

Improving Performance of CDCL SAT Solvers by Automated Design of Variable Selection Heuristics

Marketa Illetskova*, Alex R. Bertels†, Joshua M. Tuggle*, Adam Harter*, Samuel Richter*, Daniel R. Tauritz*, Samuel Mulder†, Denis Bueno†, Michelle Leger†, and William M. Siever‡

*Department of Computer Science, Missouri University of Science and Technology, Rolla, MO, U.S.A.

Email: miwdf@mst.edu, dtauritz@acm.org

†Sandia National Laboratories, Albuquerque, NM, U.S.A., Email: abertel@sandia.gov

‡Department of Computer Science and Engineering, Washington University in St. Louis, St. Louis, MO, U.S.A.

Abstract—Many real-world engineering and science problems can be mapped to Boolean satisfiability problems (SAT). CDCL SAT solvers are among the most efficient solvers. Previous work showed that instances derived from a particular problem class exhibit a unique underlying structure which impacts the effectiveness of a solver’s variable selection scheme. Thus, customizing the variable scoring heuristic of a solver to a particular problem class can significantly enhance the solver’s performance; however, manually performing such customization is very labor intensive. This paper presents a system for automating the design of variable scoring heuristics for CDCL solvers, making it feasible to tailor solvers to arbitrary problem classes. Experimental results are provided demonstrating that this system, which evolves variable scoring heuristics using an asynchronous parallel hyper-heuristics approach employing genetic programming, has the potential to create more efficient solvers for particular problem classes.

I. INTRODUCTION

The boolean satisfiability problem (SAT) is one of many problems in computer science that belongs to the NP complete complexity class, and therefore it has no efficient solution unless $P = NP$. Various real-world engineering and science problems can be modeled as instances of SAT, and searching for solutions to such real-world problems requires extensive utilization of computational resources.

Over the past few decades, SAT solver efficiency has improved dramatically, leading to even more utilization of SAT in engineering and science solutions [1]. For example, SAT has been used to minimize power consumption in an energy management system while satisfying varying requirements of its users [2]. There are many other applications of SAT in various fields such as computational biology [3], [4], software testing [5], hardware and software model checking and verification [6], [7], Internet of Things [8], and robotics [9], [10].

Encoding specific problem classes in SAT creates classes of structured logical expressions called SAT instances. Empirical evidence shows that each solver has an ideal underlying instance structure and that each class of instances has an optimal solver [11], [12], [13]. Recent work has automatically optimized SAT solver parameter configuration to target specific classes of instances [14], [15], [16]. KhudaBukhsh et al. automatically composed stochastic local search (SLS) solvers from parts of existing SLS solvers [17], and other work automatically evolved variable selection techniques for SLS

solvers [18], [19], [20], [21], [22]. However, conflict-driven clause learning (CDCL) solvers are still the most efficient SAT solvers for industrial instances. Based on Biere and Fröhlich’s demonstration of the drastic impact that restart and variable selection schemes have on CDCL solver efficiency for specific problem classes [23], [24], Bertels introduced the idea of automatic evolution of CDCL operations as a case study in [25] and as a proof-of-concept system in [26]. His work showed promising results in increasing the effectiveness of a CDCL SAT solver by targeting classes of instances with unique structure. We know of no other work that automatically evolves CDCL heuristics.

In this paper we describe the system, **ADSSEC** (Automated Design of SAT Solvers employing Evolutionary Computation), evaluate its performance on datasets and benchmarks from random and industrial tracks, and present results demonstrating its potential for producing more effective solvers targeted to arbitrary, but particular, problem classes (datasets).

II. RELATED WORK

Ideally, a single solver would be able to adapt to an application at runtime. Tools such as ParamILS [14], SMAC [15], and SpySMAC [16] automatically tailor the parameter configurations of reasonably versatile solvers to particular datasets. While the solver parameter configurations are adjustable, most of the internal methods of the solvers remain the same. The effectiveness of adapting a solver to a problem class solely through external parameter optimization is limited by the appropriateness of the solver’s architecture, such as its variable scoring heuristic, for that problem class.

In order to identify the theoretically shortest possible time to solve a given instance, one would have to generate all possible solvers and execute them for all possible configurations on that instance. As this is computationally infeasible, we must restrict ourselves to a subset of existing solvers with carefully selected parameters. This method is referred to as a portfolio approach. Xu et al. were able to predict which solvers in a portfolio were able to perform well in particular domains [27], [28]. They did this by calculating values for a set of instances, testing the portfolio on the instances, and using machine learning to relate solvers to a given instance. This pairing of solvers with instance classes allowed Xu et al. to reduce the

portfolio to the best suited solvers. Hutter et al. expanded on this work by employing these calculated values to predict the runtime of SAT solvers [29]. Portfolios of algorithms provide high flexibility in discovering the right *existing* solver for the job, assuming that the right solver is in the portfolio to begin with.

ADSSEC takes the next step by applying generative hyper-heuristics [30] to generate CDCL SAT solvers tailored for an arbitrary, but particular, problem class. *Hyper-heuristics* approaches automatically develop heuristics or algorithms; generative hyper-heuristics combine primitive heuristics, exploring a broader search space, while selective hyper-heuristics choose the best option(s) from a set of pre-defined heuristics [30]. While not quite as flexible as generating new heuristics, selecting heuristics can be beneficial if multiple components need to be matched together and effective heuristics are known. In this paper, only the variable scoring heuristic is being modified, so ADSSEC uses the generative approach. ADSSEC employs genetic programming (GP) to automatically reorganize and manipulate the algorithmic primitives constituting the variable scoring heuristic [31], [32]. These primitives can be as general as state-related variables and binary operations or as specific as carefully constructed functions with tunable inputs.

ADSSEC is a hyper-heuristics framework that employs CDCL state-based information and binary operations to automatically develop new variable scoring heuristics tailored to problem classes. ADSSEC's primitives are more granular than statements in the source code and are therefore much more versatile in developing new solver components.

We describe the ADSSEC hyper-heuristics framework in detail in the next section.

