

Clause Sharing and Partitioning for Cloud-Based SMT Solving*

Matteo Marescotti, Antti E. J. Hyvärinen, and Natasha Sharygina

Università della Svizzera Italiana, Switzerland

Abstract

Satisfiability modulo theories (SMT) allows the modeling and solving of constraint problems arising from practical domains by combining well-engineered and powerful solvers for propositional satisfiability with the expressivity of higher-order logics in a natural way. The increasing popularity of SMT as a modelling approach means that the SMT solvers need to be capable of handling increasingly complex constraint problems. This paper studies how SMT solvers can use cloud computing to scale to challenging problems through sharing of learned information in the form of clauses with approaches based on both divide-and-conquer and algorithm portfolios. Our initial experiments, executed on the OpenSMT2 solver, show that parallelization with clause sharing results in an average three-fold speed-up on our benchmark set and the speed-up is bigger on the harder instances.

1 Introduction

The *Satisfiability Modulo Theories* (SMT) [5] approach to constraint solving consists of determining whether a logical formula is satisfiable, given that some of the Boolean variables have an interpretation in *background theories*. The expressiveness of SMT makes it suitable for a vast range of application domains, including software and hardware model checking [8, 4], bioinformatics [23], and optimization [21], and has recently attained significant interest from a wide range of users. The computational cost of solving SMT instances can be very high, given that already propositional satisfiability is an NP-complete problem and the introduction of background theories can only make the problem harder. The SMT solvers tackle complexity with a tightly integrated loop where the SAT solver attempts to find a satisfying solution and queries the validity of a candidate solution from the theory solvers. In case the candidate solution is shown to be invalid the theory solvers and the SAT solver work together to extract new expressive constraints in the form of learned clauses, by combining theory specific information and resolution.

This work studies how employing parallelism and in particular cloud computing can be used in helping SMT solvers to scale to increasingly hard problems. We study two different approaches: a portfolio where several copies of a randomized SMT solver is run on a single instance; and an approach where the SMT instance is divided into several partitions that are guaranteed by construction not to share models, and each partition is solved by an SMT solver. We combine these two approaches in a natural way by having several SMT solvers work on each partition. The emphasis of the work is in how the different ways of organizing the search can co-operate to speed up the solving. We implement the co-operation by having SMT solvers working on the same partition share the clauses they learn during the execution.

To study the effects of parallelism and clause sharing we implement the approach using the SMT solver OpenSMT2 [14] and experiment with two particularly central background theories, the quantifier-free theories of uninterpreted functions and equalities [5] (QF_UF) and linear real

*This paper is presentation-only and an extended version is currently under review elsewhere.

arithmetics [6] (QF_LRA). The experimental results suggest that both the portfolio and the partitioning based approach can greatly benefit from clause sharing. Interestingly, a comparison between portfolio and partitioning reveals that the portfolio approach performs better even if the partitioning is combined with portfolio.

Related Work. The portfolio approach combined with clause sharing has been implemented using the SMT solver Z3 [22]. The implementation provides an efficient clause sharing strategy within the same machine using lockless queues that hold references to the lemmas that a solver core wants to export. The experimental evaluations show that clause sharing leads to a substantial speedup on benchmarks from the QF_IDL logic. In contrast to this work, we support two SMT theories (QF_UF and QF_LRA), and exploit the advantages of combining portfolio with search-space partitioning. Moreover our implementation is designed to run in a cluster or a cloud in addition to a single machine. Similarly to Z3, the SMT solver CVC4 [3] supports a portfolio-style parallel solving. Unlike our approach, the approach used in CVC4 is designed to run in a single computer and does not implement clause sharing.

In [13] we introduced the parallelization tree formalism for combining portfolio and search-space partitioning. The work also describes and reports results on the QF_UF logic on some instantiation of the framework. Here we extend this work by introducing clause sharing and the logic QF_LRA.

A divide-and-conquer approach for the quantifier-free bit-vector logic has been implemented on top of the SMT solver Boolector [20]. A portfolio parallelization approach for the logic of quantifier-free bit-vectors and bit-vector arrays is presented in [19]. Compared to these, our work differs in the supported theories and in that we support cloud computing and are not limited to pure divide-and-conquer or portfolio approach.

