



# Modular and Efficient Divide-and-Conquer SAT Solver on Top of the Painless Framework

Ludovic Le Frioux<sup>1,2(✉)</sup>, Souheib Baarir<sup>2,3</sup>, Julien Sopena<sup>2,4</sup>,  
and Fabrice Kordon<sup>2</sup>

<sup>1</sup> LRDE, EPITA, 94270 Le Kremlin-Bicêtre, France  
ludovic@lrde.epita.fr

<sup>2</sup> Sorbonne Université, CNRS, LIP6, UMR 7606, 75005 Paris, France  
{ludovic.le-frioux,souheib.baarir,julien.sopena,fabrice.kordon}@lip6.fr

<sup>3</sup> Université Paris Nanterre, 92000 Nanterre, France

<sup>4</sup> Inria, DELYS Team, 75005 Paris, France

**Abstract.** Over the last decade, parallel SATisfiability solving has been widely studied from both theoretical and practical aspects. There are two main approaches. First, divide-and-conquer (D&C) splits the search space, each solver being in charge of a particular subspace. The second one, portfolio launches multiple solvers in parallel, and the first to find a solution ends the computation. However although D&C based approaches seem to be the natural way to work in parallel, portfolio ones experimentally provide better performances.

An explanation resides on the difficulties to use the native formulation of the SAT problem (*i.e.*, the CNF form) to compute an *a priori* good search space partitioning (*i.e.*, all parallel solvers process their subspaces in comparable computational time). To avoid this, dynamic load balancing of the search subspaces is implemented. Unfortunately, this is difficult to compare load balancing strategies since state-of-the-art SAT solvers appropriately dealing with these aspects are hardly adaptable to various strategies than the ones they have been designed for.

This paper aims at providing a way to overcome this problem by proposing an implementation and evaluation of different types of divide-and-conquer inspired from the literature. These are relying on the **Painless** framework, which provides concurrent facilities to elaborate such parallel SAT solvers. Comparison of the various strategies are then discussed.

**Keywords:** Divide-and-conquer · Parallel satisfiability · SAT solver · Tool

## 1 Introduction

Modern SAT solvers are now able to handle complex problems involving millions of variables and billions of clauses. These tools have been used successfully

© The Author(s) 2019

T. Vojnar and L. Zhang (Eds.): TACAS 2019, Part I, LNCS 11427, pp. 135–151, 2019.

[https://doi.org/10.1007/978-3-030-17462-0\\_8](https://doi.org/10.1007/978-3-030-17462-0_8)

to solve constraints' systems issued from many contexts, such as planning decision [16], hardware and software verification [7], cryptology [23], and computational biology [20], etc.

State-of-the-art complete SAT solvers are based on the well-known Conflict-Driven Clause Learning (CDCL) algorithm [21, 28, 30]. With the emergence of many-core machines, multiple parallelisation strategies have been conducted on these solvers. Mainly, two classes of parallelisation techniques have been studied: divide-and-conquer (D&C) and portfolio. Divide-and-conquer approaches, often based on the guiding path method, decompose recursively and dynamically, the original search space in subspaces that are solved separately by sequential solvers [1, 2, 12, 14, 26, 29]. In the portfolio setting, many sequential SAT solvers compete for the solving of the whole problem [4, 5, 11]. The first to find a solution, or proving the problem to be unsatisfiable ends the computation. Although divide-and-conquer approaches seem to be the natural way to parallelise SAT solving, the outcomes of the parallel track in the annual SAT Competition show that the best state-of-the-art parallel SAT solvers are portfolio ones.

The main problem of divide-and-conquer based approaches is the search space division so that load is balanced over solvers, which is a theoretical hard problem. Since no optimal heuristics has been found, solvers compensate non optimal space division by enabling dynamic load balancing. However, state-of-the-art SAT solvers appropriately dealing with these aspects are hardly adaptable to various strategies than the ones they have been designed for [1, 2, 6]. Hence, it turns out to be very difficult to make fair comparisons between techniques (*i.e.*, using the same basic implementation). Thus, we believe it is difficult to conclude on the (non-) effectiveness of a technique with respect to another one and this may lead to premature abortion of potential good ideas.

This paper tries to solve these problems by proposing a simple, generic, and efficient divide-and-conquer component on top of the **Painless** [18] framework. This component eases the implementation and evaluation of various strategies, without any compromise on efficiency. Main contributions of this paper are the followings:

- an overview of state-of-the-art divide-and-conquer methods;
- a complete divide-and-conquer component that has been integrated to the **Painless** framework;
- a fair experimental evaluation of different types of divide-and-conquer inspired from the literature, and implemented using this component.

These implementations have often similar and sometimes better performances compared with state-of-the-art divide-and-conquer SAT solvers.

Let us outline several results of this work. First, our **Painless** framework is able to support implementation of multiple D&C strategies in parallel solvers. Moreover, we have identified “axes” for customization and adaptation of heuristics. Thus, we foresee it will be much easier to explore next D&C strategies. Second, our best implementation at this stage is comparable in terms of performance, with the best state-of-the-art D&C solvers, which shows our framework's efficiency.

This paper is organized as follows: Sect. 2 introduces useful background to deal with the SAT problem. Section 3 is dedicated to divide-and-conquer based parallel SAT solving. Section 4 explains the mechanism of divide-and-conquer we have implemented in **Painless**. Section 5 analyses the results of our experiments, and Sect. 6 concludes and gives some perspectives.