III. ADSSEC

Genetic programming (GP) utilizes a biologically inspired approach in searching for individuals adapted to a given environment by preserving genes important for success. An individual in GP may be optimized for a given input to produce desirable output in a given environment. Finding an optimal SAT solver for a given problem class is a complex and expensive task, suitable for automation using a GP-based hyper-heuristic system. Influenced by the success of Fukunaga in improving SLS runtimes by evolving specific heuristics [20], ADSSEC uses GP to evolve variable scoring heuristics to automatically tailor CDCL solvers to specific classes of structured instances.

In particular, ADSSEC utilizes Koza-style GP trees [33], which are well suited to the parse trees representing variable scoring heuristics used by CDCL solvers to select decision variables. Biere and Fröhlich's work demonstrating that the current best variable scoring heuristics are roughly equal in runtime performance when evaluated across many instance classes [23] motivated the choice to adapt a CDCL solver's variable scoring heuristics [26].

Evolutionary algorithms and SAT solving are independently computationally heavy tasks, and combining the two

TABLE I: Terminal nodes in ADSSEC

Score ($s[v]$)	The previous score of the variable.
Conflict Index (i)	The current number of conflicts encountered.
MiniSat Variable Decay Amount (f)	f is also used to derive MiniSat's Variable Increment Value (<i>MiniSat Default</i> : 0.95).
MiniSat's Variable Increment Amount (g)	$(g = (1/f)^i)$ The amount MiniSat increases a score when a variable is bumped.
Constant (C)	A constant value in $\{1, 2, 3, \dots, 10\}$ or $\{0.0, 0.1, 0.2, \dots, 0.9\}$.
Special Component (H)	Derived from the Chaff CDCL SAT solver for scaling variable scores [23]. $h_i^m = \begin{cases} 0.5 \cdot s & \text{if } m \text{ divides } i \text{ evenly} \\ s & \text{otherwise} \end{cases}$ where m is a power of 2: $\{2, 4, 8, \dots, 1024\}$.
Variables in the First UIP ($U[v]$)	The number of times a variable occurred in a first Unique Implication Point (UIP) conflict.
Eliminated Variables in Simplification Step ($E[v]$)	The number of times a variable is eliminated in the simplification of conflict clauses.

magnifies the effect. Bertels et al. showed using ADSSEC that an asynchronous parallel evolutionary algorithm (APEA) approach leads to a significant speed up in hyper-heuristic systems [25], [26]. This motivated us to use an asynchronous parallel implementation of ADSSEC.

ADSSEC creates an initial population of variable scoring heuristics and evolves the population through mutation and recombination, both explained in detail later. ADSSEC evaluates these heuristics by replacing the variable selection heuristic in MiniSat 2.2 [34], a commonly used efficient and deterministic CDCL solver with dense source code, or in Glucose 4.0 [35], one of the current state-of-the-art solvers based on MiniSat. ADSSEC employs a standard parent selection before producing offspring and a crowding survival selection specifically modified for the APEA [26]. ADSSEC returns the best heuristics from the final population after reaching the termination criteria.

Derived from currently implemented variable scoring heuristics [23], [34], [36], ADSSEC defines the terminal nodes described in Table I and basic non-terminal (binary) operator nodes: Addition (+), Subtraction (-), Multiplication (*), and Division (/). These arithmetic operators may be applied because all terminal nodes are either integer or floating-point values. These nodes allow evolution of novel schemes while still being able to represent current schemes.

A. Heuristic Representation

Mapping variable scoring heuristic functions to objects that can be easily manipulated in a GP tree is fairly straightforward. Each scoring heuristic can be represented as a parse tree where non-terminal nodes are operators and terminal nodes are state-related values. See Figure 9 for an example of a tree representation of an evolved heuristic.

ADSSEC evolves the parse tree genetic encodings. To evaluate each variable scoring heuristic, the parse tree is converted into a C++ statement. The original variable scoring heuristic in a pre-built solver is replaced by compiling and linking with the new variable scoring heuristic. The resulting solver, termed the *evolved variant*, is executed on test instances to evaluate the effectiveness of the heuristic.

B. Objectives

The evolutionary algorithm (EA) objective score represents how well an evolved variant performs on a provided training set of instances. Traditionally, the aim is to reduce the runtime needed to either find a solution or prove unsatisfiability. In Sect. IV-A, we compare performance measures, i.e., objective functions, to use in ADSSEC; our ultimate goal is to reduce the average runtime across instances in a problem class.

ADSSEC’s per-instance sub-score for an evolved variant is the ratio of the performance of the evolved solver to that of the original solver. The objective score is then simply the average of all the instance sub-scores. Thus, any evolved individual that performs identically to the original solver’s variable scoring heuristic will end up with an objective score of 1.0, and lower scores indicate better schemes.

Occasionally, the EA will construct inadequate heuristics that cause the solver to require inordinate resources to terminate on some instances. Limiting functions prevent wasting evaluation time on such heuristics. ADSSEC relies on the original solver’s performance to approximate reasonable limits for any given SAT instance. Initially, ADSSEC limits an evolved variant to three times the number of variable decisions the original solver needed to solve that instance. These generous limits allow ADSSEC to collect diverse genetic material in a population; they do not time out on all tested instances, but they are generally worse than the original solver. However, as ADSSEC progresses through the evolutionary process, interest shifts from collecting diverse heuristics to exploiting heuristics that are strictly better than the original. As such, the decision limit linearly decreases down to the exact number of decisions the original solver needed for a specific instance. For example, if ADSSEC is to complete 5000 evaluations throughout the run and the decision limit multiplier decreases from 3.0 to 1.0, then the multiplier is decremented by $((3.0 - 1.0)/5000 = 0.0004$ after each evaluation.

