Solving Difficult Instances of Boolean Satisfiability in the Presence of Symmetry

Fadi A. Aloul, Arathi Ramani, Igor L. Markov and Karem A. Sakallah

CSE-TR-463-02

September 6, 2002

THE UNIVERSITY OF MICHIGAN
Computer Science and Engineering Division
Department of Electrical Engineering and Computer Science
Ann Arbor, Michigan 48109-2122
USA

Solving Difficult Instances of Boolean Satisfiability in the Presence of Symmetry

Fadi A. Aloul, Arathi Ramani, Igor L. Markov and Karem A. Sakallah

Advanced Computer Architecture Laboratory
Department of Electrical Engineering and Computer Science
University of Michigan
Ann Arbor, Michigan 48109-2122

September 6, 2002

ABSTRACT

Research in algorithms for Boolean satisfiability (SAT) and their implementations [45, 41, 10] has recently outpaced benchmarking efforts. Most of the classic DIMACS benchmarks [21] can now be solved in seconds on commodity PCs. More recent benchmarks [54] take longer to solve due of their large size, but are still solved in minutes. Yet, small and difficult SAT instances must exist if P≠NP. To this end, our work articulates SAT instances that are unusually difficult for their size, including satisfiable instances derived from Very Large Scale Integration (VLSI) routing problems. With an efficient implementation to solve the graph automorphism problem [39, 50, 51], we show that in structured SAT instances difficulty may be associated with large numbers of symmetries.

We point out that a previously published symmetry-detection mechanism [18] based on a reduction to the graph automorphism problem often produces many spurious symmetries. Our work contributes two new reductions to graph automorphism, which detect all correct symmetries detected previously [18] as well as phase-shift symmetries not detected earlier. The correctness of our reductions is rigorously proven, and they are evaluated empirically.

We also formulate an improved construction of symmetry-breaking clauses in terms of permutation cycles and propose to use only generators of symmetries in this process. These ideas are implemented in a fully automated flow that first detects symmetries in a given SAT instance, pre-processes it by adding symmetry-breaking clauses and then calls a state-of-the-art backtrack SAT solver. Significant speed-ups are shown on many benchmarks versus direct application of the solver.

In an attempt to further improve the practicality of our approach, we propose a scheme for fast "opportunistic" symmetry detection and also show that considerations of symmetry may lead to more efficient reductions to SAT in the VLSI routing domain.

1 Introduction

Boolean satisfiability (SAT) is a pivotal problem in Computer Science with numerous applications that range from microprocessor verification [54] to FPGA layout [42]. A one-million-dollar prize is offered by the Clay Institute for Mathematical Sciences for a complete polynomial-time SAT solver or a proof that such an algorithm does not exist (the P-vs-NP problem). Additionally, industrial applications motivate intensive research in SAT algorithms that quickly solve real-life instances. The fundamental framework for state-of-the-art SAT algorithms was laid out in the 1960s, but a number of recent improvements in algorithms and implementation techniques [45, 41] have led to performance breakthroughs. Most DIMACS challenge benchmarks [21] from the early 1990s are now solved in seconds on commodity PCs. Recently posted SAT benchmarks [54] take somewhat longer to solve (minutes), but that is primarily due to their enormous size (50MB+ files, etc). With the exception of artificially constructed families of benchmarks, it appears that SAT can be solved in polynomial time "for practical purposes". It is well known that the dominant backtrack solvers, such as GRASP [45], CHAFF [41] and BerkMin [10], do not perform well on randomly-created 3-SAT instances with ≈ 4.3 clauses per variable [47]. However, such instances are not common in practical applications because they have little structure. The relative ease of structured instances from certain applications was explained [43], and generic ways to exploit certain types of structure were proposed [2].

1.1 Difficult SAT benchmarks

Our work addresses both benchmarking and algorithmic aspects of SAT research. Given the excellent performance of existing SAT solvers, there is no room for improvement on easy benchmarks, and we focus instead on difficult instances. Since the works of Haken and Urquhart [52] on lower bounds for resolution and backtracking algorithms for SAT, several instance families have been known to require exponential time for DP/DLL (Davis-Putnam [19] and Davis-Logemann-Loveland [20]) solvers and their derivatives. For example, a recent lower bound for the pigeon-hole problem is $\Omega(2^{n/20})$ [7] where n is the number of pigeons. The pigeon-hole problem can be quickly solved by induction, but the proof system behind backtrack solvers (resolution) is rather restrictive and does not allow polynomial-sized proofs for pigeon-hole instances. Short proofs without induction exist if the use of symmetry is allowed [29, 53]. Another family of difficult instances was constructed by Tseitin and Urquhart in terms of expander graphs and, unlike the pigeon-hole instances, can accommodate considerable randomness [52]. As expected, solving those instances with modern SAT solvers, such as CHAFF and BerkMin, requires long time (see Tables 4 and 5), but their relevance to application domains (e.g., Electronic Design Automation (EDA), Software Verification and Artificial Intelligence) is not clear. While lower bounds for SAT are often proven for unsatisfiable instances, it remains to be seen whether practical satisfiable instances can be difficult for the best solvers. To this end, the work in [1] contributed constructions of artificial randomly generated difficult satisfiable instances.

Our work demonstrates EDA-related SAT instances, both satisfiable and unsatisfiable, that are very difficult for their size. Observe that an easy instance of any size can be made difficult by adding a small difficult instance to it and connecting the two by inconsequential clauses to defeat partitioning.

1.2 Relevance of graph automorphism to SAT

Over many years, empirical algorithms research in many domains identified a number of fundamental problem formulations, such as Boolean satisfiability, and mustered significant efforts to solve them efficiently. State of the art is gauged by optimized solver implementations ("engines"). Performance breakthroughs are often due to novel algorithmic ideas, leaner implementations or the ability to apply a highly-optimized engine in a novel way. In this work, we observe that graph automorphism engines can be applied to the satisfiability problem in certain cases. Additionally we think that there may be significant room for future improvement given that (i) the graph automorphism problem is thought to not be NP-complete, thus potentially easier than SAT, and (ii) much less new research was done in recent years on the analysis and design of high-performance engines for graph automorphism (such works include [40, 36]). To be precise, in this work we will be dealing with the colored variant of the graph automorphism problem that can be easily extended to hypergraphs.

Besides complexity-theoretic connections between variants of Boolean satisfiability, symmetries, and the hypergraph automorphism problem [4, 34], several pre-2000 works suggested that "breaking symmetries" in CNF formulae can speed up SAT solvers [8, 12, 13, 17, 18, 36]. Symmetries of a CNF formula include clause-preserving permutations of variables. Such permutations may involve arbitrarily many variables at once, e.g., a complete cyclic shift. In this work, we do not address permutations that change the CNF formula but leave unchanged the Boolean function it represents.[1] However, if such symmetries are detected by other techniques [27], our proposed methods can process them in the same way as symmetries of the CNF formula. Similarly, many of the works we cite do not deal with symmetry detection, but rather assume that symmetries of the Boolean function are given. Using this assumption, two main directions were explored: (a) preprocessing the original CNF formula by adding symmetry-breaking clauses that do not affect satisfiability but speed up search [18], and (b) extending SAT solvers, particularly those based on backtracking, to dynamically use symmetries during the search process [13, 6, 31, 44]. In this paper we pursue the pre-processing approach due to its simplicity, but will outline how our techniques can be applied within a backtracking solver for increased efficiency.

1.3 Empirical efficiency challenges

Most prior works on symmetries in SAT predate recent breakthroughs in SAT solvers and typically use several carefully constructed instances to illustrate their approach, or do not show convincing empirical results at all. For example, Crawford et al. suggest in [18] that symmetry-based techniques allow the pigeon-hole instances to be solved in polynomial time, but their empirical data [18, Figure 3] do not support this suggestion. In the course of more recent work [31, 49], specific families of CNF formulae with extremely high numbers of symmetries were successfully attacked. Yet, it remains unclear whether the performance of leading-edge SAT solvers can be improved, via the use of symmetries, on large CNF families of practical significance. In principle, the overhead due to symmetry detection and usage may outweigh the

[1] Such permutations can be called "semantic" symmetries, in contrast with the narrower class of "syntactic" symmetries that leave the CNF formula unchanged.

benefits, and it remains to be seen that useful CNF formulae have many symmetries. Pólya (1937), Erdös and Rényi (1963) proved that a random graph on n vertices has *no symmetries* with probability $1-\binom{n}{2}2^{-n-2}(1+o(1))$ [5, p. 1461]. This claim can be extended to CNF formulae, but structured real-world instances may have richer symmetries. Indeed, Boolean functions arising in the design of hardware systems often have many symmetries [27, 9], and the overall number of functions of n variables with non-trivial symmetries grows double-exponentially. On the other hand, for a function with exponentially many symmetries, trying to explicitly use all symmetries may defeat the purpose of speeding up search [18]. Despite these pitfalls, symmetry-based approaches have been useful in model checking [26, 15, 23], hardware verification [36], software verification [11], logic synthesis [28, 9] and DSP algorithms [22]. Some researchers limited the notion of symmetry to swaps of variables [22] or subsets of variables [28] to achieve efficiency. Other works [9, 44] limited the notion of symmetry to negations of single variables or subsets of variables and referred to those restricted classes as *autosymmetries* or *phase-shift symmetries*.

1.4 Our contributions

In this work, we study and fully automate a flow that starts with a CNF formula in the DIMACS format and finds all of its symmetries within a very general class, including all permutational symmetries, variable negations, and their compositions. In this flow, all symmetries are first captured implicitly, in terms of irredundant group generators, which always guarantees exponential compression. The CNF formula is then preprocessed by adding symmetry-breaking clauses that do not affect satisfiability. A black-box SAT solver is subsequently applied to the preprocessed CNF instance to produce the final answer; any satisfying assignment to this instance is (or corresponds to) a satisfying assignment of the original instance, and if the preprocessed instance is unsatisfiable then so is the original instance. The flow is illustrated in Figure 1.

We propose new symmetry-finding techniques and empirically compare them with previously proposed constructions.

We also propose a novel construction of symmetry-breaking clauses. It is much more economical than that in [18]. Additionally, it directly applies to the compressed representation of all symmetries in the format produced by graph automorphism software [38, 39, 50, 51].

Our empirical results show significant overall performance improvements on CNF instances arising in EDA applications, as well as on highly randomized provably-difficult Urquhart benchmarks [52] that are related to Tseitin formulae used to prove lower bounds on the size of resolution proofs. Two extensions are proposed to speed up symmetry finding. One is opportunistic symmetry finding, where only some symmetries are found. The other extension pursues domain-specific symmetries and leads to improvements of SAT formulations by adding domain-specific symmetry-breaking clauses. Thus, generic symmetry finding is avoided by creating symmetry-less SAT instances that can be solved quickly.

The remaining material is organized as follows. Symmetry finding is described in Section 2 and symmetry-breaking in Section 3. Section 4 discusses constructions of SAT benchmarks. Our empirical results are presented in Section 5 and further extensions in Section 6. Section 7 concludes our work and discusses our future directions.