## 2 Background

A *propositional variable* can have two possible values  $\top$  (True), or  $\perp$  (False). A *literal*  $l$  is a propositional variable ( $x$ ) or its negation ( $\neg x$ ). A *clause*  $\omega$  is a finite disjunction of literals (noted  $\omega = \bigvee_{i=1}^k \ell_i$ ). A clause with a single literal is called *unit clause*. A *conjunctive normal form (CNF) formula*  $\varphi$  is a finite conjunction of clauses (noted  $\varphi = \bigwedge_{i=1}^k \omega_i$ ). For a given formula  $\varphi$ , the set of its variables is noted:  $V_\varphi$ . An *assignment*  $\mathcal{A}$  of variables of  $\varphi$ , is a function  $\mathcal{A} : V_\varphi \rightarrow \{\top, \perp\}$ .  $\mathcal{A}$  is total (complete) when all elements of  $V_\varphi$  have an image by  $\mathcal{A}$ , otherwise it is partial. For a given formula  $\varphi$ , and an assignment  $\mathcal{A}$ , a clause of  $\varphi$  is satisfied when it contains at least one literal evaluating to true, regarding  $\mathcal{A}$ . The formula  $\varphi$  is satisfied by  $\mathcal{A}$  iff  $\forall \omega \in \varphi, \omega$  is satisfied.  $\varphi$  is said to be SAT if there is at least one assignment that makes it satisfiable. It is defined as UNSAT otherwise.

---

**Algorithm 1.** CDCL algorithm

---

```

1 function CDCL()
2    $dl \leftarrow 0$  // Current decision level
3   while not all variables are assigned do
4      $conflict \leftarrow \text{unitPropagation}()$ 
5     if  $conflict$  then
6       if  $dl = 0$  then
7         return  $\perp$  //  $\varphi$  is UNSAT
8       end
9        $\omega \leftarrow \text{conflictAnalysis}()$ 
10       $\text{addLearntClause}(\omega)$ 
11       $dl \leftarrow \text{backjump}(\omega)$ 
12    end
13    else
14       $\text{assignDecisionLiteral}()$ 
15       $dl \leftarrow dl + 1$ 
16    end
17  end
18  return  $\top$  //  $\varphi$  is SAT

```

---

**Conflict Driven Clause Learning.** The majority of the complete state-of-the-art sequential SAT solvers are based on the Conflict Driven Clause Learning (CDCL) algorithm [21, 28, 30], that is an enhancement of the DPLL algorithm [9, 10]. The main components of a CDCL are presented in Algorithm 1.

At each step of the main loop, `unitPropagation`<sup>1</sup> (line 4) is applied on the formula. In case of conflict (line 5), two situations can be observed: the conflict is detected at decision level 0 ( $dl == 0$ ), thus the formula is declared UNSAT (lines 6–7); otherwise, a new asserting clause is derived by the conflict analysis and the algorithm backjumps to the assertion level [21] (lines 8–10). If there is no conflict (lines 11–13), a new decision literal is chosen (heuristically) and the algorithm continues its progression (adding a new decision level:  $dl \leftarrow dl + 1$ ). When all variables are assigned (line 3), the formula is said to be SAT.

**The Learning Mechanism.** The effectiveness of the CDCL lies in the *learning mechanism* (line 10). Each time a conflict is encountered, it is analyzed (`conflictAnalysis` function in Algorithm 1) in order to compute its reasons and derive a *learnt clause*. While present in the system, this clause will avoid the same mistake to be made another time, and therefore allows faster deductions (conflicts/unit propagations).

Since the number of conflicts is very huge (in avg. 5000/s [3]), controlling the size of the database storing learnt clauses is a challenge. It can dramatically affect performance of the `unitPropagation` function. Many strategies and heuristics have been proposed to manage the cleaning of the stored clauses (*e.g.*, the Literal Block Distance (LBD) [3] measure).

### 3 Divide-and-Conquer Based Parallel SAT Solvers


The divide-and-conquer strategy in parallel SAT solving is based on splitting the search space into subspaces that are submitted to different workers. If a subspace is proven SAT then the initial formula is SAT. The formula is UNSAT if all the subspaces are UNSAT. The challenging points of the divide-and-conquer mechanism are: *dividing the search space*, *balancing jobs between workers*, and *exchanging learnt clauses*.

#### 3.1 Dividing the Search Space

This section describes how to create multiple search subspaces for the studied problem, and the heuristics to balance their estimated computational costs.

**Techniques to Divide the Search Space.** To divide the search space, the most often used technique is the *guiding path* [29]. It is a conjunction of literals (called cube) that are assumed by the invoked solver (worker). Let  $\varphi$  be a formula, and  $x \in \mathcal{V}_\varphi$  a variable. Thanks to Shannon decomposition, we can rewrite  $\varphi$  as  $\varphi = (\varphi \wedge x) \vee (\varphi \wedge \neg x)$ . The two guiding paths here are reduced to a single literal:  $(x)$  and  $(\neg x)$ . This principle can be applied recursively on each subspaces to create multiple guiding paths.

<sup>1</sup> The `unitPropagation` function implements the Boolean Constraint Propagation (BCP) procedure that forces (in cascade) the values of the variables in unit clauses [9].



