(\gamma)$ is a random integer from $\{1, 2, \dots, \gamma\}$ (the role of γ is just to introduce a few stochastic noise), and $\tilde{V}(c) \subseteq V(G)$ is the set of conflicting vertices of c ($\tilde{V}(c) = \{v_i : \{v_i, v_j\} \in E(G) \rightarrow c'(v_i) = c(v_j)\}$).

Stopping criterion. TC halts whenever $f(c) = 0$ (a proper 3-coloring c has been found) or after a maximum allowed number of moves.

Given the previous components of TC, the core procedure (see the subsequent algorithm) searches for a 3-coloring $c^* \in \mathbf{C}$ (proper or not) with a minimum number of conflicting edges (with $f(c^*) = 0$ ideally, meaning that TC halts since it has found a proper 3-coloring c^*). To do so, TC iteratively moves from a 3-coloring $c \in \mathbf{C}$ to a $c' \in N(c)$. Let $M^*(c) \subset M(c)$ be the set of **best** moves (according to f) available from c and involving a conflicting vertex such that, $\forall m \in M^*(c)$, m is not tabu or m leads to a neighbor better than the best 3-coloring c^* found so far (aspiration criterion). If $M^*(c) \neq \emptyset$, m is chosen at random from $M^*(c)$ according to some probability π . Otherwise, i.e. with probability $1 - \pi$ or when $M^*(c) = \emptyset$, m is chosen at random from $M(c)$. Note that c^* is updated each time $f(c') < f(c^*)$.

TABU SEARCH ALGORITHM.

Require: A 3-colorable graph $G = (V(G), E(G))$ and a set $C = \{c_1, c_2, c_3\}$ of three colors

Require: A starting 3-coloring $c \in \mathbf{C}$ of G // Proper or not

1. $c^* \leftarrow c$ // Best 3-coloring found so far
 2. $TL(j, i) \leftarrow 0 \forall (v_j, c_i) \in V(G) \times C$ // Make the tabu list TL empty
 3. $\mu \leftarrow 0$ // Current number of moves
 4. **while** stopping criterion not met **do**
 5. $\mu \leftarrow \mu + 1$
 6. Let $M(c) = \{(v_j, c'(v_j)) : c' \in N(c)\}$
 7. Let $M^*(c) = \{(v_j, c'(v_j)) \in M(c) :$
 $\quad v_j \in \tilde{V}(c) \text{ and } \forall (v_i, c'(v_i)) \in M(c), f(c') \leq f(c'') \text{ and } (TL(j, c'(v_j)) < \mu \text{ or } f(c') < f(c^*))\}$
 8. Let r be a random real number in $[0, 1]$
 9. **if** $M^*(c) = \emptyset$ **or** $r > \pi$ **then**
 10. Randomly select a move $(v_j, c'(v_j))$ from $M(c)$
 11. **else**
 12. Randomly select a move $(v_j, c'(v_j)) \in M^*(c)$
 13. $TL(j, c(v_j)) \leftarrow \mu + TT$ // Forbid the reverse move m^{-1} at least up to iterations $\mu + TT$
 14. $c(v_j) \leftarrow c'(v_j)$ // Do the selected move
 15. **if** $f(c) < f(c^*)$ **then**
 16. $c^* \leftarrow c$
 17. **return** c^*
-

² TT is called the “tabu tenure”. We used the same dynamic TT formula than that in (Dorne & Hao, 1998) since this approach achieved effective results.

Note that selecting (lines 10 and 12 in the TC algorithm) or doing (line 14) a move in TC can be achieved efficiently, i.e. within small time complexity, using a particular data structure inspired by a technique from Fleurent & Ferland (1996b) and usually called “ δ table” in the wide tabu search literature. Basically, δ is a $n \times k$ matrix where $\delta_c(j, c'(v_j))$ stores the fitness variation (between $c \in C$ and $c' \in N(c)$) when the color assigned to $v_j \in V(G)$ changes from $c(v_j)$ to $c'(v_j)$: $\delta_c(j, c'(v_j)) = f(c') - f(c)$. δ is initialized once at the beginning of the search (before line 4, in time $O(nk)$) and updated each time a move is performed (after line 14, in time $O(nk)$ in the worst case but, in practice, only a subset of δ is updated). While selecting a move from the $M(c)$ set (line 10) takes $O(1)$ time, the evaluation of all “best” moves from the $M^*(c)$ set (line 12) is almost *incremental*: It can be achieved in $O(|\tilde{V}(c)|k)$ time in the worst case thanks to δ . Thus, each iteration takes $O(2nk)$ time at most since $|\tilde{V}(c)| \leq n$ for any 3-coloring c .

4. COMPUTATIONAL RESULTS

The computational experiments reported in Secs. 4.1–4.5 are based on the following general protocol.

Benchmark set. A collection consisting of 263 different instances is built according to Sect. 0. Recall that all these graphs are 3-colorable by construction. Their order ranges from 200 to 11 000. Note that the generator requires an integer seed for randomization initialization: We always use 5 as in (Eiben et al., 1998) to deal exactly with the same instances. Additionally, Eiben et al. (1998) noted that this parameter seems to have no great influence on results.

Reference algorithm. For reporting computational results of TC, we use the SAW evolutionary algorithm (Eiben et al., 1998) as a reference. Indeed, according to Eiben et al. (1998), SAW is effective in 3-coloring random 3-colorable graphs of large order (up to 1 500 vertices). Moreover, the authors clearly describe the graph generator employed and the seed for randomization initializations, making it possible to make direct comparisons. In all our tables shown later in the paper, “–” signals unavailable or inapplicable entries and results reported for SAW are approximated from figures in (Eiben et al., 1998). No information is given for SAW in some of our tables since it cannot be retrieved from (Eiben et al., 1998).

Performance criteria. The solution performance is assessed according to the well-known “Success Rate” measure (*SR*): It is the percentage of successful runs, i.e. in which a proper 3-coloring is found, over a given number of runs. To give an idea of the TC computational effort, we also report the mean number of moves required by TC to find a proper 3-coloring (*AMS*, for “Average number of Moves to Solution”) and its standard deviation (σ_{AMS}). Eiben et al. (1998) used a slightly different measure, namely the mean number of fitness evaluations (*AES*, for “Average number of Evaluations to Solution”). Note that *AMS* and *AES* are implementation and hardware independent measures. The mean computation time T and its standard deviation σ_T (in seconds) are also reported for successful runs of TC.

Phase transition. In some tables, the cases the closest to τ_c , τ_e , τ_g , τ_h , and τ_w are identified with the appropriate “c”, “e”, “g”, “h”, and “w” letters in the τ columns. The bold entries in Tables 1–9 (Sect. 4.1) and Tables 10–18 (Sect. 4.2) indicates which τ is the closest to the hardest cases (minimum *SR*, or maximum *AMS* or *AES*), i.e. it suggests which τ seems to be best suited to locate the phase transition.

Implementation. Our TC algorithm is coded in the C programming language (“gcc” compiler). All TC computational results were obtained on a Sun Fire V880 server with 8 Gb RAM (UltraSPARC III CPU 750 MHz).

The values of the main TC parameters were empirically determined during a few preliminary computational experiments (not shown here): $\alpha = 0.5$, $\gamma = 2$, and $k = 3$ (to compute the tabu tenure TT), and $\pi = 0.85$ (probability to select a move in M^*).

4.1. Influence of the Edge Probability p on the Problem Difficulty

Almost similarly to Eiben et al. (1998), we first limit the maximum allowed number of moves of the TC algorithm to 300 000 and vary p from 0.015 to 0.075 for $n = 200$ (step 0.005, 100 runs per p value and per graph, a total of 39 graphs), 0.006 to 0.05 for $n = 500$ (step 0.004, 50 runs, 36 instances), and 0.002 to 0.026 for $n = 1\,000$ (step 0.002, 25 runs, 45 graphs). Note that three instances were generated per p value since we consider three types of graphs (U, E, and F). Results are reported in Tables 1–9 where the two lines associated with τ (between the two dashed lines) correspond to graphs around (i.e. the closest to) the indicated phase transition thresholds while the other lines concern graphs outside of (i.e. more distant from) these thresholds.