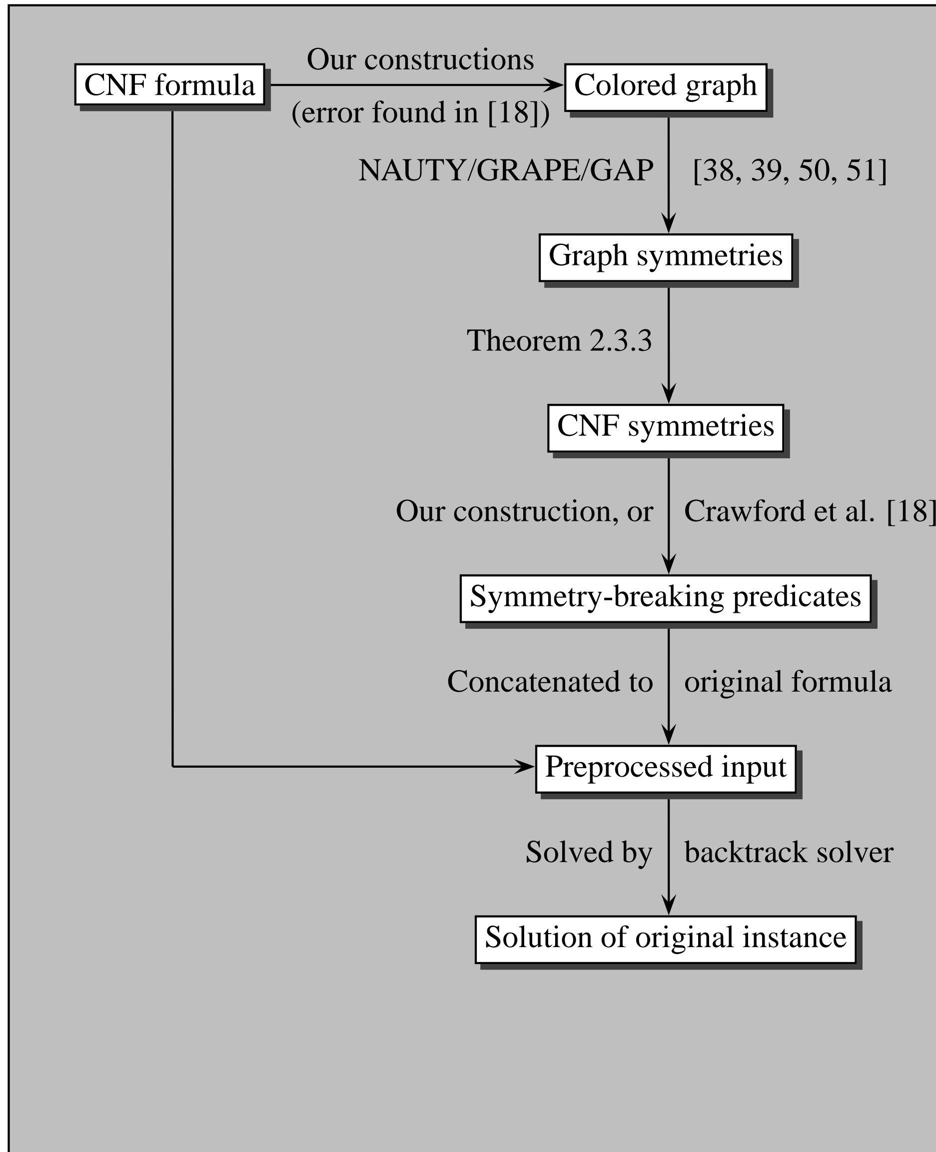

Figure 1: Preprocessing-based flow for symmetry-breaking studied in this work. Our construction of symmetry-breaking predicates improves upon that from [18].

2 Finding Symmetries

In general, a symmetry of a discrete object is a reversible transformation of its components that leaves the object unchanged. This can be taken as an informal definition, and more rigorous definitions will be given below for specific structures. Examples include permutations of graph vertices that map edges into edges, rotations of a spatial solid, e.g., a cylinder, that preserve its shape, as well as the negation of the variable a in the Boolean formula $(a+a')b$, since the formula and the function it represents are unaffected by this transformation. The discrete objects considered in our work have only finitely many symmetries. Unlike most previous works in the field, we consider, detect, represent and use several types of symmetries and their compositions, including permutational symmetries and variable negations in CNF formulae, sometimes called "phase changes" or "autosymmetries".

2.1 Representing and manipulating symmetries

Every discrete object has at least one symmetry — the "do-nothing" permutation. It is easy to see that composition of two symmetries is a symmetry, and that composition with the do-nothing permutation does not change a symmetry. The composition of symmetries is associative, and every symmetry has an inverse. However, the composition operation is often *not* commutative. An example is given by the six permutational symmetries of an equilateral triangle: (i) the do-nothing symmetry, (ii) three vertex swaps, and (iii) two cyclic rotations — counterclockwise and clockwise.

Definition 2.1.1 (from abstract algebra). *A* group *G is a set with a binary operation ("multiplication") defined on it that have the following three properties:*

- *the operation is associative, i.e.* $\forall a,b,c \in G \ (a \circ b) \circ c = a \circ (b \circ c)$;
- *there is a* unit *element* $e \in G$ *such that* $\forall a \in G \ a \circ e = e \circ a = a$;
- *for every* $a \in G$ *there is a unique* inverse $a^{-1} \in G$ *such that* $a \circ a^{-1} = a^{-1} \circ a = e$.

A subgroup *is a subset of a group that is closed under the group operation (and is therefore a group itself).*

For example, integers form a group with respect to the addition operation (0 is the unit element) and positive rationals form a group with respect to the multiplication operator ($1/1$ is the unit element). Group Theory [24] is a major branch of abstract algebra [25], and its development in the nineteenth century was motivated by groups of symmetries. Such diverse areas as the Galois theory describing solvability of polynomial equations, the periodic table of chemical elements, and Special Relativity involve analyses of groups of symmetries. In this work we will only deal with groups of symmetries whose elements can be thought of as permutations of finite sets. This obviously restricts us to finite groups. A permutation can be represented by cycles, e.g., $(23)(567)$ represents a permutation on a set of at least 7 marks (elements). This permutation swaps marks 2 and 3, cyclically permutes marks 5, 6 and 7 in that order, and leaves unchanged all other marks, e.g., 1 and 4.

Computational Group Theory (CGT), started ca 1911, is one of the oldest and most developed branches of computational algebra [48]. The flourishing of CGT began in the 1960s, and

great strides were made in the 1990s with the development of the GAP package ("Groups, Algebra and Programming") [51]. A major efficiency in computational Group Theory comes from the notion of *irredundant sets of generators* of a group.

Definition 2.1.2. *A* set of generators *consists of group elements such that any other group element can be composed of generators and their inverses. A generator is* redundant *if it can be expressed in terms of other generators. An* irredundant set of generators*, by definition, does not contain redundant generators.*

(Lagrange) Theorem 2.1.3 (from Elementary Group Theory) [24, 25]. *The size of any subgroup H of any finite group G must divide the size of G.*

Corollary 2.1.4. *For any group G with* $N > 1$ *elements, any irredundant set of generators contains* at most $\log_2 N$ *elements.*

Proof. Observe that any proper subgroup must be at least twice as small compared to the group. Given a set of n irredundant generators $x_1, \ldots, x_n$ consider a chain of subgroups G_k for $k = 1..n$, where G_k is generated by $x_1, \ldots, x_k$. By construction, G_k is a proper subgroup of G_{k+1}, and as such must be at least twice as small. Therefore, the size of $G_n = G$ must be at least 2^n. □

For example, the $k!$ permutations on k marks can be generated by (12) and $(12..k)$ or by $(12), (23), \ldots, (k-1\ k)$. Thus, representing groups by sets of generators *always ensures exponential compression*. Computational group theory provides efficient algorithms (due to Sims, Knuth, Babai and others) for manipulating groups represented by sets of generators, without decompression. Therefore, an intelligent algorithm for symmetry finding may return a small set of generators rather than list all symmetries.

Definition 2.1.5. *A mapping* $f : G_1 \to G_2$ *between two groups is a* homomorphism *iff for any* $a \in G_1$ *and* $b \in G_1$*, we have* $f(a \bullet b) = f(a) \circ f(b)$*, where* $\bullet$ *and* $\circ$ *are group operations in* G_1 *and* G_2 *respectively. A homomorphism for which an inverse mapping exists that is also a homomorphism, is called an* isomorphism*. If an isomorphism exists between* G_1 *and* G_2*, the two groups are called* isomorphic*. An isomorphism of a group with itself is called* automorphism *of that group and can be thought of as a symmetry of the group.*

Automorphisms can be composed, and form a group under this operation.

It is easy to see that if f is a homomorphism, then $f(a^{-1}) = f(a)^{-1}$. An isomorphism cannot map two different group elements to one. Additionally, the notion of isomorphism defines an equivalence relation and is useful to compare groups formally defined over different sets. In simple terms, isomorphic groups have "the same structure". Therefore, when looking for a group of symmetries of some objects, it may be convenient to find an isomorphic group instead. Since groups are often described by sets of generators, it is important to know that isomorphisms preserve such descriptions.

Theorem 2.1.6. *Any group isomorphism maps sets of generators to sets of generators, and maps irredundant sets of generators to irredundant sets of generators.*

Proof. If any element $h \in G_1$ can be written as a product of elements of a generating set or their inverses $h = g_1 \bullet g_2 \bullet \ldots \bullet g_n$, then a homomorphism $f : G_1 \to G_2$ will preserve such expressions in G_2: $f(h) = f(g_1) \circ f(g_2) \circ \ldots \circ f(g_n)$. Since every isomorphism has an inverse, any element $k \in G_2$ can be mapped back to G_1, where its pre-image can be decomposed into a product and then mapped back to G_2. This constructs a decomposition of k into a product of the images of elements of a generating set in G_1 and their inverses.

Now consider a pair of sets of generators that are mapped to each other by an isomorphism, they must have the same cardinality. Assume that one of them has a redundant element that can be expressed in terms of remaining elements. Since such an expression is preserved by an isomorphism, the image of this element must be redundant in the other set of generators. □

2.2 Colored automorphism problems

Definition 2.2.1. *Given two graphs, an* isomorphism *is a 1-to-1 mapping between the vertex sets of the two graphs that maps edges to edges. Given a graph, a* symmetry *(also called an* automorphism*) is a permutation of its vertices that maps edges to edges. In case of directed graphs, edge orientations must be preserved.*

Definition 2.2.2. *In the* graph automorphism *problem one seeks all symmetries of a given graph, e.g., in terms of group generators. The* decision version *of this problem tests for the presence of non-trivial automorphisms.*

It is known that all graphs, except for an exponentially small family, have *no symmetries* [5, p. 1461]. No general worst-case polynomial-time algorithms are known for this problem, but it is commonly believed not to be NP-complete (unless, of course, P=NP) [30]. Polynomial-time algorithms are available in many special cases [5, p. 1511], in particular for graphs of bounded degrees [33, 3]. Observe that many practical applications entail graphs of bounded degree because the objects involved (logic gates in VLSI chips, facts stored in knowledge bases, etc.) are interconnected sparsely. In contrast, Boolean Satisfiability instances of bounded degree, e.g., 3-SAT, are known to be NP-complete and 3-SAT instances may be quite difficult in practice even if every literal participates in only several clauses [47]. Generic algorithms for the graph automorphism problem [38] are based on linear-time partition refinement passes, followed by backtrack search. A simple version of partition refinement completes in three passes and does not require follow-up backtracking for all but an exponentially small family of graphs [5, p. 1513]. However, exponential worst-cases have been constructed even for very sophisticated versions [38], both theoretically and empirically [40].

The graph automorphism problem may be constrained by vertex labels — symmetries must map each vertex into a vertex with the same label. Label constraints are computationally easy and can be formally reduced to plain graph automorphism. Labels are often expressed by integers and called colors (no relation to *graph coloring*). Another extension is to colored *hypergraphs* — symmetries must map hyperedges to hyperedges (of the same cardinality because no two vertices can map to one). The colored hypergraph automorphism problem reduces to the colored graph automorphism via the bipartite graph of the hypergraph. This graph contains a vertex for each hypergraph vertex and hyperedge, and connects them with edges according to the hypergraph's incidence relation. Graph vertices in the hyperedge part are painted with a new color, and other vertices retain their original colors.

Brendan McKay implemented a practical algorithm for graph automorphism [38] in a software package called NAUTY [39], which has been continually improved for the last 20 years.[2] NAUTY has been integrated into the computational group theory system GAP [51] by means of the GRAPE package [50]. This integration enables efficient group-theoretic operations on the

[2] NAUTY version 2.0 was released in 2001.

results returned by NAUTY and facilitates some of our proposed algorithms. In 1998, Manku et al. [36] claimed speed-ups over a pre-2.0 version of NAUTY in the context of hardware verification. However, their code is not generic (built into a larger system) and is no longer supported. Finally, we observe that the runtime of existing graph automorphism programs, e.g., NAUTY, typically increases with growing numbers of vertices and symmetry generators found, but may decrease with growing numbers of vertex colors and, sometimes, graph edges.

