



PaInleSS: a Framework for Parallel SAT Solving

Ludovic Le Frioux, Souheib Baarir, Julien Sopena, Fabrice Kordon

► To cite this version:

Ludovic Le Frioux, Souheib Baarir, Julien Sopena, Fabrice Kordon. PaInleSS: a Framework for Parallel SAT Solving. The 20th International Conference on Theory and Applications of Satisfiability Testing, Aug 2017, Melbourne, Australia. pp.233-250, 10.1007/978-3-319-66263-3_15 . hal-01540785

HAL Id: hal-01540785

<https://hal.science/hal-01540785>

Submitted on 12 Mar 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

PaInleSS: A Framework for Parallel SAT Solving

Ludovic Le Frioux^{1,2,3(✉)}, Souheib Baarir^{1,2,3,4}, Julien Sopena^{2,3},
and Fabrice Kordon^{2,3}

¹ LRDE, EPITA, Kremlin-Bicêtre, France

{`ludovic.le.frioux,souheib.baarir`}@lrde.epita.fr

² Sorbonne Universités, UPMC Univ Paris 06, UMR 7606, LIP6, Paris, France

{`julien.sopena,fabrice.kordon`}@lip6.fr

³ CNRS, UMR 7606, LIP6, Paris, France

⁴ Université Paris, Nanterre, France

Abstract. Over the last decade, parallel SAT solving has been widely studied from both theoretical and practical aspects. There are now numerous solvers that differ by parallelization strategies, programming languages, concurrent programming, involved libraries, etc.

Hence, comparing the efficiency of the theoretical approaches is a challenging task. Moreover, the introduction of a new approach needs either a deep understanding of the existing solvers, or to start from scratch the implementation of a new tool.

We present **PaInleSS**: a framework to build parallel SAT solvers for many-core environments. Thanks to its genericity and modularity, it provides the implementation of basics for parallel SAT solving like clause exchanges, Portfolio and Divide-and-Conquer strategies. It also enables users to easily create their own parallel solvers based on new strategies. Our experiments show that our framework compares well with some of the best state-of-the-art solvers.

Keywords: Parallel · Satisfiability · Clause sharing · Portfolio · Cube and conquer

1 Introduction

Boolean satisfiability (SAT) has been used successfully in many contexts such as planning decision [19], hardware and software verification [6], cryptology [27] and computational biology [22], etc. This is due to the capability of modern SAT solvers to solve complex problems involving millions of variables and billions of clauses.

Most SAT solvers have long been sequential and based on the well-known DPLL algorithm [8, 9]. This initial algorithm has been dramatically enhanced by introducing sophisticated heuristics and optimizations: decision heuristics [21, 29], clauses learning [25, 32, 35], aggressive cleaning [2], lazy data structures [29], preprocessing [11, 23, 24], etc. The development of these enhancements has been greatly simplified by the introduction of **MiniSat** [10], an extensible SAT solver easing the integration of these heuristics in an efficient sequential solver.

The emergence of many-core machines opens new possibilities in this domain. Two classes of parallel techniques have been developed: competition-based (*a.k.a.*, Portfolio) and cooperation-based (*a.k.a.*, Divide-and-Conquer). In the Portfolio settings [14], many sequential SAT solvers compete for the solving of the whole problem. The first one to find a solution, or proving the problem to be unsatisfiable ends the computation. Divide-and-Conquer approaches use the guiding path method [34] to decompose, recursively and dynamically, the original problem in sub-problems that are solved separately by sequential solvers. In both approaches, sequential solvers can dynamically share learnt information. Many heuristics exist to improve this sharing by proposing trade-off between gains and overhead.

While the multiplication of strategies and heuristics provides perspectives for parallel SAT solving, it makes more complex development and evaluation of new proposals. Thus, any new contribution faces three main problems:

Problem 1: concurrent programming requires specific skills (synchronization, load balancing, data consistency, etc.). Hence, the theoretical efficiency of an heuristic may be annihilated by implementation choices.

Problem 2: most of the contributions mainly target a specific component in the solver while, to evaluate it, a complete one (either built from scratch or an enhancement of an existing one) must be available. This makes the implementation and evaluation of a contribution much harder.

Problem 3: an implementation, usually, only allows to test a single composition policy. Hence, it becomes hard to evaluate a new heuristic with different versions of the other mechanisms.

This paper presents PArallel INstantiable Sat Solver (PaInleSS)¹, a framework that simplifies the implementation and evaluation of new parallel SAT solvers for many-core environments. The components of PaInleSS can be instantiated independently to produce a new complete solver. The guiding principle is to separate the technical components dedicated to some specific aspect of concurrent programming, from the components implementing heuristics and optimizations embedded in a parallel SAT solver.

Our main contributions are the following:

- we propose a new modular and generic framework that can be used to implement new strategies with minimal effort and concurrent programming skills;
- we provide adaptors for some state-of-the-art sequential SAT solvers: **glucose** [2], **Lingeling** [5], **MiniSat** [10], and **MapleCOMSPS** [21].
- we show that it is easy to implemented strategies in PaInleSS, and provide some that are present in the classical solvers of the state-of-the-art: **glucose-syrup** [3], **Treengeling** [5], and **Hordesat** [4];
- we show the effectiveness of our modular design by instantiating, with a minimal effort, new original parallel SAT solver (by mixing strategies);

¹ painless.lrde.epita.fr.

