

Structure-Based Local Search Heuristics for Circuit-Level Boolean Satisfiability

Anton Belov^{1*} and Matti Järvisalo^{2**}

¹ Complex and Adaptive Systems Laboratory, University College Dublin, Ireland

² Department of Computer Science, University of Helsinki, Finland

Abstract. This work focuses on improving state-of-the-art in stochastic local search (SLS) for solving Boolean satisfiability (SAT) instances arising from real-world industrial SAT application domains. The recently introduced SLS method CRSAT has been shown to noticeably improve on previously suggested SLS techniques in solving such real-world instances by combining justification-based local search with limited Boolean constraint propagation on the non-clausal formula representation form of Boolean circuits. In this work, we study possibilities of further improving the performance of CRSAT by exploiting circuit-level structural knowledge for developing new search heuristics for CRSAT. To this end, we introduce and experimentally evaluate a variety of search heuristics, many of which are motivated by circuit-level heuristics originally developed in completely different contexts, e.g., for electronic design automation applications. To the best of our knowledge, most of the heuristics are novel in the context of SLS for SAT and, more generally, SLS for constraint satisfaction problems.

1 Introduction

Stochastic local search (SLS) [11] is an important paradigm which facilitates finding solutions to various kinds of hard computational problems via searching over a declarative formulation of the problem at hand. It has been recognized that one possibility to push further the efficiency of SLS techniques is to develop search techniques that exploit the *structure* of constraint satisfaction problems. Indeed, various structure-exploiting SLS methods have been developed (among others) for generic constraint satisfaction problems (CSPs; for examples see [1, 2, 18, 10]) and Boolean satisfiability (SAT; for examples see [16, 20, 21, 17, 19, 14, 13, 22, 4, 5]).

This work focuses on developing efficient structure-exploiting SLS techniques for SAT. In more detail, we study techniques that are aimed at *industrially relevant* (or, as termed in the latest 2011 SAT Competition, *application*) instance classes. The most effective methods for solving *random* SAT instances are based on SLS. Furthermore, recent advances in SLS for *crafted* SAT instances has resulted in an SLS method winning the satisfiable crafted instance category of the 2011 SAT Competition³. In contrast, on industrial instances the current SLS methods are often notably inferior to the dominant conflict-driven clause learning (CDCL) SAT solvers.

* Partially supported by SFI PI grant BEACON (09/IN.1/I2618).

** Financially supported by Academy of Finland under grant 132812.

³ Results are available at <http://satcompetition.org/2011/>.

To the best of our knowledge, currently the best performing SLS method aimed at industrial SAT instances is CRSAT [5, 6]. Instead of working on the rather low conjunctive normal form (CNF) level, CRSAT searches for a solution directly on the level of arbitrary propositional formulas, relying on the compact representation form of Boolean circuits for a succinct way of representing propositional formulas. Furthermore, instead of relying on restricting search to input variables, as often has been proposed [16, 20, 21, 17, 19], CRSAT is based on the *justification-based* circuit-level SLS approach [14, 13], searching over the whole subformula structure, and incorporates a limited form of directed circuit-level Boolean constraint propagation to further exploit structural aspects of the input formulas[5].

We have recently shown that CRSAT can be further improved by incorporating a structure-based heuristic for focusing search steps. This resulted in the *depth-based* variant of CRSAT [6]. The depth-based heuristic has interesting fundamental properties, including the fact that CRSAT remains *probabilistically approximately complete* (PAC) [12] even when focusing search via the heuristic.

Contributions The success of the depth-based search heuristic suggests that circuit-level structural properties of SAT instances can indeed be exploited to further improve SLS. Motivated by this, in this work we develop and experimentally study a wide range of novel structure-based SLS search heuristics, focusing on CRSAT. We provide a systematic large-scale study of the proposed structure-based heuristics. We relate the heuristics to the depth-based heuristic studied in detail in [6], with the aim of developing further understanding on what are the underlying properties that make the depth-based search work in practice. Furthermore, we investigate whether related (or even completely different) structural properties result in even more effective heuristics. Analysis of the experiments reveals various interesting observations on the type of structural properties of circuits result in effective search heuristics.

Finally, as a future motivation for the studied heuristics, we are interested in extending the CRSAT approach, combining justification-based search over logical combinations of constraints and exploiting limited constraint propagation, to more generic classes of constraint satisfaction problems (CSPs) for which local search is a very viable alternative [1, 2, 18, 10]. The development of good structure-based search heuristics for the circuit-level is directly applicable for the logical combinations of more high-level constraints, where the logical combinations can be viewed as circuits.

Organization Key definitions and concepts related to Boolean circuit satisfiability are reviewed as necessary preliminaries in Sect. 2. Sect. 3 is dedicated to presenting the CRSAT circuit-level SLS algorithm for which this work develops structure-based search heuristics. The heuristics are introduced in Sect. 4. Before conclusions (Sect. 6), results of an extensive empirical evaluation on the effectiveness of the structure-based heuristics are presented in Sect. 5.