Ideally, an accurate objective score would be determined by executing the evolved variable scoring heuristic against the entire training dataset of interest. This is generally too costly, therefore ADSSEC utilizes strike-based sampling to gauge the effectiveness of an evolved variant. ADSSEC randomly selects a number of instances in a user-defined range from the given training set to evaluate a variant. For each instance in this selection, ADSSEC executes the evolved variant and assigns a sub-score ratio as described before. If the variant reaches the decision limit for that instance, then the variant receives a strike and a sub-score of the current decision-limit multiplier. After a variant reaches a set number of strikes, all remaining

sub-scores are assigned the current decision-limit multiplier without evaluation.

C. Evolutionary Algorithm

1) *Population Initialization*: To avoid local optima, ADSSEC creates each individual in the initial population of variable scoring heuristics by randomly generating a parse tree from the primitives (nodes). First, ADSSEC assigns a random operator node to the root of the tree. ADSSEC then assigns two random nodes to the left and right branches of the operator node. There is a 50% chance that each node will be terminal. If the node is non-terminal, then ADSSEC repeats the process and assigns a random operator node. If the node is terminal, then there is a 50% chance that the node will be the previously assigned score s . Otherwise, ADSSEC randomly assigns one of the other terminal node options. This bias was introduced because most current schemes appear to rely heavily on the previous variable score. The maximum depth of a tree generated for the initial population is arbitrarily set to eight. Smaller depths seemed to contain much less genetic diversity while larger depths produced complex heuristics that rarely solved instances within the decision limits.

2) *Variation Operators*: As in traditional tree GP, ADSSEC uses one of two methods to develop a single offspring (variant): mutation or recombination. For mutation, ADSSEC randomly selects a node in a random individual’s parse tree and replaces it with a new subtree generated using the rules established in Population Initialization (Sect. III-C1) – without a depth limitation. For recombination, ADSSEC implements a sub-tree crossover: the system randomly selects two individuals in the population and replaces a random branch from the first parent with a random branch from the second parent. Both procedures produce a single child.

3) *Selection*: The survival selection function determines which individuals in the population continue into the next generation. In ADSSEC, genetically diverse selection is desirable so that smaller parse trees (which are generated more easily) do not flood the population. Certain small heuristics have adequate performance and, if one is discovered early on in evolution, it can overwhelm the population if diversity is not maintained.

Crowding functions are selection functions that excel at promoting genetic diversity in the population [37]. In a standard crowding function, an offspring competes with its closest parent, either replacing the parent or being dropped from the population in favor of the parent. In an APEA, however, generations are not clearly delineated and a parent can have multiple offspring being evaluated simultaneously. Bertels developed an asynchronous crowding function that allows offspring to compete with either their parents or any ‘*siblings*’ – or descendants of siblings – that replaced the parents [26], [25]. The function uses a computationally cheap distance function comparing histograms of node types (e.g., addition, constant, conflicts, etc.) to determine the closest remaining relative in the current population. As in a standard crowding function, only the offspring or its closest remaining

relative remains in the population, and the population remains at a fixed size.

We use k -Tournament parent selection, which is easily applicable to parallel systems and which provides selection pressure modification through the k parameter. We selected a low k to increase the chance that genetic material in the pool propagates to many individuals in the population while maintaining some pressure to eliminate less fit individuals.

4) *Termination*: ADSSEC terminates the evolutionary cycle after completing a user-defined number of evaluations. However, throughout the run, individuals may be replaced by randomly generated parse trees if the population converges. If the best individual has not been improved in a user-defined number of evaluations, ADSSEC introduces new material to the gene pool. Currently, all variants whose performance is worse than that of the original solver are replaced. This mechanism is useful in restarting the exploration of the variable scoring heuristic search space.

IV. EXPERIMENTS

A. Comparison of Objectives

The main motivation for this experiment was to determine a good objective for the EA employed by ADSSEC. ADSSEC evaluates several instances in parallel on the same hardware, and so runtimes for individual instances are inconsistent. For that reason, ADSSEC was originally implemented both as a single objective EA (SOEA) with the number of variable decisions as an objective and as a multi-objective EA (MOEA) with the number of variable decisions and the number of learnt clause literals (conflict literals) as objectives. An improved variable scoring heuristic should reduce the values for the number of decisions and conflict literals, potentially reducing the runtime, but lower values in CDCL solvers may not correlate well with the runtime in specific cases.

Thus, we decided to explore the performance of solvers produced by ADSSEC when using CPU time as an objective. We compare an SOEA with decisions as an objective to an SOEA with CPU time as an objective, and we compare an MOEA with decisions and conflict literals as objectives to an MOEA with decisions, conflict literals, and CPU time as objectives. We explore both SOEA and MOEA because experiments in [26] suggested that ADSSEC produced more performant evolved variants using an MOEA.

In order to achieve a statistically significant results, we ran the evolutionary process with the same settings 30 times for each of the four different classes that we compare. Since evolutionary algorithms are computationally expensive, we needed a dataset whose instances can be quickly solved. For this experiment, we chose the ibm-26 dataset, one of many subsets of the IBM - Hardware Verification set from the IBM Formal Verification Benchmark Library, from the Configurable SAT Solver Challenge (CSSC) 2014 benchmarks [38]. MiniSat can solve all instances in less than 4 seconds.

We evolved the variable scoring heuristic from the MiniSat solver, using an identical parameter configuration (Table II; more on parameter configuration in Sect. IV-B) for all SOEA

and MOEA experiments. ADSSEC evaluated each evolved variant on a randomly ordered subset of all of the instances from the ibm-26 dataset, employing the evaluation described previously.

We executed ADSSEC on the same machine for each objective function. Solvers were then compiled with the best heuristic produced by each run. To obtain more accurate CPU solve times, we averaged the CPU time collected over 30 serial runs of each evolved variant and the original MiniSat on each instance in the ibm-26 dataset. All evaluations were performed on the same machine. Results are presented in Sect. V-A.

B. Evolving Variable Scoring Heuristics

Based on preliminary results from experiments like that in Sect. IV-A, and in order to achieve results comparable to results in [26], we ran our main experiments using an MOEA with decisions and conflict literals as objectives for the training phase. The purpose of these experiments was to determine whether ADSSEC could produce more efficient solvers tailored to specific real-world problem classes.