- we evaluate our approach on the benchmark of the parallel track of the SAT Race 2015. We compare the performance of solvers instantiated using the framework with the original solvers. The results show that the genericity provided by **PaInleSS** does not impact the performances of the generated instances.

The rest of the paper is organized as follows: Sect. 2 introduces useful background to deal with sequential SAT solving. Section 3 is dedicated to parallel SAT solving. Section 4 shows the architecture of **PaInleSS**. Section 5 presents different solvers implemented using **PaInleSS**. Section 6 analyzes the results of our experiments and Sect. 7 concludes and gives some perspectives work.

2 About Sequential SAT Solving

In this section, after some preliminary definitions and notations, we introduce the most important features of modern sequential SAT solvers.

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* ω 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* φ is a finite conjunction of clauses (noted $\varphi = \bigwedge_{i=1}^k \omega_i$). For a given φ , the set of its variables is noted: V_φ . An assignment \mathcal{A} of variables of φ , is a function $\mathcal{A} : V_\varphi \rightarrow \{\top, \perp\}$. \mathcal{A} is total (complete) when all elements of V_φ have an image by \mathcal{A} , otherwise it is partial. For a given formula φ , and an assignment \mathcal{A} , a clause of φ is satisfied when it contains at least one literal evaluating to true, regarding \mathcal{A} . The formula φ is satisfied by \mathcal{A} iff $\forall \omega \in \varphi, \omega$ is satisfied. φ is said to be SAT if there is at least one assignment that makes it satisfiable. It is defined as UNSAT otherwise.

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 [25, 32, 35], that is an enhancement of the DPLL algorithm [8, 9]. The main components of a CDCL are depicted in Algorithm 1.

At each step of the main loop, **unitPropagation**² (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 [25] (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.

² The **unitPropagation** function implements the Boolean Constraint Propagation (BCP) procedure that forces (in cascade) the values of the variables in asserting clauses [8].

```

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        $\omega \leftarrow \text{conflictAnalysis}()$ 
9        $\text{addLearntClause}(\omega)$ 
10       $dl \leftarrow \text{backjump}(\omega)$ 
11    else
12       $\text{assignDecisionLiteral}()$ 
13       $dl \leftarrow dl + 1$ 
14  return  $\top$  //  $\varphi$  is SAT

```

Algorithm 1. CDCL algorithm.

The Learning Mechanism. The effectiveness of the CDCL lies in the *learning mechanism*. Each time a conflict is encountered, it is analyzed (**conflictAnalysis** function in Algorithm 1) in order to compute its reasons and to 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 [2]), 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) [2] measure).

With the two classical approaches used for parallel SAT solving: Portfolio and Divide-and-Conquer (see Sect. 3), multiple sequential solvers are used in parallel to solve the formula. With these paradigms sequential solvers can be seen as black boxes providing solving and clause sharing functionalities.

3 About Parallel SAT Solving

The arrival of many-core machines leads to new possibilities for SAT solving. Parallel SAT solving rely on two concepts: parallelization strategy and learnt clause exchanges. Two main parallelization methods have been developed: Portfolio and Divide-and-Conquer. We can also mention the hybrid approaches as alternatives, that are combinations of the first two techniques. With these parallelization strategies, it is possible to exchange learnt clauses, between the underlying sequential solvers.

3.1 Parallelization Strategies

Portfolio. The Portfolio scheme has been introduced by [14], in **ManySat**. The main idea of this approach is to run sequential solvers working in parallel on the entire formula, in a competitive way. This strategy aims at increasing the probability of finding a solution using the *diversification* [12] (also known as swarming in others contexts) principle.

The diversification can only concern the used heuristics: several solvers (workers) with different heuristics are instantiated. They differ by their decision strategies, learning schemes, the used random seed, etc.

Another type of diversification, introduced in **HordeSat** [4], uses the *phase* of the variables: before starting the search each solver receives a special phase, acting as a soft division of the search space. Solvers are invited to visit a certain part of the search space but they can move out of this region during the search.

Another technique to ensure the diversification is the *block branching* [33]: each worker focuses on a particular subset (or block) of variables. Hence, the decision variables of a worker are chosen from the block it is in charge of.

Divide-and-Conquer. The Divide-and-Conquer approach is based on splitting the search space in disjoint parts. These parts are solved independently, in parallel, by different workers. As the parts are disjunct, if one of the partitions is proven to be SAT then the initial formula is SAT. The formula is UNSAT if all the partitions are UNSAT. The challenging points of the this method are: *dividing the search space* and *balancing jobs between workers*.

To divide the search space, the most used technique is based on the Shannon's decomposition, known as the *guiding path* [34]. The guiding path is a vector of literals (*a.k.a., cube*) that are assumed by the worker when solving the formula.

Choosing the best division variables is a hard problem requiring heuristics. If some parts are too easy this will lead to repeatedly divide the search space and ask for a new job (phenomenon known as *ping-pong effect*). As all the partitions do not require the same solving time, some workers may become idle and a mechanism for load balancing is needed. Each time a solver proves that its partition is UNSAT³, it needs a new job. Another solver is chosen as target to divide its search space (*i.e.*, to extend its guiding path). The target will work on one of the new partition and the idle worker on the other one. This mechanism is often called *work stealing*.