2 Preliminaries

A *Boolean circuit* over a finite set G of *gates* is a set C of equations of the form $g = f(g_1, \dots, g_n)$, where $g, g_1, \dots, g_n \in G$ and $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is a Boolean

function, with the additional requirements that (i) each $g \in G$ appears at most once as the left hand side in the equations in C , and (ii) the underlying directed graph $\langle G, E(C) \rangle$, where $E(C) = \{\langle g', g \rangle \in G \times G \mid g = f(\dots, g', \dots) \in C\}$, is acyclic. If $\langle g', g \rangle \in E(C)$, then g' is a *child* of g and g is a *parent* of g' . For a gate g , the sets of its children (i.e., the *fanin* of g) and parents (i.e., the *fanout* of g) are denoted by $\text{fanin}(C, g)$ and $\text{fanout}(C, g)$, respectively. The *descendant* and *ancestor* relations fanin^* and fanout^* are the transitive closures of the child and parent relations, respectively. If $g = f(g_1, \dots, g_n)$ is in C , then g is an f -gate (or of type f). A gate with no children (resp. no parents) is an *input gate* (resp. an *output gate*). The sets of input gates and output gates in C are denoted by $\text{inputs}(C)$ and $\text{outputs}(C)$, respectively. A gate that is neither an input nor an output is an *internal gate*. Typical gate types include NOT ($\text{NOT}(v)$ is 1 iff v is 0) and AND ($\text{AND}(v_1, v_2)$ is 1 iff both v_1 and v_2 are 1).

An (*truth*) *assignment* for C is a (possibly partial) function $\tau : G \rightarrow \{0, 1\}$. A complete assignment τ for C is *consistent* if $\tau(g) = f(\tau(g_1), \dots, \tau(g_n))$ for each $g = f(g_1, \dots, g_n)$ in C . When convenient we write $\langle g, v \rangle \in \tau$ instead of $\tau(g) = v$. The *domain* of τ , i.e., the set of gates assigned in τ , is denoted by $\text{dom}(\tau)$. We say that two assignments, τ and τ' , *disagree* on a gate $g \in \text{dom}(\tau) \cap \text{dom}(\tau')$ if $\tau(g) \neq \tau'(g)$. For a truth assignment τ and set of gates $G \subseteq \text{dom}(\tau)$, let $\text{flip}(G, \tau)$ denote the truth assignment obtained by changing the values of the gates in G , and leaving the rest of τ unchanged.

A *constrained Boolean circuit* C^α consists of a Boolean circuit C and an assignment α for C . Each $\langle g, v \rangle \in \alpha$ is a *constraint*, and g is *constrained* to v if $\langle g, v \rangle \in \alpha$. A complete assignment τ for C *satisfies* C^α if (i) τ is consistent with C , and (ii) it respects the constraints: $\tau \supseteq \alpha$. If some assignment satisfies C^α , then C^α is *satisfiable*. A circuit that is not satisfiable is *unsatisfiable*. Without loss of generality, we assume that constraints are imposed only on output gates.

The *restriction* $\tau|_{G'}$ of an assignment τ to a set $G' \subseteq G$ of gates is defined as $\{\langle g, v \rangle \in \tau \mid g \in G'\}$. Given a gate equation $g = f(g_1, \dots, g_n)$ and a value $v \in \{0, 1\}$, a *justification* for the pair $\langle g, v \rangle$ is a partial assignment $\sigma : \{g_1, \dots, g_n\} \rightarrow \{0, 1\}$ to the children of g such that $f(\tau(g_1), \dots, \tau(g_n)) = v$ holds for all extensions $\tau \supseteq \sigma$. That is, the values assigned by σ to the children of g are enough to force g to take the consistent value v . For example, the justifications for $\langle g, 0 \rangle$, where $g = \text{AND}(g_1, g_2)$, are $\{\langle g_1, 0 \rangle\}$, $\{\langle g_2, 0 \rangle\}$, and $\{\langle g_1, 0 \rangle, \langle g_2, 0 \rangle\}$, out of which the first two are *subset-minimal*. A gate g is *justified in an assignment* τ if it is assigned, i.e. $\tau(g)$ is defined, and (i) it is an input gate, or (ii) $g = f(g_1, \dots, g_n) \in C$ and $\tau|_{\{g_1, \dots, g_n\}}$ is a justification for $\langle g, \tau(g) \rangle$. We denote the set of *unjustified* gates in an assignment τ by $\text{unjust}(C^\alpha, \tau)$.

3 CRSat: Justification-Based SLS with Forward Propagation

CRSAT is an SLS-based SAT algorithm for Boolean circuits that operates directly on circuit structure – that is, without the conversion to CNF. The algorithm was first described in [5] and was subsequently analyzed theoretically and improved in [6]. In this section we provide a high-level overview of the algorithm, and refer the reader to [5, 6] for additional details.

CRSAT is based on the *justification-based* [14, 13] approach to circuit-level SLS. In this approach, the circuit is traversed from the outputs to inputs, and the values of the *internal* gates are adjusted using local information in an attempt to eliminate all unjustified gates. CRSAT combines a weakened version of justification-based SLS with so called *limited forward propagation* – a restricted form of circuit-level Boolean constraint propagation, described in what follows.

Pseudo-code for CRSAT is presented as Algorithm 1. First, a complete extension of a random value assignment to $\text{inputs}(C^\alpha)$ is constructed, i.e., the value of each unconstrained internal gate is set consistently with the values of its children. Then, as long as $\text{unjust}(C^\alpha, \tau)$ is not empty (i.e., τ is not a satisfying assignment), the algorithm heuristically selects an unjustified gate g (line 6; we will discuss gate selection heuristics in the next section in detail). Once an unjustified gate g is chosen, the algorithm selects a justification σ for $\langle g, \tau(g) \rangle$ (lines 7- 13) and performs a search *step*. The latter consists of (i) flipping the values of gates on which σ and τ disagree (line 15), followed by (ii) propagating the consequences of the flip towards the outputs of the circuit (line 16).