**Fig. 1.** Using guiding path to divide the search space

Figure 1 illustrates such an approach where six subspaces have been created from the original formula. They are issued from the following guiding paths:  $(d \wedge b)$ ,  $(d \wedge \neg b)$ ,  $(\neg d \wedge a \wedge b)$ ,  $(\neg d \wedge a \wedge \neg b)$ ,  $(\neg d \wedge \neg a \wedge x)$ ,  $(\neg d \wedge \neg a \wedge \neg x)$ . The subspaces that have been proven UNSAT, are highlighted with red crosses. The rest of the subspaces are submitted to workers (noted  $w_i$ ).

It is worth noting that other partitioning techniques exist that were initially developed for distributed systems rather than many-cores machines. We can cite the *scattering* [13], and the *xor partitioning* [27] approaches.

**Choosing Division Variables.** Choosing the best *division variable* is a hard problem, requiring the use of heuristics. A good division heuristic should decrease the overall total solving time<sup>2</sup>. Besides, it should create balanced subspaces w.r.t. their solving time: if some subspaces are too easy to solve this will lead to repeatedly asking for new jobs and redividing the search space (phenomenon known as *ping pong effect* [15]).

Division heuristics can be classified in two categories: *look ahead* and *look back*. Look ahead heuristics rely on the possible future behaviour of the solver. Contrariwise, look back heuristics rely on statistics gathered during the past behaviour of the solver. Let us present the most important ones.

*Look Ahead.* In stochastic SAT solving (chapters 5 and 6 in [8]), look ahead heuristics are used to choose the variable implying the biggest number of unit propagations as a decision variable. When using this heuristic for the division, one tries to create the smallest possible subspaces (*i.e.*, with the least unassigned variables). The main difficulty of this technique is the generated cost of applying the unit propagation for the different variables. The so-called “cube-and-conquer” solver presented in [12] relies on such an heuristic.

*Look Back.* Since sequential solvers are based on heuristics to select their decision variables, these can naturally be used to operate the search space division. The idea is to use the variables’ VSIDS-based [25] order<sup>3</sup> to decompose the search

<sup>2</sup> Compared to the solving time using a sequential solver.

<sup>3</sup> The number of their implications in propagation conflicts.

in subspaces. Actually, when a variable is highly ranked w.r.t. to this order, then it is commonly admitted that it is a good starting point for a separate exploration [2, 13, 22].

Another explored track is the number of *flips* of the variables [1]. A flip is when a variable is propagated to the reverse of its last propagated value. Hence, ranking the variables according to the number of their flips, and choosing the highest one as a division point helps to generate search subspaces with comparable computational time. This can be used to limit the number of variables on which the look ahead propagation is applied by preselecting a predefined percentage of variables with the highest number of flips.


Another look back approach, called propagation rate (PR), tends to produce the same effect as the look ahead heuristics [26]. The PR of a variable  $v$  is the ratio between the numbers of propagations due to the branching of  $v$  divided by the number of time  $v$  has been chosen as a decision. The variable with the highest PR is chosen as division point.

### 3.2 Load Balancing

Despite all the effort to produce balanced subspaces, it is practically impossible to ensure the same difficulty for each of them. Hence, some workers often become quickly idle, thus requiring a dynamic load balancing mechanism.

A first solution to achieve dynamic load balancing is to rely on *work stealing*: each time a solver proves its subspace to be UNSAT<sup>4</sup>, it asks for a new job. A target worker is chosen to divide its search space (*e.g.*, extends its guiding path). Hence, the target is assigned to one of the new generated subspaces, while the idle solver works on the other. The most common architecture to implement this strategy is based on a master/slave organization, where slaves are solvers.

When a new division is needed, choosing the best target is a challenging problem. For example, the *Dolius* solver [1] uses a FIFO order to select targets: the next one is the worker that is working for the longest time on its search space. This strategy guarantees fairness between workers. Moreover the target has a better knowledge of its search space, resulting in a better division when using a look back heuristic.



**Fig. 2.** Load balancing through work stealing

<sup>4</sup> If the result is SAT the global resolution ends.

Let us suppose in the example of Fig. 1 that worker  $w_3$  proves its subspace to be UNSAT, and asks for a new one. Worker  $w_2$  is chosen to divide and share its subspace. In Fig. 2,  $m$  is chosen as division variable and two new guiding paths are created, one for  $w_2$  and one for  $w_3$ . Worker  $w_3$  now works on a new subspace and its new guiding path is  $(d \wedge b \wedge \neg m)$ , while the guiding path of  $w_2$  is  $(d \wedge b \wedge m)$ .

Another solution to perform dynamic load balancing is to create more search subspaces (jobs) than available parallel workers (cube-and-conquer [12]). These jobs are then managed via a work queue where workers pick new jobs. To increase the number of available jobs at runtime, a target job is selected to be divided. The strategy implemented in **Treengeling** [6] is to choose the job with the smallest number of variables; this favours SAT instances.

### 3.3 Exchanging Learnt Clauses

Dividing the search space can be subsumed to the definition of constraints on the values of some variables. Technically, there exist two manners to implement such constraints: (i) constrain the original formula; (ii) constrain the decision process initialisation of the used solver.