Hybrid Approaches. As already presented, Portfolio and Divide-and-Conquer, are the two main explored approaches to parallelize SAT solving.

The Portfolio scheme is simple to implement, and uses the principle of diversification to increase the probability of solving the problem. However, since workers can overlap their search regions, the theoretical resulting speed-up is not as good as the one of the Divide-and-Conquer approach [17]. Surprisingly, while giving

³ If the result is SAT the global solving ends.

a better theoretical speed-up, the Divide-and-Conquer approach suffers from the two challenging issues we mentioned: dividing the search space and balancing jobs between workers.

Emerging techniques, called *hybrid approaches*, propose to use simultaneously the two strategies, so that we benefit from the advantages of each, while trying to avoid their drawbacks.

A basic manner to mix the two approaches is to compose them. There are two possible strategies: Portfolio of Divide-and-Conquer (introduced by **c-sat** [30]), and Divide-and-Conquer of Portfolios (*e.g.*, **ampharos** [1] an adaptive Divide-and-Conquer that allows multiple workers on the same sub-part of the search space). Let us mention other more sophisticated ways to mix approaches like *scattering* [16,18] or *transition heuristics* based strategies [1,7,26,31].

3.2 Clauses Sharing

In all these parallelization paradigms, sharing information between workers is possible, the most important one being clauses learnt by each worker. Hence, the main questions are: *which clauses should be shared? And between which workers?* Indeed, sharing all clauses can have a bad impact on the overall behavior.

To answer the first question, many solvers rely on the standard measures, defined for sequential solvers (*i.e.*, activity, size, LBD): only clauses under a given threshold for these measures are shared. One simple way to get the threshold is to define it as constant it (*e.g.*, clauses up to size 8 are shared in **ManySat** [14]). More sophisticated approaches adapt thresholds dynamically in order to control the flow of shared clauses during the solving [4,13].

A simple solution to the second question, adopted in almost all parallel SAT solvers, is to share clauses between all workers. However, a finer (but more complex) solution is to let each worker choose its emitters [20].

As a conclusion of this section, we can say that parallel SAT solving is based on two distinct concepts. First, there exist numerous strategies to parallelize SAT solving by organizing the workers search. Secondly, with all these strategies it is possible to share clauses between the workers. These two concepts have been our intuition sources for the design of the architecture of **PaInleSS**.

4 Architecture of the Framework

There exist numerous strategies to parallelize SAT solving, their performances heavily relying on their implementation. The most difficult issues deal with concurrent programming. Languages and libraries provide abstractions to deal with these difficulties, and according to these abstractions developers have more or less control on mechanisms such as memory or threads management (*e.g.*, **Java** vs **C++**). This will affect directly the performance of the produced solver.

Therefore, it is difficult to compare the strategies without introducing technological bias. Indeed, it is difficult to integrate new strategies on top of existing

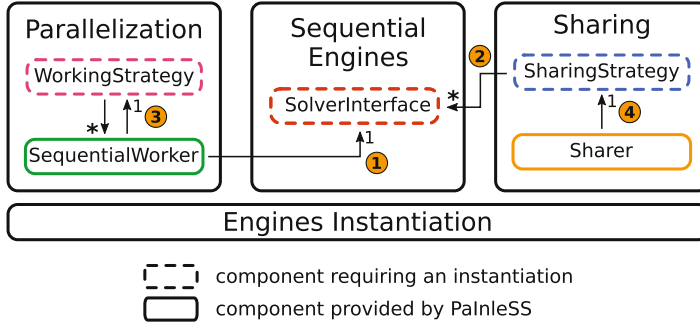


Fig. 1. Architecture of PaInleSS.

solvers, or to develop a new solver from scratch. Moreover, an implementation usually offers the possibility to modify a particular component, it is then difficult to test multiple combinations of components.

PaInleSS aims to be a solution to these problems. It is a generic, modular, and efficient framework, developed in C++11, allowing an easy implementation of parallel strategies. Taking black-boxed sequential solvers as input, it minimizes the effort to encode new parallelization and sharing strategies, thus enabling the implementation of complete SAT solvers at a reduced cost.

As mentioned earlier, a typical parallel SAT solver relies mainly on three core concepts: sequential engine(s), parallelization, and sharing. These last form the core of the **PaInleSS** architecture (see Fig. 1): the sequential engine is handled by the **SolverInterface** component. The parallelization is implemented by the **WorkingStrategy** and **SequentialWorker** components. Components **SharingStrategy** and **Sharer** are in charge of the sharing.

Sequential Engine. **SolverInterface** is an adapter for the basic functions expected from a sequential solver, it is divided in two subgroups: *solving* and *clauses export/import* (respectively represented by arrows 1 and 2 in Fig. 1). Subgroup 1 provides methods that interact with the solving process of the underlying solver. The most important methods of this interface are:

- **SatResult solve(int[*] cube)**: tries to solve the formula, with the given cube (that can be empty in case of Portfolio). This method returns SAT, UNSAT, or UNKNOWN.
- **void setInterrupt()**: stops the current search initiated using the **solve** method.
- **void setPhase(int var, bool value)**: set the phase of variable **var** to **value**.
- **void bumpVariableActivity(int var, int factor)**: bumps **factor** times the activity of variable **var**.
- **void diversify()**: adjusts internal parameters of the solver, to diversify its behaviour.