Due to the nature of evolutionary computation and SAT solving, ADSSEC requires datasets with enough instances to represent a distinct instance class for training and testing. These datasets must contain instances that ADSSEC can train on in a short period of time, and the dataset must be difficult enough to benefit from a fitted heuristic. Many available datasets do not fulfill these requirements. For these experiments, we used the unif-k5 dataset, a benchmark from the random track of CSSC 2014 consisting of 600 instances generated by Balint et al. according to the uniform generation model [39], and the ibm-18 dataset, another subset of the IBM - Hardware Verification dataset from the CSSC 2014 industrial track benchmarks [38].

We used an identical configuration for ADSSEC on both datasets (Table II), the same as used in [26]. For each experiment, ADSSEC created an initial population of 30 random individuals. The master process used several slave processes (the exact number varied based on the number of nodes available on the machine used for each experiment) to evaluate offspring populations, asynchronously creating new offspring as each node became available.

ADSSEC selected parents in a k -Tournament selection with $k = 2$ for either recombination or mutation – with a mutation probability of 0.10 and a recombination rate of 0.90 – and used an asynchronous crowding method for survival selection. Although ADSSEC terminated after 5000 evaluations, if the best individual objective score had not improved in 250 evaluations, ADSSEC replaced the worst individuals in the population with randomly generated parse trees. ADSSEC evaluated each individual on the ten SAT instances in the training set (randomly ordered); each individual was limited to four strikes against the decision limit described previously.

We used the same experimental setup to evolve variable scoring heuristics in both MiniSat and Glucose. We executed ADSSEC on several locally networked machines of varying loads. Solvers were then compiled with the best heuristics

TABLE II: ADSSEC EA parameter settings

Population (μ)	Offspring (λ)	Mutation Rate	Termination Evaluations	Restart Evaluations
30	63	0.10	5000	250

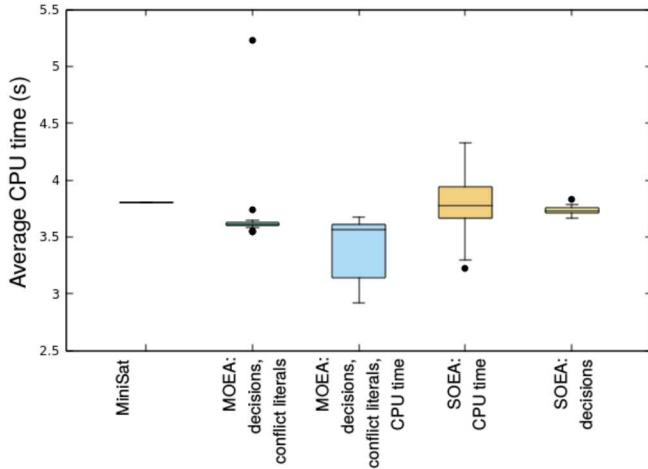


Fig. 1: Comparison of objectives on ibm-26

produced by each run. For each evolved variant and each original solver, we averaged the CPU time collected over 30 serial runs on each instance. All evaluations were performed on the same machine. Results are presented in Sect. V-B.

V. RESULTS

A. Comparison of Objectives

Figure 1 compares the CPU time that the evolved solvers produced by ADSSEC took to solve all instances in ibm-26 dataset. The SOEA with decisions as the objective produced statistically more efficient solvers than SOEA with CPU time as the objective. The MOEA with decisions, conflict literals, and CPU time produced more efficient solvers than any other combination of objectives that we compare, but these results had higher variance than those produced when not using CPU time as an objective. Our results serve as a guide for further experimentation: different EA parameter configuration or different datasets may yield different results.

B. Evolving Variable Scoring Heuristics

Figure 2 compares the CPU time required to solve instances in the ibm-18 dataset using the original MiniSat solver against that required using an evolved variant with a new variable scoring heuristic. While the minimum CPU times required are fairly close, there is a marked improvement with the new heuristic in the maximum CPU time needed. Because the dataset contains only 15 instances, ADSSEC’s training and the test instances are the same. However, in Figure 3 we show that the solver evolved on instances from the ibm-18 showed improvement in CPU time on other IBM Hardware Verification dataset subsets on which it was not trained (ibm-15 and ibm-03), especially in maximum CPU time.

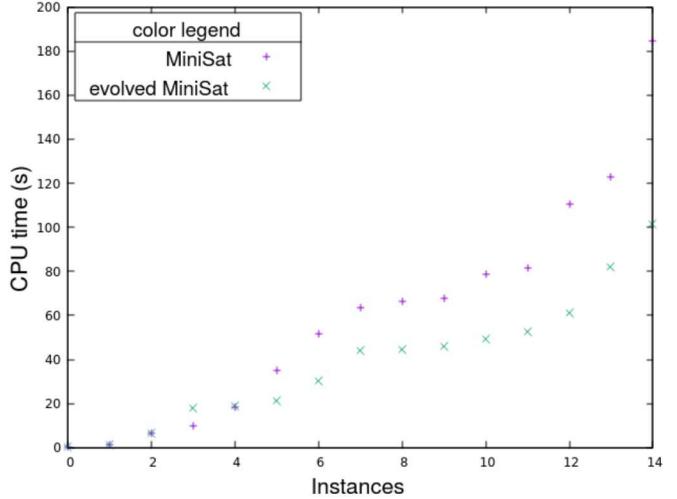


Fig. 2: ibm-18 cactus plot comparing CPU time to solve

Figure 4 compares the CPU time to solve instances from the unif-k5 dataset using the original MiniSat solver against that using MiniSat with an evolved heuristic. Significant improvement can be seen across the board using the evolved variable selection heuristic. The training set for ADSSEC was a randomly selected subset of ten instances from the dataset; both training and test instances are covered in the plot.