When the search space division is performed using (i), some learnt clauses cannot be shared between workers. This is typically the case of learnt clauses deduced from at least one clause added for space division, otherwise, correctness is not preserved. The simplest solution to preserve correctness is then to disable clause sharing [6]. Another (more complex) approach is to mark the clauses that must not be shared [17]. Clauses added for the division are initially marked. Then, the tag is propagated to each learnt clause that is deduced from at least one already marked clause.


When the search space division is performed using (ii), some decisions are forced. With this technique there is no sharing restrictions for any learnt clauses. This solution is often implemented using the assumption mechanisms [1, 2].

## 4 Implementation of a Divide-and-Conquer

This section presents the divide-and-conquer component we have built on top of the **Painless** framework. First, we recall the general architecture and operations of **Painless**. Then we describe the generic divide-and-conquer component's mechanisms. Finally we detail the different heuristics we have instantiated using this component.

### 4.1 About the Painless Framework

**Painless** [18] is a framework that aims at simplifying the implementation and evaluation of parallel SAT solvers for many-core environments. Thanks to its genericity and modularity, the components of **Painless** can be instantiated independently to produce new complete solvers.



**Fig. 3.** Architecture of Painless

The main idea of the framework is to separate the technical components (*e.g.*, those dedicated to the management of concurrent programming aspects) from those implementing heuristics and optimizations embedded in a parallel SAT solver. Hence, the developer of a (new) parallel solver concentrates his efforts on the functional aspects, namely parallelisation and sharing strategies, thus delegating implementation issues (*e.g.*, data concurrent access protection mechanisms) to the framework.

Three main components arise when treating parallel SAT solvers: *sequential engines*, *parallelisation*, and *sharing*. These form the global architecture of Painless depicted in Fig. 3.

**Sequential Engines.** The core element considered in the framework is a sequential SAT solver. This can be any CDCL state-of-the-art solver. Technically, these engines are operated through a generic interface providing basics of sequential solvers: *solve*, *interrupt*, *add clauses*, etc.

Thus, to instantiate Painless with a particular solver, one needs to implement the interface according to this engine.

**Parallelisation.** To build a parallel solver using the aforementioned engines, one needs to define and implement a parallelisation strategy. Portfolio and divide-and-conquer are the basic known ones. Also, they can be arbitrarily composed to form new strategies.

In Painless, a strategy is represented by a tree-structure of arbitrarily depth. The internal nodes of the tree represent parallelisation strategies, and leaves are core engines. Technically, the internal nodes are implemented using `WorkingStrategy` component and the leaves are instances of `SequentialWorker` component.

Hence, to develop its own parallelisation strategy, the user should create one or more strategies, and build the required tree-structure.



**Sharing.** In parallel SAT solving, the exchange of learnt clauses warrants a particular focus. Indeed, besides the theoretical aspects, a bad implementation of a good sharing strategy may dramatically impact the solver’s efficiency.

In **Painless**, solvers can export (import) clauses to (from) the others during the resolution process. Technically, this is done by using lockfree queues [24]. The sharing of these learnt clauses is dedicated to particular components called **Sharers**. Each **Sharer** in charge of sets of producers and consumers and its behaviour reduces to a loop of sleeping and exchange phases.

Hence, the only part requiring a particular implementation is the exchange phase, that is user defined.

### 4.2 The Divide-and-Conquer Component in Painless

To implement divide-and-conquer solvers with **Painless**, we define a new component. It is based on a master/slaves architecture.


Figure 4 shows the architecture of our tool. It contains several entities. The *master* is a thread executing the only **D&C** instance of the **WorkingStrategy** class. The *workers* are slave threads executing instances of the **SequentialWorker** class. An instance of the **Sharing** class allows workers to share clauses.

The master and the workers interact asynchronously by means of events. In the initialisation phase, the master may send asynchronous events to himself too.

**Master.** The master (1) initialises the **D&C** component; (2) selects targets to divide their search spaces; (3) and operates the division along with the relaunch of the associated solvers. These actions are triggered by the events **INIT**, **NEED\_JOB**, and **READY\_TO\_DIV**, respectively. In the remainder of this section we consider a configuration with  $N$  workers.

The master can be in two states: either it is sleeping, or it is currently processing an incoming event. Initially the master starts a first solver on the whole formula by sending it the **SOLVE** event. It then generates  $N - 1$  **NEED\_JOB** events to himself. This will provoke the division of the search space in  $N$  subspaces according the to implemented policy. At the end of this initialisation phase, it returns to its sleeping state. At this point, all workers are processing their subspaces.

Each time a worker needs a job, it notifies the master with a **NEED\_JOB** event.



**Fig. 4.** Architecture of the divide-and-conquer SAT solver

All over its execution, the master reacts to the `NEED_JOB` event as follows:

1. it selects a target using the current policy<sup>5</sup>, and requests this target to interrupt by sending an `INTERRUPT` event. Since this is an asynchronous communication, the master may process other events until it receives a `READY_TO_DIV` event;
2. once it receives a `READY_TO_DIV` event, the master proceeds to the effective division of the subspace of the worker which emitted the event. Both the worker which emitted the event and the one which requested a job are then invited to solve their new subspaces through the send of a `SOLVE` event.

The master may receive a `SAT` event from its workers. It means a solution has been computed and the whole execution must end. When a worker ends in an `UNSAT` situation, it makes a request for a new job (`NEED_JOB` event). When the master has no more division of the search space to perform, it states the `SAT` problem is `UNSAT`.