Algorithm 1 Generic CRSAT($C^\alpha, wp, cutoff$)

Input: C^α – constrained Boolean circuit
 wp – noise parameter ,i.e., probability of random walk
 $cutoff$ – cutoff, i.e., maximum number of steps

Output: $status$ – SAT if a satisfying assignment for C^α is found, UNKNOWN otherwise
 τ – a satisfying assignment for C^α if found, \emptyset otherwise

- 1: $\tau \leftarrow$ a complete extension of a random assignment to $\text{inputs}(C^\alpha)$
- 2: $steps \leftarrow 0$
- 3: **while** $steps < cutoff$ **do**
- 4: **if** $\text{unjust}(C^\alpha, \tau) = \emptyset$ **then**
- 5: **return** $\langle \text{SAT}, \tau \rangle$
- 6: $g \leftarrow$ a heuristically selected gate from $\text{unjust}(C^\alpha, \tau)$
- 7: $\Sigma \leftarrow$ the set of justifications for $\langle g, \tau(g) \rangle$
- 8: **with-probability** wp **do**
- 9: $\sigma \leftarrow$ random element of Σ ▷ random walk
- 10: **otherwise**
- 11: $\sigma \leftarrow$ a random justifications from those in Σ that minimize
- 12: the number of unjustified gates after the step ▷ greedy downward move
- 13: **end with-probability**
- 14: $G \leftarrow$ set of gates in σ that disagree with τ
- 15: $\tau \leftarrow \text{flip}(G, \tau)$ ▷ flip
- 16: $\tau \leftarrow \text{LBFP-FORWARD}(C^\alpha, G, \tau)$ ▷ limited forward propagation
- 17: $steps \leftarrow steps + 1$
- 18: **return** $\langle \text{UNKNOWN}, \emptyset \rangle$

The justification σ used to make a step can be selected from the set Σ of all justifications for $\langle g, \tau(g) \rangle$ either at random (with probability wp), or greedily with the objective of minimizing the number of unjustified gates after the step. Note that taking Σ to be a set of *subset-minimal* justifications results in good performance in practice; this is also how our current implementation works.

The forward propagation procedure LBCP-FORWARD is presented as Algorithm 2. It uses a priority queue Q of gates (with no duplicates) that allows to query the *smallest* gate according to a topological order in constant time⁴. Essentially, the procedure implements a circuit-level Boolean constraint propagation algorithm, except that (i) the values are propagated *only* towards the outputs of the circuit, and (ii) propagation along each path stops immediately when an unjustified gate becomes justified; hence it implements *limited forward propagation*. The addition of limited forward propagation to justification-based SLS results in multiple orders of magnitude speed-ups on industrial SAT instances [3].

Algorithm 2 LBCP-FORWARD(C^α, G, τ)

Input: C^α – constrained Boolean circuit;
 G – a set of gates whose value changes are to be propagated.
 τ – an assignment for C^α ;
Output: τ' – an assignment for C^α which is a result of limited forward propagation of the assignment $\tau|_G$.

- 1: $\tau' \leftarrow \tau$
- 2: Q .ENQUEUE(G)
- 3: **while** $\neg Q$.EMPTY **do**
- 4: $g \leftarrow Q$.POP_FRONT
- 5: **if** $g \in G$ **then** $\triangleright g$ is one of the original gates
- 6: Q .ENQUEUE(fanout(g))
- 7: **else**
- 8: **if** $g \in \text{unjust}(C^\alpha, \tau') \setminus \text{dom}(\alpha)$ **then** $\triangleright g$ unconstrained and unjustified
- 9: $\tau' \leftarrow \text{flip}(\{g\}, \tau')$
- 10: Q .ENQUEUE(fanout(g))
- 11: **return** τ'

It comes as no surprise that the effectiveness of CRSAT depends critically on the way the gates are selected for justification during the search (Line 6 of Algorithm 1). A good selection heuristic focuses search to the most important gates in terms of satisfiability. On the other hand, if a too deterministic (focused) selection procedure is used, the search may not converge into a satisfying assignment. In [6] we showed that the efficiency of CRSAT can be significantly improved by focusing the search using a structure-based gate selection heuristic which takes into account the *depth* of the selected gates. In the next section we describe a number of additional structural properties of gates and propose a number of gate selection heuristics based on these properties.

4 Structure-Based Search Heuristics for CRSAT

In this section we introduce a number of heuristics for selecting of the unjustified gate to justify at each search step in the main loop of CRSAT (line 6 of Algorithm 1). The underlying idea is that these heuristics should be able to take into account the structural

⁴ Recall that a topological order on the set of gates in a circuit is any strict total order $<$ that respects the condition “if $g_1 \in \text{fanin}(g_2)$, then $g_1 < g_2$ ”

properties of the constrained Boolean circuit at hand, and focus the search on the gates that are deemed important based on these properties. Additionally, we must aim at *efficiently* computable heuristics, as the main loop may be executed millions of times in a typical run of the algorithm (although, in contrast to typical SLS algorithms, most of the computation effort in CRSAT is attributed to the execution of forward propagation, and hence we can afford slightly more expensive computations than usual SLS heuristics).

