

Conflict Optimization for Binary CSP Applied to Minimum Partition into Plane Subgraphs and Graph Coloring

LOÏC CROMBEZ, LIMOS, Université Clermont Auvergne, France

GUILHERME D. DA FONSECA, LIS, Aix-Marseille Université, France

FLORIAN FONTAN, Independent Researcher, France

YAN GERARD, LIMOS, Université Clermont Auvergne, France

ALDO GONZALEZ-LORENZO, LIS, Aix-Marseille Université, France

PASCAL LAFOURCADE, LIMOS, Université Clermont Auvergne, France

LUC LIBRALESSO, LIMOS, Université Clermont Auvergne, France

BENJAMIN MOMÈGE, Independent Researcher, France

JACK SPALDING-JAMIESON, David R. Cheriton School of Computer Science, University of Waterloo, Canada

BRANDON ZHANG, Independent Researcher, Canada

DA WEI ZHENG, Department of Computer Science, University of Illinois at Urbana-Champaign, USA

CG:SHOP is an annual geometric optimization challenge and the 2023 edition proposed the problem of coloring a certain geometric graph defined by line segments. Surprisingly, the top three teams used the same technique, called conflict optimization. This technique has been introduced in the 2020 edition of the challenge, to solve a coordinated motion planning problem. In this paper, we present the technique in the more general framework of binary constraint satisfaction problems (binary CSP). Then, the top three teams describe their different implementations of the same underlying strategy. We evaluate the performance of those implementations to vertex color not only geometric graphs, but also other types of graphs.

CCS Concepts: • **Theory of computation** → **Design and analysis of algorithms**; *Computational Geometry*.

Additional Key Words and Phrases: CG:SHOP 2022, Constraint Satisfaction Problems, Graph Coloring, DIMACS

1 INTRODUCTION

The CG:SHOP challenge (Computational Geometry: Solving Hard Optimization Problems) is an annual geometric optimization competition, whose first edition took place in 2019. The 2022 edition proposed a problem called *minimum partition into plane subgraphs*. The input is a graph G embedded in the plane with edges drawn as straight line segments, and the goal is to partition the set of edges into a small number of plane graphs (Fig. 1) [5]. This goal can be formulated as a vertex coloring problem on a graph G' defined as follows. The vertices of G' are the segments defining the edges of G , and the edges of G' correspond to pairs of *crossing* segments (segments that intersect only at a common endpoint are not considered crossing).

Authors' addresses: Loïc Crombez, , LIMOS, Université Clermont Auvergne, Clermont-Fd, France; Guilherme D. da Fonseca, guilherme.fonseca@lis-lab.fr, LIS, Aix-Marseille Université, , France; Florian Fontan, dev@florian-fontan.fr, Independent Researcher, Paris, France; Yan Gerard, yan.gerard@uca.fr, LIMOS, Université Clermont Auvergne, Clermont-Fd, France; Aldo Gonzalez-Lorenzo, aldo.gonzalez-lorenzo@univ-amu.fr, LIS, Aix-Marseille Université, , France; Pascal Lafourcade, pascal.lafourcade@uca.fr, LIMOS, Université Clermont Auvergne, Clermont-Fd, France; Luc Libralesso, luc.libralesso@uca.fr, LIMOS, Université Clermont Auvergne, Clermont-Fd, France; Benjamin Momège, ..., Independent Researcher, Clermont-Ferrand, France; Jack Spalding-Jamieson, , David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, Canada; Brandon Zhang, , Independent Researcher, Vancouver, Canada; Da Wei Zheng, , Department of Computer Science, University of Illinois at Urbana-Champaign, IL, USA.