**Slaves.** A slave may be in three different states: *idle*, *work*, and *work\_interrupt\_requested*. Initially, it is *idle* until it receives a `SOLVE` request from the master. Then, it moves to the *work* state and starts to process its assigned subspace. It may:

- find a solution, then emit a `SAT` event to the master, and move back to *idle*;
- end processing of the subspace, with an `UNSAT` result, then it emits a `NEED_JOB` event to the master, and move back to *idle*;
- receive an `INTERRUPT` event from the master, then, it moves to the *work\_interrupt\_requested* state and continues its processing until it reaches a stable state<sup>6</sup> according to the underlying sequential engine implementation. Then, it sends a `READY_TO_DIV` event to the master prior to move back to *idle*.

### 4.3 Implemented Heuristics

The divide-and-conquer component presented in the previous section should be generic enough to allow the implementation of any of the state-of-the-art strategies presented in Sect. 3. We selected some strategies to be implemented, keeping in mind that at least one of each family should be retained:

1. Techniques to Divide the Search Space (Sect. 3.1):  
we have implemented the guiding path method based on the use of assumptions. Since we want to be as generic as possible, we have not considered techniques adding constraints to the formula (because they require tagging mechanisms complex to implement to enable clause sharing).

<sup>5</sup> This policy may change dynamically over the execution of the solver.

<sup>6</sup> For example, in `Minisat`-based solvers, a stable state could correspond to the configuration of the solver after a restart.

2. Choosing Division Variables (Sect. 3.1):  
the different division heuristics we have implemented in the `MapleCOMSPS` solver<sup>7</sup>, are: VSIDS, number of flips, and propagation rate.
3. Load Balancing (Sect. 3.2):  
a work-stealing mechanism was implemented to operate dynamic load balancing. The master selects targets using a FIFO policy (as in `Dolius`) moderated by a minimum computation time (2s) for the workers in order to let these acquire a sufficient knowledge of the subspace.

The exchange of learnt clauses (Sect. 3.3) on any of the strategies we implemented is not restricted. This allows to reuse any of the already off-the-shelf strategies provided by the `Painless` framework.

Another important issue deals with the way new subspaces are allocated to workers. We provide two strategies:

- **Reuse:** the worker reuses the same object-solver all over its execution and the master feeds it with guiding paths;
- **Clone:** each time a new subspace is assigned to a worker, the master clones the object-solver from the target and provides the copy to the idle worker. Thus, the idle worker will benefit from the knowledge (VSIDS, locally learned clauses, etc.) of the target worker.

Hence, our `Painless`-based **D&C** component can thus be instantiated to produces solvers over six orthogonal axes: (1) technique to divide the search space; (2) technique to choose the division variables; (3) load balancing strategy; (4) the sharing strategy; (5) the subspace allocation technique; (6) and the used underlying sequential solver.

By lack of time for experimentations, we select for this paper 6 solvers: all based on `MapleCOMSPS`, and sharing all learnt clauses which  $LBD \leq 4$  (this value has been experimentally deduced). Table 1 summarizes the implemented **D&C** solvers we have used for our experiments in the next section.

**Table 1.** The **D&C** solvers we use for experiments in this paper

	VSIDS	Number of flips	Propagation Rate
Reuse	P-REUSE-VSIDS	P-REUSE-FLIPS	P-REUSE-PR
Clone	P-CLONE-VSIDS	P-CLONE-FLIPS	P-CLONE-PR

## 5 Evaluation

This section presents the results of experiments done with the six **D&C** solvers we presented in Sect. 4.3. We also did comparative experiments with state-of-art **D&C** solvers (`Treengeling` [6] and `MapleAmpharos` [26]).

<sup>7</sup> We used the version that won the main track of the SAT Competition in 2016 [19].

**Treengeling** is a cube-and-conquer solver based on the **Lingeling** sequential solver. **MapleAmpharos** is an adaptive divide-and-conquer based on the solver **Ampharos** [2], and using **MapleCOMSPS** as sequential solver. Comparing our new solvers with state-of-the-art ones (*e.g.*, not implemented on **Painless**) is a way to assess if our solution is competitive despite the genericity introduced by **Painless** and ad-hoc optimizations implemented in other solvers.

All experiments were executed on a multi-core machine with 12 physical cores (2 x Intel Xeon E5645 @ 2.40 GHz), and 64 GB of memory. Hyper-threading has been activated porting the number of logical cores to 24. We used the 400 instances of the parallel track of the SAT Competition 2018<sup>8</sup>. All experiments have been conducted using the following settings:

- each solver has been run once on each instance with a time-out of 5000 s (as in the SAT Competition);
- the number of used cores is limited to 23 (the remaining core is booked to the operating system);
- instances that were trivially solved by a solver (at the preprocessing phase) were removed, indeed in this case the D&C component of solvers is not enabled, these instances are then irrelevant for our case study.

Results of these experiences are summarized in Table 2. The different columns represent respectively: the total number of solved instances, the number of UNSAT solved instances, the number of SAT solved instances, and the PAR-2 score<sup>9</sup>.