We now give a listing of the initial set of gate properties, with intuition on why these properties may be interesting. We then describe the corresponding gate selection heuristics, and, in the next section, present the results of the preliminary empirical evaluation of these heuristics. The analysis of the results will lead us to the development of additional heuristics, which will be described and analyzed in Sect. 5.

Depth: $\text{depth}(C, g)$, where the *depth* of a gate g in C is

$$\text{depth}(C, g) = \begin{cases} 0 & \text{if } g \in \text{outputs}(C^\alpha) \\ 1 + \max\{\text{depth}(C, g') \mid g' \in \text{fanout}(C, g)\} & \text{otherwise.} \end{cases}$$

The importance of gate depth for CRSAT was justified theoretically and confirmed empirically in [6]. The key aspect is that selection of gates with high depth drives the algorithm close to the inputs of the circuit, thus allowing the algorithm to explore the space of assignments to input gates faster.⁵ The depth of all gates in C can be computed in $O(|C|)$ time (where $|C|$ denotes the number of gates in C), and stored for constant time retrieval.

FO: $|\text{fanout}(C, g)|$

Gates with large *fanout size* are in a sense more influential than the rest. Intuitively, by forcing CRSAT to justify these gates, the truth values of these critical parts of the circuit are fixed first, which may result in many of the other gates' values to be set by forward propagation. The fanout size of a gate is retrieved in constant time.

TFO: $|\text{fanout}^*(C, g)|$

This is also a measure of the influence of the gate in the circuit: intuitively, the larger *the size of the transitive fanout*, the more influence the gate's value has on transitively justifying the output constraints of the circuit via forward propagation. The computation of the size of the transitive fanout of a gate requires $O(|C|)$ in the worst-case (although typically only a fraction of gates in C have to be evaluated).

TFI: $|\text{fanin}^*(C, g)|$

The *size of the transitive-fanin* of a gate g can be considered an estimate of the number of search steps required to justify all gates in the sub-circuit rooted at g . This measure is also related to the size of the *interest set* used as an objective function in justification-based SLS algorithm BC SLS [14, 13]. The computation of the size of the transitive fanin of a gate requires $O(|C|)$ in the worst-case.

CC: $\text{CC}(C, g, \tau(g))$, where the *SCOAP (Sandia Controllability and Observability Analysis Program) combinational controllability measure* [9] CC is defined as follows:

⁵ Here one should notice that driving the search towards input gates in justification-based search is different from the idea of restricting *the flips* to input gates as in [16, 20, 21, 17, 19] due to the conceptual differences of these approaches.

$$\text{CC}(C, g, 0) = \begin{cases} 1 & \text{if } g \in \text{inputs}(C) \\ 1 + \min_{g' \in \text{fanin}(C, g)} \text{CC}(C, g', 0) & \text{if } g \text{ is an AND-gate,} \end{cases}$$

$$\text{CC}(C, g, 1) = \begin{cases} 1 & \text{if } g \in \text{inputs}(C) \\ 1 + \sum_{g' \in \text{fanin}(C, g)} \text{CC}(C, g', 1) & \text{if } g \text{ is an AND-gate.} \end{cases}$$

Given a gate g and its current value v_g , SCOAP aims to provide a measure of how difficult it is to satisfy the sub-circuit rooted at g given that g is constrained to v_g (i.e., to *control* the value v_g at g). Originally, SCOAP was used as a combinational testability measure. For our purposes, SCOAP intuitively provides a measure of how difficult it is to transitively justify the output constraints of a circuit. Due to the fact that we apply And-Inverter graphs (AIGs) as benchmark instances in this paper, the definition is restricted to AND-gates only. However, the definition can be naturally extended to other gate types.

Here one should notice the original definition of SCOAP assigns for NOT-gates (negations) the value of the gate's child *incremented by one*. In contrast, here we do no increment such values, but instead *implicitly skip* NOTs in the following sense. In case $g = \text{NOT}(g')$, all gates in fanout of g' are included in fanout of g' instead of g . This is due to the fact that negations (inverters) are handled implicitly in the justification steps and forward propagation performed by CRSAT, and hence the CC value assigned to each NOT-gate equals the value assigned to the gate's child. Note that SCOAP controllability measures for all gates in C can be computed in $O(|C|)$ time.

CO: $\text{CO}(C, g)$, where the *SCOAP combinational observability measure* [9] is defined as follows:

$$\text{CO}(C, g) = \begin{cases} 0 & \text{if } g \in \text{outputs}(C) \\ 1 + \min_{g' \in \text{fanout}(C, g)} \text{CCO}(C, g', g) & \text{otherwise,} \end{cases}$$

where for an AND-gate we have

$$\text{CCO}(C, g', g) = \text{CO}(C, g') + \sum_{g'' \in \text{fanin}(C, g') \setminus \{g\}} \text{CC}(C, g'', 1).$$

As in CC, we implicitly skip NOTs in the definition. This measure attempts to capture how difficult it is to *observe* a specific value for a gate given the output constraints; in other words, how likely is it that the value is part of a minimal justification that is transitively consistent with the output constraints. The measure can be computed for all gates in C in $O(|C|)$ time.