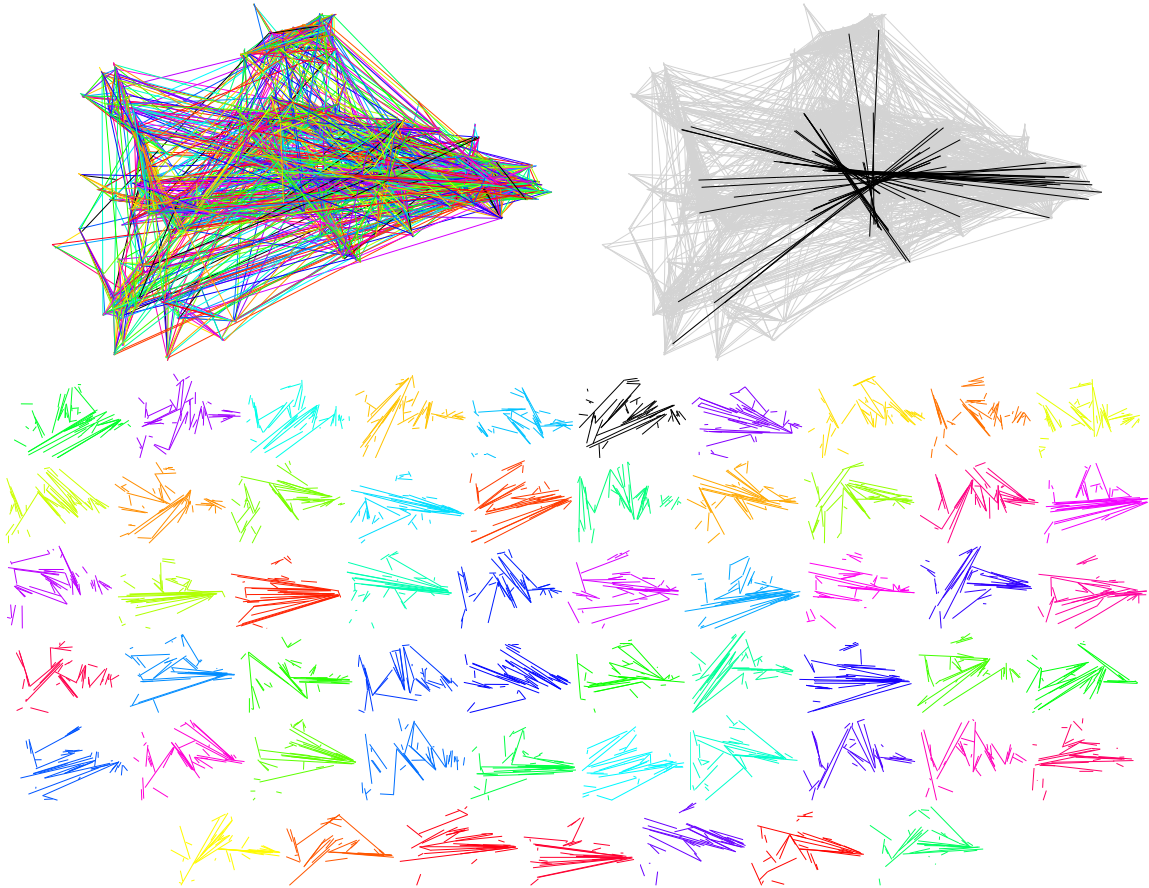


Fig. 1. A partition of the input graph of the CG:SHOP2022 instance vispecn2518 into 57 plane graphs. It is the smallest instance of the challenge with 2518 segments. On top left, you see all 57 colors together. On top right, you see a clique of size 57, hence the solution is optimal. Each of the 57 colors is then presented in small figures.

The three top-ranking teams (Lasa, Gitastrophe, and Shadoks) on the CG:SHOP 2022 challenge all used a common approach called *conflict optimization*. Conflict optimization is a technique used by Shadoks to obtain the first place in the CG:SHOP 2021 challenge for low-makespan coordinated motion planning [3], and the main ideas of the technique lent themselves well to the 2022 challenge. Next, we describe the conflict optimizer as a metaheuristic to solve constraint satisfaction problems (CSP) [27]. We start by describing a CSP.

A CSP is a triple of

- variables $X = (x_1, \dots, x_n)$,
- domains $\mathcal{D} = (D_1, \dots, D_n)$, and
- constraints \mathcal{R} .

Each variable x_i must be assigned a *value* in the corresponding domain D_i such that all constraints are satisfied. In general, the constraints may forbid arbitrary subsets of values. We restrict our attention to a particular type of constraints (*binary CSP*), which only involve pairs of assignments. A *partial evaluation* is an assignment of a subset of the variables,

called *evaluated*, with the remaining variables called *non-evaluated*. All constraints involving a non-evaluated variable are satisfied by default. We only consider assignments and partial assignments that satisfy all constraints.

The conflict optimizer iteratively modifies a partial evaluation with the goal of emptying the set S of non-evaluated variables, at which point it stops. At each step, a variable x_i is removed from S . If there exists a value $x \in D_i$ that satisfies all constraints, then we assign the value x to the variable x_i . Otherwise, we proceed as follows. For each possible value $x \in D_i$, we consider the set $K(i, x)$ of variables (other than x_i) that are part of constraints violated by the assignment $x_i = x$. We assign to x_i the value x that minimizes

$$\sum_{x_j \in K(i, x)} w(j),$$

where $w(j)$ is a weight function to be described later. The variables $x_j \in K(i, x)$ become non-evaluated and added to S .

The weight function should be such that $w(j)$ increases each time x_j is added to S , in order to avoid loops that keep moving the same variables back and forth from S . Let $q(j)$ be the number of times x_j became non-evaluated. A possible weight function is $w(j) = q(j)$. More generally, we can have $w(j) = q(j)^p$ for some exponent p (typically between 1 and 2). Of course, several details of the conflict optimizer are left open. For example, which element to choose from S , whether some random noise should be added to w , and the decision to restart the procedure from scratch after a certain time.

The CSP as is, does not apply to optimization problems. We can however, impose a maximum value of the objective function in order to obtain a CSP. In particular, if all constraints of the optimization problem are binary and the problem is min max, we obtain a binary CSP. Furthermore, in the min max case, we can start from a feasible solution that we want to improve and to target objective value k and obtain a partial assignment to CSP by setting all variables x_i with cost greater than k as non-evaluated. Also, the domain D_i of each variable x_i is reduced to contain only values whose cost is at most k .