2.3 CNF symmetries via graph automorphism

The problem of finding symmetries of a CNF formula is reduced to the colored graph automorphism problem. The main idea behind such reductions is to find a colored graph whose symmetry group is isomorphic to the symmetry group of the CNF formula. Related constructions are described in [17] and [18] for permutational symmetries, and we draw upon them in our work. We now consider a CNF formula with V variables and C clauses, of which C_2 are binary and C_x have two or more literals (clauses with fewer than two literals can be removed by preprocessing). In quotations, the word “theories” refers to CNF formulae. From [17, p.3]:

> Now consider reducing symmetry detection to graph isomorphism. We show the mapping for propositional theories (...). First note that we can “type” the nodes in the graphs, and only allow isomorphisms which preserve type (...), without increasing the difficulty of the isomorphism problem. We use five types of nodes: nodes for positive literals, nodes for negative literals, *inverse* nodes, nodes for clauses and *goal* nodes. We first link (the node for) each literal p to an inverse node and then link this inverse node to (the node for) $\neg p$. These links ensure that any graph isomorphism preserves negation. We then create a node for each clause and link it to the literals appearing in the clause. These links force graph isomorphisms to map clauses to clauses. Finally, recall that we are required to find a θ which maps p to q. To force this we create two copies of the graph for the theory. In the first we give p the type *goal* and in the second we give q the type *goal*. This typing forces any isomorphism between the two graphs to map p to q. One can then show that an isomorphism between the graphs exists iff the theory contains a simple symmetry mapping p to q.

The author then concludes that the *decision version* of the CNF symmetry detection problem is polynomial-time solvable if the length of the longest clause and the number of occurrences of the most common literal are bounded by a constant. That is because the degree of graph vertices is bounded by that constant, in which case the graph automorphism problem is poly-time solvable [33, 5]. If applied literally, the proposed construction only addresses symmetries that map p to q for particular p and q, rather than arbitrary symmetries. In order to find even a single non-trivial symmetry, one may need to traverse all pairs of variables. Thus, no isomorphism of symmetry groups is claimed in [17], and no empirical results are reported.

Additionally, we observe that for a formula with V variables and C clauses, this construction entails a graph with $6V+C$ vertices. Given that runtime of graph automorphism programs, e.g., NAUTY, grows super-linearly in terms of the number of vertices, more economical constructions (see below) can significantly reduce runtime.

Despite being impractical, the construction from [17] was apparently the first to introduce fundamental elements now used by more competitive constructions, including ours. We emphasize as particularly important

- the modeling of variables by pairs of positive-literal and negative-literal vertices,
- the modeling of each clause by a vertex connected to respective literal vertices by edges,
- connecting positive-literal and negative-literal vertices to enforce Boolean consistency.

Additional useful elements were introduced in [18, p.7]:

> The input theory is converted into a graph such that the automorphisms of the graph are exactly the symmetries of the theory. This is done using the construction in [Crawford, 1992]. There are three "colors" of vertices in this graph, the vertices representing positive literals, those representing negative literals, and those representing clauses. Graph automorphisms are constrained to always map nodes to other nodes of the same color. We also add edges from each literal to each clause that it appears in. These edges (together with the node colorings) guarantee that automorphisms of the graph are the symmetries of the theory. *Footnote 5*: For efficiency we special-case binary clauses by representing $x \wedge y$ with a link directly from x to y (instead of creating a node for the binary clause and linking x and y to it). This is important because some of the instances we consider have a huge number of binary clauses and some of the algorithms that follow are quadratic, or worse, in the number of nodes.

The reference [Crawford, 1992] in this quotation is the same as reference [17] in our paper, but the construction appears different from that cited above.[3] In fact, this formulation seems to inadvertently omit the enforcement of Boolean consistency, which leads to the generation of many spurious symmetries. For example, the formula $(a+b)$ has two symmetries: (i) the do-nothing symmetry, and (ii) the transposition (ab). The graph built by the above procedure has two positive-literal vertices, two negative-literal vertices and one clausal vertex connected to the positive-literal vertices by two edges. Since no negative literals are used, the respective vertices are disconnected and can be mapped to each other even if positive-literal vertices are fixed. There are four symmetries. One of them is the swap (transposition) of $\bar{a}$ and $\bar{b}$ with a and b fixed. It violates Boolean consistency. Notably, in [18] this construction is described in Section 7 on empirical results, next to a discussion of pigeon-hole and n-queens benchmarks. However, it produces spurious symmetries even when applied to pigeon-hole benchmarks, starting with `hole-2`.

On the positive side, this construction entails a graph with $2V + C_{\times}$ vertices — a marked improvement over [17]. We also found very useful in practice the idea to model each binary clause by one edge rather than by one vertex and two edges. Importantly, the proposed construction can be corrected by adding, for each variable, a vertex of color 4 and connecting it to the positive-literal and negative-literal vertices for the same variable (these nodes were called *inverse* nodes in [17]). We implemented this corrected version, and report empirical results for it. Similarly to [17], the reduction from [18] and its corrected version cannot find phase-shift symmetries.

In this work we propose several reductions of CNF symmetry-finding to graph automorphism, all of which allow finding phase-shift symmetries and their compositions with permutational symmetries. One of our constructions entails $2V + C_{\times}$ vertices and never finds

[3] Both papers [17] and [18] are downloadable from `http://citeseer.nj.nec.com/cs` and also from `http://www.cirl.uoregon.edu/crawford/papers/papers.html`.

spurious symmetries, but requires double edges that are not supported by the graph automorphism software NAUTY[39] used in our experiments. Another proposed construction entails $2V + C_\times + \min\{C_2, V\}$ vertices and never finds spurious symmetries. The third construction entails $2V + C_\times$ vertices, is implementable with NAUTY, produces no spurious symmetries on our benchmarks and allows a trivial check for spurious symmetries in general. Since this construction is often the fastest in practice, we characterize CNF formulae on which it produces spurious symmetries and show how spurious symmetries can be removed.

We first preprocess a given CNF formula to remove any clauses with fewer than two literals. If there is an empty clause, the formula is immediately declared unsatisfiable and the search for the symmetries of the formula becomes pointless (because every transformation of variables is a semantic symmetry as it does not affect the value of the function). If there are one-literal clauses, they can be eliminated in linear time by repeatedly (i) recording implied truth assignments (clause (a) implies $a = 1$, clause $(\bar{a})$ implies $a = 0$), (ii) eliminating the one-literal clauses, (iii) substituting the implied values of relevant variables, thus eliminating the variables, (iv) simplifying each affected clause independently. This process may prove that the original formula is satisfiable or unsatisfiable, or otherwise result in a smaller formula where every clause has at least two literals.

Given a CNF formula where every clause contains at least two literals, we represent every variable by two vertices that correspond to its positive and negative literals. We represent every non-binary clause by a single vertex, and connect that vertex, using bipartite edges, to the vertices representing literals in that clause. Binary clauses are represented by double edges connecting their respective literals. Clausal vertices are painted with color 1 and literal vertices are painted with color 2. Because vertices representing positive and negative literals in our graph are of the same color, we need to ensure Boolean consistency and mate vertices of opposite literals by single edges. Observe that no symmetry can map a single edge to a double edge, thus there is no risk of mapping a Boolean consistency edge to a binary-clause edge. This construction results in a graph with $2V + C_\times$ vertices. It corrects the reduction from [18] without increasing vertex counts and has the added advantage of detecting phase-shift symmetries (subsets of negated variables, e.g., $a \mapsto \bar{a}$) and their compositions with permutational symmetries. We refer to this construction as 2xEDGES.

Unfortunately, the graph automorphism program NAUTY [39] used in our experiments cannot represent double edges. Therefore we must seek another mechanism to distinguish Boolean consistency edges from binary-clause edges. A straightforward solution is to split every Boolean consistency edge into two edges by an added vertex of color 3 (one per edge). Alternatively, we can split binary-clause edges, which in some cases may be a better option. In fact, we can split the less numerous of the two types of edges, which yields $2V + C_\times + \min\{C_2, V\}$ vertices. Because three colors are used, this construction is referred to as MIN3C.

A far less obvious solution is *not to make explicit distinction* between the two types of edges, but represent both Boolean consistency and binary clauses by single edges. Since we first described this construction at the 2002 Design Automation Conference, we refer to it as DAC02. Figure 2 shows an example. In general, there are $2V + C_\times$ vertices, but the analysis of this construction is far more complex than that of the constructions described above. However, our efforts are justified by the often-superior empirical performance of this construction. Before

Clauses:
A $(\bar{1}+2+3)$
B $(1+\bar{2}+\bar{3})$
C $(\bar{2}+3)$

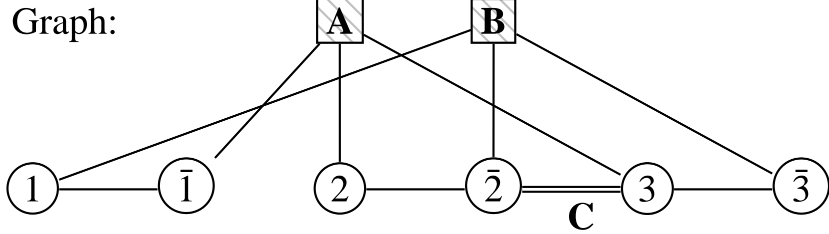

Symmetry:
$(1\bar{1})(2\bar{3})(\bar{2}3)$**(AB)**

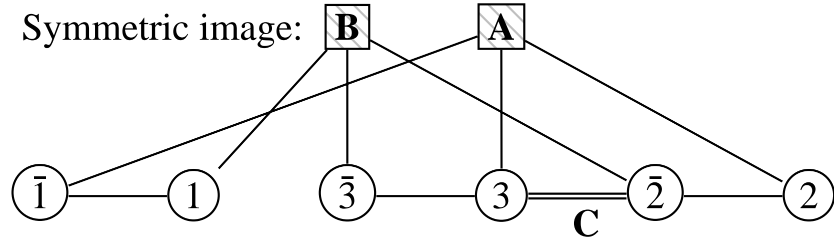

Figure 2: A CNF formula with three clauses — A, B and C, — and three variables is converted into a bi-colored graph for symmetry detection purposes. The two-literal clause C is represented by one edge (double-line) while larger clauses A and B are represented, each, by a vertex and three edges. Any symmetry must map C $\mapsto$ C, therefore this instance has only one non-trivial symmetry $(1\bar{1})(2\bar{3})(\bar{2}3)$(AB).

we proceed with formal results, let us articulate the correspondence between (i) variables and clauses in a given CNF formula, and (ii) vertices and edges of the bi-colored graph we build. Every variable corresponds to exactly two vertices of color 2. Every vertex of color 2 corresponds to a variable, and every vertex of color 1 corresponds to a clause. Every clause with more than two literals corresponds to a vertex of color 1, and every two-literal clause corresponds to an edge between two vertices of color 2. There are no edges connecting vertices of color 1, but every vertex of color 2 is connected to that of its complement literal by an edge, and there can be edges connecting pairs of vertices of different colors.

Definition 2.3.1. *A* circular chain of implications *over the variables* $x_1, x_2, \ldots, x_N$ *is a set of N binary clauses equivalent to* $(y_1 \Rightarrow y_2)(y_2 \Rightarrow y_3)\ldots(y_{N-1} \Rightarrow y_N)(y_N \Rightarrow y_1)$*, where for each k from* $1..N$*,* $y_k = x_k$ *or* $y_k = \bar{x}_k$*.*

Observe that the clause $(\bar{y}_k + y_{k+1})$ is equivalent to $(y_k \Rightarrow y_{k+1})$ and also to $(\bar{y}_{k+1} \Rightarrow \bar{y}_k)$. In terms of specific values, we have $(y_k = 1) \Rightarrow (y_{k+1} = 1)$ and $(y_{k+1} = 0) \Rightarrow (y_k = 0)$. For each k, one of the two possible values of y_k triggers an implication sequence, and thus unambiguously determines the values of all literals involved. In the remaining case, none of the variables assumes the value that triggers an implication in the circular chain. Therefore, a circular chain of implications allows only two satisfying solutions.

Theorem 2.3.2. *Assume that a given CNF formula does not contain a circular chain of implications over any subset of its variables. Then, with respect to the proposed construction of the colored graph from a CNF formula, the symmetries of the formula correspond one-to-one to the symmetries of the graph.*

The practicality of the assumption is discussed after Corollary 2.3.4 below.