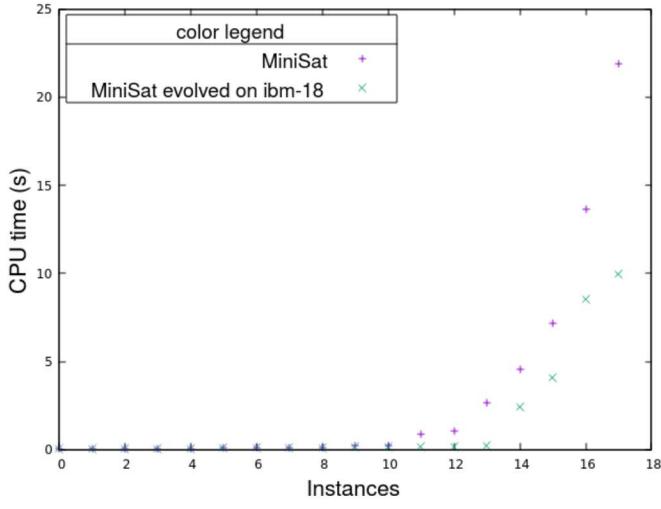
We show in Figure 5 that the solver evolved for the ibm-18 dataset did not show improvement in CPU time on instances from the unif-k5 dataset, which is from a different track. This supports the hypothesis that improved performance is obtained by targeting the specific structure of a problem class.

Figure 6 compares the CPU time to solve instances from the ibm-18 dataset using the original Glucose solver against that using Glucose with an evolved heuristic. Evolving a variable scoring heuristic in Glucose to target problems from the ibm-18 dataset did not show as much improvement as evolving the variable scoring heuristic in MiniSat. Still, some improvement was seen: the maximum CPU time to solve improved, and the average CPU time to solve all instances also improved slightly.

Figure 7 compares the CPU time to solve instances from the unif-k5 dataset using the original Glucose solver against that using Glucose with an evolved heuristic. As with MiniSat, the new heuristic for the unif-k5 dataset in Glucose significantly improved the solver for this dataset. The training set for ADSSEC was a randomly selected subset of ten instances from the dataset; both training and test instances are covered in the plot.

VI. DISCUSSION

In Figure 8 we compare solvers with evolved heuristics to some of the top state-of-the-art CDCL solvers, specifically Glucose 4.1 [35] and the SAT 16 competition version of Lingeling [40], on instances from the ibm-18 and unif-k5 datasets. We also show MiniSat for comparison. On the unif-k5 dataset, the evolved solvers significantly outperformed both Glucose and Lingeling. On ibm-18, the solvers did better than



(a) ibm-15

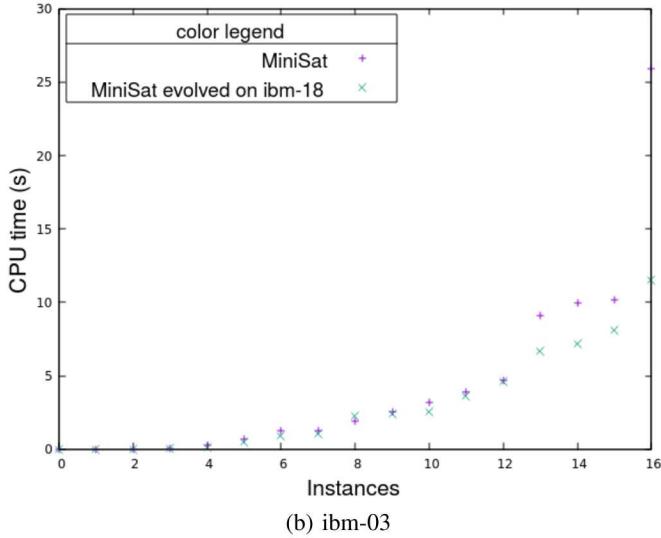


Fig. 3: Cactus plot comparing CPU time to solve of evolved solver trained on ibm-18 dataset and tested on other datasets from IBM Hardware Verification dataset

Lingeling. The evolved Glucose solver also did better than Glucose in overall runtime. Again, we averaged the CPU time collected over 30 runs of each solver on same machine in order to obtain more accurate CPU times.

In Figure 9, we present the tree representation of the best heuristic evolved by ADSSEC for the unif-k5 dataset using MiniSat. Interestingly, it is heavily based on the current number of conflicts encountered (i).

VII. CONCLUSIONS

Our results show that ADSSEC is capable of evolving variable scoring heuristics that are able to outperform default MiniSat or Glucose on a specific problem class, as detailed in Section VI. Even better, the evolved heuristics seem to outperform some of the state-of-the-art solvers.