The conflict optimizer was introduced in a coordinated motion planning setting. In that setting, the variables are the robots, the domains are their paths (of length at most k) and the constraints forbid collisions between two paths. In the graph coloring setting, the domains are the k colors of the vertices and the constraints forbid adjacent vertices from having the same color. The conflict optimizer can be adapted to non-binary CSP, but in that case multiple variables may be unassigned for a single violated constraint. The strategy has some resemblance to the similarly named *min-conflicts algorithm* [22], but notable differences are that a partial evaluation is kept instead of an invalid evaluation and the weight function that changes over time.

While the conflict optimization strategy is simple, there are different ways to apply it to the graph coloring problem. The goal of the paper is to present how the top three teams applied it or complemented it with additional strategies. We compare the relative benefits of each variant instances given in the CG:SHOP 2022 challenge. We also compare them to baselines on some instances issued from graph coloring benchmarks.

The paper is organized as follows. Section 2 presents the details of the conflict optimization strategy applied to graph coloring. In the three sections that follow, the three teams Lasa, Gitastrophe, and Shadoks present the different parameters and modified strategies that they used to make the algorithm more efficient for the CG:SHOP 2022 challenge. The last section is devoted to the experimental results.

1.1 Literature Review

The study of graph coloring goes back to the 4-color problem (1852) and it has been intensively studied since the 1970s (see [14, 17] for surveys). Many heuristics have been proposed [8, 12, 19, 21], as well as exact algorithms [4, 10, 18]. We briefly present two classes of algorithms: greedy algorithms and exact algorithms.

Greedy algorithms. These algorithms are used to find good quality initial solutions in a short amount of time. The classic greedy heuristic considers the vertices in arbitrary order and colors each vertex with the smallest non-conflicting color. The two most famous modern greedy heuristics are *DSATUR* [2] and *Recursive Largest First (RLF)* [16]. At each step (until all vertices are colored), *DSATUR* selects the vertex v that has the largest number of different colors in its neighbourhood. Ties are broken by selecting the vertex with maximum degree. The vertex v is colored with the smallest non-conflicting color. *RLF* searches for a large independent set I , assigns the vertices I the same color, removes I from G' , and repeats until all vertices are colored.

Exact algorithms. Some exact methods use a branch-and-bound strategy, for example extending the *DSATUR* heuristic by allowing it to backtrack [6, 24]. Another type of exact method (branch-and-cut-and-price) decomposes the vertex coloring problem into an iterative resolution of two sub-problems [7, 11, 20]. The “master problem” maintains a small set of valid colors using a set-covering formulation. The “pricing problem” finds a new valid coloring that is promising by solving a maximum weight independent set problem. Exact algorithms are usually able to find the optimal coloring for graphs with a few hundred vertices. However, even the smallest CG:SHOP 2022 competition instances involve at least a few thousands vertices.

2 CONFLICT OPTIMIZATION FOR GRAPH COLORING

Henceforth, we will only refer to the intersection conflict graph G' induced by the instance. Vertices will refer to the vertices $V(G')$, and edges will refer to the edges $E(G')$. Our goal is to partition the vertices using a minimum set of k color classes $C = \{C_1, \dots, C_k\}$, where no two vertices in the same color class C_i are incident to a common edge.

2.1 Conflict Optimization

We consider the classical problem of coloring the vertices of a graph $G' = (V(G'), E(G'))$. We assume that an initial solution $C = \{C_1, \dots, C_k\}$ has been previously computed (the choice of the initial solution does not seem to impact the quality of the final solution produced by the conflict optimizer). The goal of the conflict optimizer is to reduce the number of colors of C by one unit. When (and if) the conflict optimizer terminates, it will give such a solution. However, after a certain amount of time or a certain situation arrives, we may decide to abort the execution of the conflict optimizer without any solution, and perhaps try again.

Throughout the execution, we maintain a partial coloring, which is a valid coloring for a subset of the vertices. The complementary subset of uncolored vertices is called the *conflict set* and denoted S . The conflict optimizer proceeds as follows:

- (1) Pick a color class C_i to be eliminated. Uncolor all vertices in C_i and make $S \leftarrow C_i$. A valid vertex-coloring is maintained for the set $V(G') \setminus S$. If S is empty, we have a valid vertex coloring of G' which uses one fewer color.

- (2) Pick and remove an element v from S . For each color class, compute the *conflict score* with v . The conflict score of a color class C_j is

$$score(C_j) = f(C_j) \sum_{\substack{u \in C_j \\ (u,v) \in E(G')}} w(u) \quad (1)$$

where the weight $w(u)$ is a variable depending on the the number of times that u has been removed from the conflict set S in previous iterations, and where $f(C_j)$ is a random variable adding randomness in the process.