Flow: $\text{flow}(C, g)$, where the *output flow value* of a gate g in C is

$$\text{flow}(C, g) = \begin{cases} 1 & \text{if } g \in \text{outputs}(C) \\ \sum_{g' \in \text{fanout}(C, g)} \frac{\text{flow}(C, g')}{|\text{fanout}(C, g')|} & \text{otherwise.} \end{cases}$$

In other words, we compute a total flow value for each gate by pouring a unit quantity flow down from the output gates of the circuit. Here it is important to

notice that the definition of flow implicitly skips NOT-gates. This flow-based idea was first evaluated in [15] as a heuristic for restricting the set of decision variables in CDCL solvers. Our intuition is that, if a large total flow passes through a particular gate, the gate is *globally* very connected with the constraints in τ , approximating in a sense the number of possible paths for forward propagation, and thus g would have an important role in the satisfiability of the circuit.

Each of the structural properties presented above gives rise to a pair of gate selection heuristics: for a given property $f(C^\alpha, g, \tau)$, one heuristic selects at random a gate from

$$\operatorname{argmax}_{g \in \operatorname{unjust}(C^\alpha, \tau)} f(C^\alpha, g, \tau).$$

We will refer to this heuristic as a *max-variant*, ***f-max***, of the heuristic based on f . And, a dual heuristic, the *min-variant*, ***f-min***, selects at random a gate from

$$\operatorname{argmin}_{g \in \operatorname{unjust}(C^\alpha, \tau)} f(C^\alpha, g, \tau).$$

Thus, we have 7 pairs of dual heuristics, and the baseline heuristic **Rand** that simply selects a random gate from $\operatorname{unjust}(C^\alpha, \tau)$ – this is the heuristic used in the original paper on CRSAT [5].

We now note that some of the presented structural measures of gates are in parts correlated (either positively or negatively) with gate depth (these are **TFO**, **TFI**, **CC**, **CO**), while others (**FO**, **Flow**) are not. The reason that we pay a particular attention to the depth is that we know that the **Depth-max** heuristic is very effective [6]. As such, when we evaluate the heuristics based on the properties that are positively correlated with depth (*depth-friendly* heuristics) we are interested in detecting improvements over **Depth-max**. Such an improvement would suggest that another, perhaps more fundamental property, is at play in CRSAT-style circuit SLS. Furthermore, the duals of depth-friendly heuristics are expected a priori to perform poorly. In evaluating the heuristics that are not correlated with depth (*depth-agnostic* heuristics), we are also interested in detecting significant differences in performance on some classes, or even on particular problem instances. Such differences would suggest that depth-agnostic heuristics might be used as secondary heuristics in CRSAT (e.g. for tie-breaking).

To summarize, the following heuristics are the primary targets of the empirical evaluation and analysis presented in the next section:

- *Baseline*: **Rand** and also **Depth-max**.
- *Depth-friendly*: **TFO-max**, **TFI-min**, **CC-min** (small controllability value means the gate is *easy* to control, and hence intuitively close to inputs), **CO-max** (large observability value means the gate is *difficult* to observe, and hence intuitively far from outputs).
- *Depth-agnostic*: **FO-min**, **FO-max**, **Flow-min**, **Flow-max**.

5 Evaluation

In order to provide an objective empirical comparison of SLS solvers, the well-known SLS textbook by Hoos and Stütze [11] suggests a procedure for finding near-optimal

noise (the setting of the parameter wp in Alg. 1) by essentially binary searching for the noise values for each individual instance and solver to be evaluated. While full binary search is computationally infeasible given the vast number of benchmark instances used in our experiments and, on the other hand, the computational resources available to us, we applied the following approximation of the Hoos-Stützle scheme. Noise was optimized for each solver and instance individually based on 25 tries using a timeout of 200 seconds per try (with no limit on the number of steps), at noise values 0.05, 0.1, 0.2, 0.3, 0.4, 0.5. The noise with highest success rate (primary criterion) and best median time (secondary criterion) was selected. In cases where there were two or more options ranked best using both of these criteria, a random candidate among those options was picked. Note that the benchmark-class based noise optimization, which is computationally cheaper, is often insufficient on industrial application benchmarks. For example, among 61 solved instances of one of the benchmark classes described below (**sss-sat-1.0**) we found 10 instances to have a near-optimal noise value, wp_{no} , of 0.05, 10 with $wp_{no} = 0.1$, 14 instances with $wp_{no} = 0.2$, 9 instances with $wp_{no} = 0.3$, 9 instances with $wp_{no} = 0.4$ and 9 instances with $wp_{no} = 0.5$.

The reported CPU times and number of steps for each instance are the median CPU time (in seconds) and the median number of search steps with the best noise setting over 25 tries for the experiments summarized in Fig. 1 and 3, and over 100 tries for the experiments summarized in Fig. 2.⁶ The experiments were performed on an HPC cluster, each node of which runs on a dual quad-core Xeon E5450 3-GHz with 32 GB of memory.

5.1 Benchmark Families

As benchmarks, we considered over 650 And-Inverted circuits (AIGs, that is, constrained Boolean circuits in which gate types AND and NOT are used) from five different industrial application benchmark classes. We obtained the AIGs as described in the following.

hwmcc08-sat 204 satisfiable AIGs obtained from the Hardware Model Checking Competition 2008 problems⁷. The original sequential circuits were unfolded using the `aigtobmc` tool (part of the AIGer package⁸) The step bound $k = 45$ was used for the time frame expansion.

smtqfbv-sat 61 satisfiable AIGs generated by using the Boolector SMT solver⁹ [8] to bit-blast QF_BV (theory of bit-vectors) instances of the SMT Competition 2009¹⁰ into AIGs.