Subgroup 2 provides methods to add/fetch learnt clauses to/from the solver:

- `void addClause(Clause cls)`: adds a permanent clause to the solver.
- `void addLearntClause(Clause cls)`: adds a learnt clause to the solver.
- `Clause getLearntClause()`: gets the oldest produced learnt clause from the solver.

The interface also provides methods to manipulate sets of clauses. The clauses produced or to be consumed by the solver, are stored in local *lockfree* queues (based on algorithm of [28]).

Technically, to integrate a new solver in PaInleSS, one needs to create a new class inheriting from `SolverInterface` and implement the required methods (*i.e.*, wrapping the methods of the API offered by the underlying solver). The framework currently provides some basic adaptors for `Lingeling` [5], `glucose` [2], `Minisat` [10], and `MapleCOMSPS` [21].

Parallelization. Basic parallelization strategies, such as those introduced in Sect. 3, must be implemented easily. We also aim at creating new strategies and mixing them.

A tree-structured (of arbitrary depth) composition mechanism enables the mix of strategies: internal nodes represent parallelization strategies, and leaves represent solvers. As an example (see Fig. 2(a)), a Divide-and-Conquer of Portfolios is represented by a tree of depth 3: the root corresponds to the Divide-and-Conquer having children representing the Portfolios acting on several solvers (the leaves of the tree).




Fig. 2. Example of a composed parallelization strategy.

PaInleSS implements nodes using the `WorkingStrategy` class, and leaves with the `SequentialWorker` class. This last is a subclass of `WorkingStrategy` that integrates an execution flow (a thread) operating the associated solver.

The overall solving workflow within this tree is user defined and guaranteed by the two main methods of the `WorkingStrategy` (arrows 3 in Fig. 1):

- `void solve(int[*] cube)`: according to the strategy implemented, this method manages the organization of the search by giving orders to the children strategies.
- `void join(SatResult res, int[*] model)`: used to notify the parent strategy of the solving end. If the result is SAT, `model` will contain an assignment that satisfies the sub-formula treated by this node.

It is worth noting that the workflow must start by a call to the root's `solve` method and eventually ends by a call to the root's `join` method. The propagation of solving orders from a parent to one of its child nodes, is done by a call to the `solve` method of this last. The results are propagated back from a child to its parent by a call to the `join` method of this last. The solving can not be effective without a call to the leaves' `solve` methods.

Back to the example of Fig. 2(a). Consider the execution represented in Fig. 2(b). The solving order starts by a call to the root's (DC node) `solve` method. It is relayed through the tree structure to the leaves (SW nodes). Here, once its problem is found SAT by one of the SW, it propagates back the result to its PF node parent via a call to the `join` method. According to the strategy of the PF, the DC's `join` method is called and ends the global solving.

Hence, to develop its own parallelization strategy, the user should create one or more subclass of `WorkingStrategy` and to build the tree structure.

Sharing. In parallel SAT solving, we must pay a particular attention to the exchange of learnt clauses. Indeed, beside the theoretical aspects, a bad implementation of the sharing can dramatically impact the efficiency of the solver (*e.g.*, improper use of locks, synchronization problems). We now present how sharing is organized in PaInleSS.

When a solver learns a clause, it can share it according to a filtering policy such as the size or the LBD of the clause. To do so it puts the clause in a special buffer (`buff_exp` in Fig. 3). The sharing of the learnt clauses is realized by dedicated thread(s): **Sharer(s)**. Each one is in charge of a set of producers and consumers (these are references to `SolverInterface`). Its behaviour reduces to

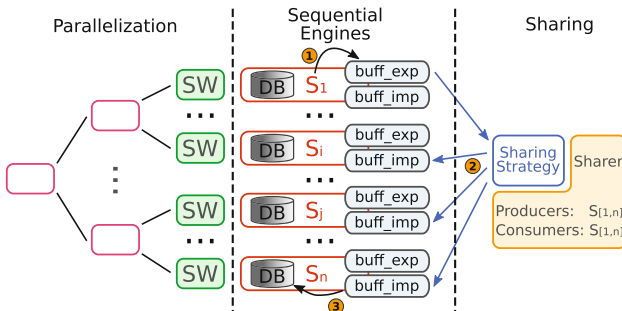


Fig. 3. Sharing mechanism implemented in PaInleSS.

a loop of sleeping and exchange phases. This last is done by calling the interface of **SharingStrategy** class (arrow 4 in Fig. 1). The main method of this class is the following:

- `void doSharing(SolverInterface[*] producers, SolverInterface[*] consumers)`: according to the underlying strategy, this method gets clauses from the producers and add them to the consumers.

In the example of Fig. 3, the **Sharer** uses a given strategy, and all the solvers (S_i) are producers and consumers. The use of dedicated workflows (*i.e.*, threads) allows CPU greedy strategies to be run on a dedicated core, thus not interfering with the solving workers. Moreover, sharing phase can be done manipulating groups of clauses, allowing the use of more sophisticated heuristics. Finally, during its search a solver can get clauses from its import buffer (`buff_imp` in Fig. 3) to integrate them in its local database.