There is a substantial body of highly relevant research on parallel SAT solving, overviewed for instance in [16]. Among relevant works are the portfolio combined with clause sharing implemented in the ManySAT solver [7], and more recently with a higher degree of parallelism by HordeSAT [2]. The combination of search-space partitioning and clause sharing has been recently studied in [1].

In this paper we first introduce the basic concepts required for interpreting our results in Sec. 2, and then describe our implementation framework where we combine the concepts in Sec. 3. Some experimental results obtained with the implementation are presented in Sec. 4, and conclusions are drawn in Sec. 5.

2 Background

A literal is a Boolean variable x or its negation $\neg x$. A clause is a set of literals and a propositional formula in conjunctive normal form (CNF) is a set of clauses. Throughout the text we use both a set of literals and disjunction, and a set of clauses and a conjunction, interchangeably. An assignment σ is a set of literals such that for no variable x , both $x \in \sigma$ and $\neg x \in \sigma$. A variable x is *assigned* if either $x \in \sigma$ or $\neg x \in \sigma$. An assignment σ satisfies a clause c if $\sigma \cap c \neq \emptyset$ and a formula F if it satisfies all its clauses.

Most SMT solvers are based on the DPLL(T) framework [18] which takes as input a problem instance presented as a propositional formula where some of the Boolean variables have an interpretation as Boolean relations, such as equalities, disequalities, and inequalities, in a theory T . A DPLL(T) solver consists of a solver for the propositional satisfiability problem (SAT) and one or more theory solvers that are capable of reasoning on a conjunction of Boolean relations over the theory T . In the pre-processing phase the input formula is converted into an

equisatisfiable propositional formula F in CNF while maintaining the special T -interpretations of the Boolean variables.

The SMT solving process is driven by a SAT solver maintaining a set of clauses which initially consists of the formula F . During the search the SAT solver builds an assignment σ and alternates between two phases.

- In the *propagation* phase the solver identifies clauses $c = l_1 \vee \dots \vee l_n$ such that (i) there is a single unassigned literal $l_i \in c$, and (ii) σ falsifies the literals $l_j, 1 \leq j \leq n, j \neq i$. Any such literal l_i is added to σ until no new clauses satisfying (i) and (ii) can be found.
- In the *decision* phase the solver chooses a literal l_i unassigned in σ and adds it to σ .

A *conflict* occurs if during the propagation phase the SAT solver detects a clause with all literals falsified. The decisions and propagations are stored in the *implication graph* [15], a directed graph having as nodes the literals of the assignment σ and as edges the arches $\{(-l_j, l_i) \mid 1 \leq j \leq n, j \neq i\}$ obtained in the propagation phase.¹

A SAT solver *learns* a clause c by performing essentially resolution steps directed by the implication graph when it finds a conflict. The learned clauses are by construction guaranteed to be logical consequences of F , and are both used in guiding the search and added temporarily to the clause database to reduce the number of assignments the solver needs to cover during the search. Finally the solver makes with a decreasing frequency a *restart* where the assignment σ is cleared and the search is continued without otherwise changing the state of the solver.

The SAT solver queries periodically whether the conjunction of the theory atoms in σ is consistent with the theory. In case a theory solver determines an inconsistency it identifies a subset σ' of σ that causes the inconsistency and returns the clause $c_T := \{-l \mid l \in \sigma'\}$ to the SAT solver. Minimizing σ' is critical for the good performance of the SMT solvers (see, e.g., [17, 6]). The clause c_T is used together with the implication graph to learn a clause c in the way described above for clause learning. The solving process terminates when either the clause database becomes unsatisfiable or a satisfying assignment consistent with the theory T is found.

Parallel algorithm portfolios. An algorithm portfolio [9] is a set of algorithms that compete in finding a solution of a given problem. The decision phase employs a heuristic for choosing l_i and introduces in a natural way nondeterminism into the solver. Such changes to the heuristic often result in significant variation in the run time of the solver [11] and therefore a natural algorithm portfolio can be obtained by seeding differently the pseudo-random-number generator of the SMT solver and running several solvers in parallel.