⁶ Based on our experience, given the large number of instances, 25 tries is enough to detect the main trends. The experiments described in Fig. 2 require more precision.

⁷ Original instances available at <http://fmv.jku.at/hwmcc08/>.

⁸ Available at <http://fmv.jku.at/aiger/>

⁹ <http://fmv.jku.at/boolector/>

¹⁰ <http://www.smtcomp.org/2009/>

- sss-sat-1.0** 98 satisfiable AIGs from “formal verification of buggy variants of a dual-issue superscalar microprocessor”¹¹ [23]. These circuits, originally in the ISCAS format, were converted to AIG using the ABC system¹² [7].
- vliw-sat-1.1** 98 satisfiable AIGs from “formal verification of buggy variants of a VLIW microprocessor”, originating from the same place and converted to AIG in a similar fashion as sss-sat-1.0 instances.
- sat-race** Satisfiable AIGs filtered from a total of 200 instances used in the final round of *structural SAT track* of the SAT Race 2008 and 2010 competitions.¹³

In order to be able carry out the experiments in practice, we picked a selection of a total of 300 instances from these benchmark classes as follows. Based on the good performance reported in [6] for the **Depth-max** heuristic, we filtered out trivial instances for **Depth-max** (instances for which the median number of steps was < 730). From the remaining ones, in order to we picked those instances that we consider *solved* by **Depth-max** (i.e., instances for which the success rate for **Depth-max** was $\geq 50\%$)¹⁴. This resulted in the following distribution of instances: hwmcc08 – 95, smtqfbv – 46, sss-sat-1.0 – 61, vliw-sat-1.1 – 96, and sat-race – 2.

5.2 Results and Analysis

Fig. 1 presents a “cactus” plot, i.e., the number of instances that can be solved within a given time¹⁵, summarizing the comparative performance of the 15 structure-based gate selection heuristics described in Sect. 4. The following conclusions can be drawn.

First, we note that whenever a heuristic outperforms the baseline **Rand** heuristic, its dual performs worse than **Rand**, and vice versa. In fact, we see that in many cases the better the performance of a heuristic, the worse is the performance of its dual. This suggests that the properties proposed in Sect. 4 are meaningful in the context of CRSAT. One exception to the nice “symmetric” picture is the pair based on SCOAP combinational controllability **CC**, where the worse of the duals, **CC-max**, performs surprisingly close to the baseline **Rand** – we will discuss this point later. An additional observation is that the depth-friendly heuristics **TFO-max**, **TFI-min**, **CC-min** and **CO-max** always perform significantly better than their duals, and, furthermore, form most of the best performing heuristics. This corroborates the hypothesis that there is an important underlying property correlated with the depth of gates.

Second, we observe surprisingly good performance from the depth-agnostic **Flow-min**. Recall that, intuitively, gates with high flow are those that have large influence on other gates in the circuit. Thus, on the surface, this result casts doubt on the role of the

¹¹ Available at http://www.miroslav-velev.com/sat_benchmarks.html



¹² <http://www.eecs.berkeley.edu/~alanmi/abc/>

¹³ Available at <http://baldur.iti.uka.de/sat-race-2010/downloads.html>

¹⁴ This allowed us to perform these extensive experiments in practice within the given time frame. We hope to extend the experiments also to those instances unsolved by **Depth-max**.

¹⁵ The median CPU times were used for the plot. The median number of search steps would also be an appropriate measure for comparing the quality of search heuristics. However, the relative performance differences based on time and on number steps are very close in this case, and the cactus plot using running times is easier to read.

Fig. 1. A comparison of the performance of 15 gate-selection heuristics described in Sect. 4 as a cactus plot, i.e., the number of those instances that can each *solved* within a given time limit. An instance is considered *solved* if a success rate over the 25 tries is $\geq 50\%$. The CPU time of a solved instance is the median CPU time for the instance over all runs (including the unsuccessful ones).



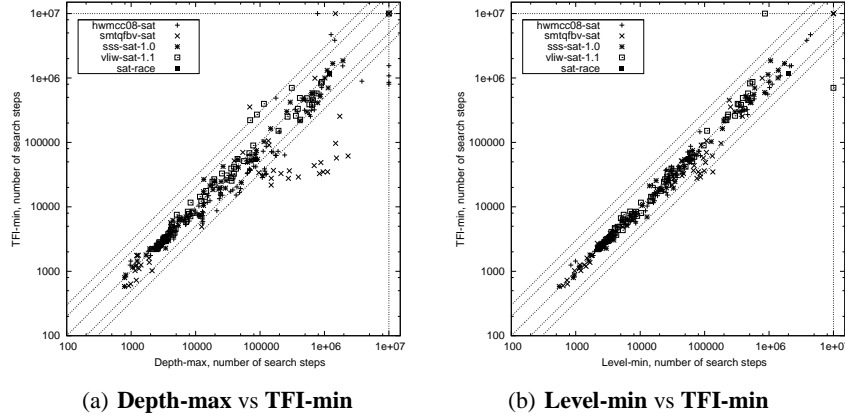


Fig. 2. Scatter plots that compare the performances of selected heuristics in terms of the median number of steps, over 100 tries. Timed-out instances are plotted with the number of steps set to 10^7 , on the vertical and horizontal lines.