**Table 2.** Results of the different solvers

Solver	ALL (360)	UNSAT	SAT	PAR-2
P-CLONE-FLIPS	198	87	111	1732696.65
P-CLONE-PR	183	73	110	1871614.48
P-CLONE-VSIDS	183	77	106	1880281.54
P-REUSE-FLIPS	190	83	107	1796426.72
P-REUSE-PR	180	72	108	1938621.48
P-REUSE-VSIDS	184	75	109	1868619.43
MapleAmpharos	153	29	124	2190680.55
Treengeling	200	84	116	1810471.56

## 5.1 Comparing the Implemented Divide-and-Conquer Solvers

Figure 5 presents the cactus plot of the performances of the different D&C solvers. These differ in two orthogonal axes: the used subspace allocation technique, and the used division heuristic. We analyse here each axe separately.

<sup>8</sup> <http://sat2018.forsyte.tuwien.ac.at/benchmarks/Main.zip>.

<sup>9</sup> The used measure in the annual SAT Competition.

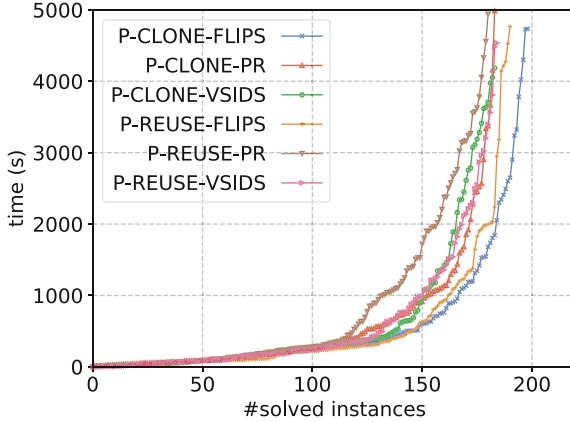


Fig. 5. Cactus plot of the different divide-and-conquer based solvers

When considering the allocation technique (clone *vs.* reuse), we can observe that the cloning based strategy is globally more efficient, even if it has a supplementary cost (due to the cloning phase). The scatter plots of Fig. 6 confirm this observation (most plots are below the diagonal, showing evidence of a better average performance). We believe this is due to the local knowledge that is implicitly shared between the (cloned) workers.

When considering the division heuristics (VSIDS, number of flips, and propagation rate), we observe that number of flips based approach is better than the two others. Both, by the number of solved instances and the PAR-2 measure. This is particularly true when considering the cloning based strategy. VSIDS and propagation rate based solvers are almost identical.

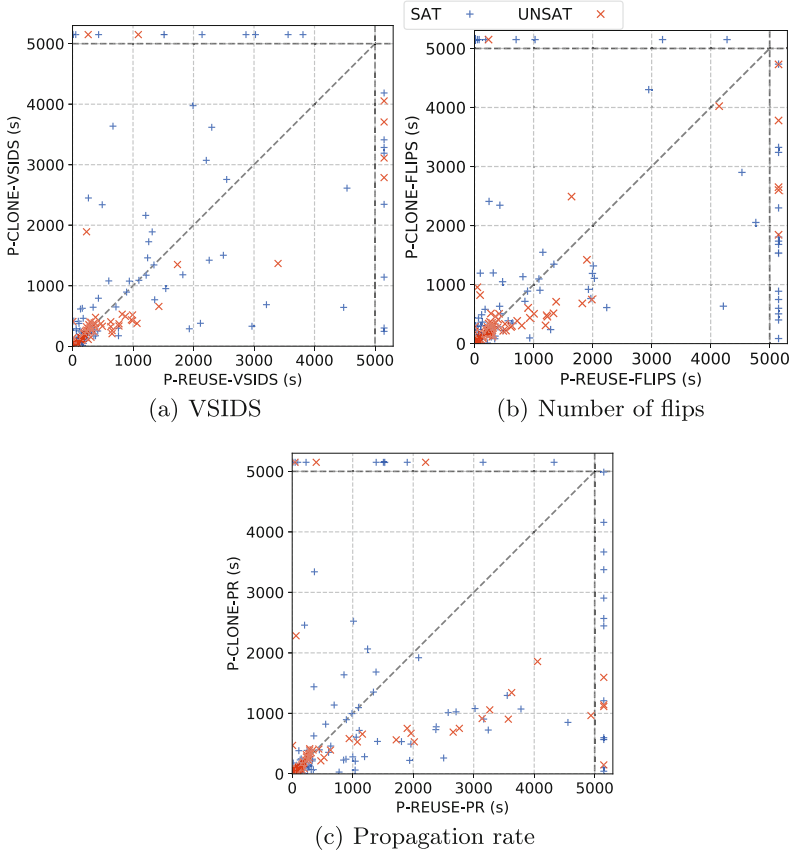
### 5.2 Comparison with State-of-the-Art Divide-and-Conquer

Figure 7 shows a cactus plot comparing our best divide-and-conquer (*i.e.*, P-CLONE-FLIPS) against **Treengeling** and **MapleAmpharos**.

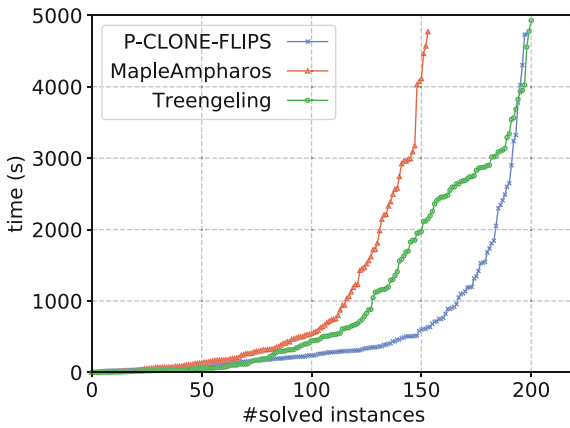
The **MapleAmpharos** solver seems to be less efficient than our tool, and solves less instances. When considering only the 123 instances that both solvers were able to solve, we can calculate the cumulative execution time of this intersection (CTI) for **MapleAmpharos** and P-CLONE-FLIPS: it is, respectively, 24h33min and 14h34min.