To define a particular sharing strategy the user only needs to provide a subclass of **SharingStrategy**. With our mechanism it is possible to have several groups of sharing each one manage by a **Sharer**. Moreover, solvers can be dynamically added/deleted from/to the producers and/or customers sets of a **Sharer**.

Engine Instantiation. To create a particular instance of **PaInleSS**, the user has to adapt the main function presented by Algorithm 2. The role of this function is to instantiate and bind all the components correctly. This actions are simplified by the use of parameters.

First, the concrete solver classes (inheriting from **SolverInterface**) are instantiated (line 2). Then the **WorkingStrategy** (including **SequentialWorker**) tree is implemented (line 3). This operation links **SequentialWorker** to their **SolverInterface**. Finally, the **Sharer**(s) and their concrete **SharingStrategy**(s) are created; the producers and consumers sets are initialized (line 4).

```

1 function main-PaInleSS (args: the program arguments)
2   solvers ← Create SolverInterface
3   root ← Create WorkingStrategy tree (solvers)
4   sharers ← Create SharingStrategy and Sharer (solvers)
5   root.solve()
6   while timeout or stop do
7     sleep(...)
8   print(root.getResult())
9   if root.getResult() == SAT then
10    print(root.getModel())

```

Algorithm 2. The main function of **PaInleSS**.

The solving starts by the call to the `solve` method of the root `WorkingStrategy` tree. The main thread will execute a loop, where it sleeps for an amount of time, and then checks if either the timeout has been reached or the solving ended (lines 6–7). It prints the final result (line 8), plus the model in case of a SAT instance (lines 9–10).

5 Implementing and Combining Existing Strategies

To validate the generic aspect of our approach, we selected three efficient state-of-the-art parallel SAT solvers: `glucose-syrup` [3], `Treengeling` [5], and `Horde-sat` [4]. For each selected solver, we implemented a solver that mimics the original one using `PaInleSS`. To show the modularity of `PaInleSS`, we used the already developed components to instantiate two new original solvers that combine existing strategies.

Solver “à la Glucose-Syrup”. The `glucose-syrup`⁴ solver is the winner of the parallel track of the SAT Race 2015. It is a Portfolio based on the sequential solver `glucose` [2]. The sharing strategy exchanges all the exported clauses between all the workers. Beside, the workers have customized settings in order to diversify their search.

Hence, implementing a solver “à la `glucose-syrup`”, namely `PaInleSS-breakglucose-syrup`, required the following components: `Glucose` an adaptor to use the `glucose` solver; `Portfolio` a simple `WorkingStrategy` that implements a Portfolio strategy; `SimpleSharing` a `SharingStrategy` that exchanges all the exported clauses from the producers to all the consumers with no filtering.

The implementation of `PaInleSS-glucose-syrup` required 355 lines of code (LoC) for the adaptor, 95 LoC for the Portfolio, and 44 LoC for the sharing strategy.

Solver “à la Treengeling”. The `Treengeling`⁵ solver is the winner of the parallel track of the SAT Competition 2016. It is based on the sequential engine `Lingeling` [5]. Its parallelization strategy is a variant of Divide-and-Conquer called Cube-and-Conquer [15]. The solving is organized in rounds. Some workers search for a given number of conflicts. When the limit is reached, some are selected to split their sub-spaces using a lookahead heuristic. The sharing is restricted to the exchange of unit clauses from a special worker. This last is also in charge of the solving of the whole formula during all the execution.

To implement a solver “à la `Treengeling`”, namely `PaInleSS-treengeling`, we needed the following components: `Lingeling`, an adaptor of the sequential solver `Lingeling`; `CubeAndConquer` a `WorkingStrategy`, that implements a

⁴ www.labri.fr/perso/lSimon/downloads/softwares/glucose-syrup.tgz.

⁵ www.fmv.jku.at/lingeling/lingeling-bbc-9230380-160707.tar.gz.

Cube-and-Conquer [15]; SimpleSharing already used to define for the `glucose-syrup` like solver. In this case, the underlying sequential solvers are parametrized to export only unit clauses, and only the special worker is a producer.

For the `CubeAndConquer` we choose time to manage rounds because it allows, once one worker has encountered an UNSAT situation, to restart the worker with another guiding-path. In the original implementation, rounds are managed using numbers of conflicts, this makes the reuse of idle CPU much harder.

The implementation of `PaInleSS-treengeling` needed 377 LoC for the adaptor and 249 LoC for `CubeAndConquer`.

Solver “à la Hordesat”. `Hordesat`⁶ is a Portfolio-based solver with a modular design. `Hordesat` uses as sequential engine either `Minisat` [10] or `Lingeling`. It is a Portfolio where the sharing is realized by controlling the flow of exported clauses. Every second, 1500 literals (*i.e.*, sum of the size of the clauses) are exported from each sequential engine. Moreover, we used the `Lingeling` solver and the native diversification of `Plingeling` [5] (a Portfolio solver of `Lingeling`) combined to the random sparse diversification (presented as the best combination by [4]).

The solver “à la Hordesat”, namely `PaInleSS-hordesat`, required the following components: `Lingeling` and `Portfolio` that have been implemented earlier; `HordesatSharing` a `SharingStrategy` that implements the `Hordesat` sharing strategy. This last required only 148 LoC.