On the set of small-order instances ($n = 200$, see Tables 1–3), TC always succeeds in all runs (SR is always 1) for all the graphs within the time limit of 300 000 moves, but needs more moves to find a solution for a graph at the phase transitions (when $p = 0.035$) than outside of the thresholds. Note that the initialization procedure DSATUR alone always finds a proper 3-coloring whenever $p = 0.015$ and for the $F_{200, 0.02}$ graph ($AMS = 0.0$ means that TC performs no move at all). DSATUR also obtains proper 3-colorings in some runs for $p \in \{0.02, 0.025\}$ in each class.

At $n = 500$ (Tables 4–6), while more computational effort (AMS) is sometimes needed by TC, the problem is still easy for TC outside of τ_g (SR is always 1). At τ_g , TC is always competitive in terms of SR , especially on the U graph where $SR = 0.9$ (see Table 4). However, the problem is here slightly harder than the $n = 200$ cases for TC. This is particularly true on the F and E graphs where the SR achieved by TC at τ_g falls, respectively, to 0.72 and 0.56 (see Tables 6 and 5). DSATUR continues to produce proper 3-colorings for $n = 500$ in each class, in all runs when $p = 0.006$ and sometimes for $F_{500, 0.01}$.

On large-order graphs ($n = 1\,000$, Tables 7–9), TC finds proper 3-colorings in all the 25 runs for each class whenever p is outside of τ_h . In these cases, mean computing times are still short. At τ_h , TC succeeds in all runs, but only on U and E graphs, see Tables 7–8 respectively. Indeed, it achieves $SR = 0.04$ for the F instance (Table 9). Here again, the DSATUR algorithm directly identifies proper 3-colorings in all runs whenever $p = 0.002$ and for $E_{1\,000, 0.004}$ and $F_{1\,000, 0.004}$, and in some runs for $U_{1\,000, 0.004}$.

Now, we turn our attention to the performance of the reference algorithm SAW. At $n = 200$, SAW obtained interesting SR values on U and E graphs, see Tables 1–2 where SR is always 1 except when $p = 0.035$ ($SR \approx 0.9$ and $SR \approx 0.85$, respectively). For F graphs (Table 3), while SAW still verifies $SR = 1$ outside of τ , it achieves a lower SR around τ : $SR \approx 0.65$ for $p = 0.04$ and $SR \approx 0.37$ when $p = 0.035$. This confirms the well known fact that F graphs may be harder than U and E instances, even on small-order graphs. For medium-order graphs (see Tables 4–6), the SR of SAW is always 1 outside of τ_g except on $F_{500, 0.022}$ ($SR \approx 0.94$) and $F_{500, 0.018}$ ($SR \approx 0.54$). SAW starts to have (great) difficulties in finding proper 3-colorings at τ_g when $n = 500$. Indeed, $SR \approx 0.1$ on the U graph and $SR \approx 0.08$ for the F instance. Furthermore, it seems to fail on the E instance ($SR \approx 0$). At $n = 1\,000$ (Tables 7–9), SAW always finds proper 3-colorings whenever p is outside of τ_h except on two E graphs ($SR \approx 0.96$ for $p \in \{0.006, 0.01\}$) and two F graphs ($SR \approx 0.88$ for $p = 0.012$ and $SR \approx 0.48$ for $p = 0.01$). SAW dramatically fails at τ_h : $SR \approx 0.04$ for the U instance and SAW seems to never solve E and F graphs ($SR \approx 0$). Consequently, one can conclude that TC reaches always the same or higher success rate than SAW on all the graphs.

Table 1. Small-order U graphs ($n = 200$): Influence of the edge probability p (100 runs).

p	τ	TC (300 000 moves)					SAW	
		SR	AMS	σ_{AMS}	T (s)	σ_T (s)	SR	AES
0.015		1	0.0	0.0	< 1	< 1	1	0
0.02		1	0.6	2.6	< 1	< 1	1	0
0.025		1	124.2	264.6	< 1	< 1	1	0
0.03		1	3 376.5	2 982.0	< 1	< 1	1	10 000
0.035	g, h	1	14 423.6	13 371.8	< 1	< 1	0.90	75 000
0.04	c, e, h, w	1	2 851.6	2 140.5	< 1	< 1	1	10 000
0.045		1	840.9	618.1	< 1	< 1	1	4 000
0.05		1	1 150.6	661.3	< 1	< 1	1	4 000
0.055		1	869.2	520.2	< 1	< 1	1	2 000
0.06		1	1 242.1	1 390.7	< 1	< 1	1	2 000
0.065		1	731.5	611.3	< 1	< 1	1	1 000
0.07		1	720.6	405.2	< 1	< 1	1	1 000
0.075		1	519.3	309.8	< 1	< 1	1	500

Table 2. Small-order E graphs ($n = 200$): Influence of the edge probability p (100 runs).

p	τ	TC (300 000 moves)					SAW	
		SR	AMS	σ_{AMS}	T (s)	σ_T (s)	SR	AES
0.015		1	0.0	0.0	< 1	< 1	1	0
0.02		1	0.5	2.8	< 1	< 1	1	0
0.025		1	65.3	103.9	< 1	< 1	1	0
0.03		1	4 540.7	3 984.8	< 1	< 1	1	13 000
0.035	g, h	1	11 865.1	9 946.7	< 1	< 1	0.85	68 000
0.04	c, e, h, w	1	3 699.8	2 993.3	< 1	< 1	1	68 000
0.045		1	998.3	709.6	< 1	< 1	1	9 000
0.05		1	766.3	398.3	< 1	< 1	1	9 000
0.055		1	1 019.4	858.6	< 1	< 1	1	4 500
0.06		1	1 786.9	1 418.6	< 1	< 1	1	4 500
0.065		1	971.5	1 440.8	< 1	< 1	1	2 000
0.07		1	510.4	291.5	< 1	< 1	1	2 000
0.075		1	248.9	221.0	< 1	< 1	1	1 000

Table 3. Small-order F graphs ($n = 200$): Influence of the edge probability p (100 runs).

p	τ	TC (300 000 moves)					SAW	
		SR	AMS	σ_{AMS}	T (s)	σ_T (s)	SR	AES
0.015		1	0.0	0.0	< 1	< 1	1	0
0.02		1	0.0	0.0	< 1	< 1	1	0
0.025		1	6.7	25.0	< 1	< 1	1	0
0.03		1	720.5	718.2	< 1	< 1	1	8 000
0.035	g, h	1	58 636.4	47 428.0	< 1	< 1	0.37	110 000
0.04	c, e, h, w	1	14 226.4	13 675.2	< 1	< 1	0.65	75 000
0.045		1	2 749.3	1 779.0	< 1	< 1	1	13 500
0.05		1	1 053.1	960.2	< 1	< 1	1	12 500
0.055		1	1 146.8	652.6	< 1	< 1	1	6 000
0.06		1	2 785.1	2 929.8	< 1	< 1	1	6 000
0.065		1	941.1	748.7	< 1	< 1	1	3 000
0.07		1	931.8	774.2	< 1	< 1	1	3 000
0.075		1	398.2	280.2	< 1	< 1	1	3 000

Table 4. Medium-order U graphs ($n = 500$): Influence of the edge probability p (50 runs).

p	τ	TC (300 000 moves)					SAW	
		SR	AMS	σ_{AMS}	T (s)	σ_T (s)	SR	AES
0.006		1	0.0	0.0	< 1	< 1	1	0
0.01		1	286.9	381.9	< 1	< 1	1	8 000
0.014	c, g, h, w	0.9	98 080.4	74 802.3	1.6	1.1	0.1	90 000
0.018	e, w	1	4 754.1	2 405.9	< 1	< 1	1	25 000
0.022		1	5 113.1	2 852.4	< 1	< 1	1	8 000
0.026		1	5 235.8	3 378.6	< 1	< 1	1	8 000
0.03		1	1 769.5	744.1	< 1	< 1	1	8 000
0.034		1	2 504.7	1 937.4	< 1	< 1	1	8 000
0.038		1	956.8	796.2	< 1	< 1	1	8 000
0.042		1	935.4	480.5	< 1	< 1	1	8 000
0.046		1	1 380.1	4 739.4	< 1	< 1	1	8 000
0.05		1	874.7	556.0	< 1	< 1	1	8 000

Table 5. Medium-order E graphs ($n = 500$): Influence of the edge probability p (50 runs).

