

Seven Challenges in Parallel SAT Solving

Youssef Hamadi and Christoph M. Wintersteiger

■ *This article provides a broad overview of the state of the parallel SAT solving field. A set of challenges to researchers is presented that, we believe, must be met to ensure the practical applicability of parallel SAT solvers in the future. All these challenges are described informally but put into perspective with related research results, and a (subjective) grading of difficulty for each of them is provided.*

*Parallelism is the wave of the future ...
and always will be.*

This is a famous saying in the parallel computing community. It conveys a general sentiment that the coming of parallel architectures would forever be delayed. This was indeed true at a time when clock-speed growth seemed always possible, allowing sequential code seamlessly to become faster. This remained true until the thermal wall¹ stopped this free lunch scenario. Chip makers had only one way to escape: packing multiple processing units on a single processor in order to provide support for parallelism. The future was there, and that's when problems started for programmers.

Parallelizing code is not straightforward and beyond mere conceptual difficulties; for example, which part should be parallelized?, it includes low-level technicalities like race conditions, deadlocks, starvation, and nondeterminism, all of which must be taken into consideration in parallel algorithm design and implementation.

Historically, the parallel computing community quickly adopted combinatorial search as a playground for applications. Search algorithms have the advantage of being conceptually simple (think of the most basic backtrack-style algorithm) and computationally demanding due to the (usually) exponential size of the search space. In contrast, the search community did not really focus its research on parallelizing. The lack of proper

SAT

SAT is the problem of determining whether the variables of a given Boolean formula can be assigned in such a way as to make the formula evaluate to true. If no such assignment exists, the function expressed by the formula always evaluates to false. In this latter case, the formula is called unsatisfiable; otherwise it is called satisfiable. SAT was the first known example of an NP-complete problem (Garey and Johnson 1979). Briefly, this means that there is no known algorithm that efficiently solves all instances of SAT and it is generally believed (but not proven; see P versus NP problem) that no such algorithm can exist.

infrastructure and, for many, the feeling that sequential algorithms were still full of research opportunities can go toward explaining that. In that community, parallelism was often only put in the perspectives of papers with no real perspectives. This led to a situation where parallel search algorithms were designed by people with only one part of the required skills.

Most computational problems solved on a computer have a deterministic nature. Sometimes, these problems can be large, and divide-and-conquer parallelism is a suitable approach. In that context, if the overhead of dividing is well controlled, linear or close to linear speedups are possible. When parallel computing researchers started to address search, they reused their main concept and tried the most efficient way to apply divide-and-conquer techniques. Research was often about crafting the best load-balancing strategies in order to avoid the starvation problem, while minimizing the overhead.

Search problems are intrinsically nondeterministic, and this very particular nature was indeed discovered by the aforementioned community. They encountered this fact in the form of observing superlinear speed-ups, which was so unusual to them that they called them speedup anomalies (Pruul and Nemhauser 1988, Rao and Kumar 1993).

In divide-and-conquer parallel search, superlinear speedups are indeed possible when the sequential algorithm is poorly driven by its heuristics and when the division of the search space artificially brings solutions to the beginning of a subspace. This means that a sequential search algorithm does not need to exhaust the search space to find a solution or often even when proving that a problem has no solution, as is the case with conflict-driven solvers (Moskewicz, Madigan et al. 2001).

By 2005, it was apparent that the thermal wall had been hit and that processor speed would not continue to increase as before. This gradually prompted the interest of search researchers who then started seriously to consider parallelism as a path into the future.

Boolean satisfiability (SAT), that is, the problem of determining whether a Boolean formula can evaluate to true, benefits from very mature and advanced algorithms with large practical impact. Application and research domains like software and hardware verification, automated planning, computational biology, and many others benefit from modern SAT solvers. These domains have large and difficult instances that provide the SAT community with meaningful benchmarks.

Most of the following challenges are general in such a way that the questions they raise should positively affect not only research in parallel SAT but in parallel search in general. We first present the current situation in sequential and parallel SAT solving and then give a set of challenges. Each of these challenges comes with an overly optimistic estimate of its inherent difficulty represented as black circles, where we would estimate that every black circle represents, roughly, about two years of research.