influential gates in the context of CRSAT. On the other hand, between the two duals based on the size of the fanout of gates, it is the **FO-max** that performs well, rather than **FO-min**. A closer look at some of our instances resolves this apparent contradiction – the flow is *not* depth-agnostic, but, in fact, is negatively correlated with depth. The reason for this is that most of our benchmark circuits have significantly more inputs than outputs, and thus gates that are close to inputs tend to have small flow values. At the same time, we did not detect any interesting relationships between **Depth-max** and **FO-min**, most likely due to the fact that the latter is much more a local property than the former. This suggests that to further study the effects of “influence” of gates in the context of CRSAT, alternative measures are needed, e.g., ones that are based on graph-theoretic centrality measures. This conclusion is corroborated by the fact that, although the depth-friendly heuristics capture high influence — gates with large depth often have large transitive fanout and thus have high influence through forward propagation — the results show that **TFO-max** is not the best performing heuristic.

Finally, we observe that the SCOAP-based heuristics **CC-min** and **CO-max**, as well as the **TFO-max** heuristic based on the size of the transitive fanout of gates, do not perform as well as **Depth-max**. However, in contrast, the heuristic that prefers gates with *small transitive fanin*, **TFI-min**, appears to perform noticeably better than **Depth-max**. The scatter plot in Fig. 2(a) which compares the performance of these two heuristics in terms of the number of search steps demonstrates that the size of the transitive fanin of gates can provide a better guidance to CRSAT than the depth of the gate.

Note that gates with small transitive fanin are very likely to be close to the inputs. Based on the theoretical analysis of CRSAT in [3] and [6] the performance of the algorithm should improve if it arrives to the input level frequently. Hence, to get insight into the reasons of the good performance of **TFI-min**, we need to understand whether the heuristic is effective simply because it brings the algorithm faster to the input level, or whether there is another mechanism at play. One way to investigate the answer to this

question is to compare the performance of **TFI-min** with a heuristic that is based on a measure that disregards the number of gates in the sub-circuit rooted at the gate, and takes into account only the distance from the gate to the input level. Such a measure, well known in EDA literature, is called the *level* of a gate, and is defined as follows:

Level: $\text{level}(C, g)$, where the *level* of a gate g in C is

$$\text{level}(C, g) = \begin{cases} 0 & \text{if } g \in \text{inputs}(C) \\ 1 + \max\{\text{level}(C, g') \mid g' \in \text{fanin}(C, g)\} & \text{otherwise.} \end{cases}$$

Thus, $\text{level}(C, g)$ is simply the maximum distance from the gate g to an input gate, and so the depth-friendly heuristic based on level, **Level-min**, would control the search solely based on the distance to the inputs.

The comparative performance of **TFI-min** and **Level-min** is presented in the scatter plot in Fig. 2(b). We observe that performances of the two heuristics are highly correlated. As such, this comparison does not give a definitive answer to the question of which measure is more fundamental for CRSAT. To gain some insight, we can introduce heuristics that go for the input gates more aggressively than **Level-min**. Such heuristics can, for instance, be based on the following measures:

LLevel: $\text{llevel}(C, g)$, where the “low” level of a gate g in C is

$$\text{llevel}(C, g) = \begin{cases} 0 & \text{if } g \in \text{inputs}(C) \\ 1 + \min\{\text{llevel}(C, g') \mid g' \in \text{fanin}(C, g)\} & \text{otherwise.} \end{cases}$$

Alevel: $\text{alevel}(C, g)$, where the “average” level of a gate g in C is

$$\text{alevel}(C, g) = \begin{cases} 0 & \text{if } g \in \text{inputs}(C) \\ 1 + \sum_{g' \in \text{fanin}(C, g)} \text{level}(C, g') / |\text{fanin}(C, g)| & \text{otherwise.} \end{cases}$$

Thus, the “low” level of g is the shortest distance from g to some input gate, while the “average” level of g is somewhere in between the level and the “low” level; that is, we always have $\text{level}(C, g) \geq \text{alevel}(C, g) \geq \text{llevel}(C, g)$. As such, the **LLevel-min** heuristics will drive the search to the input gates extremely aggressively, while the **Alevel-min** heuristic represents a middle ground between **Level-min** and **LLevel-min**.

The cactus plot in Fig. 3 summarizes the comparative performance in terms of CPU time of the three level-based heuristics described above and **TFI-min**. We note that the performance of level-based heuristics degrades as the heuristics attempt to drive the search towards the inputs more aggressively. This observation provides partial evidence to the hypothesis that the size of transitive fanin of a gate, which provides an estimate of the amount of work needed to justify a sub-circuit rooted at the gate, is a more fundamental structural property in the context of CRSAT. However, in order to evaluate this hypothesis properly, we need to discover classes of problems where the measures **Level** and **TFI** are not correlated. Finally, due to the fact that on the instances in our benchmark set the two measures appear to be correlated, we note that since **Level** is a cheaper-to-compute measure, in practical applications one might want to consider using **Level-min**, rather than **TFI-min**, as a gate-selection heuristic.




Fig. 3. A comparison of the performance of **TFI-min** with various level-based gate-selection heuristics as a cactus plot, i.e. the number of those instances that can each *solved* within a given time limit. An instance is considered *solved* if a success rate over the 25 tries is $\geq 50\%$. The CPU time of a solved instance is the median CPU time for the instance over all runs (including the unsuccessful ones).