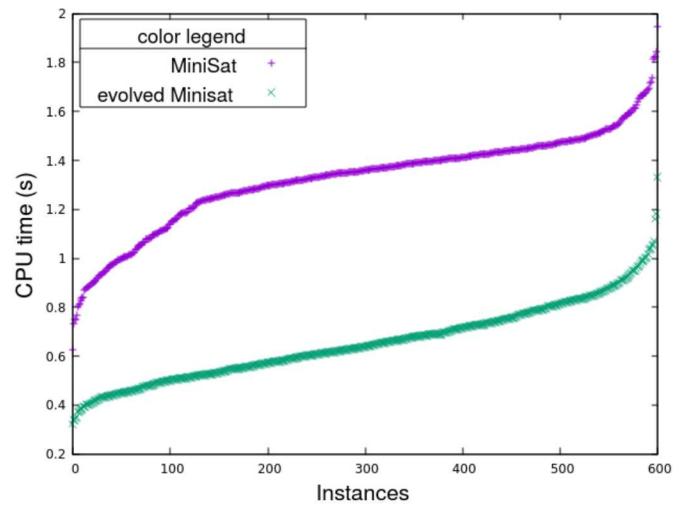


Fig. 4: unif-k5 cactus plot comparing CPU time to solve

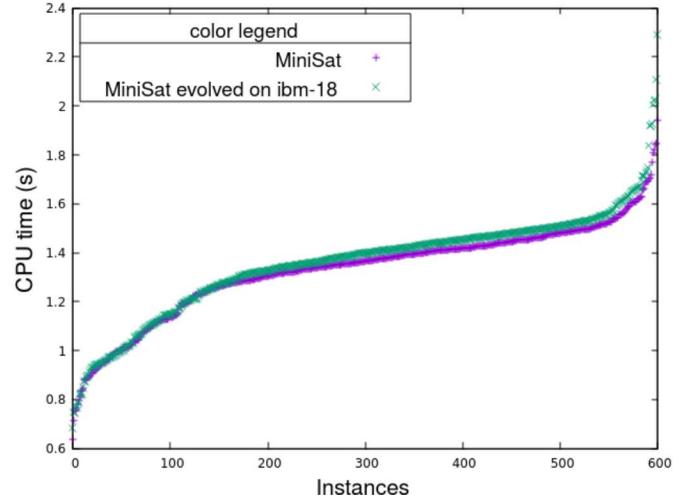


Fig. 5: unif-k5 cactus plot comparing CPU time to solve, evolved solver trained on ibm-18 dataset

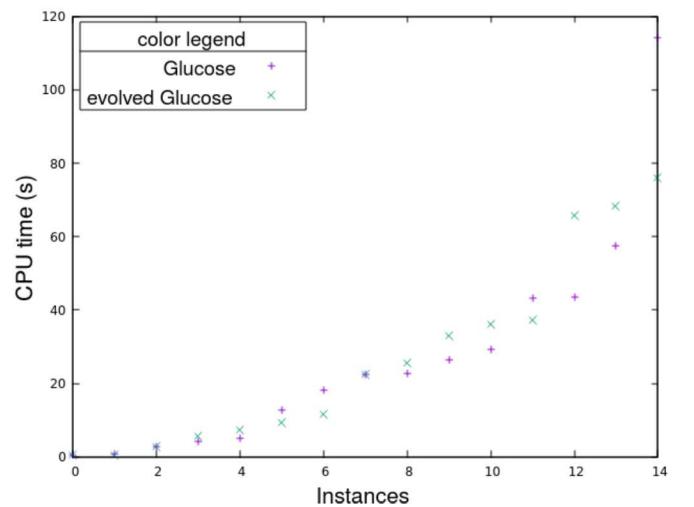


Fig. 6: ibm-18 cactus plot comparing CPU time to solve

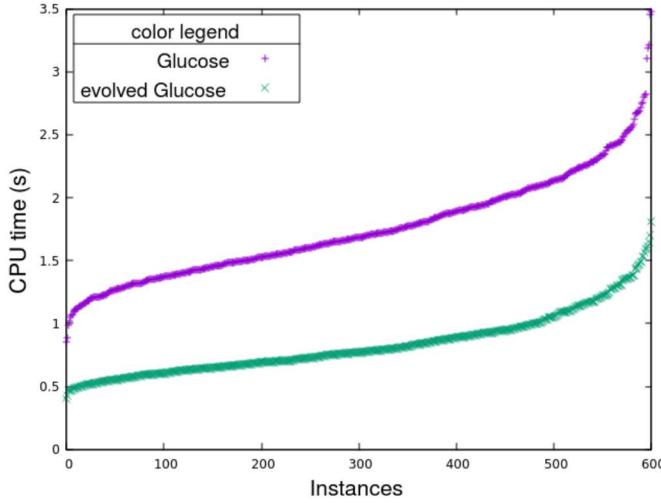
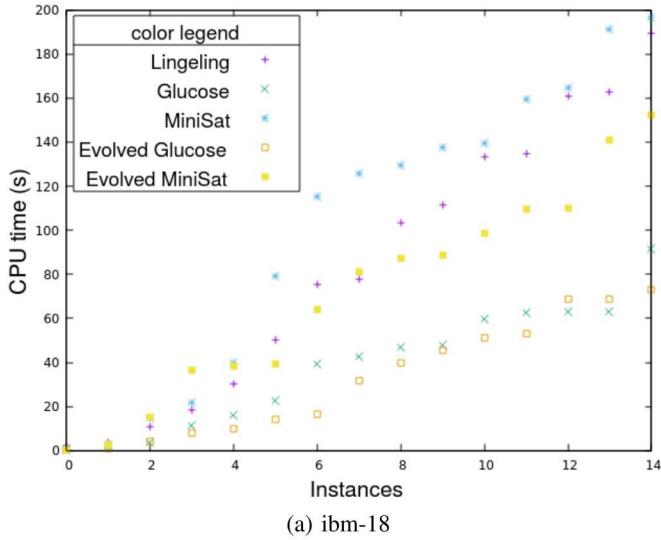


Fig. 7: unif-k5 cactus plot comparing CPU time to solve



(a) ibm-18

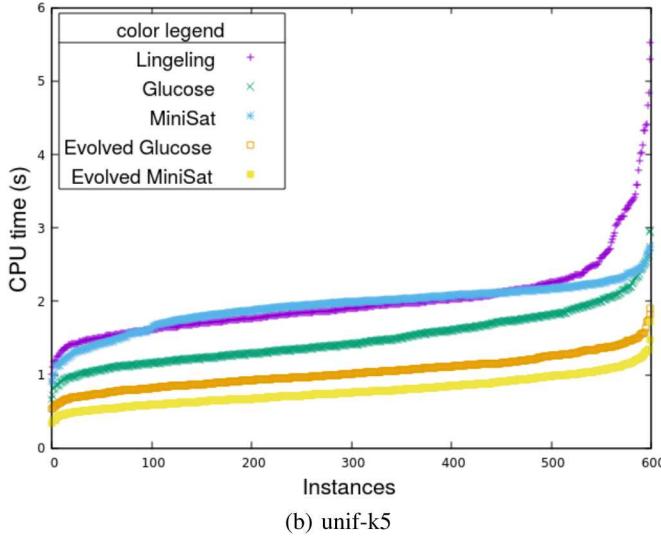


Fig. 8: Cactus plot comparing CPU time to solve of evolved solvers and top solvers