- (3) Pick the color class C_j with the lowest conflict score. Uncolor all vertices in C_j which are adjacent to v and add those vertices to S . (Notice that if the BDFS option detailed in the later is activated, the algorithm does not put in the conflict set S all the vertices in conflict with S . Some of them are recolored easily so that they do not pass by the conflict set S). Insert v into C_j .
- (4) Repeat steps 2 and 3 until the set S is empty.

The three teams provided different variants of the algorithm by playing with different options of the optimizer.

- (a) The first option is the choice of the initial color C_i to be eliminated at the first step of the loop. It is random for Gitastrophe, and the greedy coloring for Shadoks and Lasa variants.
- (b) The second option is the way to choose the element v from S in step 2. Random for Gitastrophe, a fifo queue for Shadoks, and the element that provides the least total conflict score after its removal for Lasa.
- (c) The third option is the choice of the weight function $w(\cdot)$ defined on the vertices. Different functions have been used, all depending on the parameter $q(u)$ that is defined as the number of times that a vertex u has been removed from S . Lasa used $w(u) = 1 + q(u)$. Gitastrophe used $w(u) = 1 + q(u)^2$. Shadoks used $w(u) = 1 + q(u)^p$ with $p \in [1, 2]$. Shadoks also added a threshold q_{\max} with $w(u) = \infty$ if $q(u) > q_{\max}$. Gitastrophe also had such a threshold, but instead used it as a heuristic to abort the execution and start again.
- (d) The fourth option is the choice of $f(C_i)$. Lasa and Gitastrophe simply set $f(C_i) = 1$, while Shadoks used a Gaussian random variable with average 1 for $f(C_i)$. The right amount of randomness, controlled by the variance σ , has a significant impact on the search time.
- (e) The fifth option is that Shadoks added a Bounded Depth-First Search (BDFS) option which detects vertices that can be recolored easily. These vertices are recolored immediately, instead of entering S , and consequently does not suffer an increase in the value of $q(\cdot)$.

Some extra options are useful in order to drive the computation.

- Restart: The computation is restarted from step 2 if the size of the conflict set S becomes too large because the coloring of $V(G') \setminus S$ has deteriorated too much to come back to a valid coloring.
- Multistart: Shadoks also used a multistart option to restart from step 1.

The different parameters, options and complementary strategies used by each team are precised in the three next sections.

3 LASA TEAM

3.1 Finding Initial Solutions

Lasa team used two approaches to find initial solutions:

- (1) **DSATUR** is the classical graph coloring algorithm presented in Section 1.

- (2) **Orientation greedy** is almost the unique algorithm where the geometry of the segments is used. If segments are almost parallel, it is likely that they do not intersect (thus forming an independent set). This greedy algorithm first sorts the segments by orientation, ranging from $-\frac{\pi}{2}$ to $\frac{\pi}{2}$. For each segment in this order, the algorithm tries to color it using the first available color. If no color has been found, a new color is created for coloring the considered segment. This algorithm is efficient, produces interesting initial solutions and takes into account the specificities of the competition.

3.2 Conflict Optimization

TABUCOL inspired neighbourhood. One classical approach for the vertex coloring involves allowing solutions with conflicting vertices (two adjacent vertices with the same color). It was introduced in 1987 [13] and called TABUCOL. It starts with an initial solution, removes a color (usually the one with the least number of vertices), and assigns uncolored vertices with a new color among the remaining ones. This is likely to lead to some conflicts (*i.e.* two adjacent vertices sharing a same color). The local search scheme selects a conflicting vertex, and tries to swap its color, choosing the new coloring that minimises the number of conflicts. If it reaches a state with no conflict, we obtain a new best-known solution. The process is repeated until the stopping criterion is met. While the original TABUCOL algorithm includes this “tabu-list” mechanism to avoid cycling, it is not always sufficient, and requires some hyper-parameter tuning in order to obtain a good performance on a large variety of instances. To overcome this issue, this neighbourhood is embedded within a conflict optimizer scheme as presented above.

PARTIALCOL inspired neighbourhood. Another local search algorithm solving the vertex coloring problem was introduced in 2008 (PARTIALCOL). This algorithm proposes a new local search scheme that allows partial coloring (thus allowing uncolored vertices). The goal is to minimize the number of uncolored vertices. Similarly to TABUCOL, PARTIALCOL starts with an initial solution, removes one color (unassigning its vertices), and performs local search iterations until no vertex is left uncolored. When coloring a vertex, the adjacent conflicting vertices are uncolored. Then, the algorithm repeats the process until all vertices are colored, or the stopping criterion is met. This neighbourhood was also introduced alongside a tabu-search procedure. The tabu-search scheme is replaced with a conflict-optimization scheme. Note that this neighbourhood was predominantly used by the other teams.