Context: Sequential SAT Solvers

State-of-the-art solvers extend the original Davis, Putnam, Logemann, and Loveland (DPLL) procedure (Davis, Logemann, and Loveland 1962) with conflict analysis (Zhang et al. 2001). The general architecture of such conflict-directed clause-learning solvers (CDCL) is presented in figure 1. These procedures include an optional preprocessing step 0, which performs variable elimination and clause subsumption checks in order to reduce the size of the formula and improve the performance of the search process (Eén and Biere 2005). The search then repeatedly creates tree nodes by setting the truth value of a literal (a Boolean variable or its negation). This assignment is used to trigger an inference step 1 that deduces and propagates some forced unit literal assignments. This is recorded in the implication graph, a central data structure, which stores the partial assignment together with its implications. This branching process is repeated until it finds a model or reaches a conflict. In the first case, the formula is answered to be satisfiable, and the model is reported. In the second case, a conflict clause is generated. This is performed by the conflict analysis component through a bottom-up traversal of the implication graph and resolution of clauses encountered during this traversal (step 2). It stops when a conflict clause containing only one literal from the current decision level is generated. Such a conflict clause (or

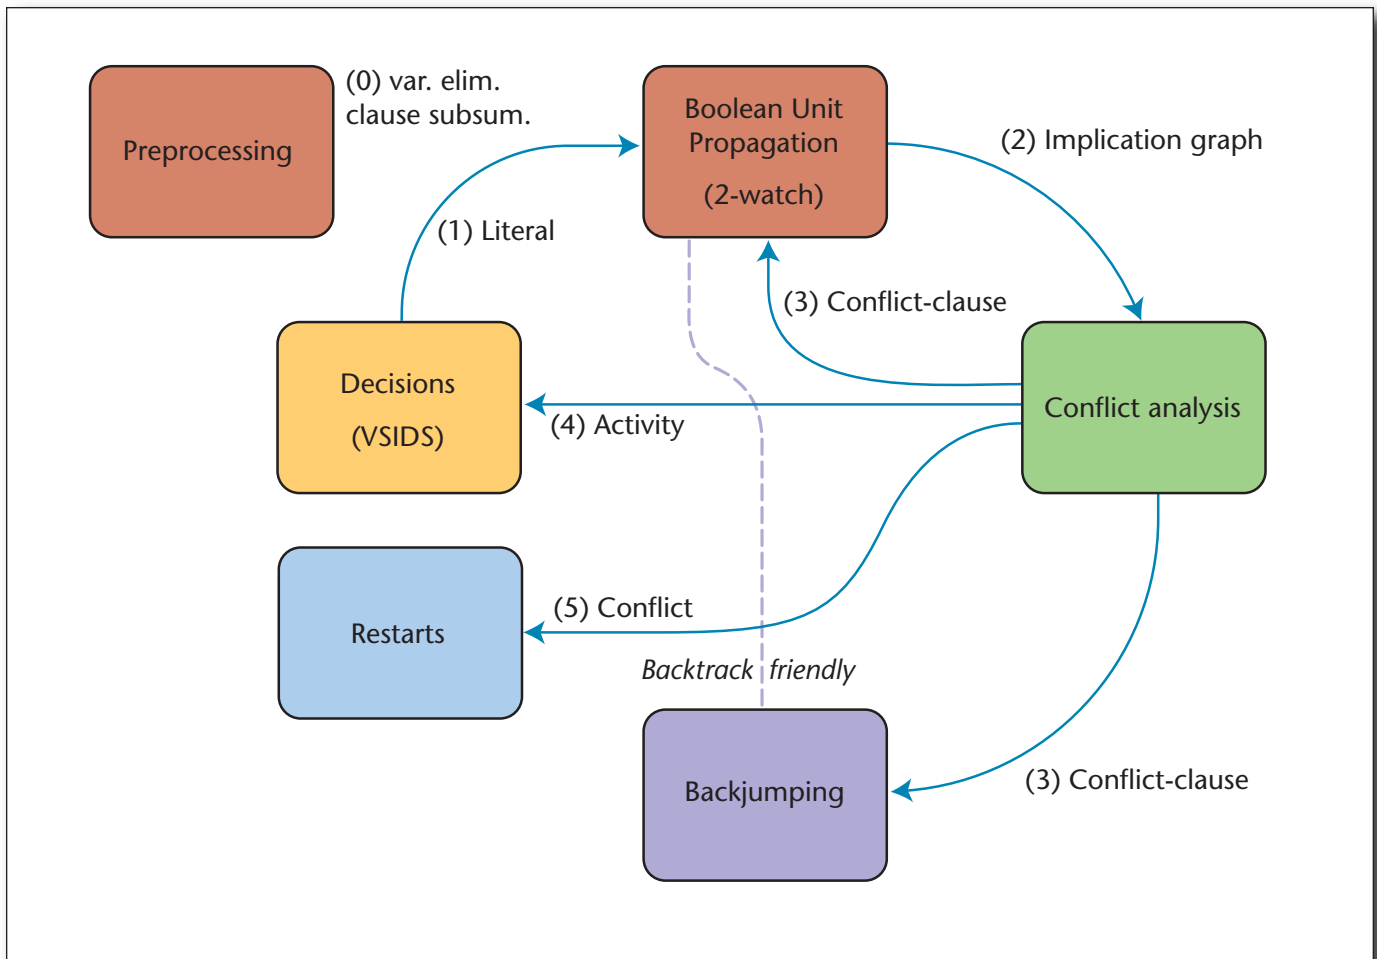


Figure 1. The General Architecture of a Sequential SAT Solver.