Clause sharing. In clause sharing the clauses learned by an SMT solver while solving a formula F are distributed among the solvers in the parallel portfolio. Since clause learning plays an important part of the SMT solving process this sharing can speed up the parallel solving process. For example the shared clauses make it easier to produce the required clauses in case of unsatisfiability, and reduce the number of assignments the solver covers before finding a satisfying assignment.

Search space partitioning. The SMT solver bases its search on the SAT solver, and therefore a natural way of dividing the work and avoiding overlap for the solvers is to constrain the SMT formula F into partitions F_1, \dots, F_n such that the original formula is satisfiable if and

¹We equate x and $\neg\neg x$.

only if one of F_1, \dots, F_n is satisfiable. This can be done through adding additional constraints C_i to the formula F , resulting in the partition $F_i := F \wedge C_i$. While in principle, for example conjoining a single literal $\{l\}$ to the formula F halves the search space, this happens rarely in practice. Often the resulting partitions F_i will have overlap in their search due to the heuristics of the solver and therefore the observed speedup will be less dramatic. The situation is made worse by the unpredictability of the SMT solver run time in case the instance is unsatisfiable. Assuming the shape of the run-time distribution is the same for both the instance F and the partition F_i , it can be shown that independent of the number of partitions there are distributions for which the expected run time increases when partitioning is done as described above [12]. To lessen this effect several more complex parallelization algorithms combining elements from search-space partitioning and algorithm portfolios have been suggested [13]. For example running a parallel portfolio for each partition makes it less likely that one of the partitions will require excessive time for being solved.

Constructing partitions. We use in this work an approach for constructing partitions called *scattering*, initially introduced in [10]. Given a formula F , a number of partitions to be created n , and a sequence d_1, \dots, d_{n-1} the partitions are obtained following the iteration

$$\begin{aligned} F_1 &:= F \wedge l_1^1 \wedge \dots \wedge l_{d_1}^1 \\ F_k &:= F \wedge (\neg l_1^1 \vee \dots \vee \neg l_{d_1}^1) \wedge \dots \wedge (\neg l_1^{k-1} \vee \dots \vee \neg l_{d_{k-1}}^{k-1}) \wedge l_1^k \wedge \dots \wedge l_{d_k}^k \\ F_n &:= F \wedge (\neg l_1^1 \vee \dots \vee \neg l_{d_1}^1) \wedge \dots \wedge (\neg l_1^{n-1} \vee \dots \vee \neg l_{d_{n-1}}^{n-1}) \end{aligned}$$

The goal of using the sequence d_i is to make the search space of each F_i as close as possible to $1/n$ of the search space of the instance F . We obtain the sequence d_i by assuming that conjoining a disjunction of k literals with a formula F reduces the size of the search space by factor of $(1 - 1/2^k)$. For example for $n = 2$ partitions this gives $d_1 = 1$, while for $n = 8$ we get the sequence $d_1 = 3, d_2 = 3, d_3 = 3, d_4 = 3, d_5 = 2, d_6 = 2$, and $d_7 = 1$. Note that the number of constructed partitions in this method does not have to be a power of 2. Finally, the literals l_i^j are chosen using the same heuristic the SMT solver uses during the search.

3 The parallelization Framework

In this section we present a framework that we use in the experiments of the paper for studying the effect of parallelization approaches and clause sharing. We give an overview of the framework in Fig. 1. The framework is built on the assumption that it will be run in a cluster of machines or a cloud, even though it is also possible to run the system on a single computer. The design follows a client-server approach in which the server, acting as a front-end to the user, receives input instances in the SMTLIB2 format², and at the same time handles the connection with the clients, managing client failures and asynchronous connection of new clients gracefully. The clients are implemented as SMT solvers wrapped by a network layer that handles the connection with the server.

The framework is designed highly modular in order to allow to study the effects of the components. For example the system can handle several different policies for scheduling the partitions among the possibly high number of solvers in the cluster. This allows to research on how to best combine portfolio and search-space partitioning.