4 GITASTROPHE

4.1 Solution Initialization

The gitastrophe team used the traditional greedy algorithm of Welsh and Powell [28] to obtain initial solutions: order the vertices in decreasing order of degree, and assign each vertex the minimum-label color not used by its neighbors. They attempted to use different orderings for the greedy algorithm, such as sorting by the slope of the line segment associated with each vertex (as the orientation greedy initialization presented in Section 3), and also tried numerous other strategies. Ultimately, after running the solution optimizer for approximately the same amount of time, all initializations resulted in equal number of colors.

4.2 Modifications to the Conflict Optimizer

Taking inspiration from memetic algorithms, which alternate between an intensification and a diversification stage, the algorithm continually switched between a phase using the above conflict score, and one minimizing only the number of

conflicts (i.e. the conflict score of (1) is replaced with $\sum_{u \in C_i, (u,v) \in E(G')} 1$). Each phase lasted for 10^5 iterations. Adding the conflict-minimization phase gave minor improvements to some of the challenge instances.

5 SHADOKS

In this section, we describe the choices used by the Shadoks team for the options described in Section 2.1.

Option (a). We generally chose to eliminate the color with the smallest number of elements. However, if the multistart option is toggled on, then a random color is used each time.

Option (b). The conflict set S is stored in a queue. We tried other strategies, but found that the queue gives the best results.

Option (c). The weight function used is $w(u) = 1 + q(u)^p$, mostly with $p = 1.2$. The effect of the parameter p is drawn in Fig. 2. Notice that in all figures, the number of colors shown is the average of multiple executions of the code using different random seeds.




Fig. 2. Number of colors over time for the instance `vispecn13806` using different values p . The algorithm uses $\sigma = 0.15$, easy vertices, $q_{\max} = 59022$, but does not use the BDFS nor any clique.

If $q(u)$ is larger than a threshold q_{\max} , we fix $w(u) = \infty$ so that the vertex u never reenters S . If at some point we need to color a vertex v that is adjacent to a vertex u of infinite weight in every color class, then we restart the conflict optimizer. When restarting, the coloring is shuffled by moving vertices that fit into multiple color classes.

Looking at Fig. 3, the value of q_{\max} does not seem to have much influence as long as its value is not too small. Throughout the challenge Shadoks almost exclusively used $q_{\max} = 2000 \cdot (75000/m)^2$, where m is the number of vertices. This value roughly ensures a restart every few hours.

If the clique option is toggled on, each vertex u in the largest known clique has $w(u) = \infty$. The impact of the clique option on the computation is shown in Fig. 4. The idea is that since each vertex of the clique must have a different color, it is useless to change their color, we can change the color of the other vertices instead. We used several methods to produce large cliques, including simulated annealing and mixed integer programming.

Option (d). We used the function f as a Gaussian random variable of mean 1 and variance σ . A good default value is $\sigma = 0.15$. The effect of the variance is drawn in Fig. 5. Notice that setting $\sigma = 0$ gives much worse results.




Fig. 3. Number of colors over time with different values of q_{\max} obtained on the instance vispecn13806. Parameters are $\sigma = 0.15$, $p = 1.2$, no clique knowledge, and no BDFS.




Fig. 4. Number of colors over time with and without clique knowledge and BDFS obtained on the instance vispecn13806. Parameters are $\sigma = 0.15$, $p = 1.2$, and $q_{\max} = 1500000$.




Fig. 5. Number of colors over time for the instance vispecn13806 for different values of σ . In both figures the algorithm uses $p = 1.2$, easy vertices, $q_{\max} = 59022$, but does not use the BDFS nor any clique. For $\sigma \geq 0.25$, no solution better than 248 colors is found.

Option (e). The goal of BDFS is to further optimize very good solutions that the conflict optimizer is not able to improve otherwise. Fig. 4 shows the influence of BDFS. While on this figure, the advantages of BDFS cannot be noticed, its use near the end of the challenge improved about 30 solutions.

The *bounded depth-first search* (BDFS) algorithm tries to improve the dequeuing process. The goal is to prevent a vertex from being queued by locally recoloring a bounded number of vertices in the current partial solution. To do so, we perform a local search into the tree of possible ways to color the vertices.

The BDFS algorithm has two parameters: *crossing bound* c_{\max} and *depth* d . In order to recolor a vertex v , BDFS gets the set C of color classes with at most c_{\max} neighbors of v . If a class of C has no neighbor of v , we assign v to C . Otherwise, for each class $C \in C$, BDFS tries to recolor the vertices in C that cross v by recursively calling itself with depth $d - 1$. At depth $d = 0$ the algorithm stops trying to color the vertices.

During the challenge Shadoks used BDFS with parameters $c_{\max} = 3$ and $d = 3$. The depth was increased to 5 (resp. 7) when the number of vertices in the queue was 2 (resp. 1).