p	τ	TC (300 000 moves)					SAW	
		SR	AMS	σ_{AMS}	T (s)	σ_T (s)	SR	AES
0.006		1	0.0	0.0	< 1	< 1	1	0
0.01		1	263.7	252.2	< 1	< 1	1	8 000
0.014	c, g, h, w	0.56	180 950.9	63 173.1	2.8	< 1	0	-
0.018	e, w	1	6 913.4	7 106.6	< 1	< 1	1	30 000
0.022		1	4 678.5	2 060.6	< 1	< 1	1	20 000
0.026		1	9 008.2	19 218.4	< 1	< 1	1	12 500
0.03		1	1 855.7	1 363.8	< 1	< 1	1	12 500
0.034		1	1 205.9	1 628.4	< 1	< 1	1	12 500
0.038		1	2 021.1	1 149.9	< 1	< 1	1	8 000
0.042		1	1 415.1	4 915.3	< 1	< 1	1	8 000
0.046		1	5 756.4	27 617.0	< 1	< 1	1	8 000
0.05		1	469.4	609.2	< 1	< 1	1	8 000

Table 6. Medium-order F graphs ($n = 500$): Influence of the edge probability p (50 runs).

p	τ	TC (300 000 moves)					SAW	
		SR	AMS	σ_{AMS}	T (s)	σ_T (s)	SR	AES
0.006		1	0.0	0.0	< 1	< 1	1	0
0.01		1	26.7	60.2	< 1	< 1	1	500
0.014	g, h, w	0.72	133 391.3	69 861.7	2.1	1.1	0.08	115 000
0.018	c, e, w	1	26 981.2	28 508.1	< 1	< 1	0.54	85 000
0.022		1	7 931.2	4 684.5	< 1	< 1	0.94	55 000
0.026		1	17 668.1	30 915.4	< 1	< 1	1	16 500
0.03		1	1 732.8	962.6	< 1	< 1	1	12 500
0.034		1	3 757.1	1 727.3	< 1	< 1	1	4 000
0.038		1	2 247.9	1 975.9	< 1	< 1	1	4 000
0.042		1	288.5	250.8	< 1	< 1	1	4 000
0.046		1	1 289.3	1 001.6	< 1	< 1	1	4 000
0.05		1	1 019.7	856.1	< 1	< 1	1	4 000

Table 7. Large-order U graphs ($n = 1\ 000$): Influence of the edge probability p (25 runs).

p	τ	TC (300 000 moves)					SAW	
		SR	AMS	σ_{AMS}	T (s)	σ_T (s)	SR	AES
0.002		1	0.0	0.0	< 1	< 1	1	0
0.004		1	82.9	203.7	< 1	< 1	1	4 000
0.006	g	1	113 985.8	78 314.6	3.6	2.4	1	95 000
0.008	c, e, h, w	1	117 190.8	73 166.9	5.7	3.6	0.04	135 000
0.01		1	13 928.6	5 770.6	< 1	< 1	1	60 000
0.012		1	35 546.6	32 219.9	1.4	1.1	1	35 000
0.014		1	4 972.1	2 014.5	< 1	< 1	1	20 000
0.016		1	11 020.4	6 874.4	< 1	< 1	1	20 000
0.018		1	8 920.1	4 135.9	< 1	< 1	1	20 000
0.02		1	3 220.8	1 130.9	< 1	< 1	1	10 000
0.022		1	4 874.2	3 029.0	< 1	< 1	1	10 000
0.024		1	4 685.3	1 895.9	< 1	< 1	1	10 000
0.026		1	1 652.3	711.4	< 1	< 1	1	10 000

Table 8. Large-order E graphs ($n = 1\ 000$): Influence of the edge probability p (25 runs).

p	τ	TC (300 000 moves)					SAW	
		SR	AMS	σ_{AMS}	T (s)	σ_T (s)	SR	AES
0.002		1	0.0	0.0	< 1	< 1	1	0
0.004		1	0.0	0.0	< 1	< 1	1	8 000
0.006	g	1	132 932.7	74 624.1	4.2	2.3	0.96	120 000
0.008	c, e, h, w	1	102 510.6	73 566.2	5.1	3.8	0	-
0.01		1	12 648.1	6 684.0	< 1	< 1	0.96	85 000
0.012		1	81 719.4	43 227.4	2.9	1.5	1	40 000
0.014		1	9 683.8	4 300.5	< 1	< 1	1	30 000
0.016		1	12 685.9	6 086.4	< 1	< 1	1	20 000
0.018		1	14 078.4	9 242.3	< 1	< 1	1	16 500
0.02		1	4 762.6	1 995.0	< 1	< 1	1	16 500
0.022		1	7 353.6	3 467.2	< 1	< 1	1	16 500
0.024		1	6 206.8	3 096.5	< 1	< 1	1	16 500
0.026		1	3 523.1	1 168.8	< 1	< 1	1	16 500

Table 9. Large-order F graphs ($n = 1\ 000$): Influence of the edge probability p (25 runs).

p	τ	TC (300 000 moves)					SAW	
		SR	AMS	σ_{AMS}	T (s)	σ_T (s)	SR	AES
0.002		1	0.0	0.0	< 1	< 1	1	0
0.004		1	0.0	0.0	< 1	< 1	1	0
0.006	g	1	22 952.4	14 409.0	< 1	< 1	1	50 000
0.008	c, e, h, w	0.04	102 504.0	0.0	5.7	0.0	0	-
0.01		1	38 349.3	30 597.6	3.1	2.5	0.48	132 500
0.012		1	83 314.1	55 334.2	3.0	1.9	0.88	100 000
0.014		1	19 289.2	12 835.6	1.2	< 1	1	50 000
0.016		1	10 703.1	3 251.3	< 1	< 1	1	22 500
0.018		1	9 633.8	4 490.6	< 1	< 1	1	22 500
0.02		1	5 937.3	2 540.5	< 1	< 1	1	10 000
0.022		1	6 327.8	2 685.7	< 1	< 1	1	10 000
0.024		1	4 107.8	2 385.4	< 1	< 1	1	10 000
0.026		1	3 523.8	1 568.8	< 1	< 1	1	10 000

4.2. Deeper experiments around the phase transitions

Tables 1–9 disclose that 3-COL is typically harder at τ_h than at τ_c , τ_e , τ_g , or τ_w , i.e. that τ_h may be more effective at identifying the hardest instances. To try to verify this observation, we report deeper experiments with TC in Tables 10–18 for more detailed p values around τ . Note that this section include 21 new graphs not considered in Sect. 4.1 (they appear in italic typeface).

Small-order graphs ($n = 200$) are still easy, even at τ , see Tables 10–12. Indeed, SR is always 1 except on $F_{200, 0.0375}$ where $SR = 0.82$. Furthermore, mean computing time of TC is always smaller than a second. Medium-order graphs ($n = 500$, Tables 13–15) also seem to be quite easy for TC, even at τ . Indeed, SR is always 1 except on $U_{500, 0.014}$ (0.9), $E_{500, 0.014}$ (0.56), $F_{500, 0.014}$ (0.72), and $F_{500, 0.016}$ (0.64). Some large-order graphs ($n = 1\,000$, Tables 16–18) are especially difficult for TC within the time limit of 300 000 moves. This is particularly true at τ_g for all the instances (since $SR \leq 0.08$) and outside of τ_g for one F graph ($SR = 0.04$ when $p = 0.008$). Furthermore, the difficulty also holds outside of τ in one case, when $p = 0.009$ for the F instance ($SR = 0.48$).