The server works in two modes: depending on the configuration it either splits the instance into several partitions using the scattering approach described in Sec. 2, or runs in a pure

²<http://smtlib.cs.uiowa.edu>

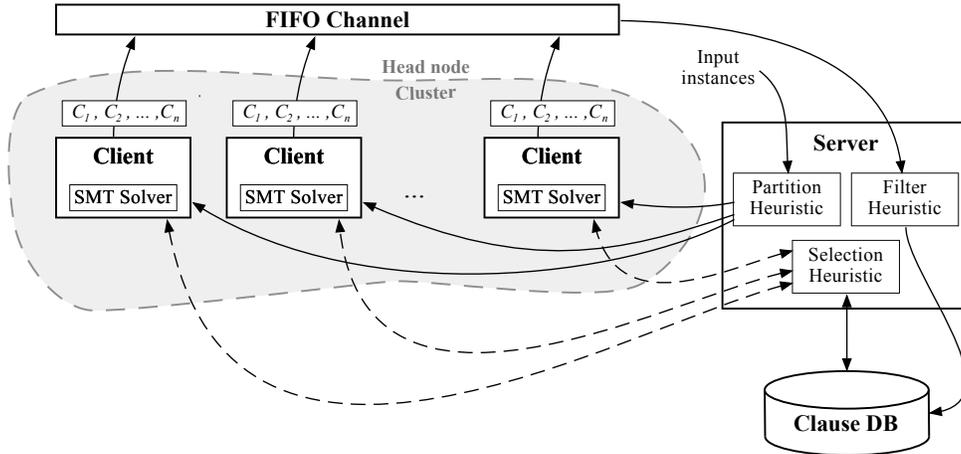


Figure 1: Parallel SMT solver framework with clause sharing

portfolio mode without splitting. Once the initial phase is over, the server distributes either the partitions or the original formula to the solvers in all the available clients.

During the solving phase, each client contacts the server both to publish newly learned clauses and to request new interesting clauses published by the clients. To avoid problems with the high throughput we use a *FIFO Channel* that allows multiple clients to push clauses to the server without turnaround delays. Once a batch of clauses is received, the server uses the *filter heuristic* to choose potential clauses for merging them with the previously received clauses into the clause database (*Clause DB* in Fig. 1). Hence, at any given time, that database will contain all the learned clauses sent by the clients and that have passed the filter heuristic of the server. Inside the clause database the clauses are divided by partitions and each client has only access to the clauses published by the solvers working on the same partition.

Upon a restart each client will ask for clauses from the server. However, the request cannot be replied by sending the entire clause set of the partition that being solved, because it usually consists of prohibitively many clauses. A high number of clauses slows down the client solver since the overhead related to the growth of the internal data structure is higher than the speedup obtained from the clauses. To address the problem the framework allows the use of the *selection heuristic* which attempts to choose clauses that are particularly promising for the problem at hand. The connection between the clients and the selection heuristic is the only one that does not use the push mechanism, but instead needs a pull request. This choice was done since the schedule at which the clients can receive clauses is unpredictable. To indicate the pull request the interaction in Fig. 1 is drawn with dashed lines.

The modular approach also makes it easy to study which clauses are worth sharing. This can be done through experimenting with different filter and selection heuristics inside the server. The current version implements naïve heuristics for both filtering and selection. The filtering heuristic is based on the number of literals inside each clause: clauses with more than a fixed number of literals will be discarded. The selection heuristic works by randomly selecting a fixed number of clauses from the database, and each new set will replace the old one inside the solver.

For the implementation we have had to handle several challenges. The goal of the implementation is to provide scalability, fault-tolerance, and low latencies during data transfer, as

well as ease of use and portability from a cluster of machines to a single machine with many cores / CPUs. In the rest of this section we will present some details about our choices during the implementation.

The SMT solver. While the abstract framework supports any DPLL(T) solver, our current implementation of the framework uses the SMT solver OpenSMT2 [14]. OpenSMT2 is a light-weight implementation that currently supports the quantifier-free theories of uninterpreted functions with equalities (QF_UF) and linear real arithmetics (QF_LRA). The solver is written in C++ and has been developed in the University of Lugano since 2008. The code is easily approachable because of its limited size and the object-oriented architecture. It is released under the MIT license³ and most recently competed in the SMT competitions in 2014 and 2015, performing in the mid-range in the competition. The code base is relatively efficient featuring efficient low-level memory management and a cache-friendly design for many of the central algorithms.