Proof. It is not hard to see that every permutational symmetry of the initial formula (i.e., a permutation of variables that maps clauses to clauses) corresponds to a colored symmetry of

the bi-colored graph we built. Such a graph symmetry will map vertices to vertices of the same color and edges to edges. In particular, if a maps to b, then $\bar{a}$ maps to $\bar{b}$ and the edge $a\bar{a}$ maps to the edge $b\bar{b}$. Edges between vertices of color 2 will always map to edges between vertices of color 2, and the same can be said about edges between vertices of different colors. Phase-shift symmetries of the original formula also correspond to colored graph symmetries. For example $a \mapsto \bar{a}$ will induce a swap between the vertices a and $\bar{a}$, leaving the edge $a\bar{a}$ in place and swapping any existing edges ac and $\bar{a}c$ for a clausal vertex c. An immediate consequence is that every composition of permutational and phase-shift symmetries of the original formula correspond to a colored graph symmetry. For example, if a is symmetric to $\bar{b}$, then $a \mapsto \bar{b}$ and $\bar{a} \mapsto b$ so that the edge $a\bar{b}$ maps to $\bar{a}b$.

Our next observation is that given a colored graph symmetry that corresponds to some CNF symmetry, we can always uniquely reconstruct the CNF symmetry as long as the correspondence between variables and vertices of color 2 is available. This is also shown by first considering purely permutational symmetries, then phase shift symmetries and then their compositions. A graph symmetry that corresponds to a permutational CNF symmetry must map positive-literal vertices to other such. Therefore we can restrict the graph symmetry to this subset of vertices, thus producing a permutation of CNF variables. A graph symmetry that corresponds to a phase-shift CNF symmetry must either preserve a given literal vertex or map it to the complement-literal vertex, preserving the edge between them. Therefore, a list of positive-literal vertices that are not preserved uniquely identifies a phase-shift CNF symmetry. To reconstruct a CNF symmetry that is a composition of permutations and phase-shifts, we distinguish (i) positive-literal vertices that map to positive-literal vertices, from (ii) positive-literal vertices that map to negative-literal vertices. In each case, a given CNF variable is mapped to another variable, possibly with a follow-up negation. By ignoring the follow-up negations, we reconstruct the purely permutational component of the CNF symmetry. The phase-shift component, i.e., variables to be negated before the permutation is applied, can be reconstructed by listing positive-literal vertices that map to negative-literal vertices.[4]

Perhaps, the least trivial property of the proposed reduction to graph automorphism is that every colored symmetry of the graph corresponds to a symmetry of the original formula. To prove this, we show that the reconstruction procedure from the previous paragraph can be successfully applied to any colored graph symmetry. A vertex permutation is a colored symmetry if and only if (i) vertices are mapped to vertices of the same color, and (ii) edges are mapped to edges. This is consistent with CNF symmetries' mapping variables to variables and clauses with more than two literals to such clauses. However, it is more difficult to prove Boolean consistency, i.e., $\forall\, a, b\ (a \mapsto b) \Rightarrow (\bar{a} \mapsto \bar{b})$, where a and b are *literals*. This is easy in the absence of 2-literal clauses because all edges connecting vertices of color 2 are Boolean consistency edges of the form $\bar{a}a$. Since every such edge can only map to another such edge, $(a \mapsto b)$ leave no choice for $a\bar{a}$ but to map to $b\bar{b}$ because $b\bar{b}$ is the only edge that connects b to another vertex of color 2. This simple proof also applies if the two-literal clauses are represented by vertices, rather than by edges as in Figure 2.

[4]If one should perform negations *after* the permutation is applied, then listed should be the negative-literal vertices that map to positive-literal vertices.

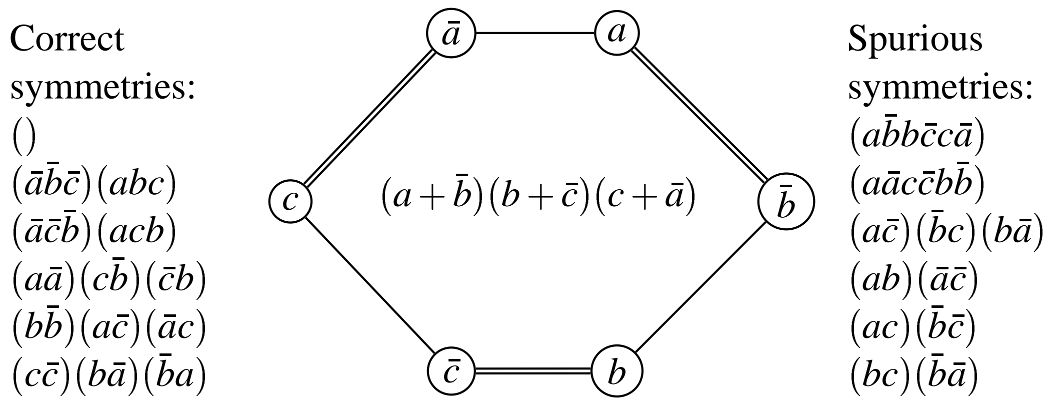

Figure 3: An illustration of spurious symmetries: a CNF formula and its graph. Boolean consistency edges are shown by double-lines, but are indistinguishable from other edges by graph automorphism software NAUTY which cannot handle double-edges. Therefore the graph has 12 symmetries: 6 rotations and 6 axial flips. Only 6 of them — 3 rotations and 3 flips, — preserve Boolean consistency edges and correspond to symmetries of the CNF formula. The remaining 6 symmetries are spurious (the first three spurious symmetries shown are rotations, and the remaining three are axial flips).

The difficulty in the general case is due to our modeling of two-literal clauses by edges that connect vertices of color 2. Such edges may potentially map to Boolean consistency edges, and our task is to prove that impossible. This is done in the Appendix. □

The 2xEDGE reduction circumvents this difficulty by connecting positive-literal vertices to negative-literal vertices with double-edges.

Theorem 2.3.3. *Under the assumption of Theorem 2.3.2, the symmetry groups of the CNF formula and the bi-colored graph are isomorphic.*

The proof consists of the straightforward verification that the one-to-one mapping constructed in the proof of Theorem 2.3.2 is a homomorphism.

Corollary 2.3.4. *Under the assumption of Theorem 2.3.2, sets of symmetry generators of the bi-colored graph correspond one-to-one to sets of symmetry generators of the CNF formula.*

To evaluate the practicality of the assumption in Theorems 2.3.2 and 2.3.3. observe that the failure of this assumption implies that in every satisfying assignment, the variables involved in the circular chain of implications can have one of two different sets of values (models). This can be illustrated by the CNF formula $(a+\bar{b})(b+\bar{c})(c+\bar{a})$, which allows only two models (000 and 111) but has six symmetries (do-nothing, two three-cycles and three variable swaps combined with negation of all variables). However, the graph produced by our construction is a hexagon having 12 symmetries (the so-called *dihedral group* D_6 [24, 25]). Half of those are spurious as explained in Figure 3.

From the practical standpoint, we note that

- Circular chains of implications do not arise in standard SAT models from many application domains. For example, they do not appear in equivalence checking of combinational circuits because combinational circuits are directed acyclic graphs.

Reduction type	#Colors	#Vertices	Detects phase-shifts?	Finds spurious symmetries ?	Practical with NAUTY[39]?
[17]	5	$6V+C$	No	No	No
[18]	3	$2V+C_\times$	No	Many+often	No
2xEDGES	2	$2V+C_\times$	Yes	No	No
MIN3C	3	$2V+C_\times$ $+\min\{C_2,V\}$	Yes	No	Yes
DAC02	2	$2V+C_\times$	Yes	In rare cases + trivial check $\exists$	Yes

Table 1: Comparing reductions of CNF symmetry detection to graph automorphism. V is the number of variables in the original CNF instance, C is the number of clauses, C_2 is the number of binary clauses, $C_\times = C - C_2$. The 2xEDGES reduction is not practical with NAUTY because NAUTY does not support double edges in graphs. CNF instances for which the DAC02 reduction finds spurious symmetries are characterized in Theorem 2.3.2.

- The presence of circular chains of implications does not invalidate our construction. As can be seen from the proof of Theorem 2.3.2, the only potential problem is spurious graph symmetries that do not correspond to any CNF symmetries. Since any application of Theorem 2.3.2 must convert symmetry generators returned by a graph automorphism problem into CNF symmetries, any spurious symmetry generators can be identified with minimal computational effort and minimal programming overhead.

- If some, but not all, symmetry generators are spurious, the non-spurious generators are still useful for symmetry-breaking, while spurious generators can be discarded (but that is not necessarily the best approach).

- Since the product of non-spurious symmetries cannot be spurious, there can be no spurious symmetries at all if none of symmetry generators are spurious. In other words, if spurious symmetries exist, at least one generator must be spurious.

- Once a spurious symmetry generator is found, a circular chain of implications can be identified in linear time along the lines of analysis in the proof of Theorem 2.3.2 in the Appendix. Since every circular chain of implications implies two sets of values for variables involved, circular chains of implications can be *removed* by introducing one Boolean variable to represent the two sets of values (old variables get eliminated).

- In applications where many spurious symmetries are expected and can slow down symmetry detection, circular chains of implications can be identified in linear time *before symmetry detection*, using Depth-First Search on a directed graph of binary clauses.

While the correctness of representing binary clauses with edges (Theorem 2.3.2) appears much harder to prove compared to the correctness of graph reductions proposed earlier, our construction reduces the number of vertices in the graph by the number of binary clauses in the CNF instance. Application-derived CNF instances typically have a significant proportion

of binary clauses, and our construction DAC02 leads to non-trivial run time savings in practice. Table 1 summarizes the main properties of various reductions of CNF symmetry finding to graph automorphism. Additionally, we empirically compare MIN3C, DAC02, the reduction from [18] and a corrected version of that reduction. In the corrected version, to ensure Boolean consistency, we add one extra node of color 4 for each variable and two edges connecting that node to the positive and negative literals of that variable.

Our testbed includes five sets of difficult benchmarks with non-trivial symmetries.

1. The `hole-n` benchmark set, available within the DIMACS collection [21].
2. Randomized benchmarks `Urq` proposed by Urquhart[52], based on parity checks and expander graphs.
3. Randomized benchmarks `grout` derived in this work in the context of global grid-based routing for VLSI.
4. Benchmarks `fpga` derived in this work in the context of detailed routing for Field-Programmable Gate Arrays.
5. Recent benchmarks from the micro-processor verification domain [54].

Descriptions of all benchmark sets except for the `Urq` and microprocessor verification sets are given in Section 4. Our implementation of symmetry-finding uses the program NAUTY [39] version 2.0, shipped with the GAP package [51] version 4 release 3. Table 2 compares sizes of graphs produced by four constructions. We make the following observations:

- Because all of our benchmarks contain more binary clauses than variables, MIN3C generates exactly as many vertices and edges as the corrected version of the reduction from [18]. However, MIN3C entails one color less and detects phase-shift symmetries.
- Graphs produced by MIN3C always have more vertices than those produced by DAC02.
- DAC02 and [18] produce graphs with the same numbers of vertices, but DAC02 generates more edges because it ensures Boolean consistency.

Table 3 compares symmetry-finding runtime and rounded number of symmetries (sizes of symmetry groups) discovered with each reduction. All runtimes are recorded on a Linux workstation with 1.2GHz AMD Athlon and 1Gb or DDRAM.

Several entries of the table with sizes of symmetry groups can be verified independently. For example, the number of symmetries in `hole-n` benchmarks is $n!(n+1)!$ because the symmetry group is the Cartesian product of S_n (holes can be permuted arbitrarily) and S_{n+1} (pigeons can be permuted arbitrarily). For $n = 7$ this yields 203212800, which rounds off to $2.03e8$. Furthermore, we make the following observations:

- Except for the second (`Urq`) and the last (microprocessor verification) benchmark sets, the reduction from [18] entails more symmetries than other reductions. This is because Boolean consistency is not enforced by that reduction, and spurious symmetries are found. `Urq` benchmarks do not have permutational symmetries, as checked by the corrected version of [18]. The reduction from [18] cannot detect phase-shift symmetries.