learnt clause) expresses that the last literal is implied at a previous level (it is “asserting”). The solver then jumps back to this decision level and assigns the literal to true in step 3. When an empty conflict clause is generated, the literal is implied at level 0, and the original formula can be reported as unsatisfiable. In addition to this basic scheme, modern solvers use additional components such as literal selection heuristics and a restart policy. For instance, the rank or activity of each Boolean variable encountered during the previous resolution process is increased (step 4). The variable with greatest activity is selected to be assigned as the next decision. This corresponds to the so called VSIDS variable branching heuristic (Zhang et al. 2001). When branching, after a certain amount of conflicts, a cutoff limit is reached and the search is restarted (step 5).

Context: Parallel SAT Solvers

There are two main approaches to parallel SAT solving. The first one implements the historical

divide-and-conquer idea, which incrementally divides the search space into subspaces, successively allocated to sequential CDCL workers. Workers cooperate through some load-balancing strategy, which performs the dynamic transfer of subspaces to idle workers, and through the exchange of conflict clauses.

The parallel portfolio approach was introduced in 2008 (Hamadi, Jabbour, and Sais 2008; Wintersteiger, Hamadi, and de Moura 2009; Guo et al. 2010). It exploits the complementarity of different sequential DPLL strategies to let them compete and cooperate on the same formula. Since each worker addresses the whole formula, there is no need to introduce load-balancing overheads, and cooperation is only achieved through the exchange of conflict clauses. With this approach, the crafting of the strategies is important, especially with only a few workers. The objective is to cover the space of good search strategies in the best possible way.

In general, the interleaving of computation can lead to the previously mentioned problem of non-determinism. This is true for solvers that use a