Networking. The server and the clients communicate using our custom-built protocol through TCP/IP sockets, making the solution light-weight, portable, easy to modify and easy to use. Moreover it achieves our goal of portability by being compatible with clusters, cloud computing and single computers.

A binary format for SMT. In order to use clause sharing we must ensure that the internal clausal representation of each instance is the same in every client solver. This property cannot be guaranteed by the SMTLIB2 language since small changes in the input formula might result in subtle optimizations that will dramatically change the CNF structure seen by the SAT solver embedded in the SMT solver. For that reason we designed a binary format for SMT, representing the internal state of OpenSMT2. This format is used for data transfers between each client and the server but also results in us being forced to limit ourselves to a specific SMT solver.

FIFO channel and clause DB. For these challenges we use REDIS⁴, an open source in-memory data structure store, used as database, cache, and message broker. In order to get a scalable, fast, and fault-tolerant push connection from multiple sources we use the *Publisher / Subscriber messaging paradigm* of REDIS. The clauses are stored using the REDIS *set feature* that automatically handles cases where a clause would be added to clause database that is syntactically equal to an already present clause. The set feature is used by both the filter and the selection heuristics.

4 Experiments

This section gives a brief experimental evaluation of the implementation, available at <http://verify.inf.usi.ch/opensmt>. For lack of space we only present results on the effects of clause sharing and the difference in portfolio and the divide-and-conquer approach.

Our hardware configuration is kept the same in all experiments we run. The experiments were run in a cluster where we used eight compute nodes for the clients and the head node for the server. Each compute node is equipped with two CPU Quad-Core AMD Opteron 2384

³<http://opensource.org/licenses/MIT>

⁴<http://redis.io>

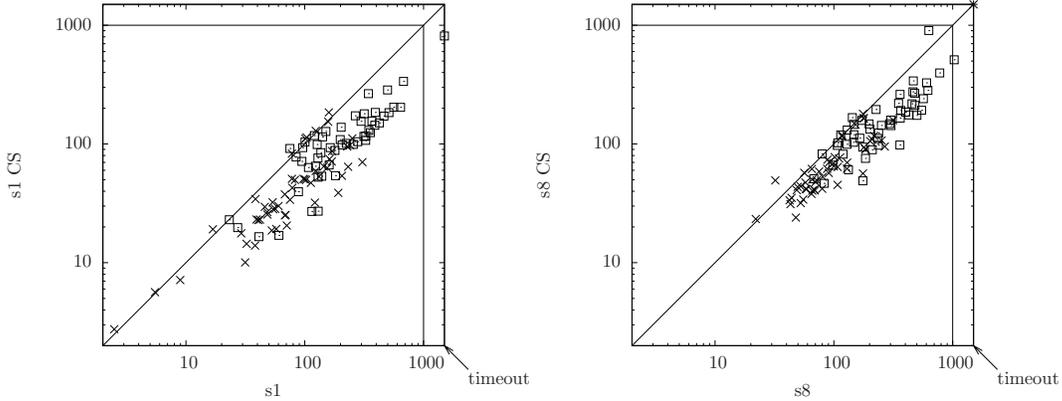


Figure 2: Using clause sharing against not using clause sharing with pure portfolio and 8 partitions. Framework run with 64 solvers on the QF_LRA benchmark set