Combining Existing Strategies. Based on the implemented solvers, we reused the obtained components to quickly build two new original solvers.

`PaInleSS-treengeling-hordesat`: it is a `PaInleSS-treengeling`-based solver that shares clauses using the strategy of `Hordesat`. The implementation of this solver reuses the `Lingeling`, `CubeAndConquer`, and `HordesatSharing` classes. To instantiate this solver we only needed a special parametrization. Beside, the modularity aspects, by this instantiation, we aimed to investigate the impact of a different sharing strategy on the overall performances of `PaInleSS-treengeling`.

`PaInleSS-treengeling-glucose`: it is a Portfolio solver that mixes `Cube-and-Conquer` of `Lingeling`, and a Portfolio of `Glucose` solvers. Here, `Glucose` workers export unit and *glue* clauses [2] (*i.e.*, clauses with LBD equals to 2) to the other solvers. This last solver reuses the following components: `Lingeling`, `Glucose`, `Portfolio`, `CubeAndConquer`, `SimpleSharing`. Only 15 LoC are required to build the parallelization strategy tree. By the instantiation of this solver, we aimed to study the effect of mixing some parallelisation strategies.

⁶ baldur.iti.kit.edu/hordesat/files/hordesat.zip.

6 Numerical Results

This section presents the results of experiments we realized using the solvers described in Sect. 5: `PaInleSS-glucose-syrup`, `PaInleSS-treengeling`, `PaInleSS-hordesat`, `PaInleSS-treengeling-hordesat`, and `PaInleSS-treengeling-glucose`. The goal here is to show that the introduction of genericity does not add an overhead *w.r.t.* the original solvers.

All the experiments have been executed on a parallel machine with 40 processors Intel Xeon CPU E7- 2860 @ 2.27 GHz, and 500 Go of memory. We used the 100 instances of the parallel track of the SAT Race 2015⁷. All experiments have been conducted using the following parametrisations: each solver has been run once on each instance, with a time-out of 5000 s (as in the SAT Race). We limited the number of involved CPUs to 36.

Table 1. Results of the different solvers. The different columns represent: the number of UNSAT solved instances, SAT solved instances, total solved instances, and the cumulative time spent solving the instances solved by the two solvers.

Solver	UNSAT	SAT	Total	Cum. Time Inter.
<code>glucose-syrup</code>	30	41	71	15 h37
<code>PaInleSS-glucose</code>	32	46	78	13 h18
<code>Treengeling</code>	32	50	82	20 h55
<code>PaInleSS-treengeling</code>	32	50	82	14 h12
<code>Hordesat</code>	31	44	75	15 h05
<code>PaInleSS-hordesat</code>	31	43	74	14 h19

The number of solved instances per solver are reported in Table 1. Globally, these primary results show that our solvers compare well to the studied state-of-the-art solvers. We can deduce that the genericity offered by `PaInleSS` does not impact the global performances. Moreover, on instances solved by both, the original solver and our implementation, the cumulative solving time is in our favor (see column Cum. Time. Inter. in Table 1). A more detailed analysis is given for each solver in the rest of the section.

PaInleSS-glucose-syrup vs. glucose-syrup. Our implementation of the `glucose-syrup` parallelization strategy was able to solve 7 more instances compared to `glucose-syrup`. This concerns both SAT and UNSAT instances as shown in the scatter plot of Fig. 4(a) and, in the cactus plots of Figs. 5(a) and 6(b). This gain is due to our careful design of the sharing mechanism that is decentralized and uses lock-free buffers. Indeed in `glucose-syrup` a global buffer is used to exchange clauses, which requires import/export to use a unique lock, thus introducing a bottleneck. The absence of bottleneck in our implementation increases the parallel all over the execution, explaining our better performances.

⁷ baldur.itl.kit.edu/sat-race-2015/downloads/sr15bench-hard.zip.

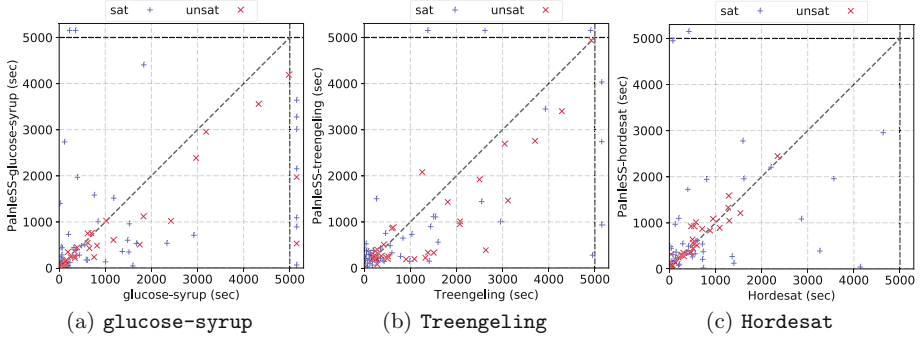


Fig. 4. Scatter plots of PaInleSS's solvers against state-of-the-art ones.




Fig. 5. Cactus plots of SAT instances of PaInleSS's solvers against state-of-the-art ones.