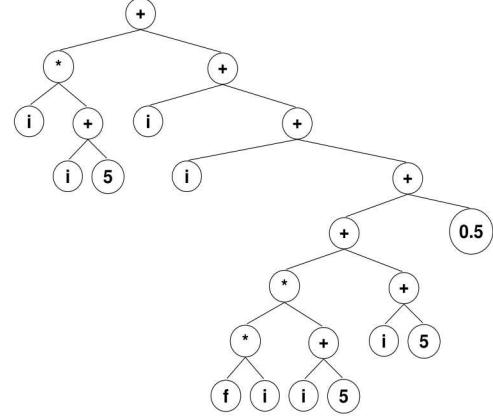


Fig. 9: Evolved heuristic for MiniSat for the unif-k5 dataset

We found that even when not using CPU time as an objective during the training phase, we were still able to show major improvement in CPU times of evolved solvers. As our results showed, a multi-objective evolutionary algorithm using CPU time as one of the objectives might lead to even better results than the ones we achieved.

Previous research showed that an automatically optimized solver that needs less CPU time to solve a set of instances is more energy efficient than an original solver [41]. We showed that ADSSEC found solvers that required less CPU time to solve instances from some benchmarks than some of the most energy efficient state-of-the-art solvers; thus, we believe that our system can produce more energy efficient solutions to real-world problems.

REFERENCES

- [1] J. Marques-Silva, “Practical applications of Boolean Satisfiability,” in *2008 9th International Workshop on Discrete Event Systems*, May 2008, pp. 74–80.
- [2] F. Corno and F. Razzak, “Intelligent Energy Optimization for User Intelligent Goals in Smart Home Environments,” *IEEE Transactions on Smart Grid*, vol. 3, no. 4, pp. 2128–2135, Dec. 2012.
- [3] G. Chin, D. Chavarra-Miranda, G. Nakamura, and H. Sofia, “BioGraphE: High-performance bionetwork analysis using the Biological Graph Environment,” *BMC bioinformatics*, vol. 9, p. S6, May 2008.
- [4] R. Soto, H. Kjellerstrand, O. Durn, B. Crawford, E. Monfroy, and F. Paredes, “Cell formation in group technology using constraint programming and Boolean satisfiability,” *Expert Systems with Applications*, vol. 39, no. 13, pp. 11423–11427, 2012.
- [5] S. Khurshid and D. Marinov, “TestEra: Specification-Based Testing of Java Programs Using SAT,” *Automated Software Engineering*, vol. 11, no. 4, pp. 403–434, Oct. 2004.
- [6] M. Soeken, R. Wille, M. Kuhlmann, M. Gogolla, and R. Drechsler, “Verifying UML/OCL Models Using Boolean Satisfiability,” in *Proceedings of the Conference on Design, Automation and Test in Europe*. 3001 Leuven, Belgium, Belgium: European Design and Automation Association, 2010, pp. 1341–1344.
- [7] Y. Vizel, G. Weissenbacher, and S. Malik, “Boolean Satisfiability Solvers and Their Applications in Model Checking,” *Proceedings of the IEEE*, vol. 103, no. 11, pp. 2021–2035, Nov. 2015.
- [8] C.-J. M. Liang, B. F. Karlsson, N. D. Lane, F. Zhao, J. Zhang, Z. Pan, Z. Li, and Y. Yu, “SIFT: Building an Internet of Safe Things,” in *Proceedings of the 14th International Conference on Information Processing in Sensor Networks*, ser. IPSN ’15. New York, NY, USA: ACM, 2015, pp. 298–309.

[9] F. Imeson and S. L. Smith, "A language for robot path planning in discrete environments: The TSP with Boolean satisfiability constraints," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, May 2014, pp. 5772–5777.

[10] ——, "Multi-robot task planning and sequencing using the SAT-TSP language," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, May 2015, pp. 5397–5402.

[11] E. Zarpas, "Benchmarking SAT Solvers for Bounded Model Checking," in *Theory and Applications of Satisfiability Testing—SAT 2005*, ser. Lecture Notes in Computer Science, vol. 3569. St Andrews, UK: Springer Berlin Heidelberg, 2005, pp. 340–354.

[12] ——, "Back to the SAT05 Competition: an a Posteriori Analysis of Solver Performance on Industrial Benchmarks," *Journal on Satisfiability, Boolean Modeling and Computation*, vol. 2, pp. 229–236, Jan. 2006.

[13] A. Van Gelder, "Another Look at Graph Coloring via Propositional Satisfiability," *Discrete Applied Mathematics*, vol. 156, no. 2, pp. 230–243, Jan. 2008.

[14] F. Hutter, H. H. Hoos, K. Leyton-Brown, and T. Stützle, "ParamILS: An Automatic Algorithm Configuration Framework," *Journal of Artificial Intelligence Research*, vol. 36, no. 1, pp. 267–306, Sep. 2009.

[15] F. Hutter, H. H. Hoos, and K. Leyton-Brown, "Sequential model-based optimization for general algorithm configuration," in *Learning and Intelligent Optimization*, ser. Lecture Notes in Computer Science. Rome, Italy: Springer Berlin Heidelberg, Jan. 2011, vol. 6683, pp. 507–523.

[16] S. Falkner, M. Lindauer, and F. Hutter, "SpySMAC: Automated Configuration and Performance Analysis of SAT Solvers," in *Theory and Applications of Satisfiability Testing—SAT 2015*, ser. Lecture Notes in Computer Science. Austin, TX, USA: Springer International Publishing, Sep. 2015, vol. 9340, pp. 215–222.

[17] A. R. KhudaBukhsh, L. Xu, H. H. Hoos, and K. Leyton-Brown, "SATenstein: Automatically Building Local Search SAT Solvers from Components," in *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, ser. IJCAI'09. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2009, pp. 517–524.

[18] M. Bader-El-Den and R. Poli, "Generating SAT Local-Search Heuristics Using a GP Hyper-Heuristic Framework," in *Artificial Evolution*, ser. Lecture Notes in Computer Science, vol. 4926. Tours, France: Springer Berlin Heidelberg, Oct. 2008, pp. 37–49.