Table 10. Small-order U graphs ($n = 200$): Deeper experiments with TC around τ (100 runs, 300 000 moves).

p	τ	SR	AMS	σ_{AMS}	T (s)	σ_T (s)
0.03		1	3 376.5	2 982.0	< 1	< 1
<i>0.0325</i>	<i>g</i>	1	8 256.4	6 963.3	< 1	< 1
0.035	<i>h</i>	1	14 423.6	13 371.8	< 1	< 1
<i>0.0375</i>	<i>c, h</i>	1	5 849.7	3 597.3	< 1	< 1
0.04	<i>c, e, h, w</i>	1	2 851.6	2 140.5	< 1	< 1
<i>0.0425</i>		1	1 916.3	1 436.7	< 1	< 1

Table 11. Small-order E graphs ($n = 200$): Deeper experiments with TC around τ (100 runs, 300 000 moves).

p	τ	SR	AMS	σ_{AMS}	T (s)	σ_T (s)
0.03		1	4 540.7	3 984.8	< 1	< 1
0.0325	<i>g</i>	1	16 016.6	12 633.1	< 1	< 1
0.035	<i>h</i>	1	11 865.1	9 946.7	< 1	< 1
<i>0.0375</i>	<i>c, h</i>	1	5 518.3	5 319.6	< 1	< 1
0.04	<i>e, h, w</i>	1	3 699.8	2 993.3	< 1	< 1
<i>0.0425</i>		1	1 399.6	1 279.6	< 1	< 1

Table 12. Small-order F graphs ($n = 200$): Deeper experiments with TC around τ (100 runs, 300 000 moves).

p	τ	SR	AMS	σ_{AMS}	T (s)	σ_T (s)
0.03		1	720.5	718.2	< 1	< 1
<i>0.0325</i>		1	6 018.3	5 727.3	< 1	< 1
0.035	<i>g, h</i>	1	58 636.4	47 428.0	< 1	< 1
0.0375	<i>h</i>	0.82	78 110.1	76 498.8	< 1	< 1
0.04	<i>c, e, h, w</i>	1	14 226.4	13 675.2	< 1	< 1
<i>0.0425</i>		1	14 091.3	12 096.2	< 1	< 1

Table 13. Medium-order U graphs ($n = 500$): Deeper experiments with TC around τ (50 runs, 300 000 moves).

p	τ	SR	AMS	σ_{AMS}	T (s)	σ_T (s)
<i>0.012</i>		1	12 126.5	9 510.6	< 1	< 1
0.014	<i>g, h</i>	0.9	98 080.4	74 802.4	1.6	1.1
<i>0.016</i>	<i>c, e, h, w</i>	1	16 411.1	14 708.0	< 1	< 1
0.018		1	4 754.1	2 405.9	< 1	< 1

Table 14. Medium-order E graphs ($n = 500$): Deeper experiments with TC around τ (50 runs, 300 000 moves).

p	τ	SR	AMS	σ_{AMS}	T (s)	σ_T (s)
0.012		1	7 852.9	4 524.5	< 1	< 1
0.014	c, g, h	0.56	180 950.9	63 173.1	2.8	< 1
0.016	e, h, w	1	24 717.1	20 670.7	< 1	< 1
0.018		1	6 913.4	7 106.6	< 1	< 1

Table 15. Medium-order F graphs ($n = 500$): Deeper experiments with TC around τ (50 runs, 300 000 moves).

p	τ	SR	AMS	σ_{AMS}	T (s)	σ_T (s)
0.012		1	2 735.7	1 780.9	< 1	< 1
0.014	g, h	0.72	133 391.3	69 861.7	2.1	1.1
0.016	c, e, h, w	0.64	148 273.7	76 569.5	3.3	1.7
0.018		1	26 981.2	28 508.1	< 1	< 1

Table 16. Large-order U graphs ($n = 1 000$): Deeper experiments with TC around τ (25 runs, 300 000 moves).

p	τ	SR	AMS	σ_{AMS}	T (s)	σ_T (s)
0.006		1	113 985.8	78 314.6	3.6	2.4
0.007	g, h	0.04	236 891.0	0.0	8.3	0.0
0.008	c, e, h, w	1	117 190.8	73 166.9	5.7	3.6
0.009		1	23 644.5	6 916.5	1.5	< 1

Table 17. Large-order E graphs ($n = 1 000$): Deeper experiments with TC around τ (25 runs, 300 000 moves).

p	τ	SR	AMS	σ_{AMS}	T (s)	σ_T (s)
0.006		1	132 932.7	74 624.1	4.2	2.3
0.007	g, h	0.08	215 655.5	31 855.5	8.1	1.0
0.008	c, e, h, w	1	102 510.7	73 566.2	5.1	3.8
0.009		1	26 270.6	14 541.4	1.7	< 1

Table 18. Large-order F graphs ($n = 1 000$): Deeper experiments with TC around τ (25 runs, 300 000 moves).

p	τ	SR	AMS	σ_{AMS}	T (s)	σ_T (s)
0.006		1	22 952.4	14 409.0	< 1	< 1
0.007	g, h	0	-	-	-	-
0.008	c, e, h, w	0.04	102 504.0	0.0	5.7	0.0
0.009		0.48	139 301.5	86 264.1	9.5	5.8

Table 19 recalls the most effective τ measure from Tables 10–18 depending on n and the class of graphs. The last three columns (respectively lines) also propose a ranking of τ_c , τ_e , τ_g , τ_h , and τ_w for a particular n value (respectively for a particular graph class). For instance, τ_h is classified as “Best” when $n = 200$ since “h” appears more than the other thresholds on the “ $n = 200$ ” line. Similarly, τ_c , τ_e , and τ_w are categorized as “Worst” for $n = 200$ since they are missing on the “ $n = 200$ ” line.

From Table 19, one can observe that τ_h is (almost) always the most effective τ measure whatever the value of n or the graph class. Indeed, if we define the overall score Σ (for all n values and all graphs) of a τ measure as the number of times it appears in the inner table (intersection of lines 3–5 and columns U–F), we obtain $\Sigma_h > \Sigma_g > \Sigma_c > \Sigma_{e, w}$ (since $8 > 6 > 3 > 2$). One can then establish the following overall τ ranking: $\tau_h >_{\Sigma} \tau_g >_{\Sigma} \tau_c >_{\Sigma} \tau_{e, w}$, where “ $>_{\Sigma}$ ” means “more effective than”. Consequently, we will mainly use τ_h as the phase transition threshold in the rest of the paper.

Table 19. Which τ measure is the best to identify hard 3-COL instances?

n	Graph class			Best	Worst
	U	E	F		
200	h	g	h	h	g
500	g, h	c, g, h	c, e, h, w	h	c, g
1 000	g, h	g, h	c, e, g, h, w	g, h	c, e, w
Best	h	g	h		
	g	c, h	c, e, w		
Worst	c, e, w	e, w	g		

4.3. Influence of the problem size n on the problem difficulty

The scalability of TC, i.e. how its performance changes with growing problem size, can be observed in Tables 20–24 (27 new instances), on graphs respectively outside of τ_h (within 500 000 moves for TC) and around τ_h (1 000 000 moves), for various n values in [250, 1 500] (see also Sect. 4.5, where we use much larger graph with n up to 11 000 to test the limit of TC).

Tables 20–21 show that graphs of these sizes outside of τ_h are really easy for TC since SR is always 1. Around τ_h (Tables 22–24), the U and E graphs are still easy for TC ($SR = 1$) but the F instances become harder when $n \geq 1\,000$ ($SR \leq 0.04$).

Table 20. E graphs: Influence of the problem size outside of τ_h ($\rho = 10/n$, 50 runs).

n	TC (500 000 moves)					SAW	
	SR	AMS	σ_{AMS}	T (s)	σ_T (s)	SR	AES
250	1	1 487.3	862.1	< 1	< 1	1	12 500
500	1	3 568.5	2 091.2	< 1	< 1	1	37 500
750	1	5 932.9	2 500.8	< 1	< 1	1	57 000
1 000	1	10 239.2	5 300.0	1.5	< 1	1	100 000
1 250	1	13 254.8	6 376.8	2.3	1.1	0.9	150 000
1 500	1	21 103.1	9 217.5	4.4	1.8	0.9	185 500

Table 21. U and F graphs: Influence of the problem size on TC outside of τ_h ($\rho = 10/n$, 50 runs, 500 000 moves).

n	U					F				
	SR	AMS	σ_{AMS}	T (s)	σ_T (s)	SR	AMS	σ_{AMS}	T (s)	σ_T (s)
250	1	1 186.9	612.5	< 1	< 1	1	2 238.3	1 272.3	< 1	< 1
500	1	2 885.5	1 357.3	< 1	< 1	1	7 915.5	5 158.7	< 1	< 1
750	1	8 110.6	4 363.4	< 1	< 1	1	17 802.9	13 352.0	1.8	1.3
1 000	1	9 727.2	4 187.5	1.4	< 1	1	33 667.9	25 020.5	4.9	3.6
1 250	1	9 696.3	4 253.2	1.7	< 1	1	68 762.2	65 591.5	2.1	2.0
1 500	1	19 528.4	9 281.9	3.9	1.9	1	70 217.6	48 963.0	2.6	1.8