Degeneracy order. Given an target number of colors k , we call *easy vertices* a set of vertices Y such that, if the remainder of the vertices of G' are colored using k colors, then we are guaranteed to be able to color all vertices of G' with k colors. This is obtained using the degeneracy order Y . To obtain Y we iteratively remove from the graph a vertex v that has at most $k - 1$ neighbors, appending v to the end of Y . We repeat until no other vertex can be added to Y . Notice that, once we color the remainder of the graph with at least k colors, we can use a greedy coloring for Y in order from last to first without increasing the number of colors used. Removing the easy vertices reduces the total number of vertices, making the conflict optimizer more effective. Shadoks always toggled this option on (the challenge instances contains from 0 to 23% easy vertices).

6 RESULTS

We provide the results of the experiments performed with the code from the three teams on two classes of instances. First, we present the results on some selected CG:SHOP 2022 instances. These instances are intersection graphs of line segments. Second, we execute the code on graphs that are not intersection graphs, namely the classic DIMACS graphs [15], comparing the results of our conflict optimizer implementations to previous solutions.

6.1 CG:SHOP 2022 Instances

We selected 14 instances (out of 225) covering the different types of instances given in the CG:SHOP 2022 challenge. The results are presented in Table 1. At the end of the challenge, 8 colorings computed by Lasa, 11 colorings computed by Gitastrophe, and 23 colorings computed by Shadoks over 225 instances have been proved optimal (their number of colors is equal to the size of a clique).

In order to compare the efficiency of the algorithms, we executed the different implementations on the CG:SHOP instance `vispecn13806`. The edge density of this graph is 19%, the largest clique that we found has 177 vertices and the best coloring found during the challenge uses 218 colors. Notice that `vispecn13806` is the same instance used in other Shadoks experiments in Section 5. We ran the three implementations on three different servers and compared the results shown in Figure 6. For each implementation, the x coordinate is the running time in hours, while the y coordinate is the smallest number of colors found at that time.

6.2 Results on DIMACS Graphs

We tested the implementation of each team on the DIMACS instances [15] to gauge the performance of the conflict optimizer on other classes of graphs. We compared our results to the best known bounds and to the state of the art coloring algorithms HEAD [23] and QACOL [25, 26].

Table 1. Several CG:SHOP 2022 results. We compare the size of the largest known clique to the smallest coloring found by each team on a selection of 14 CG:SHOP 2022 instances.

Instance	Clique	Best	Lasa	Gitastrophe	Shadoks
rvisp5013	46	49	50	49	49
rsqrpecn8051	173	175	177	176	175
vispecn13806	77	218	224	221	218
rsqrp14364	134	136	137	137	136
vispecn19370	169	192	197	194	192
rvisp24116	97	104	110	105	104
visp26405	78	81	83	81	81
sqrp28863	190	190	191	191	190
visp38574	118	133	138	134	133
sqrpecn45700	460	462	465	465	462
reecn51526	308	310	315	312	310
vispecn58391	305	367	380	369	367
vispecn65831	357	439	453	440	439
sqrp72075	264	269	272	271	269

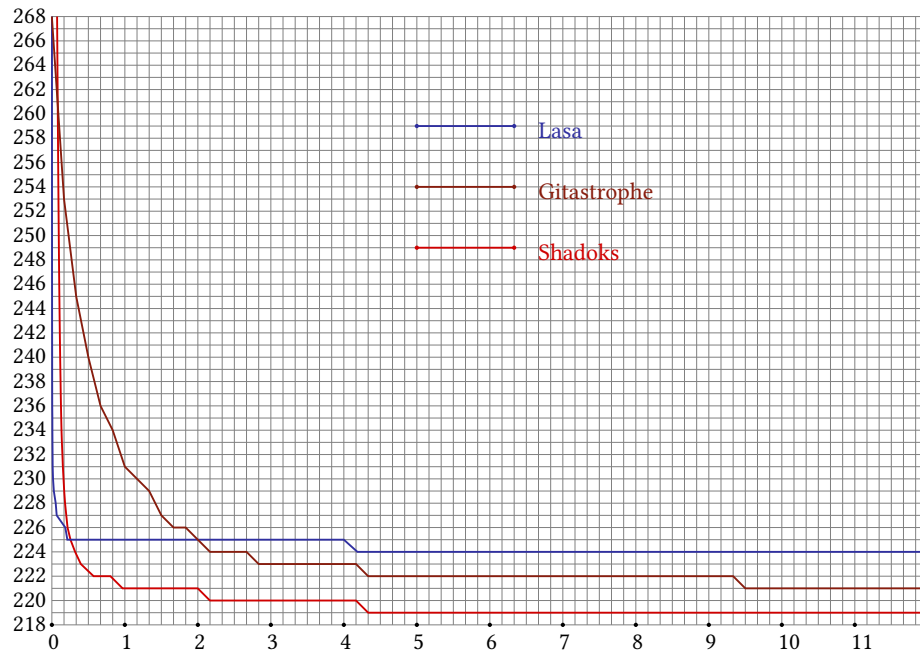


Fig. 6. Number of colors over time (in hours) for the instance vispecn13806.