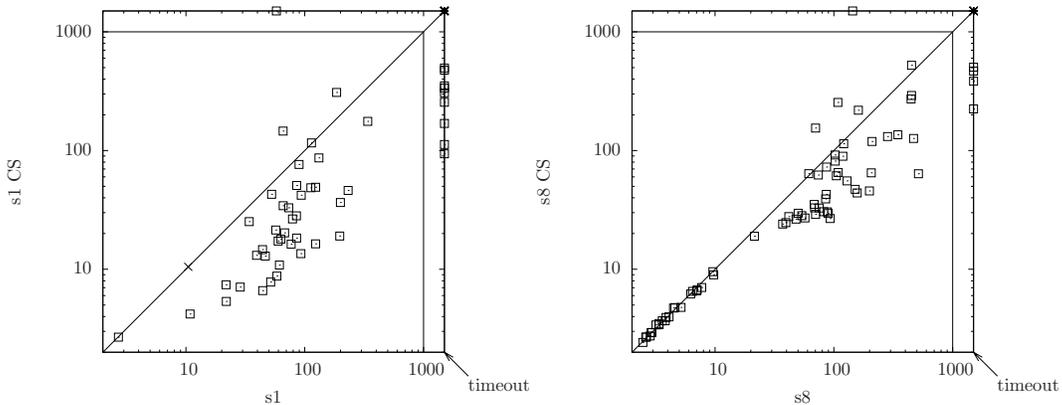


Figure 3: Using clause sharing against not using clause sharing with pure portfolio and 8 partitions. Framework run with 64 solver on the QF_UF benchmark set

and 16GB of RAM. During the experiments each node had eight client processes implementing the SMT solver OpenSMT2, resulting in total of 64 solvers in the entire cluster. The timeout is fixed everywhere to 1000 seconds. In the experiments running the search-space partitioning heuristic we used the scattering approach described in Sec. 2.

We used a fixed benchmark set obtained from the SMTLIB2 benchmarks repository⁵ and the QF_LRA and QF_UF theories. The set from the QF_LRA theory was created by selecting the instances with an average sequential execution time between 100 and 1000 seconds (including those in timeout) using OpenSMT2; the set consists of 106 instances in total. The benchmark set for the QF_UF theory consists of 254 instances. This set includes all instances that could be solved with sequential OpenSMT2 between 100 and 1000 seconds; 11 instances which are

⁵<http://smtlib.cs.uiowa.edu/benchmarks.shtml>

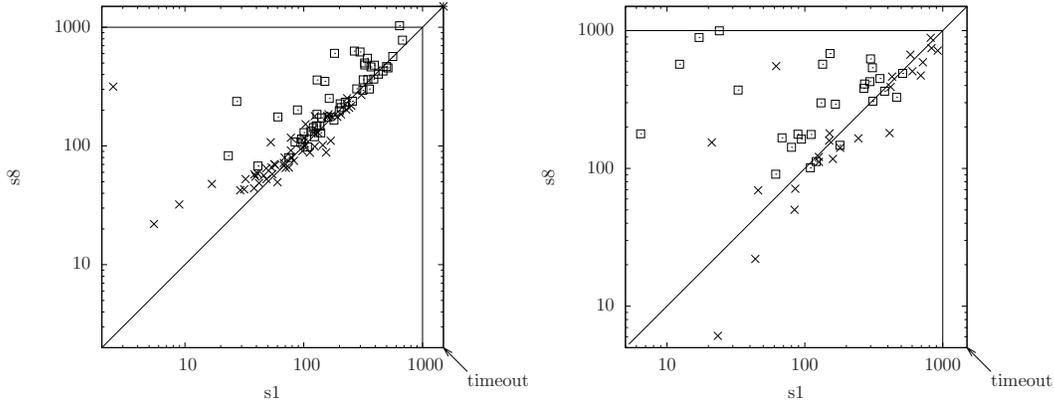


Figure 4: Comparing pure portfolio of 64 solvers against the partitioning in 8 part and using 8 portfolio solver on each part on the QF_LRA benchmark set. The graph on the left shows the results using the framework of Sec. 3, the graph on the right shows a controlled experiment where the partitions are constructed offline and the portfolio size for each partition is fixed. The benchmark sets are different on the left and on the right.

known to be difficult for OpenSMT2 and time out in 1000 seconds; and 200 randomly chosen instances of which half are guaranteed to be satisfiable and the other half unsatisfiable.

In the figures we use the labels $S1$, $S2$ and $S8$ to indicate the number of partitions created from the input instance. Therefore the label $S1$ indicates the pure portfolio approach. The label CS indicates that clause sharing is used. We use the convention that the symbol \times denotes a satisfiable instance and the symbol \square denotes an unsatisfiable instance.