Instance	vari-ables	clau-ses	Previous work				Our work			
			[18]		[18] corrected		MIN3C		DAC02	
			#vert	#edges	#vert	#edges	#vert	#edges	#vert	#edges
hole07	56	204	120	252	164	364	164	364	120	308
hole08	72	297	153	360	225	504	225	504	153	431
hole09	90	415	190	465	280	675	280	675	190	585
hole10	110	561	231	660	341	880	341	880	231	770
hole11	132	738	276	858	408	1122	408	1122	276	990
hole12	156	949	325	1092	481	1404	481	1404	325	1248
Urq3_5	46	470	560	2910	606	3002	606	3002	560	2956
Urq4_5	74	674	840	4270	914	4418	914	4418	840	4434
Urq5_5	121	1210	1450	7546	1571	7788	1571	7788	1450	7667
Urq6_5	180	1756	2122	10772	2292	11132	2292	11132	2112	10952
Urq7_5	240	2194	2672	13194	2912	13674	2912	13674	2672	13434
grout3.3-01	864	7592	2306	9024	3170	10752	3170	10752	2306	9888
grout3.3-03	960	9156	2558	10740	3518	12660	3518	12660	2558	11700
grout3.3-04	912	8356	2432	9864	3344	11688	3344	11688	2432	10776
grout3.3-08	912	8356	2432	9864	3344	11688	3344	11688	2432	10776
grout3.3-10	1056	10.8k	2796	12564	3852	14676	3852	14676	2796	13620
fpga10.08	120	448	328	840	448	1080	448	1080	328	960
fpga10.09	135	549	369	1035	503	1305	504	1305	369	1170
fpga12.11	198	968	539	1815	737	2211	737	2211	539	2013
fpga12.12	216	1128	588	2124	804	2556	804	2556	588	2340
fpga12.08	144	560	392	1032	536	1320	536	1320	392	1176
fpga12.09	162	684	441	1269	603	1593	603	1593	441	1431
fpga13.09	176	759	478	1398	654	1750	654	1750	478	1574
fpga13.10	195	905	530	1675	725	2065	725	2065	530	1870
fpga13.12	234	1242	636	2322	870	2790	870	2790	636	2556
2pipe_1.ooo	834	7026	3851	14925	4685	16593	4685	16593	3851	15759
2pipe_2.ooo	925	8212	4133	17231	5058	19081	5058	19081	4133	18156
2pipe	861	6695	2621	12841	3482	14563	3482	14563	2621	13702
3pipe	2392	27533	7428	53620	9820	58404	9820	58404	7428	56012

Table 2: Comparison of reductions in terms of sizes of graphs produced.

- Except for the second and the last benchmark sets, the reductions MIN3C, DAC02 and corrected [18] find the same numbers of symmetries. In particular, those three reductions produce correct numbers of symmetries for `hole-n` instances. This is consistent with the reduction from [18] being erroneous, as it discovers many spurious symmetries.
- Except for the second (`Urq`) benchmark set, the runtimes of MIN3C and corrected [18] are comparable. This is expected because they generate equal numbers of vertices and edges, differing only in the number of colors. The runtimes for `Urq` benchmarks are different because MIN3C leads to the discovery of more symmetries.

Instance	vari-ables	clau-ses	Previous work				Our work			
			[18]		[18] corrected		MIN3C		DAC02	
			time	#symm	time	#symm	time	#symm	time	#symm
hole07	56	204	0.25	1.47e70	0.0	2.03e8	0 .0	2.03e8	0.0	2.03e8
hole08	72	297	0.35	1.24e77	0.47	1460	0.1	1460	0.0	1460
hole09	90	415	0.73	5.69e77	0.23	1.32e12	0.0	1.32e12	0.05	1.32e12
hole10	110	561	2.0	4.06e77	0.12	1.45e14	0.15	1.45e14	0.08	1.45e14
hole11	132	738	3.38	1.53e78	0.32	1.91e16	0.19	1.91e16	0.12	1.91e16
hole12	156	949	6.66	1.61e78	0.24	2.98e18	0.39	2.98e18	0.13	2.98e18
Urq3_5	46	470	0.05	1	0.21	1	0.51	5.37e8	0.39	5.37e8
Urq4_5	74	674	0.0	1	0.15	1	1.56	8.8e12	1.6	8.8e12
Urq5_5	121	1210	0.12	1	0.15	1	14.16	4.72e21	13.73	4.72e21
Urq6_5	180	1756	0.53	1	1.2	1	70.29	6.49e32	63.37	6.49e32
Urq7_5	240	2194	1.06	1	1.62	1	189.0	1.12e43	175.99	1.12e43
grout3.3-01	864	7592	109.1	8.28e77	16.36	8.71e9	15.44	8.71e9	4.86	8.71e9
grout3.3-03	960	9156	173.1	4.67e77	32.2	6.97e10	28.02	6.97e10	9.07	6.97e10
grout3.3-04	912	8356	143.5	1.10e78	21.85	2.61e10	19.65	2.61e10	7.01	2.61e10
grout3.3-08	912	8356	143.6	1.80e78	26.04	3.48e10	22.23	3.48e10	7.09	3.48e10
grout3.3-10	1056	10.8k	263.6	1.03e78	42.54	3.48e10	33.03	3.48e10	10.73	3.48e10
fpga10.08	120	448	2.24	5.49e77	0	6.69e11	0.21	6.69e11	0.18	6.69e11
fpga10.09	135	549	3.92	1.44e78	0.32	1.50e13	0.35	1.50e13	0.07	1.50e13
fpga12.11	198	968	15.05	1.99e78	1.17	1.79e18	1.0	1.79e18	0.47	1.79e18
fpga12.12	216	1128	24.2	2.81e78	1.73	2.57e20	1.68	2.57e20	0.64	2.57e20
fpga12.08	144	560	4.96	1.01e78	0.52	2.41e13	0.51	2.41e13	0.18	2.41e13
fpga12.09	162	684	7.38	8.18e77	0.55	5.42e14	0.53	5.42e14	0.28	5.42e14
fpga13.09	176	759	10.16	8.50e77	0.76	3.79e15	0.66	3.79e15	0.25	3.79e15
fpga13.10	195	905	16.64	2.13e78	1.15	1.90e17	0.94	1.90e17	0.44	1.90e17
fpga13.12	234	1242	31.81	9.29e77	1.89	9.01e20	1.8	9.01e20	0.73	9.01e20
2pipe_1.ooo	834	7026	9.61	2	13.19	2	15.95	8	9.14	8
2pipe_2.ooo	925	8212	12.26	2	21.59	2	20.17	32	11.15	32
2pipe	861	6695	3.19	32	7.6	8	7.28	128	3.21	128
3pipe	2392	27.5k	72.09	32	165.75	8	163.2	512	70.95	512

Table 3: Comparison of reductions in terms of symmetry-finding runtime and rounded numbers of discovered symmetries. Runtimes are in seconds on a 1.2GHz AMD Athlon running Linux.

- DAC02 is generally the fastest reduction. No other reduction generates fewer vertices, and DAC02 does not discover any spurious symmetries on given benchmarks as its results always agree with MIN3C.

In addition, we explicitly verified that the symmetries discovered by MIN3C and DAC02, but not by the two versions of [18], are phase-shift symmetries and their compositions with permutational symmetries. An implementation of the DAC02 reduction is available in our software package Šatter that targets symmetry-finding and symmetry-breaking for SAT. This package can be downloaded from `http://gigascale.org/bookshelf/Slots/shatter/`.

3 Symmetry-breaking

Symmetries induce equivalence classes on the set of truth assignments (in group theory, they are called *orbits*). Specifically, given a satisfying truth assignment, all other truth assignments to which it can be mapped by symmetries, must also be satisfying. Similarly, symmetries always map unsatisfying assignments to unsatisfying assignments. Therefore, for a complete SAT solver it suffices to reason about one representative from each such class. This restriction can be implemented by selecting unique representatives from every equivalence class and adding clauses that are only satisfied by those representatives. An earlier construction of such symmetry-breaking clauses [18] is based on a given ordering of variables. Its main idea is (i) to order all elements from the solution space lexicographically, and (ii) to select the lexicographically smallest element from each equivalence class as its representative.

3.1 Previous work

The lex-leader symmetry-breaking predicates described by Crawford et al. in [18] are built for a given group of permutational symmetries. Such predicates are conjunctions of smaller predicates for individual symmetries. Below, let n be the number of variables and $LL(G)$ be the lex-leader symmetry-breaking predicate for the group G. Boolean variables x_k are traversed according to the original ordering.

$$LL(G) = \bigwedge_{\pi \in G} LL(\pi) \tag{1}$$

$$LL(\pi) = \bigwedge_{1 \le i \le n} C(\pi, i) \tag{2}$$

$$C(\pi, i) = [\bigwedge_{1 \le j < i} (x_j = x_j^{\pi})] \Rightarrow (x_i \le x_i^{\pi}) \tag{3}$$

Theorem 3.1.1 [18]. *For a group G acting on truth assignments, the truth assignments that satisfy $LL(G)$ are the lexicographically smallest representatives from each class of truth assignments that can be mapped to each other by symmetries from G.*

Each $C(\pi, i)$ is then expressed in the CNF form using $i-1$ auxiliary variables $e_j = (x_j = x_j^{\pi})$:

$$C(\pi, i) = (e_1 e_2 \dots e_{i-1} \Rightarrow (x_j \le x_j^{\pi})) = (\bar{e}_1 + \bar{e}_2 + \dots + \bar{e}_{i-1} + \bar{x}_i + x_i^{\pi}) \tag{4}$$

Due to clauses of growing size, CNF expressions for each $LL(\pi)$ have $\Theta(n^2)$ literals, which may be prohibitively expensive even for one permutation π with, say, 9,000 variables (see Table 4). Additionally, $LL(\pi)$ for different π may contain redundant clauses. To prune redundant clauses, the authors propose the concept of a symmetry tree, but it does not always prevent redundant clauses and is itself not always prunable to polynomial size.[18].[5]

[5] In the special case of the symmetry group S_n, according to [18], the symmetry-breaking predicate produced using a symmetry tree has size $\Theta(n^2)$. Techniques proposed in our work generate a linear-sized predicate.

The need for more efficient, and also partial symmetry-breaking has been understood for some time [18, 37, 35], but no satisfactory generic approaches have been proposed that can be fully automated. In a recent work [35] Luks and Roy show that, even for an Abelian (commutative) symmetry group and a given ordering of variables, full lex-leader symmetry-breaking predicates can be exponentially large. This drawback can be circumvented by reordering variables, which facilitates polynomial-sized full lex-leader symmetry-breaking predicates for Abelian symmetry groups. However, the construction in [35] is not practical and is rather used for an existence proof. Also, it does not address non-Abelian groups.

3.2 Using symmetry generators

In this work we explore partial symmetry-breaking, i.e., we do not require that symmetry-breaking predicates be satisfied by lex-leaders only (but we still require that all lex-leaders satisfy symmetry-breaking predicates). As other authors, we compute symmetry-breaking clauses on a per-symmetry basis, but consider only irredundant sets of symmetry generators (returned by graph automorphism programs) instead of the entire symmetry group G. By breaking generator symmetries only, one does not necessarily break all symmetries. However, we believe that by breaking only generator symmetries, one can often achieve significant pruning because an irredundant set of generators contains "maximally independent" symmetries — none of them can be expressed in terms of others. The following example suggested to us by Eugene Goldberg of Cadence Berkeley Labs demonstrates that symmetry-breaking by generators is not complete in some cases.

Consider a formula with four Boolean variables x_1, x_2, x_3 and x_4 that can be permuted arbitrarily, e.g., $(x_1+x_2+x_3+x_4)$. The symmetry group, S_4, can be given by the two generators: $g_1 = (12)$ and $g_2 = (1234)$. Let us assume that in each equivalence class of truth assignments under those symmetries we select the lexicographically smallest element with respect to the original order of variables, i.e., x_1 is the most significant bit. Note that Boolean cube is split into 5 equivalence classes by the action of S_4 because the number of 1's in truth assignments is invariant under permutational symmetries. In particular, the equivalence class of the truth assignment 0101 has six elements, and the smallest element is 0011. However, if we build symmetry-breaking predicates using g_1 and g_2 only, 0101 will satisfy them because $g_1(0101) = 1001 > 0101$ and $g_2(0101) = 1010 > 0101$. Thus, such symmetry-breaking predicates select more than one representative from some equivalence classes. Moreover, additionally conjoining symmetry-breaking predicates for *powers* of generators does not help in this case because $g_1^2 = ()$, $g_2^4 = ()$, $g_2^2(0101) = 0101$ and $g_2^3(0101) = 1010 > 0101$.