The time limit for Lasa’s algorithms is 1 hour. CWLS is the Lasa’s conflict optimizer with the neighbourhood presented in TABUCOL [13], while PWLS is the optimizer with the neighbourhood presented in PARTIALCOL [1]. Gitastrophe algorithm ran 10 minutes after which the number of colors no longer decreases. Shadoks algorithm ran for 1 hour without the BDFS option (results with BDFS are worse).

Results are presented in Table 2. We only kept the difficult DIMACS instances. For the other instances, all the results match the best known bounds. The DIMACS instances had comparatively few edges (on the order of thousands or millions); the largest intersection graphs considered in the CG:SHOP challenge had over 1.5 billion edges.

Table 2. Comparison of our method with state-of-the-art graph coloring algorithms. The conflict optimizer underperforms except on the geometric graphs *r** and *dsjr**.

Instance	Best	HEAD [23]	QACOL [25, 26]	Lasa CWLS	Lasa PWLS	Gitastrophe	Shadoks
dsjc250.5	28	28	28	28	29	29	28
dsjc500.1	12	12	12	13	13	13	13
dsjc500.5	47	47	48	49	51	52	50
dsjc500.9	126	126	126	126	130	130	128
dsjc1000.1	20	20	20	21	22	21	21
dsjc1000.5	82	82	82	89	94	93	91
dsjc1000.9	222	222	222	223	240	235	231
r250.5	65	65	65	65	65	65	65
r1000.1c	98	98	98	98	98	98	98
r1000.5	234	245	238	234	234	234	237
dsjr500.1c	84	85	85	85	85	85	85
dsjr500.5	122	-	122	122	122	122	122
le450_25c	25	25	25	26	26	26	26
le450_25d	25	25	25	26	26	26	26
flat300_28_0	28	31	31	31	32	33	32
flat1000_50_0	50	50	-	50	50	91	54
flat1000_60_0	60	60	-	60	92	93	90
flat1000_76_0	81	81	81	88	93	92	90
C2000.5	145	146	145	165	173	173	168
C4000.5	260	266	259	311	320	317	312

We noticed that the conflict optimizer works extremely poorly on random graphs, but it is fast and appears to perform well on geometric graphs (r250.5, r1000.1c, r1000.5, dsjr500.1c and dsjr500.5), matching the best-known results [9]. Interestingly, these geometric graphs are not intersection graphs as in the CG:SHOP challenge, but are generated based on a distance threshold. On the DIMACS graphs, Lasa implementation shows better performance than the other implementations.

7 ACKNOWLEDGMENTS

Shadoks would like to thank H el ene Toussaint, Rapha el Amato, Boris Lonjon, and William Guyot-L enat from LIMOS, as well as the Qarma and TALEP teams and Manuel Bertrand from LIS, who continue to make the computational resources of the LIMOS and LIS clusters available to our research. We would also like to thank the challenge organizers and other competitors for their time, feedback, and making this whole event possible.

The work of Lo ic Crombez has been sponsored by the French government research program “Investissements d’Avenir” through the IDEX-ISITE initiative 16-IDEX-0001 (CAP 20-25). The work of Guilherme D. da Fonseca is supported by the French ANR PRC grant ADDS (ANR-19-CE48-0005). The work of Yan Gerard is supported by the French ANR PRC grants ADDS (ANR-19-CE48-0005), ACTIVmap (ANR-19-CE19-0005) and by the French government IDEX-ISITE initiative 16-IDEX-0001 (CAP 20-25). The work of Aldo Gonzalez-Lorenzo is supported by the French ANR PRC grant COHERENCE4D (ANR-20-CE10-0002). The work of Pascal Lafourcade is supported by the French ANR

Loïc Crombez, Guilherme D. da Fonseca, Florian Fontan, Yan Gerard, Aldo Gonzalez-Lorenzo, Pascal Lafourcade, Luc Libralesso, Benjamin Momège, Jack Spalding-Jamieson, Brandon Zhang, and Da Wei Zheng

PRC grant MobiS5 (ANR-18-CE39-0019), DECRYPT (ANR-18-CE39-0007), SEVERITAS (ANR-20-CE39-0005) and by the French government IDEX-ISITE initiative 16-IDEX-0001 (CAP 20-25). The work of Luc Libralesso is supported by the French ANR PRC grant DECRYPT (ANR-18-CE39-0007).