[19] R. H. Kibria and Y. Li, "Optimizing the Initialization of Dynamic Decision Heuristics in DPLL SAT Solvers Using Genetic Programming," in *Genetic Programming*, ser. Lecture Notes in Computer Science. Budapest, Hungary: Springer Berlin Heidelberg, Apr. 2006, vol. 3905, pp. 331–340.

[20] A. S. Fukunaga, "Evolving Local Search Heuristics for SAT Using Genetic Programming," in *Genetic and Evolutionary Computation—GECCO 2004*, ser. Lecture Notes in Computer Science, vol. 3103. Seattle, WA, USA: Springer Berlin Heidelberg, Jun. 2004, pp. 483–494.

[21] ——, "Automated Discovery of Local Search Heuristics for Satisfiability Testing," *Evolutionary Computation*, vol. 16, no. 1, pp. 31–61, Apr. 2008.

[22] ——, "Massively Parallel Evolution of SAT Heuristics," in *Proceedings of the Eleventh Conference on Congress on Evolutionary Computation*, ser. CEC'09. Piscataway, NJ, USA: IEEE Press, 2009, pp. 1478–1485.

[23] A. B. andAndreas Fröhlich, "Evaluating CDCL Variable Scoring Schemes," in *SAT*, ser. Lecture Notes in Computer Science, vol. 9340. Springer, 2015, pp. 405–422.

[24] A. Biere and A. Fröhlich, "Evaluating CDCL Restart Schemes," in *Proceedings of the International Workshop on Pragmatics of SAT (POS'15)*, Austin, TX, Sep. 2015.

[25] A. R. Bertels and D. R. Tauritz, "Why Asynchronous Parallel Evolution is the Future of Hyper-heuristics: A CDCL SAT Solver Case Study," in *Proceedings of the 18th Annual Conference Companion on Genetic and Evolutionary Computation (GECCO '16)*, Denver, Colorado, USA, Jul. 2016, pp. 1359–1365.

[26] A. R. Bertels, "Automated Design of Boolean Satisfiability Solvers Employing Evolutionary Computation," Master's thesis, Missouri University of Science and Technology, Rolla, Missouri, USA, 2016.

[27] L. Xu, F. Hutter, H. H. Hoos, and K. Leyton-Brown, "SATzilla: Portfolio-Based Algorithm Selection for SAT," *Journal of Artificial Intelligence Research*, vol. 32, no. 1, pp. 565–606, May 2008.

[28] ——, "SATzilla2009: An Automatic Algorithm Portfolio for SAT," in *SAT 2009 Competitive Events Booklet*, Sep. 2009, pp. 53–55.

[29] F. Hutter, L. Xu, H. H. Hoos, and K. Leyton-Brown, "Algorithm Runtime Prediction: Methods & Evaluation," *Artificial Intelligence*, vol. 206, pp. 79–111, Jan. 2014.

[30] E. K. Burke, M. Gendreau, M. Hyde, G. Kendall, G. Ochoa, E. Özcan, and R. Qu, "Hyper-heuristics: a survey of the state of the art," *Journal of the Operational Research Society*, vol. 64, no. 12, pp. 1695–1724, Dec. 2013.

[31] M. A. Martin and D. R. Tauritz, "A Problem Configuration Study of the Robustness of a Black-Box Search Algorithm Hyper-Heuristic," in *Proceedings of the 16th Annual Conference Companion on Genetic and Evolutionary Computation (GECCO '14)*. Vancouver, BC, Canada: ACM, Jul. 2014, pp. 1389–1396.

[32] ——, "Hyper-Heuristics: A Study On Increasing Primitive-Space," in *Proceedings of the 17th Annual Conference Companion on Genetic and Evolutionary Computation (GECCO '15)*. Madrid, Spain: ACM, Jul. 2015, pp. 1051–1058.

[33] R. Poli, W. B. Langdon, N. F. McPhee, and J. R. Koza, *A Field Guide to Genetic Programming*. Lulu Enterprises, UK Ltd, Mar. 2008.

[34] N. Eén and N. Sörensson, "An Extensible SAT-solver," in *Theory and Applications of Satisfiability Testing—SAT 2003*, ser. Lecture Notes in Computer Science, vol. 2919. Santa Margherita Ligure, Italy: Springer Berlin Heidelberg, May 2003, pp. 502–518.

[35] G. Audemard and L. Simon, "Glucose and Syrup in the SAT Race 2015," in *Reports on the SAT 2015 Competition*, Austin, TX, USA, Sep. 2015.

[36] R. Tichy and T. Glase, "Clause Learning in SAT," http://www.cs.uni-potsdam.de/wv/lehre/05WS/05-AutoProbSolv/slides/learning_a.pdf, Apr. 2006, accessed: 2017-09-13.

[37] A. E. Eiben and J. E. Smith, *Introduction to Evolutionary Computing*, 2nd ed. Springer, 2015.

[38] R. G. on Learning, Optimization, and A. A. D. at Freiburg University, "Configurable SAT Solver Challenge (CSSC) 2014 - Benchmarks," <http://aclib.net/cssc2014/benchmarks.html>, [Online; accessed 02-April-2017].

[39] A. Balint, A. Belov, M. Jyrvisalo, and C. Sinz, "SAT Challenge 2012 Random SAT Track: Description of Benchmark Generation," in *Proceedings of SAT Challenge 2012: Solver and Benchmark Descriptions*, Department of Computer Science Series of Publications B, University of Helsinki, Finland, 2012, pp. 72–73.

[40] A. Biere, "Lingeling and Friends Entering the SAT Race 2015," FMV Reports Series, Institute for Formal Models and Verification, Johannes Kepler University, Altenbergerstr. 69, 4040 Linz, Austria, Tech. Rep., 2015.

[41] B. R. Bruce, J. Petke, and M. Harman, "Reducing Energy Consumption Using Genetic Improvement," in *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, ser. GECCO '15. New York, NY, USA: ACM, 2015, pp. 1327–1334.