Fig. 6. Cactus plots of UNSAT instances of PaInleSS's solvers against state-of-the-art ones.

PaInleSS-treengeling vs. Treengeling. Concerning **Treengeling**, our implementation has comparable results. Figure 4(b) shows that the average solving time of SAT instances is quite similar, while for the UNSAT instances, our implementation is in average faster. This is corroborated by the cactus plot depicted in Fig. 6(b). This speed up is due to our fine implementation of the Cube-and-Conquer strategy, thus increasing the real parallelism all over the execution and explaining our better performances on UNSAT instances.

PaInleSS-hordesat vs. Hordesat. Although **Hordesat** was able to solve 1 more instance than our tool, results are comparable. Moreover scatter plot of Fig. 4(c), and cactus plots of Figs. 5(c) and 6(c) exhibit quit similar results for the two tools. For instances solved by both tools, our tool was a beat faster and used almost 3000 seconds less as pointed out in Table 1. As the sharing strategy of **Hordesat** is mainly based on two parameters, namely the number of exchanged literals per round, and the sleeping time of sharer by round, we think that a finer tuning of this couple of parameters for our implementation could improve the performances of our tool.




Fig. 7. Scatter plots of the composed solvers against **PaInleSS-treengeling**.

Results of the Composed Solvers. **PaInleSS-treengeling-hordesat** solved 81 instances (49 SAT and 32 UNSAT), and **PaInleSS-treengeling-glucose** solved 81 instances (48 SAT and 33 UNSAT). The scatter plot of the two strategies (Fig. 7), show that these strategies are almost equivalent w.r.t. the original ones. These results allow us to conclude that the introduced strategies do not add any value to the original one.

7 Conclusion

Testing and implementing new strategies for parallel SAT solving has become a challenging issue. Any new contribution in the domain faces the following problems: concurrent programming requires specific skills, testing new strategies required a prohibitive development of a complete solver (either built from scratch or an enhancement of an existing one), an implementation often allows to test only a single composition policy and avoids the evaluation of a new heuristic with different versions of the other mechanisms.

To tackle these problems we proposed **PaInleSS**, a modular, generic and efficient framework for parallel SAT solving. We claimed that its modularity

and genericity allow the implementation of basic strategies, as well as new ones and their combination with a minimal effort and concurrent programming skills.

We have proven our claims, first, by the implementation of strategies present in some state-of-the-art solvers: **glucose-syrup**, **Treengeling**, and **Hordesat**. Second, we reused the developed complements to derive, easily, new solvers that mix strategies. We also show that the instantiated solvers are as efficient as the original one (and even better), by conducting a set experiments using benchmarks of the SAT Race 2015.

As perspectives, we plan to adapt our framework for mutli-machine environments. We also would like to enhance **PaInleSS** with helpful tools to monitor algorithm metrics (*e.g.*, number of shared clauses), system metrics (*e.g.*, synchronization time, load balancing), and to facilitate the debugging work. Another interesting point is the simplification of the instantiation mechanism by providing a domain specific language (DSL).

References

1. 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). doi:[10.1007/978-3-319-44953-1_3](https://doi.org/10.1007/978-3-319-44953-1_3)
2. Audemard, G., Simon, L.: Predicting learnt clauses quality in modern SAT solvers. In: IJCAI, vol. 9, pp. 399–404 (2009)
3. 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). doi:[10.1007/978-3-319-09284-3_15](https://doi.org/10.1007/978-3-319-09284-3_15)
4. 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). doi:[10.1007/978-3-319-24318-4_12](https://doi.org/10.1007/978-3-319-24318-4_12)
5. Biere, A.: Splat, lingeling, plingeling, treengeling, yalsat entering the SAT competition 2016. SAT COMPETITION 2016, p. 44 (2016)
6. 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). doi:[10.1007/3-540-49059-0_14](https://doi.org/10.1007/3-540-49059-0_14)
7. Blochinger, W.: Towards robustness in parallel SAT solving. In: PARCO, pp. 301–308 (2005)
8. Davis, M., Logemann, G., Loveland, D.: A machine program for theorem-proving. Commun. ACM **5**(7), 394–397 (1962)
9. Davis, M., Putnam, H.: A computing procedure for quantification theory. J. ACM **7**(3), 201–215 (1960)
10. Eén, N., Sörensson, N.: An extensible SAT-solver. In: Giunchiglia, E., Tacchella, A. (eds.) SAT 2003. LNCS, vol. 2919, pp. 502–518. Springer, Heidelberg (2004). doi:[10.1007/978-3-540-24605-3_37](https://doi.org/10.1007/978-3-540-24605-3_37)
11. Eén, N., Biere, A.: Effective preprocessing in SAT through variable and clause elimination. In: Bacchus, F., Walsh, T. (eds.) SAT 2005. LNCS, vol. 3569, pp. 61–75. Springer, Heidelberg (2005). doi:[10.1007/11499107_5](https://doi.org/10.1007/11499107_5)
12. Guo, L., Hamadi, Y., Jabbour, S., Sais, L.: Diversification and intensification in parallel SAT solving. In: International Conference on Principles and Practice of Constraint Programming, pp. 252–265. Springer (2010)