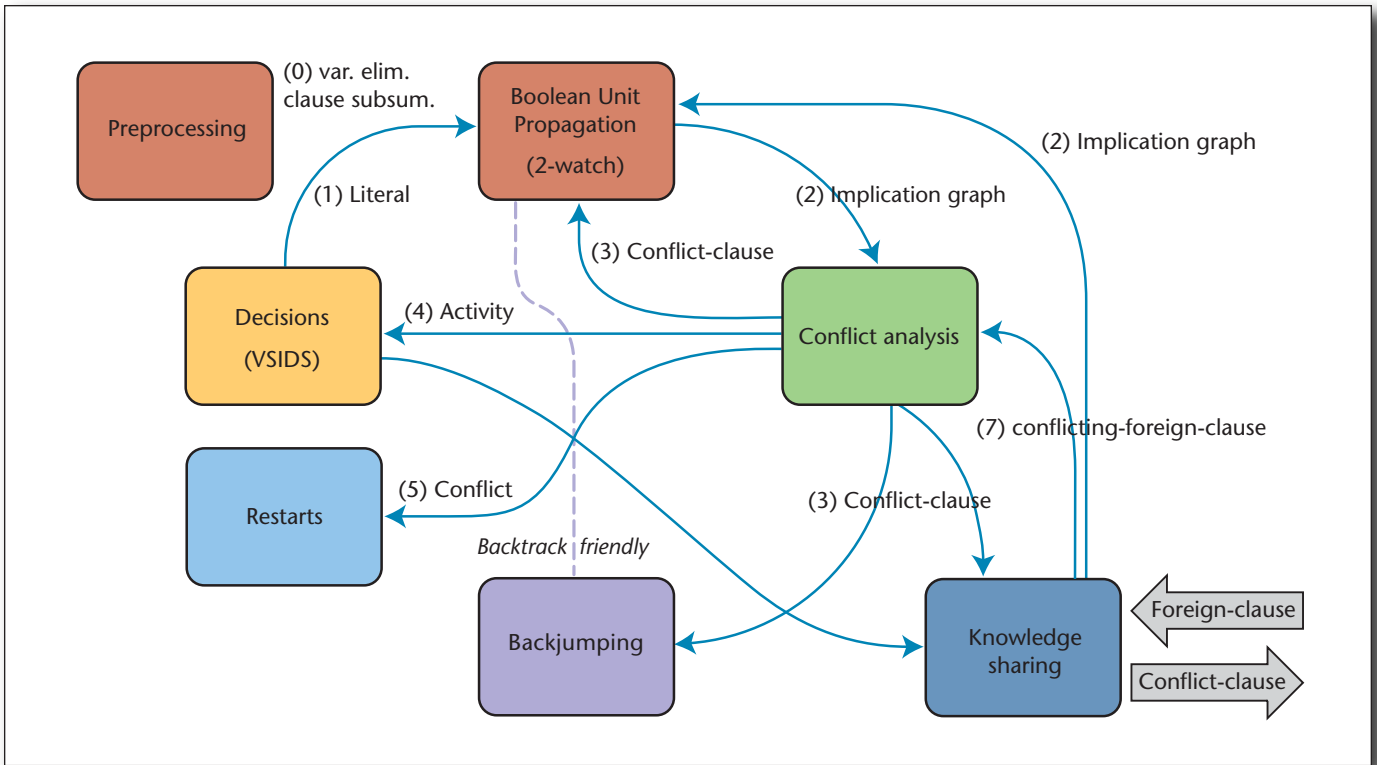


Figure 2. The General Architecture of a Parallel SAT Solver.

divide-and-conquer or a portfolio approach. Hamadi, Jabbour, Piette, and Sais (2011) propose a new technique to efficiently ensure the determinization of any parallel portfolio algorithm. Their method performs dynamic synchronization, which minimizes waiting time at barriers. This allows a parallel SAT portfolio always to return the same solution (or proof of unsatisfiability) in about the same run time, while preserving performance.

In figure 2 we present the CDCL architecture of a typical worker. It extends the original architecture presented in figure 1 with a knowledge-sharing component that exports and imports conflict clauses. Clauses are exported during step 3. If an imported conflict clause (in step 6) contradicts the current branch, the aforementioned conflict analysis and backjumping steps are executed. Finally, the knowledge-sharing unit can use the information gained by the VSIDS heuristic to filter out incoming information. This is figured through a link between the decisions and knowledge sharing components and will be detailed in Challenge 4. For a broader overview of parallel SAT solving see Martins, Manquinho, and Lynce (2012).

Context: Performance Evaluation

We suggest that performance evaluation of parallel SAT solvers be conducted on practically relevant

benchmark sets as is currently done in the biannual SAT competitions. We consider randomly generated benchmarks of mostly theoretical interest, but not necessarily as an indicator of the performance of a parallel SAT solver in practice. Especially non-deterministic solvers may benefit from an evenly distributed set of benchmarks, which may translate into performance figures that are only achievable in theory but not in practice.

Usually, the speedup of a parallel solver over a sequential one is defined as

$$S = \frac{T_s}{T_p}$$

such that a parallel solver that runs in time T_p exhibits a speedup S over a sequential solver that runs in time T_s . In practice, there are two different categories of applications for parallel SAT solvers that have different objectives: efficiency or effectiveness. The speedup by itself is not considered an indicative measure of performance for either of these categories. Instead, in the first category of applications, the run-time efficiency

$$E = \frac{S}{r} = \frac{T_s}{r \cdot T_p}$$

where r is the number of resources available to the solver, is of the greatest interest. For example, in

applications where energy consumption is an issue, a solver that performs at lower efficiency may be considered inferior to a solver that performs efficiently, even if its speedup figure is smaller. We expect this will be the case for many software and hardware verification applications in the near future, where limited size clusters are used to verify designs overnight. In the second category of applications, the absolute wall-clock time required to solve a problem is of paramount importance; we call this the run-time effectiveness of the solver, which we consider a better measure of performance in applications where energy consumption is of little or no importance. For example in cryptographic applications, especially for code breaking, we may assume that energy consumption and the available size of the cluster are irrelevant.

In general, the trade-off between efficiency and effectiveness highly depends on the application, and it is ultimately a decision that the community of SAT solver developers cannot make for the end user. We therefore suggest providing both measures of efficiency and effectiveness in a performance evaluation of parallel SAT solvers.