Interestingly, for the symmetry group S_4, GAP/GRAPE/NAUTY do not produce the two generators used in the above example. They produce the following set of three generators: (12), (23) and (34). Our construction proposed below generates the symmetry-breaking clauses $(x_1 \leq x_2)$, $(x_2 \leq x_3)$ and $(x_3 \leq x_4)$, which admit only five truth assignments: 0000, 0001, 0011, 0111 and 1111 — one from each equivalence class under S_4. This analysis shows that the particular choice of irredundant generating sets is important for symmetry-breaking. In our experience, GAP/GRAPE/NAUTY often produce "lucky" sets of generators that lead to fuller

symmetry-breaking. Our future research will attempt to explain why that is happening.

As shown by our experiments in Section 5 below, symmetry-breaking by generators offers an attractive trade-off between effective pruning and small overhead. However, we would like to articulate an important pitfall in this direction. Firstly, adding symmetry-breaking predicates should not change the satisfiability of the original CNF instance. In our and previous work this is ensured by the fact that symmetry-breaking predicates are satisfied by at least one truth assignment from each class of symmetric truth assignments. The lex-leader predicates described above are satisfied by lexicographically smallest truth assignments because all $LL(\pi)$ are. The pitfall lies in the possibility to conjoin symmetry-breaking predicates that are satisfied by non-lex-leader representatives of classes of symmetric truth assignments. A conjunction of such predicates may be unsatisfiable and thus unusable as a symmetry-breaking predicate. Therefore, in this work, we adhere to lex-leader predicates.

3.3 Using cycles of permutations

Our construction is formulated in terms of cycles of a permutation, which is convenient because the output of graph automorphism programs is expressed in cycle notation. We observe that in overwhelmingly many instances all generators have two-cycles only. Even in rare cases when three-cycles were present, two-cycles dominate by far. Another important observation about the output of graph automorphism programs is that collections of two-cycles returned on the output are sorted according to the given variable ordering. Therefore, we can apply the Crawford construction in Equations 2 and 3 to individual cycles and further optimize it for two-cycles. In particular, for the variable swap (ab) the construction in [18] entails one additional variable and six symmetry-breaking clauses. Our construction below entails only one clause.

Single cycles. First observe that if the cycle (ab) is a symmetry, whenever there is a satisfying assignment with $a = 0, b = 1$, there should be a symmetric (equivalent) satisfying assignment with $a = 1, b = 0$ and other variables unchanged. To allow only the first assignment, we add the symmetry-breaking clause $(\bar{a} + b)$, which can also be interpreted as $(a \leq b)$. Similarly, to “break” a cycle of length three (abc), we add $(\bar{a} + b)(\bar{b} + c)$, i.e., $(a \leq b)(b \leq c)$. To make sure that lexicographically smallest representatives of symmetric truth assignments satisfy our predicates, one has to choose an ordering of all variables at the beginning, and always use the $\leq$ sign consistently with that ordering. When $a = \bar{b}$, we get the cycle $(\bar{a}a)$, and it can be broken in two ways. In terms of the original CNF instance, the value of a can be fixed arbitrarily, and this can be expressed by a single one-literal symmetry-breaking clause: (a) or $(\bar{a})$. The construction in [18] does not address such phase-shift symmetries and never results in one-literal clauses.

In general, longer cycles require more complex symmetry-breaking clauses, but apparently one can always improve on the construction from [18]. A particular difficulty with cycles of length > 3 is that they cannot, in general, be ordered according to a given ordering of variables. For example, the cycle (1324) can be written as (3241), (2413) or (4132), but none of these representations are ordered. Therefore we are not considering longer cycles in this work (and they do not appear useful for symmetry-breaking on our benchmarks).

Multiple cycles. While single-literal symmetry-breaking clauses are most efficient (they

reduce the solution space by 50%), they are associated with variables whose values do not affect satisfiability. After such variables are found and eliminated, other symmetries may remain. Indeed, we can produce symmetry-breaking clauses from any one two-cycle or three-cycle of any symmetry. However, clauses of the form $(\bar{a}+b)$ achieve no pruning in areas of the solution space where the variables involved have identical values. A key idea in that case, similar to that in [18], is to process another cycle, but only if $a=b$. In fact, this is similar to Equation 3, except that we now operate on cycles and do not need to involve *all* variables, which can dramatically reduce the size of symmetry-breaking clauses. Specifically, when building a symmetry-breaking predicate for the symmetry $(ab)(cd)(ef)...$, we first add $(\bar{a}+b)$, then $(a=b)\Rightarrow(c\leq d)$, then $((a=b)(c=d))\Rightarrow(e\leq f)$, etc. In the spirit of Equation 4, we introduce one additional variable per cycle to indicate the equality of all variables in the cycle. A sample clause with new variables looks like $(\bar{x}_{a=b}+\bar{x}_{c=d}+\bar{e}+f)$. This construction is given only for permutations with two-cycles and three-cycles.

It can be seen that both Equation 3 and our construction entail a lexicographic comparison between the tested truth assignment and its symmetric image: the former operates on bits and the latter operates on cycles. In practice, this often leads to very large reductions in the number of generated clauses. As a result of the bitwise comparison, lexicographically smallest truth assignments are identified if single-bit comparisons are performed according to the global ordering of variables. However, in the context of cyclewise comparison, the situation is more complex. We only assert that a lexicographic comparison is performed when (i) each cycle is a two-cycle, (ii) each cycle is ordered according to the global ordering of variables, and (iii) cycles are ordered lexicographically (which is equivalent to ordering them by the first element since they must be disjoint). Any chain of two-cycles can be brought to this form by sorting.

Theorem 3.3.1. *Consider an arbitrary single permutation consisting of two-cycles only. Apply the proposed construction of symmetry-breaking predicates, including the sorting of cycles and elements within each cycle. All resulting CNF clauses are satisfied by lexicographically smallest representatives of classes of truth assignments that are symmetric under the given permutation. No other truth assignments satisfy all of those clauses.*

Proof. We begin by noting that variables not involved in any cycles can be skipped during a lexicographic comparison of a truth assignment to its image under the given permutation. Our construction, indeed, skips those variables. The remaining part of the proof employs induction on the number n_c of cycles. In the base case $n_c=0$, the lexicographic comparison always returns true, and no clauses are generated. For an added cycle (ab) where a precedes b, we note that the clause $(\bar{a}+b)$, also known as $(a\leq b)$, lexicographically compares the partial assignments ab and ba. In other words, the test $(\bar{a}+b)$ checks that the value of a in the current truth assignment be $\leq$ to the value of a in the symmetric assignment. If those values are different, the overall comparison is finished. Otherwise, the comparison shifts its focus to the least variable unseen before (which may be ordered before or after b) and its image under the permutation. This corresponds to considering the next two-cycle. We would like to articulate that our construction does not require variables in two-cycles to be pairwise-adjacent in the variable ordering.

Since the square of the permutation is the identity, the classes of symmetric truth assign-

ments consist of one or two elements only. The clauses we consider are satisfied by a given assignment if and only if (by construction) the image of this assignment is not lexicographically smaller than the assignment itself. Therefore, all clauses are satisfied by (i) one-element classes, and (ii) the smaller elements of all two-element classes. □

In our experiments, most generators returned by graph automorphism software consist of two-cycles only. For rare benchmarks, some generators have small numbers of cycles of other lengths, typically three-cycles. It turns out that three-cycles can be ignored without violating the correctness of the symmetry-breaking procedure.

Theorem 3.3.2. *Consider a single permutation having (i) cycles of length two, (ii) cycles of odd lengths, and no other cycles. If the proposed construction of symmetry-breaking clauses is applied to two-cycles only, the resulting clauses must be satisfied by all lex-leader truth assignments, and potentially other truth assignments.*

Proof. Consider the product (or the least common multiple) p of all odd cycle lengths. The p-th power of the given permutation has the same two-cycles, but no other cycles. Since it is also a symmetry, Theorem 3.3.1 applies. Moreover, any lex-leader truth assignment with respect to the original permutation (i.e., cannot be improved by applying the permutation or its powers) is also a lex-leader with respect to the p-th power. □

3.4 Further improvements

In practice, the process of constructing symmetry-breaking clauses is often dwarfed by the symmetry-detection time. However, with every cycle processed, we add larger and larger symmetry-breaking clauses. Since large clauses that do not affect satisfiability rarely improve run time of SAT solvers, we optionally limit symmetry-breaking clauses to the first 10 cycles of every symmetry. For the price of incomplete symmetry-finding, this technique considerably reduces the overhead of symmetry-breaking clauses. Based on Theorem 3.3.1, we make the following observation.

Observation 3.4.1. *Consider a variant of the proposed construction of symmetry-breaking predicates for permutations with two-cycles only. After cycles are sorted, only the first k cycles are considered and the remaining cycles ignored. The clauses produced by this reduced construction are all satisfied by lex-leader truth assignments, but other truth assignments may satisfy those clauses.*

While the reduced construction does not achieve as much pruning as the full construction involving all cycles, its overhead is smaller. In our experiments the reduced construction often performed better.

Another potential improvement is to modify the source code of a backtrack SAT solver to dynamically check conditions of the form $((a = b)(c = d)...(u = v))$. We do not pursue this option in the current work because our discussion is limited to pre-processing.

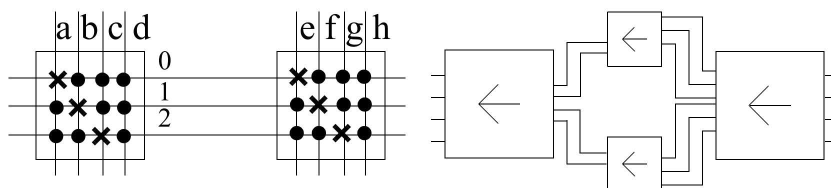

Figure 4: Construction of difficult SAT instances: (left) two switchboxes in common FPGA architectures, (right) similar N-by-M switchboxes are used to build hard satisfiable instances. Four connections are sought between (i) a, b, c and d, and (ii) e, f, g and h. Crosses correspond to input connections mated to channels, and every solid dot indicates the absence of a link.

4 Difficult SAT Instances

The *pigeon-hole principle* asserts that $n+1$ pigeons cannot be assigned to n holes as long as (i) two pigeons are not assigned to the same hole, and (ii) every pigeon must be assigned to one hole. These constraints can be expressed in terms of $n(n+1)$ Boolean variables: x_{ij} is interpreted as the indicator of assignment of pigeon j to hole i. The first family of clauses consists of $n^2(n+1)/2$ mutual exclusions $(\overline{x_{ij_1}} + \overline{x_{ij_2}})$, $j_1 \neq j_2$. The second family consists of $n+1$ n-literal clauses $(\sum_{i=1}^{n} x_{ij})$ — one for every pigeon j. The pigeon-hole principle then asserts that those two families of clauses cannot be satisfied simultaneously. However, its easy proof by induction is beyond the capabilities of backtrack SAT-solvers that typically operate within the resolution proof system.

The pigeon-hole instances `hole-n` described above are provably difficult for backtrack SAT solvers tied to resolution [7] and empirically difficult for the leading-edge implementation CHAFF as shown in Table 4. However, they are often treated as artificial in the EDA literature. Below, we derive equivalent instances `chnl[N]x[N+1]` from the domain of detailed routing for Field-Programmable Gate Arrays (FPGAs) and generalize them in two ways: `chnl[N]x[M]` (unsatisfiable) and `fpga[N].[M]` (satisfiable). We also give randomized constructions of difficult global routing instances `grout`.