Table 22. E graphs: Influence of the problem size around τ_h ($\rho = 8/n$, 25 runs).

n	TC (1 000 000 moves)					SAW	
	SR	AMS	σ_{AMS}	T (s)	σ_T (s)	SR	AES
250	1	5 256.3	3 524.6	< 1	< 1	1	28 500
500	1	20 774.4	13 021.1	< 1	< 1	0.88	200 000
750	1	44 542.8	34 333.0	3.2	2.4	0.52	300 000
1 000	1	102 510.7	73 566.2	5.1	3.8	0.16	418 500
1 250	1	130 037.1	184 316.4	15.4	21.3	0.20	400 000
1 500	1	172 020.7	154 432.5	25.2	22.5	0.08	771 900

Table 23. U graphs: Influence of the problem size on TC around τ_h ($p = 8/n$, 25 runs, 1 000 000 moves).

n	SR	AMS	σ_{AMS}	T (s)	σ_T (s)
250	1	2 432.7	1 376.1	< 1	< 1
500	1	17 637.4	17 818.6	< 1	< 1
750	1	51 467.3	34 610.7	3.6	2.4
1 000	1	117 190.8	73 166.9	5.7	3.6
1 250	1	118 455.1	88 368.2	13.9	10.4
1 500	1	177 317.3	179 939.3	3.4	3.4

Table 24. F graphs: Influence of the problem size on TC around τ_h ($p = 8/n$, 25 runs, 1 000 000 moves).

n	SR	AMS	σ_{AMS}	T (s)	σ_T (s)
250	1	321 279.1	241 163.2	6.2	4.7
500	1	306 117.2	234 178.1	2.2	1.7
750	0.24	219 788.5	171 377.6	2.1	1.7
1 000	0.04	102 504.0	0.0	5.7	0.0
1 250	0	–	–	–	–
1 500	0	–	–	–	–

SAW was checked for scalability only on E graphs in (Eiben et al., 1998). While it reached good SR values outside of τ_h (see Table 20), its performance dramatically falls around τ_h when $n \geq 1\,000$ (Table 22).

4.4. Impact of longer runs on the solution performance

We just observed that, in some or all runs, TC fails to find a proper 3-coloring for some graphs within 300 000 moves (see Tables 12–18 in Sect. 4.2) or 1 000 000 moves (Table 24 in Sect. 4.3). We study here the effect of giving more search time to TC, i.e. if longer runs can increase its success rates for solving these instances. So, we first extend the maximum number of moves per run to 1 000 000 for graphs in Sect. 4.2 and rerun TC whenever $SR < 1$ for TC in Tables 12–18. In Table 25, SR_s again lists the SR achieved by TC in Tables 12–18 (short runs with 300 000 moves). Similarly, SR_l , AMS_l , and T_l are for 25 long runs (i.e. within 1 000 000 moves).

Table 25. Long TC runs on the hardest instances from Tables 12–18 where $SR < 1$ (25 runs, 1 000 000 moves).

Graph	τ	SR_s	SR_l	AMS_l	σ_{AMS_l}	T_l (s)	σ_{T_l} (s)
F _{200,0.0375}	h	0.82	1	147 017.4	171 001.0	1.7	2.0
U _{500,0.014}	g, h	0.90	1	196 277.7	156 315.5	3.0	2.3
E _{500,0.014}	c, g, h	0.56	0.96	304 047.6	235 099.6	4.7	3.6
F _{500,0.014}	g, h	0.72	1	293 927.9	171 530.3	4.5	2.6
F _{500,0.016}	c, e, h, w	0.64	0.96	329 983.9	270 385.5	7.1	5.8
U _{1000,0.007}	g, h	0.04	0.28	601 305.1	130 523.4	21.9	4.5
E _{1000,0.007}	g, h	0.08	0.60	619 195.7	254 241.8	22.0	9.0
F _{1000,0.007}	g, h	0	0.40	470 637.1	320 741.1	3.7	2.4
F _{1000,0.008}	c, e, h, w	0.04	0.04	102 504.0	0.0	5.7	0.0
F _{1000,0.009}		0.48	0.72	410 530.3	254 225.0	12.1	7.2

Table 25 confirms that small and medium-order graphs ($n \leq 500$) are easily solved now by TC, even around τ_h ($SR_l \geq 0.96$). Significant improvements can also be observed on large-order U and E graphs ($n = 1\,000$). Nevertheless, the U instance is still quite challenging ($SR_l = 0.28$). The large-order F graphs remain difficult to color, even if some improvements are sometimes observed. Indeed, no improvement at all was possible when $p = 0.008$ ($SR_l = SR_s$).

Note that Eiben et al. (1998) reported one similar experiment using only one graph ($E_{1\,000, 0.008}$): The SR of SAW increased from 0 within 300 000 evaluations to 0.44 within 1 000 000 evaluations ($AES = 407\,283$)³.

Since TC still fails to reach $SR = 1$ within 1 000 000 moves for 10 instances (4 in Table 24 and 7 in Table 25, but $F_{1\,000, 0.008}$ is considered in both tables), we remove this limit and allow TC to run until it finds a proper 3-coloring. Results are summarized in Tables 26–27⁴. “MAXINT” entries in Table 27 indicate values larger than the maximal integer authorized by the system (i.e. 4 294 967 295). In these cases, T_∞ indicates the minimum time needed to reach a proper 3-coloring.

Table 26. Achieving $SR = 1$ with TC on the hardest instances from Table 25 where $SR_l < 1$ (5 runs, without time limit).

Graph	τ	SR_l	SR_∞	AMS_∞	σ_{AMS_∞}	T_∞ (s)	σ_{T_∞} (s)
$E_{500,0.014}$	c, g, h	0.96	1	693 830.4	641 679.3	2.2	2.1
$F_{500,0.016}$	c, e, h, w	0.96	1	450 009.4	430 902.2	3.4	3.2
$U_{1000,0.007}$	g, h	0.28	1	2 904 052.2	3 138 476.6	28.5	31.0
$E_{1000,0.007}$	g, h	0.60	1	1 161 061.6	1 209 402.4	11.6	11.7
$F_{1000,0.007}$	g, h	0.40	1	1 888 195.4	1 413 201.3	11.2	8.2
$F_{1000,0.008}$	c, e, h, w	0.04	1	298 129 024.1	165 232 840.6	3 983.7	2 205.2
$F_{1000,0.009}$		0.72	1	633 880.2	523 436.4	11.6	9.5

Table 27. Achieving $SR = 1$ with TC around $\tau_h(p = 8/n)$ on the hardest F instances from Table 24 where $SR_l < 1$ (5 runs, without time limit).

n	SR_l	SR_∞	AMS_∞	σ_{AMS_∞}	T_∞ (s)	σ_{T_∞} (s)
750	0.24	1	11 933 517.1	9 686 691.4	114.5	92.3
1 000	0.04	1	298 129 024.1	165 232 840.6	3 983.7	2 205.2
1 250	0	1	MAXINT	–	> 3 674.4	–
1 500	0	1	MAXINT	–	> 454 662.9	–

Two main observations can be made from Tables 26–27. First, all graphs are quite easy for TC whenever $p \neq 8/n$, see Table 26 where $AMS_\infty \leq 2\,904\,052$ in this case. Second, only the large-order F instances constitute a real challenge for TC whenever $p = 8/n$, see Table 27 where $AMS_\infty \geq 298\,129\,024$ for $n \geq 1\,000$.

4.5. How far can we go with TC?

The scalability of TC was studied in Sect. 4.3 for graphs with up to 1 500 vertices (see also Sect. 4.4 for longer runs, with or without time limit), as in (Eiben et al., 1998) for SAW. In this section, we report additional results for TC in Tables 28–36 for some n values in $[2\,000, 11\,000]$ ⁵ around and outside of the threshold τ_h to try to determine the limits of TC (95 new graphs).

³ However, note that “0.44” is contradictory with Fig. 14 in (Eiben et al., 1998). Indeed, the plot rather suggests 0.16 as already indicated in Table 22.