We wish to remark upon the number in efficiency computations. In many evaluations as well as the theoretical analysis of algorithms, this number is simply taken to be the number of computing elements available to the parallel solver. This is fully justified for theoretical purposes. In practice, this is not realistic, especially for multicore machines (compare with for example, Wintersteiger, Hamadi, and de Moura, [2009]). It is sometimes assumed that an n -core machine is able to perform n times the work of the corresponding single-core machine, which is simply not true due to memory and cache congestion issues, but also because modern processors change their behavior when multiple cores are under load, for example, by reducing the clock speed to avoid overheating. We therefore propose to compute the efficiency of a parallel multicore SAT solver with respect to its true capacity, which is to be measured in a prior calibration experiment. For example, this may be estimated by running copies of a sequential SAT solver in parallel with an observed runtime of $T_{ns'}$ which will be greater than T_s (Arbelaez and Hamadi 2011). To compute the efficiency of a parallel n -core solver we propose to use

$$r = n \cdot \frac{T_s}{T_{ns'}}$$

which we consider more realistic. In what follows we refer only to the general performance of a solver. Depending on the intended application, this is to be taken as either the efficiency or the effectiveness of the solver.

The Challenges

As presented in the introductory paragraphs, a divide-and-conquer approach can be lucky. A run can benefit from a good split, which brings a solution at the beginning of some subspace and allows for an early stop. In contrast, a different division can decrease performance. What is interesting here is that adding resources can decrease the performance since it can produce more demanding subspaces.

Even if portfolio-based approaches are less prone to this problem, extending the size of a portfolio can still be detrimental to its performance. In general, this increases the overhead, due to more frequent and broader clause sharing, and worsens cache-congestion issues. A priori, the question of deciding the most effective number of resources to use against a given formula is a difficult one.

One possible direction of research is to extend automatic tuning techniques. These approaches use machine learning to craft a predictive function that relates the features of an instance and the parameters of a given solver to its expected run time. This function can be learned and tested offline against a large set of representative instances and used at run time to configure a solver and maximize its performance. This offline approach assumes that a large and representative set of instances is available beforehand (Xu et al. 2008). A more recent approach avoids this problem by learning the function online (Arbelaez, Hamadi, and Sebag 2010). We believe that the previous offline and online approaches could be extended to consider the number of resources as an additional parameter of the solver.

Challenge 1. Generalize automatic tuning techniques to decide, among other solver parameters, the best amount of computational resource.

●○○○○

Decomposition

In the area of parallel algorithms it is natural to think of decomposition of the problem into a number of smaller subproblems. Most parallel SAT solvers are based on search algorithms, and we identify two inherently different types of decomposition for search algorithms: search-space decompositions and instance decompositions.

In the first category, the search space of the problem is decomposed; that is, the nodes or processes explore different (potentially overlapping) parts of the search space of the problem. In the case of SAT, the simplest way of achieving this is by duplication of the problem and assignment of a variable to contradicting values in the two branches. The set of assigned literals in any of the leaves of such a decomposition tree is then called a guiding path (Zhang, Bonacina, and Hsiang

1996). As we have seen with the previous challenge, finding a good decomposition prior to solving the formula is a hard problem as it is hard to predict the hardness of each of the subproblems.

In the second category of decompositions, the instance itself is decomposed such that none of the computing elements has knowledge of the whole problem instance. This type of decomposition is especially important when large formulas are considered²; for example, deep BMC unwindings in hardware verification (Ganai et al. 2006). On the one hand, finding an optimal decomposition that balances the size of the subproblems is easy for SAT problems, but the resulting subproblems are usually not balanced with respect to their hardness. On the other hand, finding a good instance decomposition that minimizes the number of shared variables is a hard problem in itself, and for this reason approximations may result in better overall performance. Recently, it has been shown that it is possible to recover from very crude approximations quickly through the use of Craig interpolation procedures, which are techniques to synthesize implied facts (called interpolants) from unsatisfiable implications $A \rightarrow B$, such that $A \rightarrow I \rightarrow B$ and I uses only variables common to A and B . It has been demonstrated not only that the incorporation of such techniques into the SAT solver presents a large number of opportunities for parallel solvers, but also that dynamic instance decompositions may even improve the performance of a sequential SAT solver when combined with well-chosen interpolation methods (Hamadi, Marques-Silva, and Wintersteiger 2011).

Clearly, for both types of decomposition, the state of the art is unsatisfactory and further research is needed to find good decompositions and recover methods that perform well in practice, both for large search spaces and for large problem instances.

Challenge 2. Design for either of the two classes of decomposition a dynamic decomposition technique that is efficiently computable and results in decompositions that enable solvers to consistently outperform currently known methods.

●●●○○

Preprocessing

In the recent past, preprocessing for SAT formulas has received increased attention, and it has been shown that some types of preprocessing have a great effect on the performance of sequential SAT solvers, for example, Eén and Biere (2005). We believe that in the context of parallel SAT solving, new preprocessing techniques are required. For instance, it may not be necessary (or even beneficial) aggressively to reduce the number of clauses in a problem before it is split or distributed to the computing elements.