The effect of clause sharing. Our first experiments show how sharing clauses affects the solving time using different partitioning methods for QF_LRA (Fig. 2) and QF_UF (Fig. 3). For both figures the graph on the left shows how the parallelization algorithm based on pure portfolio benefits most from clause sharing: with both theories it gives a 2.05 times speedup, as well as 1 more QF_LRA instance and 9 more QF_UF instances solved within the timeout compared to not using clause sharing. The effect of clause sharing is slightly less dramatical for the approach combining portfolio and search-space partitioning: the speedup due to clause sharing is 1.67 for partitioning in 8 parts and solving each partition with a portfolio of 8 solvers ($S8$). A possible reason for this is that the number of learned clauses available inside the clause database for a single portfolio is bigger when there are more solvers running in the portfolios, and therefore also the quality of clauses that the heuristic picks is higher.

The effect of partitioning Figure 4 (*left*) compares the portfolio approach against the approach where an instance is split into 8 partitions in the framework, from the QF_LRA benchmark set. Interestingly the portfolio approach is almost consistently better than the approach using partitioning, in particular for the unsatisfiable instances. To study this effect in more detail and to rule out effects such as network delays or time used in constructing the partitions, we designed a second experiment in more controlled setting (Fig. 4 (*right*)). For this experiment we chose a set of instances that require more than 1000 seconds to solve using the

sequential version of OpenSMT2. The instances were split off-line into 8 partitions and each partition was solved with a portfolio of 8 instances of OpenSMT2. The time reported in Fig. 4 (*right*) is the maximum of the fastest solver for each portfolio for the unsatisfiable instances and the fastest solver for the satisfiable instances, mimicking the implementation of the framework. Similarly the portfolio approach was run with 64 solvers and the reported time does not include the time required to transfer the instance to the computing node.

The more controlled experiment verifies the phenomenon that an approach based on partitioning performs worse in particular in the unsatisfiable instances, while the results seem to be better for many of the satisfiable instances. This behavior is often observed when the shape of the distribution of an unsatisfiable instance has a significant probability mass at the end of the scale but still a probability mass at very low run times [12]. In such cases the effect of partitioning that tends to make instances easier to solve is not enough to compensate the benefit from a pure portfolio approach.

5 Conclusions

SMT solving in cloud environments is a challenging topic which so far has received relatively little attention from the community developing and researching SMT solvers. This paper addresses the challenges related to integrating one of the key components of a modern SMT solver, the sharing of learned clauses, to parallel SMT solving algorithms. We provide a generic framework for clause sharing in a cloud computing environment and implement a system that supports clause sharing with parallelization algorithms based on both a portfolio and splitting the input formula into partitions.

The framework and the parallelization algorithms are agnostic to the underlying theories used by the SMT solver. We provide results on two fundamental theories: the quantifier-free theories of uninterpreted functions and equality (QF_UF), and linear real arithmetics (QF_LRA). The results show that both theories can benefit significantly from clause sharing. In the experiments we also observe that the partitioning approach, while working relatively well for QF_UF, performs somewhat worse for QF_LRA in the benchmark set we study on this paper. We conjecture that this results from the partitioning heuristic behaving in an unexpected way where the problems sometimes get more difficult to solve, in combination with a run-time increasing phenomenon observed with unsatisfiable instances.

Future work. The framework we set in this paper opens several interesting research directions. In particular we point out two central open questions we plan to address in the future: how to construct a good heuristic for (i) partitioning for QF_LRA and (ii) for filtering and selecting the clauses to be shared in a portfolio.

References

- [1] G. Audemard, B. Hoessen, S. Jabbour, and C. Piette. DoliUS: A distributed parallel SAT solving framework. In D. L. Berre, editor, *POS-14*, volume 27 of *EPiC Series*, pages 1–11. EasyChair, 2014.
- [2] T. Balyo, P. Sanders, and C. Sinz. HordeSat: A massively parallel portfolio SAT solver. In *Proc. SAT 2015*, volume 9340 of *LNCS*, pages 156–172. Springer, 2015.
- [3] C. Barrett, C. L. Conway, M. Deters, L. Hadarean, D. Jovanovic, T. King, A. Reynolds, and C. Tinelli. CVC4. In *Proc. CAV 2011*, volume 6806 of *LNCS*, pages 171 – 177. Springer, 2011.