⁴ For runs without time limit, we only report (mean) values based on 5 executions since no significant differences were observed (on easy instances) with a larger number of runs.

⁵ The graph generator employed to build the graphs is restricted to $n \leq 5\,000$. So, we just modified two constants of the generator to generate instances with $n > 5\,000$.

Tables 28–30 show computational results outside of the phase transition with a time limit of 500 000 moves. All **U** and **E** instances, and **F** graphs where $n \leq 2\,500$, are really easy for TC (since $SR = 1$ in this cases). Note that TC also performs well for $F_{3\,000, 10/n}$ since $SR = 0.68$. The problem becomes harder only on **F** instances from $n = 3\,500$ since the best SR achieved by TC when $n \geq 3\,500$ falls to 0.30. So, Table 30 clearly confirms that **F** graphs are harder than **U** and **E** instances, even outside of τ .

Table 28. **U** graphs: The limits of TC outside of τ_n ($p = 10/n$, 50 runs, 500 000 moves).

n	SR	AMS	σ_{AMS}	T (s)	σ_T (s)
2 000	1	28 506.6	12 760.3	1.3	< 1
2 500	1	26 098.0	9 887.6	1.6	< 1
3 000	1	43 744.6	13 434.0	3.2	< 1
3 500	1	60 434.0	23 081.1	2.0	< 1
4 000	1	69 174.7	25 391.7	2.8	1.0
4 500	1	67 340.3	26 422.2	3.3	1.3
5 000	1	82 123.8	29 815.0	4.5	1.5
5 500	1	88 009.3	26 729.8	4.6	1.4
6 000	1	104 856.0	32 455.8	6.9	2.1
6 500	1	122 111.1	42 108.0	8.3	2.8
7 000	1	123 161.9	41 666.6	8.4	2.8
7 500	1	167 213.7	57 518.7	13.1	4.5
8 000	1	168 917.2	56 505.6	13.2	4.4
8 500	1	170 589.2	46 071.5	13.9	3.5
9 000	1	216 444.0	70 894.6	19.3	6.2
9 500	1	221 415.8	71 346.4	22.0	6.8
10 000	1	199 860.9	69 193.1	19.5	6.7
10 500	1	223 878.2	68 928.2	22.8	6.9
11 000	1	264 433.7	78 143.3	28.5	8.3

Table 29. **E** graphs: The limits of TC outside of τ_n ($p = 10/n$, 50 runs, 500 000 moves).

n	SR	AMS	σ_{AMS}	T (s)	σ_T (s)
2 000	1	31 544.8	17 327.4	1.5	< 1
2 500	1	37 282.5	17 758.8	2.2	1.1
3 000	1	41 050.8	16 714.3	2.8	1.2
3 500	1	59 544.4	19 657.7	2.4	< 1
4 000	1	66 063.5	26 473.9	2.6	< 1
4 500	1	69 276.8	25 403.4	3.2	1.1
5 000	1	101 027.9	34 619.6	5.4	1.8
5 500	1	99 081.7	33 051.4	5.5	1.8
6 000	1	109 455.7	44 881.8	6.6	2.6
6 500	1	121 805.9	36 992.5	7.7	2.2
7 000	1	123 962.3	43 498.6	8.6	2.9
7 500	1	123 982.2	45 346.8	9.1	3.1
8 000	1	145 698.6	45 759.8	11.8	3.6
8 500	1	172 399.4	54 661.4	14.7	4.5
9 000	1	185 468.3	53 877.2	16.9	4.8
9 500	1	215 814.2	69 888.2	20.7	6.3
10 000	1	211 838.6	71 073.1	21.5	7.0
10 500	1	218 459.6	59 538.2	22.1	6.1
11 000	1	268 026.0	95 549.2	29.8	10.3

Table 30. F graphs: The limits of TC outside of τ_h ($p = 10/n$, 50 runs, 500 000 moves).

n	SR	AMS	σ_{AMS}	T (s)	σ_T (s)
2 000	1	144 239.7	89 756.7	11.2	6.9
2 500	1	155 953.6	91 872.9	15.2	8.9
3 000	0.68	270 358.6	129 986.0	7.7	3.7
3 500	0.28	262 949.6	110 083.5	9.2	3.8
4 000	0.28	311 626.9	89 354.3	11.3	3.2
4 500	0.30	312 356.9	106 744.1	15.8	5.4
5 000	0.22	326 451.3	111 325.0	16.8	5.4
5 500	0.28	340 731.7	112 833.7	20.0	6.6
6 000	0.20	388 958.3	83 763.7	24.1	5.1
6 500	0.04	359 221.5	90 857.5	25.0	5.8
7 000	0.04	375 040.5	25 217.5	24.6	1.4
7 500	0.04	456 035.5	24 982.5	36.0	1.9
8 000	0.06	439 747.3	56 787.3	35.6	4.4
8 500	0.04	334 201.0	78 082.0	27.6	6.4
9 000	0	–	–	–	–
9 500	0	–	–	–	–
10 000	0	–	–	–	–
10 500	0	–	–	–	–
11 000	0	–	–	–	–

Tables 31–32 shows results for “longer” runs, with a time limit of 1 000 000 moves (Table 31) or without time limit (Table 32), to achieve $SR = 1$ on the hardest F instances from Table 30. One observes that a solution is always found but, contrary to U and E instances, the computation effort required for 3-coloring large F graphs properly can be very high (up to more than 59 million moves in average).

Table 31. Long TC runs outside of τ_h ($p = 10/n$) on the hardest F instances from Table 30 where $SR < 1$ (25 runs, 1 000 000 moves).

n	SR_s	SR_l	AMS_l	σ_{AMS_l}	T_l (s)	σ_{T_l} (s)
3 000	0.68	1	395 950.3	272 423.2	31.6	21.7
3 500	0.28	0.44	425 406.7	241 345.6	14.9	8.3
4 000	0.28	0.44	568 688.4	269 936.2	23.7	11.2
4 500	0.30	0.80	518 964.3	174 691.5	24.7	8.3
5 000	0.22	0.48	589 879.2	266 224.1	31.2	14.1
5 500	0.28	0.48	606 378.3	217 482.9	34.0	12.1
6 000	0.20	0.52	623 241.4	194 497.9	38.3	12.1
6 500	0.04	0.12	630 536.3	139 949.1	40.6	8.8
7 000	0.04	0.24	789 536.0	159 775.2	55.1	10.8
7 500	0.04	0.12	623 690.0	189 473.5	45.3	13.7
8 000	0.06	0.24	612 446.7	214 622.9	48.1	16.8
8 500	0.04	0.16	703 282.3	133 815.8	59.1	11.0
9 000	0	0.08	684 127.5	216 838.5	64.5	20.5
9 500	0	0	–	–	–	–
10 000	0	0	–	–	–	–
10 500	0	0.04	787 970.0	0.0	81.7	0.0
11 000	0	0	–	–	–	–

Table 32. Achieving $SR = 1$ with TC on the hardest F instances from Table 31 where $SR_l < 1$ (5 runs, without time limit).

n	SR_l	SR_∞	AMS_∞	σ_{AMS_∞}	T_∞ (s)	σ_{T_∞} (s)
3 500	0.44	1	1 423 879.8	672 368.9	52.1	24.5
4 000	0.44	1	1 161 122.8	801 233.9	46.0	31.8
4 500	0.8	1	700 484.6	380 628.5	35.2	19.1
5 000	0.48	1	1 315 440.6	766 691.4	65.5	37.9
5 500	0.48	1	1 011 681.4	933 814.3	55.7	49.8
6 000	0.52	1	1 468 846.8	759 937.0	90.0	46.6
6 500	0.12	1	4 705 684.0	2 950 583.4	291.6	182.9
7 000	0.24	1	3 781 609.5	1 832 451.7	259.3	125.5
7 500	0.12	1	7 628 363.0	8 251 686.1	583.0	630.0
8 000	0.24	1	1 522 375.0	721 937.9	122.0	58.1
8 500	0.16	1	2 118 416.3	1 432 638.0	182.6	123.1
9 000	0.08	1	3 428 184.8	2 060 651.2	301.4	179.8
9 500	0	1	12 454 689.0	4 959 205.0	1 160.5	461.9
10 000	0	1	59 920 576.0	50 207 203.5	5 870.1	4 909.1
10 500	0.04	1	6 780 762.5	875 675.5	690.2	90.5
11 000	0	1	10 497 934.0	5 181 142.3	1 103.2	546.2

Tables 33–34 show computational results around the phase transition for U and E instances within a time limit of 1 000 000 moves. Note that no result is reported here (i.e. around τ_h) for the F graphs since, as already showed in Table 24 (Sect. 4.3), TC cannot solve such instances once $n \geq 1 250$ within the time limit of 1 000 000 moves. Indeed, Table 27 (Sect. 4.4) indicates that TC needs more than 4 billion moves (about 126 hours) to solve $F_{1500, 8/n}$. This seems to indicate that, for F graphs around τ_h , $F_{1500, 8/n}$ would be the largest graph that can be colored by TC.