Furthermore, preprocessing in the context of parallel SAT should take into account the nature of the parallelization approach, especially the type of decomposition that is used, that is, search-space or instance decomposition. Depending on the type of decomposition, different preprocessing techniques may have the best effect on the performance of the solver. For example, in instance decompositions it may be much more effective to minimize the set of overlapping variables between subproblems than to minimize the overall size of the formula.

For very large formulas, it may be infeasible to preprocess a whole problem instance before solving it. We therefore consider it worthwhile to investigate parallel preprocessing algorithms as well.

Challenge 3. Devise new parallel preprocessing techniques that, with knowledge of the type of decomposition being used, simplify a problem instance such that the overall performance of the solver is increased. ●●●○○

Improved Knowledge Sharing

Modern SAT solvers generate conflict clauses to prevent the reoccurrence of a conflict and to back-jump effectively in the list of decisions. Recent parallel solvers have leveraged these clauses by sharing them. Since search can generate a large (exponential) number of new clauses, strategies were defined to limit the overhead of communication.

The most basic strategy limits the size of the shared clauses up to some fixed limit. This has two advantages. It restricts the overhead and focuses the cooperation to powerful clauses.

However, the static-size strategy can miss situations where more cooperation would help, for instance, when two strategies explore the same subspace. Also, it might maintain useless exchanges between strategies which focus on independent subproblems.

To alleviate these problems, Hamadi, Jabbar, and Sais (2009) have introduced a dynamic strategy that uses control theory techniques to increase or reduce automatically the quantity of clauses shared between two search efforts. Their technique estimates the quality of incoming clauses as the observed performance and uses this information to extend or restrict the cooperation.

Assessing the quality of a clause with respect to its local impact is difficult and a generalization of the clause deletion problem in modern CDCL solvers. We think that the community should spend some effort to define better quality measures in order to leverage the benefits of clause sharing, and we therefore propose the following challenge.

Challenge 4. Drive the cooperation through better estimates of the local quality of incoming clauses.

●●○○○

Integer Factorization

We believe that it is beneficial to the community to contemplate solving challenging problems from related areas for which SAT solvers may ultimately present an effective solution. Recently there has been an increased interest in solving problems related to security applications in the SAT community. One problem that is particularly challenging, and of utmost importance in practical security applications, is the (decision version of the) integer factorization problem IF.

This problem is known to be in NP and there exists a trivial encoding to SAT, for example, through the Boolean encoding of a multiplier circuit, but the performance of current SAT technology on such formulas is not competitive with that of dedicated, subexponential algorithms like the quadratic and general number field sieve (for an introduction see, for example, Crandall and Pomerance [2001]). It is typical for these dedicated algorithms to require a large number of resources for a long time. For instance, the recent success in factoring a 768-bit integer through a distributed number field sieve kept many hundreds of machines busy for almost two years, a total equivalent of 1500 years of computation on a single-core processor (Kleinjung et al. 2010).

We consider IF a prime example of a challenging problem for parallel SAT solving, not only for its potential practical implications, but also because advances in this direction would shed more light on the structure of NP. Currently, IF is believed not to be NP-complete, but also to lie outside of P. It is a candidate for the NP-intermediate complexity class (Ladner 1975), which, currently, very little is known about. Finding practically efficient parallel algorithms for problems in this class would have a great impact not only in practice but also for the theory of SAT and parallel algorithms in general.

Challenge 5. Design an encoding of IF instances and a parallel SAT solver that performs competitively with dedicated algorithms for IF. ●●●●

Specific Encodings

As a sixth challenge, we suggest to investigate new encodings of the SAT

problem. Most SAT solvers support only the solving of formulas in CNF form and it is possible that this encoding, while convenient, poses a limitation for parallel solvers. For example, it is conceivable that, when many processors are employed, a pipelined evaluation of assignments on deep circuits could perform better than a CNF encoding with clauses held in the usual watchlists, simply because the locking/synchronization overhead on the watchlists grows too quickly as the number of processors is increased.

Challenge 6. Devise a new encoding of SAT problems specifically for parallel solvers. ●●●●

Starting from Scratch

Much of the ongoing research in parallel SAT is focused on parallelizing existing algorithms and implementations, many of them based on CDCL solvers. We believe that parallelizing existing procedures is not the best way to obtain a truly well performing parallel SAT algorithm. Instead we propose to start from scratch and to investigate completely new algorithms and data structures for parallel SAT or to revisit techniques that were deemed inefficient in the past.