- [4] A. Biere, A. Cimatti, E. M. Clarke, and Y. Zhu. Symbolic model checking without BDDs. In *Proc. TACAS '99*, volume 1579 of *LNCS*, pages 193–207. Springer, 1999.
- [5] D. Detlefs, G. Nelson, and J. B. Saxe. Simplify: a theorem prover for program checking. *Journal of the ACM*, 52(3):365–473, 2005.
- [6] B. Dutertre and L. M. de Moura. A fast linear-arithmetic solver for DPLL(T). In *Proc. CAV 2006*, volume 4144 of *LNCS*, pages 81–94. Springer, 2006.
- [7] Y. Hamadi, S. Jabbour, and L. Sais. ManySAT: a parallel SAT solver. *Journal on Satisfiability Boolean Modeling and Computation*, 6(4):245 – 262, 2009.
- [8] T. A. Henzinger, R. Jhala, R. Majumdar, and G. Sutre. Software verification with Blast. In *Tenth International Workshop on Model Checking of Software (SPIN)*, volume 2648 of *LNCS*, pages 235–239. Springer, 2003.
- [9] B. A. Huberman, R. M. Lukose, and T. Hogg. An economics approach to hard computational problems. *Science*, 275(5296):51–54, 1997.
- [10] A. E. J. Hyvärinen, T. Junttila, and I. Niemelä. A distribution method for solving SAT in grids. In *Proc. SAT 2006*, volume 4121 of *LNCS*, pages 430–435. Springer, 2006.
- [11] A. E. J. Hyvärinen, T. A. Junttila, and I. Niemelä. Incorporating clause learning in grid-based randomized SAT solving. *Journal on Satisfiability Boolean Modeling and Computation*, 6(4):223–244, 2009.
- [12] A. E. J. Hyvärinen, T. A. Junttila, and I. Niemelä. Partitioning search spaces of a randomized search. *Fundamenta Informaticae*, 107(2-3):289–311, 2011.
- [13] A. E. J. Hyvärinen, M. Marescotti, and N. Sharygina. Search-space partitioning for parallelizing SMT solvers. In *Proc. SAT 2015*, volume 9340 of *LNCS*, pages 369–386. Springer, 2015.
- [14] A. E. J. Hyvärinen, M. Marescotti, and N. Sharygina. OpenSMT2: An SMT solver for multi-core and cloud computing. In *Proc. SAT 2016*, LNCS. Springer, 2016. To appear.
- [15] J. P. Marques-Silva and K. A. Sakallah. GRASP: A search algorithm for propositional satisfiability. *IEEE Transactions on Computers*, 48(5):506–521, 1999.
- [16] R. Martins, V. M. Manquinho, and I. Lynce. An overview of parallel SAT solving. *Constraints*, 17(3):304–347, 2012.
- [17] R. Nieuwenhuis and A. Oliveras. Proof-producing congruence closure. In *Proc. RTA 2005*, volume 3467 of *LNCS*, pages 453 – 468. Springer, 2005.
- [18] R. Nieuwenhuis, A. Oliveras, and C. Tinelli. Solving SAT and SAT modulo theories: From an abstract Davis–Putnam–Logemann–Loveland procedure to DPLL(T). *Journal of the ACM*, 53(6):937 – 977, 2006.
- [19] H. Palikareva and C. Cadar. Multi-solver support in symbolic execution. In *Proc. CAV 2013*, volume 8044 of *LNCS*, pages 53–68. Springer, 2013.
- [20] C. Reisenberger. PBoolector: a parallel SMT solver for QF_BV by combining bit-blasting with look-ahead. Master’s thesis, Johannes Kepler Univesität Linz, Linz, Austria, 2014.
- [21] R. Sebastiani and S. Tomasi. Optimization modulo theories with linear rational costs. *ACM Transactions on Computational Logic*, 16(2):12:1–12:43, 2015.
- [22] C. M. Wintersteiger, Y. Hamadi, and L. M. de Moura. A concurrent portfolio approach to SMT solving. In *Proc. CAV 2009*, volume 5643 of *LNCS*, pages 715–720. Springer, 2009.
- [23] B. Yordanov, C. M. Wintersteiger, Y. Hamadi, and H. Kugler. SMT-based analysis of biological computation. In *Proc. NFM 2013*, volume 7871 of *LNCS*, pages 78–92. Springer, 2013.