13. Hamadi, Y., Jabbour, S., Sais, J.: Control-based clause sharing in parallel SAT solving. In: Hamadi, Y., Monfroy, E., Saubion, F. (eds.) *Auton. Search*, pp. 245–267. Springer, Heidelberg (2011)
14. Hamadi, Y., Jabbour, S., Sais, L.: ManySAT: a parallel SAT solver. *J. Satisfiability Boolean Model. Comput.* **6**, 245–262 (2009)
15. 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). doi:[10.1007/978-3-642-34188-5_8](https://doi.org/10.1007/978-3-642-34188-5_8)
16. 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). doi:[10.1007/11814948_39](https://doi.org/10.1007/11814948_39)
17. Hyvärinen, A.E.J., Junttila, T., Niemelä, I.: Partitioning search spaces of a randomized search. In: Serra, R., Cucchiara, R. (eds.) *AI*IA 2009. LNCS*, vol. 5883, pp. 243–252. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-10291-2_25](https://doi.org/10.1007/978-3-642-10291-2_25)
18. 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). doi:[10.1007/978-3-642-31612-8_17](https://doi.org/10.1007/978-3-642-31612-8_17)
19. Kautz, H.A., Selman, B., et al.: Planning as satisfiability. In: *ECAI*, vol. 92, pp. 359–363 (1992)
20. Lazaar, N., Hamadi, Y., Jabbour, S., Sebag, M.: Cooperation control in Parallel SAT Solving: a Multi-armed Bandit Approach. Research Report RR-8070, INRIA, September 2012
21. Liang, J.H., Ganesh, V., Poupart, P., Czarnecki, K.: Learning rate based branching heuristic for SAT solvers. In: Creignou, N., Le Berre, D. (eds.) *SAT 2016. LNCS*, vol. 9710, pp. 123–140. Springer, Cham (2016). doi:[10.1007/978-3-319-40970-2_9](https://doi.org/10.1007/978-3-319-40970-2_9)
22. 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). doi:[10.1007/11814948_16](https://doi.org/10.1007/11814948_16)
23. Manthey, N.: Coprocessor 2.0 – a flexible CNF simplifier. In: Cimatti, A., Sebastiani, R. (eds.) *SAT 2012. LNCS*, vol. 7317, pp. 436–441. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-31612-8_34](https://doi.org/10.1007/978-3-642-31612-8_34)
24. Manthey, N.: Coprocessor – a standalone SAT preprocessor. In: Tompits, H., Abreu, S., Oetsch, J., Pührer, J., Seipel, D., Umeda, M., Wolf, A. (eds.) *INAP/WLP -2011. LNCS*, vol. 7773, pp. 297–304. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-41524-1_18](https://doi.org/10.1007/978-3-642-41524-1_18)
25. Marques-Silva, J.P., Sakallah, K., et al.: GRASP: a search algorithm for propositional satisfiability. *IEEE Trans. Comput.* **48**(5), 506–521 (1999)
26. Martins, R., Manquinho, V., Lynce, I.: Improving search space splitting for parallel SAT solving. In: *IEEE International Conference on Tools with Artificial Intelligence*, vol. 1, pp. 336–343. IEEE (2010)
27. Massacci, F., Marraro, L.: Logical cryptanalysis as a SAT problem. *J. Autom. Reasoning* **24**(1), 165–203 (2000)
28. Michael, M.M., Scott, M.L.: Simple, fast, and practical non-blocking and blocking concurrent queue algorithms. In: *Proceedings of the Fifteenth Annual ACM Symposium on Principles of Distributed Computing*, pp. 267–275. ACM (1996)
29. Moskewicz, M.W., Madigan, C.F., Zhao, Y., Zhang, L., Malik, S.: Chaff: engineering an efficient SAT solver. In: *38th annual Design Automation Conference*, pp. 530–535. ACM (2001)

30. Ohmura, K., Ueda, K.: **c-sat**: A parallel SAT solver for clusters. In: Kullmann, O. (ed.) SAT 2009. LNCS, vol. 5584, pp. 524–537. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-02777-2_47](https://doi.org/10.1007/978-3-642-02777-2_47)
31. Schulz, S., Blochinger, W.: Cooperate and compete! a hybrid solving strategy for task-parallel SAT solving on peer-to-peer desktop grids. In: High Performance Computing and Simulation, pp. 314–323. IEEE (2010)
32. Silva, J.P.M., Sakallah, K.A.: GRASP—a new search algorithm for satisfiability. In: IEEE/ACM International Conference on Computer-Aided Design, pp. 220–227. IEEE (1997)
33. Sonobe, T., Inaba, M.: Portfolio with block branching for parallel SAT solvers. In: Nicosia, G., Pardalos, P. (eds.) LION 2013. LNCS, vol. 7997, pp. 247–252. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-44973-4_25](https://doi.org/10.1007/978-3-642-44973-4_25)
34. Zhang, H., Bonacina, M.P., Hsiang, J.: PSATO: a distributed propositional prover and its application to quasigroup problems. *J. Symbolic Comput.* **21**(4), 543–560 (1996)
35. Zhang, L., Madigan, C.F., Moskewicz, M.H., Malik, S.: Efficient conflict driven learning in a boolean satisfiability solver. In: IEEE/ACM International Conference on Computer-Aided Design, pp. 279–285. IEEE Press (2001)