4.1 FPGA routing instances

The pigeon-hole principle is directly related to routing because it can be interpreted as the impossibility of routing $n+1$ connections through n channels. As one can imagine, trying to make m connections through n channels is typical for FPGA routing, and in some cases $m > k$. We encode such instances in terms of $m \times k$ FPGA switchboxes that mate m input connections to k channels. A switchbox can connect any given input to any one channel, but no two inputs can be connected to the same channel, and every input must be connected to some channel. The state of an FPGA switchbox is described by an $m \times k$ matrix of binary variables and, similarly to the encoding of the pigeon-hole principle above, is subject to two families of constraints. These constraints are violated iff there are fewer channels than inputs. We put two $m \times k$ switchboxes on both sides of a batch of k channels, which entails $2mk$ variables (see [42] for details of SAT formulations). Figure 4(left), which illustrates our construction. It shows two

4×3 FPGA switchboxes connected to three horizontal channels. Four connections are sought between (i) a, b, c and d on the left, and (ii) e, f, g and h on the right. Crosses represent input connections mated to channels, and every dot indicates the absence of a link. Empirical results in Table 4 are shown for six routing configurations (`chnl`) in which one tries to route (a) 11, 12 or 13 connections through 10 tracks, and (b) 12, 13 or 20 connections through 11 tracks. These instances are extremely difficult for the leading-edge SAT solver CHAFF [41] and also have many symmetries. They can appear as sub-instances in larger routing instances, and such sub-instances may be difficult to find.

From the benchmarking point of view, it is natural to expect *unsatisfiable* instances among the most difficult to solve. Indeed, randomized restarts used by CHAFF [41] typically allow it to avoid difficult regions of the search space and to quickly find satisfying solutions if they exist. However, our second construction is designed to create difficult *satisfiable* instances that trap even the best solvers in hopeless regions of their solution space for a long time before a satisfying solution can be found. The main idea is to create a satisfiable instance with a large number of hard-to-avoid unsatisfiable sub-instances. If the number of unsatisfiable branches is much larger than the number of satisfiable branches, then random restart will keep on jumping from one unsatisfiable branch to another for a long time. Solvers without random restarts will, too, need to prove the unsatisfiability of many branches.

Our second construction entails routing a number of wires through four FPGA switchboxes of the type used in the first construction. The rightmost switchbox in the configuration in Figure 4(right) has several redundant outgoing tracks that are divided into two channels. Each channel is connected to a smaller switchbox with an insufficient number of outgoing tracks. The two groups of tracks that leave the smaller switchboxes are connected to the leftmost switchbox. When routing connections through tracks right-to-left, connections must be split between switchboxes subject to the throughput constraints of switchboxes. However, to a SAT solver, the throughput constraints are obscured by the pigeon-hole principle. SAT solvers first partition the connections between the two channels and backtrack from every partition that does not lead to a satisfying assignment. If the capacities of the two channels leading to the smaller switchboxes are greater than the throughput of those switchboxes, an overwhelming majority of partitions will lead to unsatisfiable pigeon-hole instances. On average, at least several such instances must be solved before a good partition is found. Empirical results for these satisfiable instances (`fpga`) in Table 4 show that they are difficult for CHAFF. We observe that these instances become more harder when the difference between the throughput of the small switchboxes and the capacities of the channels that lead to them is increased. This is consistent with our observations for the unsatisfiable `chnl` instances. Conceivably, some SAT-solvers may order variables related to the leftmost switchbox first and find satisfying assignments faster than CHAFF. This is consistent with our empirical data for the BerkMin solver [10] in Table 5. However, the configuration of switchboxes in Figure 4 (right) can be further modified to generate more difficult benchmarks. Specifically, one can add three new switchboxes on the left which are copies of existing three switchboxes on the right. The overall configuration will then be symmetric about the vertical axis passing through the currently leftmost switchbox in Figure 4 (right).

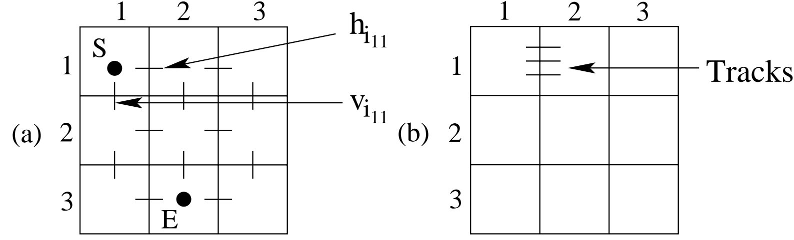

Figure 5: Construction of difficult SAT instances (global routing).

4.2 Global routing instances

We propose a new construction of difficult randomized *satisfiable* instances unrelated to pigeonholes. They express routing two-pin connections in a grid with edge capacity constraints. To ensure that an instance is satisfiable but difficult, we use *randomized flooding*. Namely, we create a routing configuration by adding shortest possible routes while unused routing resources (edge capacities) remain. Shortest routes are created by breadth-first-search between pairs of randomly chosen grid cells or, if that fails, by finding a maximal shortest route starting at a given grid cell with unused routing resources. After a routing configuration is created, routes are erased and their end-points are used to formulate a SAT instance.

Our SAT encoding of routing instances has two components. One deals with *route definition* and captures possible ways to route each connection. The other addresses *capacity constraints* and restricts the number of connections that can be routed across a grid cell boundary.

Route definition constraints. Routes are specified in terms of edges across cell boundaries in a grid. For each connection, we consider routing tracks across each cell boundary on the grid. In the SAT formulation, each track (for a given connection) is treated as a variable. Figure 5 (a) illustrates routing tracks in a 3-by-3 grid. Consider a two-terminal connection from S to E. Horizontal tracks for connection i are labeled $h_{i_{r,c}}$, where r and c are the row and column indices of the cell whose boundary the track crosses. Vertical tracks are labeled $v_{i_{r,c}}$. In Figure 5 (a), let the points marked S and E be the terminals of some two-terminal connection i. The SAT formulation proceeds as follows. For every connection, we add groups of clauses corresponding to individual grid cells.

For each of the two terminals, we add a clause consisting of positive literals of variables of all tracks to which the terminal can connect. For example, we add the clause $(h_{i_{1,1}} + v_{i_{1,1}})$ for terminal marked S in Figure 5(a) because any route for this connection must pass through $h_{i_{1,1}}$ or $v_{i_{1,1}}$. In the general case, we also need to add [binary] mutual exclusion clauses ensuring that only one of the incident tracks is actually taken. For the terminal S, this entails only one clause $(\overline{h_{i_{1,1}}} + \overline{v_{i_{1,1}}})$. For the terminal E, this entails three clauses $(\overline{h_{i_{3,1}}} + \overline{v_{i_{2,2}}})(\overline{h_{i_{3,1}}} + \overline{h_{i_{3,2}}})(\overline{v_{i_{2,2}}} + \overline{h_{i_{3,2}}})$.

We now consider every grid cell other than the terminals. Either *none* or *two* of its boundary edges must be selected. This is enforced as follows. Observe that a given cell may have two, three or four boundaries with tracks passing through them. Only two track variables, label them x_1 and x_2, are involved when "corner" grid cells are considered. In this case, we add clauses $(x_1 + \overline{x_2})(\overline{x_1} + x_2)$. In the case of three or four track variables ("border" grid cells or "internal"

grid cells respectively), we add two types of clauses. First, for every variable x_i, we add the constraint $(x_i \Rightarrow \sum_{j \neq i} x_j)$, which can be captured by one clause $(\overline{x_i} + \sum_{j \neq i} x_j)$ and says that if one boundary edge is selected, then another must be selected as well. The second type of clauses prohibits selecting three or four boundary edges. In the case of three variables x_1, x_2 and x_3, we add $(\overline{x_1} + \overline{x_2} + \overline{x_3})$. For four variables x_1, x_2, x_3 and x_4, we add

$$(\overline{x_1} + \overline{x_2} + \overline{x_3})(\overline{x_1} + \overline{x_2} + \overline{x_4})(\overline{x_1} + \overline{x_3} + \overline{x_4})(\overline{x_2} + \overline{x_3} + \overline{x_4})$$

As an illustration, we apply this procedure to the grid cell (1,2) in Figure 5(a) and produce

$$(\overline{h_{i_{1,1}}} + h_{i_{1,2}} + v_{i_{1,2}})(h_{i_{1,1}} + \overline{h_{i_{1,2}}} + v_{i_{1,2}})(h_{i_{1,1}} + h_{i_{1,2}} + \overline{v_{i_{1,2}}})(\overline{h_{i_{1,1}}} + \overline{h_{i_{1,2}}} + \overline{v_{i_{1,2}}})$$

The correctness of the general construction can be proven by following argument. First, it can be seen that any given connection, interpreted as a truth assignment, satisfies those constraints. Now assume an arbitrary satisfying assignment and show that, topologically, it is a valid connection. Start at a terminal. Exactly one track must be taken to a neighboring grid cell. If that grid cell is a terminal, we are done. Otherwise, exactly one track must be taken to a grid cell not visited before. The same argument shows that if a partial route is not completed, it can be extended by one track. Since there is only a finite number of grid cells, the route must be completed sooner or later.

When the layout is not *obstructed*, the above construction can be applied to all grid cells in an arbitrary order. However, if some tracks are removed or if certain grid cells are not available for routing, some grid cells may be unreachable from the terminals. Since no routes can pass through unreachable grid cells, they can be ignored when a SAT instance is constructed. We perform this optimization by traversing grid cells by a breadth-first search. Once a terminal is enqueued, our algorithm enters a loop that dequeues one grid cell, marks it visited, adds relevant clauses and enqueues unvisited adjacent grid cells. The algorithm finishes when the queue is empty. If the other terminal was not visited in the process, no routes connect the two terminals.

Capacity constraints. Each edge of a grid cell boundary has a capacity associated with it to restrict the number of connections that can be routed through it. The capacity limits are intended to prevent routing congestion. If C is the capacity limit for an edge of a grid cell, we include C variables per edge for each connection. In other words, each connection can be routed through one of C tracks across a cell boundary as shown in Figure 5 (b).

Consider two connections i and j. Consider horizontal route tracks for each connections, $h_{i_{r,c}}$, and $h_{j_{r,c}}$ for some row r and column c. Let $i_{r,c_1}, i_{r,c_2}, \ldots i_{r,c_C}$ and $j_{r,c_1}, j_{r,c_2}, \ldots j_{r,c_C}$ be the C extra variables introduced in the SAT formulation for the horizontal track in question. Then clearly, for any $i_{r,c_k}, 1 \leq k \leq C$, $i_{r,c_k} \Rightarrow h_{i_{r,c}}$, and also $h_{i_{r,c}} \Rightarrow (i_{r,c_1} + \ldots + i_{r,c_C})$. Clauses of this form are added to the SAT instance. Another restriction is that a route cannot pass through two tracks in the same channel (edge of a grid cell), i.e., if for some $k, 1 \leq k \leq C$, if i_{r,c_k} is true, then for all $l, 1 \leq l \leq C, l \neq k, (i_{r,c_k} \Rightarrow \overline{i_{r,c_l}})$. These clauses are also added. Finally, two connections cannot be routed through the same track, i.e. for all $k, 1 \leq k \leq C$, $(i_{r,c_k} \Rightarrow \overline{j_{r,c_k}})$ for all $j \neq i$, where j represents another connection.

We created ten routing configurations by randomly flooding a 3-by-3 routing grid with connections subject to edge capacity constraints of 3. Then we applied the SAT encoding above. The difficulty of these randomly generated benchmarks varies, and we only report empirical results for the five most difficult instances (`grout` in Table 4).

5 The Effect of Breaking Symmetries

Our computational experiments were performed on PCs with AMD Athlon processors @1.2GHz and 1GB of RAM. All codes were compiled with `g++ 2.95.4 -O3` and ran on Debian Linux. In addition to the instances described in Section 4 (`chnl` and `fpga`) and (`grout`), Table 4 lists six standard pigeon-hole instances (`hole`), five families of artificially constructed randomized Urquhart benchmarks (`Urq`) [52] and seven recent benchmarks from the micro-processor verification domain [54].

CHAFF run times in Table 4 are averages of (up to) 20 independent starts because CHAFF uses randomization internally and results of different runs may vary significantly. All runs not completed in 1000 seconds were aborted and did not contribute to averages. The percent of time-outs is shown for each instance.

To detect symmetries in CNF formulae, we convert them into colored graphs as outlined in Section 2. Those graphs are subsequently processed by the NAUTY program [38, 39]. For each run, the result is a list of permutation generators of the group of symmetries, specified by their cycles. For each SAT instance, Table 4 lists NAUTY run time in seconds excluding I/O, the total number of symmetries and the number of permutation generators. Those symmetry detection implementations are deterministic and not affected by re-ordering of vertices in the input graph. For some benchmarks we built symmetry-breaking clauses only for ten cycles per symmetry. The first ten cycles typically capture most of the speed-up provided by "breaking" a given symmetry. After new clauses were added, the preprocessed CNF instance was solved with CHAFF. Table 4 lists CHAFF run times for each instance. Because CHAFF run time on a given instance fluctuates from run to run, we report averages of 20 independent runs for each instance. Pre-processed CNFs never timed out in our experiments.