Table 33. U graphs: The limits of TC around τ_h ($\rho = 8/n$, 25 runs, 1 000 000 moves).

n	SR	AMS	σ_{AMS}	T (s)	σ_T (s)
2 000	1	312 539.2	160 692.6	8.0	3.9
2 500	1	474 737.1	204 351.5	24.0	9.9
3 000	1	328 232.3	137 527.3	12.8	5.2
3 500	0.72	556 514.9	230 236.4	11.2	4.7
4 000	0.68	689 607.9	182 578.8	16.3	4.2
4 500	0.6	610 822.6	205 490.8	17.6	6.0
5 000	0.52	680 168.1	244 471.0	23.6	8.5
5 500	0.68	604 993.5	200 738.5	18.8	6.2
6 000	0.36	739 870.9	120 713.8	28.3	4.6
6 500	0.28	854 454.0	104 911.0	30.4	4.0
7 000	0.12	762 356.7	152 977.1	30.0	5.9
7 500	0.04	947 253.0	0.0	37.7	0.0
8 000	0.08	897 777.0	2 250.0	44.3	< 1
8 500	0.32	811 344.4	152 395.1	41.5	8.0
9 000	0.08	858 772.5	70 563.5	45.1	3.1
9 500	0.04	872 204.0	0.0	50.8	0.0
10 000	0.04	790 561.0	0.0	44.1	0.0
10 500	0.04	915 827.0	0.0	54.3	0.0
11 000	0	–	–	–	–

Table 34. E graphs: The limits of TC around τ_h ($p = 8/n$, 25 runs, 1 000 000 moves).

n	SR	AMS	σ_{AMS}	T (s)	σ_T (s)
2 000	0.96	504 763.3	212 903.3	12.2	5.2
2 500	0.92	426 471.0	203 155.8	12.9	6.3
3 000	0.80	493 869.6	212 989.6	10.1	4.4
3 500	0.56	626 114.4	214 233.8	15.4	5.2
4 000	0.60	540 508.9	169 444.7	15.1	5.0
4 500	0.64	569 527.4	187 684.6	18.1	5.7
5 000	0.28	638 217.7	206 569.0	22.7	7.6
5 500	0.64	584 719.5	169 439.6	18.8	5.5
6 000	0.36	755 428.4	179 760.4	26.7	6.1
6 500	0.16	759 410.3	38 527.2	27.8	1.5
7 000	0.16	764 425.5	170 939.2	30.1	7.0
7 500	0.16	689 463.3	78 321.7	27.3	4.0
8 000	0.16	840 018.0	154 616.0	39.5	7.9
8 500	0.08	884 325.0	10 851.0	43.2	< 1
9 000	0.12	761 693.7	155 293.8	45.2	6.9
9 500	0.04	947 668.0	0.0	46.7	0.0
10 000	0.04	857 340.0	0.0	49.0	0.0
10 500	0	–	–	–	–
11 000	0	–	–	–	–

According to Table 33, TC still always solves easily U graphs around τ_h up to $n = 3\,000$ since $SR = 1$ in these cases. Furthermore, TC also performs quite well on larger U instances since $SR \geq 0.52$ for n up to 5 500. E graphs (see Table 34) start here to be a little bit harder than U instances since TC never reached $SR = 1$ but it performs well up to $n = 5\,500$ ($SR \geq 0.56$ except for $E_{5\,000, 8/n}$). The performance of TC falls below 0.5 only for the largest graphs ($n \geq 6\,000$ and for $E_{5\,000, 8/n}$).

Table 35. Achieving $SR = 1$ with TC around τ_h ($p = 8/n$) on the hardest U instances from Table 33 (5 runs, without time limit).

n	SR_l	SR_∞	AMS_∞	σ_{AMS_∞}	T_∞ (s)	σ_{T_∞} (s)
3 500	0.72	1	674 481.0	524 132.0	16.3	12.5
4 000	0.68	1	718 282.2	577 635.5	17.4	13.0
4 500	0.60	1	735 476.4	395 813.7	20.3	10.9
5 000	0.52	1	1 299 003.3	1 050 951.6	29.2	17.2
5 500	0.68	1	1 377 980.4	406 716.0	40.5	12.1
6 000	0.36	1	1 639 610.8	554 939.4	55.5	21.0
6 500	0.28	1	1 887 605.3	929 657.3	69.1	32.7
7 000	0.12	1	1 958 313.0	753 376.6	73.4	27.1
7 500	0.04	1	3 541 162.0	2 309 180.9	126.6	71.1
8 000	0.08	1	2 359 020.8	1 947 452.6	101.0	78.0
8 500	0.32	1	2 543 023.5	1 329 579.5	124.8	63.8
9 000	0.08	1	2 937 435.0	1 129 824.7	149.7	59.1
9 500	0.04	1	2 407 969.5	975 149.0	129.9	53.7
10 000	0.04	1	2 969 634.0	1 495 455.2	180.6	91.7
10 500	0.04	1	4 426 329.0	3 903 536.4	246.4	207.3
11 000	0	1	4 877 196.0	2 224 861.0	295.1	130.7

Table 36. Achieving $SR = 1$ with TC around $\tau_n(p = 8/n)$ on the hardest E instances from Table 34 (5 runs, without time limit).

n	SR_l	SR_∞	AMS_∞	σ_{AMS_∞}	T_∞ (s)	σ_{T_∞} (s)
2 000	0.96	1	511 059.2	277 824.1	13.2	6.9
2 500	0.92	1	464 184.2	343 839.3	14.8	10.8
3 000	0.80	1	687 144.4	310 374.7	27.1	12.0
3 500	0.56	1	1 032 754.4	823 348.6	25.3	19.6
4 000	0.60	1	868 927.4	378 278.6	25.2	10.9
4 500	0.64	1	844 836.2	481 748.7	26.3	14.5
5 000	0.28	1	2 097 527.5	1 047 561.4	72.5	36.1
5 500	0.64	1	2 100 852.5	894 367.8	57.0	22.7
6 000	0.36	1	1 144 047.0	245 704.9	40.1	9.1
6 500	0.16	1	2 123 158.3	1 409 654.3	84.1	56.6
7 000	0.16	1	1 969 999.4	1 149 751.0	81.9	46.0
7 500	0.16	1	2 247 856.3	763 401.7	91.9	30.0
8 000	0.16	1	1 997 386.0	1 203 102.3	86.5	52.9
8 500	0.08	1	3 118 057.0	1 693 211.6	151.4	81.2
9 000	0.12	1	3 243 706.3	3 819 214.4	175.5	204.6
9 500	0.04	1	3 269 792.5	527 271.4	165.7	24.3
10 000	0.04	1	3 582 580.8	1 845 239.8	196.5	98.6
10 500	0	1	4 844 833.0	1 346 013.9	299.5	82.8
11 000	0	1	4 904 942.0	1 222 402.3	294.1	80.7

Tables 35–36 show results for runs without time limit on the graphs from Tables 33–34 where $SR < 1$. One observes that a solution is always found for each run of TC, even for the largest instances with 11 000 vertices. This indicates that TC is probably able to color U and E graphs with much larger n , even around the phase transition.

5. CONCLUSIONS

We present an experimental investigation of a simple tabu search algorithm for coloring random 3-colorable graphs, studying three well-known classes of graphs (Uniform, Equipartite, and Flat) outside of or around the phase transition thresholds. The main findings of this study can be summarized as follows.