Although our tool solves 2 less instances as **Treengeling**, it has better PAR-2 measure. The CTI calculated on the 169 instances solved by both solvers, is 49h14min and 22h23min, respectively for **Treengeling** and P-CLONE-FLIPS. We can say that even if both solve almost the same number of instances, our D&C solver is faster. We clearly observe this phenomenon in Fig. 7.

Thus, in addition to highlight the performance of our instantiation, this shows the effectiveness of the flip-based approach with respect to the well-proven cube-and-conquer strategies.



**Fig. 6.** Scatter plots of divide-and-conquer reusing *vs.* cloning solvers



**Fig. 7.** Cactus plot of the best instantiated divide-and-conquer and state-of-the-art solvers

## 6 Conclusion

This paper proposed an optimal implementation of several parallel SAT solvers using the divide-and-conquer (D&C) strategy that handle parallelisms by performing successive divisions of the search space.

Such an implementation was performed on top of the *Painless* framework that allows to easily deal with variants of strategies. Our *Painless*-based implementation can be customized and adapted over six orthogonal axes: (1) technique to divide the search space; (2) technique to choose the division variables; (3) load balancing strategy; (4) the sharing strategy; (5) the subspace allocation technique; (6) and the used underlying sequential solver.

This work shows that we have now a modular and efficient framework to explore new D&C strategies along these six axes. We were then able to make a fair comparison between numerous strategies.

Among the numerous solvers we have available, we selected six of them for performance evaluation. Charts are provided to show how they competed, but also how they cope face to natively implemented D&C state-of-the-art solvers.

This study shows that the flip-based approach in association with the clone policy outperforms the other strategies whatever the used standard metrics is. Moreover, when compared with the state-of-the-art D&C-based solvers, our best solver shows to be very efficient and allows us to conclude the effectiveness of our modular platform based approach with respect to the well-competitive D&C solvers.

In the near future, we want to conduct more massive experiments to measure the impact of clauses sharing strategies in the D&C context, and evaluate the scalability of the various D&C approaches.

## References

1. Audemard, G., Hoessen, B., Jabbour, S., Piette, C.: Dolius: a distributed parallel SAT solving framework. In: *Pragmatics of SAT International Workshop (POS) at SAT*, pp. 1–11. Citeseer (2014)
2. Audemard, G., Lagniez, J.-M., Szczepanski, N., Tabary, S.: An adaptive parallel SAT solver. In: Rueher, M. (ed.) *CP 2016*. LNCS, vol. 9892, pp. 30–48. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-44953-1\\_3](https://doi.org/10.1007/978-3-319-44953-1_3)
3. Audemard, G., Simon, L.: Predicting learnt clauses quality in modern SAT solvers. In: *Proceedings of the 21st International Joint Conferences on Artificial Intelligence (IJCAI)*, pp. 399–404. AAAI Press (2009)
4. Audemard, G., Simon, L.: Lazy clause exchange policy for parallel SAT solvers. In: Sinz, C., Egly, U. (eds.) *SAT 2014*. LNCS, vol. 8561, pp. 197–205. Springer, Cham (2014). [https://doi.org/10.1007/978-3-319-09284-3\\_15](https://doi.org/10.1007/978-3-319-09284-3_15)
5. Balyo, T., Sanders, P., Sinz, C.: HordeSat: a massively parallel portfolio SAT solver. In: Heule, M., Weaver, S. (eds.) *SAT 2015*. LNCS, vol. 9340, pp. 156–172. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-24318-4\\_12](https://doi.org/10.1007/978-3-319-24318-4_12)
6. Biere, A.: CaDiCaL, Lingeling, Plingeling, Treengeling and YalSAT entering the SAT competition 2018. In: *Proceedings of SAT Competition 2018: Solver and Benchmark Descriptions*, pp. 13–14. Department of Computer Science, University of Helsinki, Finland (2018)