REFERENCES

- [1] Ivo Blöchliger and Nicolas Zufferey. 2008. A graph coloring heuristic using partial solutions and a reactive tabu scheme. *Computers & Operations Research* 35, 3 (2008), 960–975.
- [2] Daniel Brélaz. 1979. New methods to color the vertices of a graph. *Commun. ACM* 22, 4 (1979), 251–256.
- [3] Loïc Crombez, Guilherme D. da Fonseca, Yan Gerard, Aldo Gonzalez-Lorenzo, Pascal Lafourcade, and Luc Libralesso. 2021. Shadoks Approach to Low-Makespan Coordinated Motion Planning (CG Challenge). In *37th International Symposium on Computational Geometry, SoCG 2021*. 63:1–63:9.
- [4] David Eppstein. 2002. Small maximal independent sets and faster exact graph coloring. *J. Graph Algorithms Appl* 7, 2 (2002), 131–140.
- [5] Sándor P. Fekete, Phillip Keldenich, Dominik Krupke, and Stefan Schirra. 2022. Minimum Partition into Plane Subgraphs: The CG: SHOP Challenge 2022. *CoRR* abs/2203.07444 (2022). arXiv:2203.07444 <https://arxiv.org/abs/2203.07444>
- [6] Fabio Furini, Virginie Gabrel, and Ian-Christopher Ternier. 2017. An Improved DSATUR-Based Branch-and-Bound Algorithm for the Vertex Coloring Problem. *Networks* 69, 1 (2017), 124–141.
- [7] Fabio Furini and Enrico Malaguti. 2012. Exact weighted vertex coloring via branch-and-price. *Discrete Optimization* 9, 2 (2012), 130–136.
- [8] Philippe Galinier and Jin-Kao Hao. 1999. Hybrid evolutionary algorithms for graph coloring. *Journal of combinatorial optimization* 3, 4 (1999), 379–397.
- [9] Olivier Goudet, Cyril Grelier, and Jin-Kao Hao. 2021. A deep learning guided memetic framework for graph coloring problems. arXiv:cs.LG/2109.05948
- [10] Stefano Gualandi and Federico Malucelli. 2012. Exact solution of graph coloring problems via constraint programming and column generation. *INFORMS Journal on Computing* 24, 1 (2012), 81–100.
- [11] Stefano Gualandi and Federico Malucelli. 2012. Exact solution of graph coloring problems via constraint programming and column generation. *INFORMS Journal on Computing* 24, 1 (2012), 81–100.
- [12] Alain Hertz and Dominique de Werra. 1987. Using tabu search techniques for graph coloring. *Computing* 39, 4 (1987), 345–351.
- [13] Alain Hertz and Dominique de Werra. 1987. Using tabu search techniques for graph coloring. *Computing* 39, 4 (1987), 345–351.
- [14] Tommy R. Jensen and Bjarne Toft. 2011. *Graph coloring problems*. John Wiley & Sons.
- [15] David S Johnson and Michael A Trick. 1996. *Cliques, coloring, and satisfiability: second DIMACS implementation challenge, October 11-13, 1993*. Vol. 26. American Mathematical Soc.
- [16] Frank Thomson Leighton. 1979. A graph coloring algorithm for large scheduling problems. *Journal of research of the national bureau of standards* 84, 6 (1979), 489.
- [17] R. M. R. Lewis. 2015. *A Guide to Graph Colouring: Algorithms and Applications* (1st ed.). Springer Publishing Company, Incorporated.
- [18] Corinne Lucet, Florence Mendes, and Aziz Moukrim. 2006. An exact method for graph coloring. *Computers & Operations Research* 33, 8 (2006), 2189–2207.
- [19] David W. Matula, George Marble, and Joel D. Isaacson. 1972. Graph coloring algorithms. In *Graph theory and computing*. Elsevier, 109–122.
- [20] Anuj Mehrotra and Michael A Trick. 1996. A column generation approach for graph coloring. *INFORMS Journal on Computing* 8, 4 (1996), 344–354.
- [21] Isabel Méndez-Díaz and Paula Zabala. 2006. A branch-and-cut algorithm for graph coloring. *Discrete Applied Mathematics* 154, 5 (2006), 826–847.
- [22] Steven Minton, Mark D Johnston, Andrew B Philips, and Philip Laird. 1992. Minimizing conflicts: a heuristic repair method for constraint satisfaction and scheduling problems. *Artificial intelligence* 58, 1-3 (1992), 161–205.
- [23] Laurent Moalic and Alexandre Gondran. 2018. Variations on memetic algorithms for graph coloring problems. *Journal of Heuristics* 24, 1 (2018), 1–24.
- [24] Pablo San Segundo. 2012. A new DSATUR-based algorithm for exact vertex coloring. *Computers & Operations Research* 39, 7 (2012), 1724–1733.
- [25] Olawale Titiloye and Alan Crispin. 2011. Quantum annealing of the graph coloring problem. *Discret. Optim.* 8 (2011), 376–384.
- [26] Olawale Titiloye and Alan Crispin. 2012. Parameter tuning patterns for random graph coloring with quantum annealing. *PLoS one* 7, 11 (2012), e50060.
- [27] Edward P. K. Tsang. 1993. Foundations of constraint satisfaction. In *Computation in cognitive science*.
- [28] D. J. A. Welsh and M. B. Powell. 1967. An upper bound for the chromatic number of a graph and its application to timetabling problems. *Comput. J.* 10, 1 (01 1967), 85–86. <https://doi.org/10.1093/comjnl/10.1.85> arXiv:<https://academic.oup.com/comjnl/article-pdf/10/1/85/1069035/100085.pdf>