The root cause of our suggestion is the fact that most modern sequential SAT solvers are ultimately based on Boolean constraint propagation (BCP), which is a P-complete problem and thus suspected to be hard to parallelize (Hamadi 2002). If we think of a CDCL solver as a dynamic decomposition of the search space (through decision variables), then most of the speedups are likely to be obtained on this higher level of decomposition and recombination (decision making, conflict analysis, and sharing), but it might ultimately remain difficult to effectively parallelize the rest of the algorithm. Further research into parallelizations of existing solvers may help to gain a better understanding of the challenges of parallelizations of P-complete problems, but we believe that it will be hard to design algorithms that perform well in practice. It is conceivable that there exist other algorithms that are much easier to parallelize. For instance, it is conceivable that an algorithm based on a reduction to a series of bounded-

The Integer Factorization Problem (IF)

Given two integers N and M such that $1 < M \leq N$, determine whether N has a nontrivial factor $d < M$.

width branching programs would be considerably easier to parallelize, since it is known that branching programs of width 5 and of polynomial length recognize exactly those languages in NC^1 (Barrington 1986), a complexity class for which algorithms are suspected to be easy to parallelize.

Challenge 7. Devise a parallel algorithm for SAT that is not based on a reduction to a (set of) P-complete problem(s) and that performs on par with or better than parallelizations of CDCL. ●●●●

Conclusion

Today, computers have multiple cores, and cloud computing provides users with an inexpensive way to rent virtual resources on which to run their applications. Still, most search researchers restrict themselves to sequential algorithms. This is paradoxical, especially when we consider the importance of search. There are two complementary explanations to this situation: The first one lies in the lack of parallel programming skills and the second comes from the difficulty of good intuition building.

The first problem is very general and can only be tackled by making progress in parallel programming languages and tools and through an increase in parallelism courses in undergraduate curricula. This is difficult, but feasible. Solving the second problem is much more challenging. It requires years of practice that can only sometimes provide the expertise and intuition required for significant contributions.

In this article, we try to address the second point. Our strategy is to share our views and understanding of the evolution of parallel search in general and parallel SAT solving in particular. From that understanding, we present a

list of important challenges. They have different goals and different inherent complexities. Our objective is not necessarily to put the community onto them, but we believe that by sharing our views we can contribute to fostering an increased interest in parallel SAT solving and parallel search in general. We hope that this will eventually result in better parallel algorithms that further increase the practical applicability of search.

Notes

1. Operating at higher clock rates consumes more electrical energy partly dissipated in the form of heat. Overcoming this heat has become technically difficult and economically inefficient. This was originally presented as “hitting a thermal wall” by chip manufacturers.
2. Note that in some cases the distribution is a given. In distributed constraint reasoning privacy can be a concern, and agents only exploit a partial view of the problem (Ringwelski and Hamadi 2005).