6 Conclusions

We presented results of experiments on the applicability of different circuit-level properties as the basis of structure-based search (gate selection) heuristics for the state-of-the-art SLS method CRSAT for industrial-related Boolean satisfiability instances. The results can be seen as first steps towards understanding the role of structural information in justification-based local search for SAT with limited Boolean propagation integrated into the search. We identified a number of easy-to-compute structural properties which appear suitable as the basis of heuristics for CRSAT, some of which can even outperform the recently introduced depth-based variant of CRSAT. The promise of the resulting heuristics was also corroborated by showing the dual properties result in extremely weakly performing heuristics.

The now presented results open up various interesting questions for further work on improving structure-based SLS for SAT. First, the observation that somewhat differently defined structural properties result in good heuristics suggests to study different ways of *combining* the resulting heuristics for achieving even better performance. This includes the question of what are the actual underlying properties to give good performance, and which the now studied easy-to-compute properties may be approximating. In addition to gate selection heuristics, we also aim to study different *objective functions* that are based on structural properties of SAT instances. Finally, we note that the development of good structure-based search heuristics for the circuit-level is directly applicable for the logical combinations of more high-level constraints (more generic CSPs), where the logical combinations can be viewed as circuits. This is one of the main research directions we are currently pursuing.

Acknowledgements We thank the anonymous referees for helpful comments.

References

1. Ågren, M., Flener, P., Pearson, J.: Generic incremental algorithms for local search. *Constraints* 12(3), 293–324 (2007)
2. Ågren, M., Flener, P., Pearson, J.: Revisiting constraint-directed search. *Inf. Comput.* 207(3), 438–457 (2009)
3. Belov, A.: Stochastic Local Search for Non-clausal and Circuit Satisfiability. Ph.D. thesis, York University, Toronto, Canada (2010)
4. Belov, A., Stachniak, Z.: Improving variable selection process in stochastic local search for propositional satisfiability. In: *Proc. SAT. LNCS*, vol. 5584, pp. 258–264. Springer (2009)
5. Belov, A., Stachniak, Z.: Improved local search for circuit satisfiability. In: *Proc. SAT. LNCS*, vol. 6175, pp. 293–299. Springer (2010)
6. Belov, A., Jarvisalo, M., Stachniak, Z.: Depth-driven circuit-level stochastic local search for SAT. In: Walsh, T. (ed.) *Proc. IJCAI*. pp. 504–509. AAAI Press (2011)
7. Brayton, R., Mishchenko, A.: ABC: An academic industrial-strength verification tool. In: *Proc. CAV. LNCS*, vol. 6174, pp. 24–40. Springer (2010)
8. Brummayer, R., Biere, A.: Boolector: An efficient SMT solver for bit-vectors and arrays. In: *Proc. TACAS. LNCS*, vol. 5505, pp. 174–177. Springer (2009)
9. Goldstein, L.: Controllability/observability analysis of digital circuits. *IEEE Transaction on Circuits and Systems* 26, 685–693 (1979)
10. Hentenryck, P.V., Michel, L.: Constraint-based local search. MIT Press (2005)
11. Hoos, H., Stützle, T.: *Stochastic Local Search: Foundations and Applications*. Elsevier / Morgan Kaufmann (2004)
12. Hoos, H.H.: On the run-time behaviour of stochastic local search algorithms for SAT. In: *Proc. AAAI*. pp. 661–666. AAAI Press (1999)
13. Jarvisalo, M., Junttila, T., Niemelä, I.: Justification-based local search with adaptive noise strategies. In: *Proc. LPAR. LNCS*, vol. 5330, pp. 31–46. Springer (2008)
14. Jarvisalo, M., Junttila, T., Niemelä, I.: Justification-based non-clausal local search for SAT. In: *Proc. ECAI*. pp. 535–539. IOS Press (2008)
15. Jarvisalo, M., Niemelä, I.: The effect of structural branching on the efficiency of clause learning SAT solving: An experimental study. *Journal of Algorithms* 63(1–3), 90–113 (2008)
16. Kautz, H., McAllester, D., Selman, B.: Exploiting variable dependency in local search. In: *IJCAI* (1997)
17. Muhammad, R., Stuckey, P.: A stochastic non-CNF SAT solver. In: *Proc. PRICAI. LNCS*, vol. 4099, pp. 120–129. Springer (2006)
18. Naveh, Y.: Guiding stochastic search by dynamic learning of the problem topography. In: *Proc. CPAIOR. LNCS*, vol. 5015, pp. 349–354. Springer (2008)
19. Pham, D., Thornton, J., Sattar, A.: Building structure into local search for SAT. In: *Proc. IJCAI*. pp. 2359–2364 (2007)
20. Sebastiani, R.: Applying GSAT to non-clausal formulas. *J. Artif. Intell. Res.* 1, 309–314 (1994)
21. Stachniak, Z.: Going non-clausal. In: *SAT (2002)*, <http://gauss.ececs.uc.edu/Conferences/SAT2002/Abstracts/stachniak.ps>
22. Stachniak, Z., Belov, A.: Speeding-up non-clausal local search for propositional satisfiability with clause learning. In: *Proc. SAT. LNCS*, vol. 4996, pp. 257–270. Springer (2008)
23. Velev, M.N., Bryant, R.E.: Superscalar processor verification using efficient reductions of the logic of equality with uninterpreted functions to propositional logic. In: *Proc. CHARME. LNCS*, vol. 1703, pp. 37–53. Springer (1999)