Outside of the phase transition thresholds

The simple tabu search algorithm can color any graph (U, E, F) with $200 \leq n \leq 11\,000$ vertices at each run. Moreover, as already observed in other studies, F graphs are more difficult to color than U and E graphs. More precisely:

- For the U and E classes, any graph with up to 11 000 vertices can very easily be colored within 500 000 moves (less than 30 seconds in average). This suggests that TC is probably able to color much larger ($n \gg 11\,000$) U and E graphs within reasonable time.
- For the F class, a solution can always be found for graphs with $n \leq 3\,000$ in average within 1 million moves (less than 60 seconds). Larger graphs with $3\,500 \leq n \leq 11\,000$ can also always be colored if more computing time is allowed. Typically this can be achieved in average with 60 millions of moves (about 1.5 hours).

Around the phase transition thresholds

The simple tabu search algorithm can color any U and E graph with $200 \leq n \leq 11\,000$ vertices at each run. E graphs are a little more difficult to color than U graphs. It is very difficult to color F graphs with more than 1 500 vertices. More precisely:

- For the U and E classes, any graph with up to 11 000 vertices can be colored in average within 5 million moves (less than 5 minutes). This suggests that TC is probably able to color still larger ($n \gg 11\,000$) U and E graphs within reasonable time.
- For the F class, with a time limit of 1 million moves (a few seconds), a proper 3-coloring can always be found for graphs with up to 500 vertices, a solution can occasionally be found for graphs with $500 < n \leq 1\,000$. F graphs with up to 1 500 vertices can also always be colored if no time limit is imposed. However, this may require up to more than 4 billion moves (about 126 hours). This suggests that F graphs larger than 1 500 vertices around the phase transition thresholds constitute a real challenge for TC, but very probably for many (k -)coloring algorithms.

Phase transition thresholds

Finally, concerning the different phase transition thresholds reported in the literature, the experimental results coincide globally well with what is predicted by these thresholds as to the relative hardness of a given graph. Nevertheless, it is observed that the threshold τ_h proposed in (Eiben, van der Hauw, & van Hemert, 1998) is better suited to locate the phase transitions compared with other τ measures. To be more precise, the lower bound of τ_h ($7/n$) seems more adequate for U and E instances while the whole interval ($7/n \leq p \leq 8/n$) remains valid for (sufficiently large) F graphs. Moreover, a ranking among these thresholds is proposed based on the computational observations in Sect. 4.2.

ACKNOWLEDGEMENTS

We would like to thank the reviewers of the paper for their useful comments. This work was partially supported by the French “Pays de la Loire” region (MILES and RadaPop projects, Chair of Excellence), the French Research Ministry, and the French “Centre National de la Recherche Scientifique” (CNRS, GdR RO project).

REFERENCES

- Alekseev, V. E., Boliac, R., Korobitsyn, D. V., & Lozin, V. V. (2007). NP-hard graph problems and boundary classes of graphs. *Theoretical Computer Science*, 389(1-2), 219–236.
- Barbosa, V. C., & Ferreira, R. G. (2004). On the phase transitions of graph coloring and independent sets. *Physica A: Statistical Mechanics and its Applications*, 343, 401–423.
- Bollobás, B. (editor) (2001). *Random graphs*. Volume 73 of Cambridge Studies in Advanced Mathematics, second edition. Cambridge: Cambridge University Press.
- Braunstein, A., Mulet, R., Pagnani, A., Weigt, M., & Zecchina, R. (2003). Polynomial iterative algorithms for coloring and analyzing random graphs. *Physical Review E*, DOI 10.1103/PhysRevE.68.036702.
- Brélaz, D. (1979). New methods to color the vertices of a graph. *Communications of the Association for Computing Machinery*, 22(4), 251–256.

- Cheeseman, P., Kanefsky, B., & Taylor, W. M. (1991). Where the really hard problems are. In J. Mylopoulos, & R. Reiter (editors), *Twelfth International Joint Conference on Artificial Intelligence* (volume 1, pp. 331–337). San Mateo: Morgan Kaufmann.
- Culberson, J. C., & Gent, I. P. (2001). Frozen development in graph coloring. *Theoretical Computer Science*, 265(1-2), 227–264.
- Dorne, R., & Hao, J. K. (1998). Tabu search for graph coloring, T-colorings and set T-colorings. In S. Voss, S. Martello, I. H. Osman, & C. Roucairol (editors), *Meta-Heuristics: Advances and Trends in Local Search Paradigms for Optimization* (pp. 77–92). Dordrecht: Kluwer.
- Dubois, O., Monasson, R., Selman, B., & Zecchina, R. (editors) (2001). *Theoretical Computer Science*, 265 (1-2).
- Eiben, A. E., van der Hauw, J. K., & van Hemert, J. I. (1998). Graph coloring with adaptive evolutionary algorithms. *Journal of Heuristics*, 4(1), 24–46.
- Erben, W. (2001). A grouping genetic algorithm for graph colouring and exam timetabling. In E. Burke, & W. Erben (editors), *Practice and Theory of Automated Timetabling III* (volume 2079 of Lecture Notes in Computer Science, pp. 132–156). Berlin: Springer.
- Fleurent, C., & Ferland, J. (1996a). Genetic and hybrid algorithms for graph coloring. *Annals of Operations Research*, 63(3), 437–461.
- Fleurent, C., & Ferland, J. (1996b). Object-oriented implementation of heuristic search methods for graph coloring, maximum clique, and satisfiability. In D. S. Johnson, & M. A. Trick (editors), *Cliques, Coloring, and Satisfiability* (volume 26 of DIMACS Series in Discrete Mathematics and Theoretical Computer Science, pp. 619–652). New Providence: American Mathematical Society.
- Garey, M. R., & Johnson, D. S. (1979). *Computers and Intractability: A Guide to the Theory of NP-Completeness*. San Francisco: W. H. Freeman and Company.
- Gent, I. P., MacIntyre, E., Prosser, P., & Walsh, T. (1996). The constrainedness of search. In W. J. Clancey, D. Weld, H. E. Shrobe, & T. E. Senator (editors), *Thirteenth National Conference on Artificial Intelligence* (volume 1, pp. 246–252). Menlo Park: AAAI Press.
- Glover, F. W., & Laguna M. (1997). *Tabu Search*. Dordrecht: Kluwer.
- Hartmann, A. K., & Weigt, M. (2005). *Phase Transitions in Combinatorial Optimization Problems: Basics, Algorithms and Statistical Mechanics*. New York: Wiley.
- Hertz, A., & de Werra, D. (1987). Using tabu search techniques for graph coloring. *Computing*, 39(4), 345–351.
- Hogg, T., Huberman, B. A., & Williams, C. P. (editors) (1996). *Artificial Intelligence*, 81(1-2).
- Karp, R. M. (1972). Reducibility among combinatorial problems. In R. E. Miller, & J. W. Thatcher (editors), *Complexity of Computer Computations* (pp. 85–103). New York: Plenum.
- Kochol, M., Lozin, V., & Randerath, B. (2003). The 3-colorability problem on graphs with maximum degree four. *SIAM Journal on Computing*, 32(5), 1128–1139.
- Krivelevich, M., & Sudakov, B. (2006). Pseudo-random graphs. In E. Gyori, G. O. H. Katona, & L. Lovász (editors), *More sets, graphs and numbers* (volume 15 of Bolyai Society Mathematical Studies, pp. 199–262). Berlin: Springer.

Krzakała, F., Pagnani, A., & Weigt, M. (2004). Threshold values, stability analysis, and high- q asymptotics for the coloring problem on random graphs. *Physical Review E*, DOI 10.1103/PhysRevE.70.046705.

Maffray, F., & Preissmann, M. (1996). On the NP-completeness of the k -colorability problem for triangle-free graphs. *Discrete Mathematics*, 162(1-3), 313–317.

Monasson, R., Zecchina, R., Kirkpatrick, S., Selman, B., & Troyansky, L. (1999). Determining computational complexity from characteristic ‘phase transitions’. *Nature*, 400, 133–137.

Petford, A. D., & Welsh, D. J. A. (1989). A randomized 3-colouring algorithm. *Discrete Mathematics*, 74(1-2), 253–261.

Zdeborová, L., & Krzakała, F. (2007). Phase transitions in the coloring of random graphs. *Physical Review E*, DOI 10.1103/PhysRevE.76.031131.