References

- Arbelaez, A., and Hamadi, Y. 2011. Improving Parallel Local Search for SAT. *Learning and Intelligent Optimization*, Lecture Notes in Computer Science Volume 6683, 46–60. Berlin: Springer.
- Arbelaez, A.; Hamadi, Y.; and Sebag, M. 2010. Continuous Search in Constraint Programming. In *Proceedings of the 22nd IEEE International Conference on Tools with Artificial Intelligence*, 53–60. Piscataway, NJ: Institute of Electrical and Electronics Engineers.
- Barrington, D. A. 1986. Bounded-Width Polynomial-Size Branching Programs Recognize Exactly Those Languages in NC^1 . In *Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing*, 1–5. New York: Association for Computing Machinery.
- Crandall, R., and Pomerance, C. 2001. *Prime numbers: A Computational Perspective*. Berlin: Springer.
- Davis, M.; Logemann, G.; and Loveland, D. W. 1962. A Machine Program for Theorem-Proving. *Communications of the ACM* 5(7): 394–397. New York: Association for Computing Machinery.
- Eén, N., and Biere, A. 2005. Effective Pre-processing in SAT Through Variable and Clause Elimination. In *Proceedings of the 8th International Conference on Theory and Applications of Satisfiability Testing*, 61–75. Berlin: Springer.
- Ganai, M.; Gupta, A.; Yang, Z.; and Ashar, P. 2006. Efficient Distributed SAT and SAT-Based Distributed Bounded Model Checking. *International Journal on Software Tools for Technology Transfer* 8(4): 387–396.
- Garey, M. R., and Johnson, D. S. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. San Francisco: W. H. Freeman.
- Guo, L.; Hamadi, Y.; Jabbour, S.; and Sais, L. 2010. Diversification and Intensification in Parallel SAT Solving. In *Principles and Practice of Constraint Programming CP 2010*, Lecture Notes in Computer Science Volume 6308, 255–265. Berlin: Springer.
- Hamadi, Y. 2002. Optimal Distributed Arc-Consistency. *Constraints* 7(3/4): 367–385.
- Hamadi, Y.; Jabbour, S.; and Sais, L. 2008. ManySAT: Solver Description. Microsoft Research MSR-TR-2008-83, Redmond, WA.
- Hamadi, Y.; Jabbour, S.; and Sais, L. 2009. Control-Based Clause Sharing in Parallel SAT Solving. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence*, 499–504. Menlo Park, CA: AAAI Press.
- Hamadi, Y.; Jabbour, S.; Piette, C.; and Sais, L. 2011. Deterministic Parallel DPLL. *Journal of Satisfiability* 7(4): 127–132.
- Hamadi, Y.; Marques-Silva, J.; and Wintersteiger, C. M. 2011. Lazy Decomposition for Distributed Decision Procedures. In *Proceedings of the Workshop on Parallel and Distributed Methods in Model Checking (PDMC)*, 43–54. Ithaca, NY: Electronic Proceedings in Theoretical Science, Cornell University.
- Kleinjung, T.; Aoki, K.; Franke, J.; Lenstra, A.; Thomé, E.; Bos, J.; Gaudry, P.; Kruppa, A.; Montgomery, P.; Osvik, D. A.; te Riele, H.; Timofeev, A.; and Zimmermann, P. 2010. Factorization of a 768-bit RSA Modulus. Deposited in Cryptology ePrint Archive, Report 2010/006. Carson City, NV: International Association for Cryptologic Research.
- Ladner, R. E. 1975. On the Structure of Polynomial Time Reducibility. *Journal of the ACM* 22(1): 155–171.
- Martins, R.; Manquinho, V. M.; and Lynce, I. 2012. An Overview of Parallel SAT Solving. *Constraints* 17(3): 304–347.
- Moskewicz, M. W.; Madigan, C. F.; Zhao, Y.; Zhang, L.; and Malik, S. 2001. Chaff: Engineering an Efficient SAT Solver. In *Proceedings of the Design Automation Conference*, 530–535. Piscataway, NJ: Institute of Electrical and Electronics Engineers.
- Pruul, E. A., and Nemhauser, G. L. 1988. Branch-and-Bound and Parallel Computation: A Historical Note. *Operations Research Letters* 7(2): 65–69.
- Rao, V. N., and Kumar, V. 1993. On the Efficiency of Parallel Backtracking. *IEEE Transactions on Parallel and Distributed Systems* 4(4): 427–437.
- Ringwelski, G., and Hamadi, Y. 2005. Boosting Distributed Constraint Satisfaction. In *Principles and Practice of Constraint Programming*, Lecture Notes in Computer Science Volume 3709, 549–562. Berlin: Springer.
- Wintersteiger, C. M.; Hamadi, Y.; and de Moura, L. 2009. A Concurrent Portfolio Approach to SMT Solving. *Computer Aided Verification*, Lecture Notes in Computer Science Volume 5643, 715–720. Berlin: Springer.
- Xu, L.; Hutter, F.; Hoos, H. H.; and Leyton-Brown, K. 2008. SATzilla: Portfolio-Based Algorithm Selection for SAT. *Journal of Artificial Intelligence Research* 32: 565–606.
- Zhang, H.; Bonacina, M. P.; and Hsiang, J. 1996. PSATO: a Distributed Propositional Prover and its Application to Quasigroup Problems. *Journal of Symbolic Computation* 21(4): 543–560.
- Zhang, L.; Madigan, C. F.; Moskewicz, M. W.; and Malik, S. 2001. Efficient Conflict Driven Learning in Boolean Satisfiability Solver. In *Proceedings of the International Conference on Computer-Aided Design*, 279–285. New York: Association for Computing Machinery.

Youssef Hamadi is a senior researcher at Microsoft Research. He holds a doctoral degree in computer science from the University of Montpellier in France and a Habilitation from the University of Paris-Sud, France. His research interests involve the practical resolution of large-scale real-life problems. His work is set at the intersection of optimization and artificial intelligence. His research considers the design of complex systems based on multiple formalisms fed by different information channels that plan ahead and perform smart decisions. His current focus is on autonomous search, parallel search, and Boolean satisfiability, with applications to environmental intelligence, business intelligence, and software verification.

Christoph M. Wintersteiger is a researcher at Microsoft Research. He holds an engineering degree in computer science from the University of Linz, Austria, and a doctoral degree in computer science from ETH Zurich, Switzerland. His research is focused on the investigation and design of automated reasoning techniques and applications thereof, especially in the field of automated software verification. He currently works on parallel and distributed methods for SAT and SMT solving.