7. Biere, A., Cimatti, A., Clarke, E., Zhu, Y.: Symbolic model checking without BDDs. In: Cleaveland, W.R. (ed.) TACAS 1999. LNCS, vol. 1579, pp. 193–207. Springer, Heidelberg (1999). [https://doi.org/10.1007/3-540-49059-0\\_14](https://doi.org/10.1007/3-540-49059-0_14)
8. Biere, A., Heule, M., van Maaren, H.: Handbook of Satisfiability, vol. 185. IOS Press, Amsterdam (2009)
9. Davis, M., Logemann, G., Loveland, D.: A machine program for theorem-proving. *Commun. ACM* **5**(7), 394–397 (1962)
10. Davis, M., Putnam, H.: A computing procedure for quantification theory. *J. ACM* **7**(3), 201–215 (1960)
11. Hamadi, Y., Jabbour, S., Sais, L.: ManySAT: a parallel SAT solver. *J. Satisfiability Boolean Model. Comput.* **6**(4), 245–262 (2009)
12. Heule, M.J.H., Kullmann, O., Wieringa, S., Biere, A.: Cube and conquer: guiding CDCL SAT solvers by lookaheads. In: Eder, K., Lourenço, J., Shehory, O. (eds.) HVC 2011. LNCS, vol. 7261, pp. 50–65. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-34188-5\\_8](https://doi.org/10.1007/978-3-642-34188-5_8)
13. Hyvärinen, A.E.J., Junttila, T., Niemelä, I.: A distribution method for solving SAT in grids. In: Biere, A., Gomes, C.P. (eds.) SAT 2006. LNCS, vol. 4121, pp. 430–435. Springer, Heidelberg (2006). [https://doi.org/10.1007/11814948\\_39](https://doi.org/10.1007/11814948_39)
14. Hyvärinen, A.E.J., Manthey, N.: Designing scalable parallel SAT solvers. In: Cimatti, A., Sebastiani, R. (eds.) SAT 2012. LNCS, vol. 7317, pp. 214–227. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-31612-8\\_17](https://doi.org/10.1007/978-3-642-31612-8_17)
15. Jurkowiak, B., Li, C.M., Utard, G.: Parallelizing satz using dynamic workload balancing. *Electron. Notes Discrete Math.* **9**, 174–189 (2001)
16. Kautz, H.A., Selman, B., et al.: Planning as satisfiability. In: Proceedings of the 10th European Conference on Artificial Intelligence (ECAI), vol. 92, pp. 359–363 (1992)
17. Lanti, D., Manthey, N.: Sharing information in parallel search with search space partitioning. In: Nicosia, G., Pardalos, P. (eds.) LION 2013. LNCS, vol. 7997, pp. 52–58. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-44973-4\\_6](https://doi.org/10.1007/978-3-642-44973-4_6)
18. Le Frioux, L., Baarir, S., Sopena, J., Kordon, F.: PaInleSS: a framework for parallel SAT solving. In: Gaspers, S., Walsh, T. (eds.) SAT 2017. LNCS, vol. 10491, pp. 233–250. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-66263-3\\_15](https://doi.org/10.1007/978-3-319-66263-3_15)
19. Liang, J.H., Oh, C., Ganesh, V., Czarnecki, K., Poupard, P.: MapleCOMSPS, mapleCOMSPS LRB, mapleCOMSPS CHB. In: Proceedings of SAT Competition 2016: Solver and Benchmark Descriptions, p. 52. Department of Computer Science, University of Helsinki, Finland (2016)
20. Lynce, I., Marques-Silva, J.: SAT in bioinformatics: making the case with haplotype inference. In: Biere, A., Gomes, C.P. (eds.) SAT 2006. LNCS, vol. 4121, pp. 136–141. Springer, Heidelberg (2006). [https://doi.org/10.1007/11814948\\_16](https://doi.org/10.1007/11814948_16)
21. Marques-Silva, J.P., Sakallah, K.: GRASP: a search algorithm for propositional satisfiability. *IEEE Trans. Comput.* **48**(5), 506–521 (1999)
22. Martins, R., Manquinho, V., Lynce, I.: Improving search space splitting for parallel SAT solving. In: Proceedings of the 22nd IEEE International Conference on Tools with Artificial Intelligence (ICTAI), vol. 1, pp. 336–343. IEEE (2010)
23. Massacci, F., Marraro, L.: Logical cryptanalysis as a SAT problem. *J. Autom. Reasoning* **24**(1), 165–203 (2000)
24. Michael, M.M., Scott, M.L.: Simple, fast, and practical non-blocking and blocking concurrent queue algorithms. In: Proceedings of the 15th ACM Symposium on Principles of Distributed Computing (PODC), pp. 267–275. ACM (1996)



25. Moskewicz, M.W., Madigan, C.F., Zhao, Y., Zhang, L., Malik, S.: Chaff: engineering an efficient SAT solver. In: Proceedings of the 38th Design Automation Conference (DAC), pp. 530–535. ACM (2001)
26. Nejati, S., et al.: A propagation rate based splitting heuristic for divide-and-conquer solvers. In: Gaspers, S., Walsh, T. (eds.) SAT 2017. LNCS, vol. 10491, pp. 251–260. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-66263-3\\_16](https://doi.org/10.1007/978-3-319-66263-3_16)
27. Plaza, S., Markov, I., Bertacco, V.: Low-latency SAT solving on multicore processors with priority scheduling and XOR partitioning. In: the 17th International Workshop on Logic and Synthesis (IWLS) at DAC (2008)
28. Silva, J.P.M., Sakallah, K.A.: GRASP—a new search algorithm for satisfiability. In: Proceedings of the 16th IEEE/ACM International Conference on Computer-Aided Design (ICCAD), pp. 220–227. IEEE (1997)
29. Zhang, H., Bonacina, M.P., Hsiang, J.: PSATO: a distributed propositional prover and its application to quasigroup problems. *J. Symb. Comput.* **21**(4), 543–560 (1996)
30. Zhang, L., Madigan, C.F., Moskewicz, M.H., Malik, S.: Efficient conflict driven learning in a boolean satisfiability solver. In: Proceedings of the 20th IEEE/ACM International Conference on Computer-Aided Design (ICCAD), pp. 279–285. IEEE (2001)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