The last column in Table 4 shows the relative speed-up ratios due to the use of symmetry-breaking clauses. For a given CNF instance, the first number is the ratio of (i) the CHAFF run time on original instance, and (ii) the total run time of symmetry detection and CHAFF on pre-processed instances. The second number is produced similarly, except that symmetry detection run time is ignored. This is the maximal possible speed-up if symmetries are detected instantaneously or provided as domain-specific knowledge. We make the following observations:

1. The proposed SAT instances are only a fraction of the size of recent micro-processor verification benchmarks [54], but are more difficult to solve.

2. Some difficult SAT instances have astronomical numbers of symmetries; this includes the randomized `Urq` and `grout` benchmarks.

3. Symmetry-breaking clauses often speed up the best available SAT solver CHAFF [41].

4. Symmetry-breaking clauses typically do not slow down CHAFF and often speed it up, even when few symmetries are present.

5. Either CHAFF or symmetry detection may be a bottleneck.

6. Among the `chnl` instances, the hardest to solve was routing of 20 connections through 11 tracks. Adding extra unrouted connections consistently increased difficulty. That is somewhat counter-intuitive.

Instance	Satis-fiable?	#vars and #clauses	Plain CHAFF sec	Time -out %	Symmetries					Speed-up: total / search only
					Finding sec	Number of	#generators cycles		CHAFF sec	
hole07	UNS	56;204	0.37	0%	0.1	2.03e8	all	13	0.01	3.32; 36.50
hole08	UNS	72;297	1.27	0%	0.07	1.46e10	all	15	0.01	15.22; 94.15
hole09	UNS	90;415	3.79	0%	0.1	1.32e12	all	17	0.02	32.0; 204.97
hole10	UNS	110;561	22.44	0%	0.15	1.45e14	all	19	0.02	132; 1122
hole11	UNS	132;738	212.73	0%	0.18	1.91e16	all	21	0.03	1.23k; 7.09k
hole12	UNS	156;949	>1000	100%	0.24	2.98e18	all	23	0.04	— ; —
Urq3_5	UNS	46;470	232.44	10%	0.48	2.32e6	all	29	0.0	484.16; —
Urq4_5	UNS	74;694	250.01	25%	1.35	2.50e6	all	43	0.0	185.18; —
Urq5_5	UNS	121;1210	>1000	100%	13.15	>1e7	all	72	0.0	—; —
Urq6_5	UNS	180;1756	>1000	100%	62.93	>1e7	all	109	0.0	—; —
Urq7_5	UNS	240;2194	>1000	100%	176.62	>1e7	all	143	0.0	—; —
grout3.3-01	SAT	864;7592	19.01	0%	4.79	8.71e9	10	26	0.67	3.48; 28.37
grout3.3-03	SAT	960;9156	44.35	0%	8.94	6.97e10	10	29	0.40	4.75; 110.9
grout3.3-04	SAT	912;8356	19.36	0%	6.81	2.61e10	10	27	0.36	2.70; 53.79
grout3.3-08	SAT	912;8356	21.30	0%	7.14	3.48e10	10	28	0.67	2.73; 31.80
grout3.3-10	SAT	1056;10.8k	28.18	0%	10.65	3.48e10	10	28	0.85	2.45; 33.15
chnl10x11	UNS	220;1122	22.17	0%	0.45	4.20e28	all	39	0.11	39.91; 210.1
chnl10x12	UNS	240;1344	81.88	0%	0.61	6.04e30	all	41	0.12	111.6; 663.0
chnl10x13	UNS	300;2130	657.61	25%	1.28	4.50e37	all	47	0.17	454.8; 3.96k
chnl11x12	UNS	264;1476	207.37	0%	0.75	7.31e32	all	43	0.15	231.3; 1.41k
chnl11x13	UNS	286;1742	788.32	20%	1.08	1.24e35	all	45	0.16	633.5; 4.79k
chnl11x20	UNS	440;4220	>1000	100%	4.4	1.89e52	all	59	0.31	—; —
fpga10.08	SAT	120;448	7.56	0%	0.63	6.00e71	all	62	0.05	11.15; 157.6
fpga10.09	SAT	135;549	3.80	0%	0.88	6.33e77	all	68	0.03	4.16; 113.4
fpga12.11	SAT	198;968	694.00	50%	3.76	7.18e77	all	95	0.06	181.6; 11.3k
fpga12.12	SAT	216;1128	80.20	0%	5.31	7.44e77	all	104	0.13	14.74; 616.9
fpga12.08	SAT	144;560	246.70	10%	1.23	8.41e77	all	72	0.08	188.4; 3.10k
fpga12.09	SAT	162;684	885.00	80%	1.7	2.25e77	all	79	0.05	504.6; 16.4k
fpga13.09	SAT	176;759	550.00	85%	2.57	2.56e77	all	84	0.06	208.8; 8594
fpga13.10	SAT	195;905	>1000	100%	4.04	5.76e77	all	93	0.08	— ; —
fpga13.12	SAT	234;1242	>1000	100%	6.9	8.85e77	all	110	0.08	— ; —
2dlx_ca_mc*	UNS	3250;24.6k	6.54	0%	38.36	4	10	2	6.30	0.15; 1.04
2pipe.cnf	UNS	892; 6695	2.08	0%	10.74	128	10	7	1.56	0.17; 1.33
2pipe_1_ooo	UNS	834; 7026	2.55	0%	9.37	8	10	3	1.80	0.23; 1.41
2pipe_2_ooo	UNS	925; 8213	3.43	0%	11.14	32	10	5	2.82	0.25; 1.22
3pipe	UNS	2468;27.5k	36.44	0%	463.57	512	10	9	19.65	0.08; 1.85
4pipe	UNS	5237;80.2k	337.61	0%	>1000	—	—	—	—	—; —
5pipe	UNS	9471;195k	325.92	0%	>1000	—	—	—	—	—; —

Table 4: CHAFF run time on original SAT instances is compared to the combined run time of symmetry detection and CHAFF on instances with symmetry-breaking clauses added. The rightmost column also shows pure search speed-up (that does not take symmetry detection into account). The full name of benchmark `2dlx_ca_mc` is `2dlx_ca_mc_ex_bp_f` . The numbers of symmetry generators and max cycles used per generator (10 or all) are shown. The benchmarks we generated for these experiments are available at `http://gigascale.org/bookshelf/Slots/SATbench`

Instance	original	with SBPs
hole10	110	0.01
Urq3_5	>1000	0.17
groute3.3-03	5.5	0.6
chnl10_11	109.95	0.02
fpga10_8	0.02	0.02
3pipe	2.69	0.3

Table 5: Runtime of the BerkMin solver [10] (version 56) on sample SAT instances: original and with symmetry-breaking predicates added.

Not to limit out results to a single SAT-solver (CHAFF), we ran similar experiments with the BerkMin solver [10] version 56. Representative results are shown in Table 5 where solver runtimes are compared with and without symmetry-breaking predicates added. BerkMin solves the `grout`, `fpga` and microprocessor verification benchmark sets faster than CHAFF, but other benchmark sets are harder for BerkMin. Symmetry-breaking reduces runtime in most cases. In similar experiments with GRASP[45], all of our benchmarks are solved faster with the help of symmetry-breaking predicates, even if symmetry-finding time is charged for.

6 Opportunistic Symmetry-Finding

The use of symmetry-breaking clauses does not require finding *all* symmetries. In fact, an algorithm that does not guarantee finding all symmetries may finish sooner. Some symmetries may be found using domain-specific knowledge, and then symmetry-breaking clauses can be added during the creation of SAT instances.

6.1 Window-based symmetry finding

We observed that a variable would sometimes be symmetric to another variable connected by a clause (one hop) or through a chain of two clauses (two hops). When this is not true for all symmetries of a CNF formula, many symmetries may be composable from permutation generators of that kind. We therefore focus on "local" symmetries that permute small subsets of variables and fix all other variables.[6] We define the subsets by sliding a window of fixed size along a given linear ordering of the variables — either the original variable ordering of the CNF formula or the connectivity-based MINCE ordering [2]. For a window, we consider the left and right cuts, as in Figure 6. To find symmetries local to a given window, the standard construction of colored graph is applied to clauses and literals that are entirely inside the window. Each *cut clause* is represented by a vertex of a unique color that is connected to literals inside the window. Vertices beyond the current window are ignored. To argue that the proposed construction is correct, i.e. does not add spurious symmetries, we consider the following *recoloring* process.

Definition 6.1.1. *Given a colored graph and a subset of its vertices, change the color of each vertex into a unique color — one new color per vertex. This process is called* recoloring

[6]The complexity of such a restricted version of the graph automorphism problem was studied in [32].

of a given set of vertices. The following lemma shows how to restrict the set of symmetries of a colored graph. This can be done, e.g., with the purpose of accelerating symmetry-finding for the price of losing some symmetries.

Lemma 6.1.2. *Given a colored graph G, consider an arbitrary recoloring of an arbitrarily-chosen subset of its vertices. Call the resulting graph G^R. Then the following claims hold.*

(a) every symmetry of G^R is a symmetry of G, and must map each recolored vertex to itself;

(b) symmetries of G^R form a subgroup in the group of symmetries G;[7]

(c) the choice of new (unique) colors does not affect the symmetry group G^R.

While reducing the number of symmetries can, in principle, be consistent with smaller symmetry-detection runtimes, most graph automorphism programs are most sensitive to the number of vertices in the input graph rather than to the number of symmetries. The following lemma shows how to reduce the vertex set of the graph in the context of Lemma 6.1.2.

Lemma 6.1.3. *Given a colored graph G, consider an arbitrary recoloring of an arbitrarily-chosen subset of its vertices. Call the recolored graph G^R. Consider a non-empty subset W of recolored vertices such that each of them is adjacent to recolored vertices only (if such a subset exists). Remove all vertices in W from G^R together with all incident edges. Then the symmetries of the remaining colored graph G^R_W are in one-to-one correspondence with the symmetries of G^R, in fact the two groups of symmetries are isomorphic.*

Proof. Every symmetry of G^R maps every vertex from W to itself by Lemma 6.1.2 (a). Therefore, every such symmetry gives rise to a symmetry of G^R_W. Vice versa, every symmetry of G^R_W can be unambiguously extended to a symmetry of G^R by mapping every vertex from W to itself. This construction restores every symmetry of G^R mapped to a symmetry of G^R_W.

Lemma 6.1.3 reduces the number of vertices under the assumption that set W exists — the larger W, the greater the reduction. Constructively finding W remains an open problem.

Lemma 6.1.4. *Given a colored graph G and an arbitrary edge-cut in it, pick one of the partitions and recolor all vertices in it. Then the set of vertices in that partition that are not incident to any edges in the cut can play the role of set W in Lemma 6.1.3.*

Observe that colored graph G^R_W from Lemma 6.1.3 may still contain a large number of recolored vertices. This may be undesirable because the total number of vertices in G^R_W is limited by the scalability of available symmetry-detection software, and non-trivial symmetries of G^R_W do not involve recolored vertices. Indeed, recolored vertices are included into the vertex set of G^R_W, thus potentially slowing down symmetry-detection programs or at least increasing memory usage.[8] Therefore, this construction can be improved by minimizing the number of vertices incident to cut edges, e.g., by minimizing the size of the cut itself.

Another concern about restricting symmetry detection along the lines of Lemmas 6.1.2-6.1.4 is that one should apply it several times, with different sets of vertices recolored. This way more symmetries can be detected. Indeed, if the size of G^R_W is limited by a constant, then the number of calls to symmetry-detection software should grow at least linearly so that every vertex in G be "given an opportunity" to map elsewhere.

The concerns mentioned above can be addressed in the context of window-based symmetry detection. We first order CNF variables by representing the CNF as a hypergraph (clauses

[7] This subgroup is the *stabilizer* [24, 25] of the set of recolored vertices in the symmetry group of G.

[8] NAUTY maintains the input graph in a dense adjacency matrix.

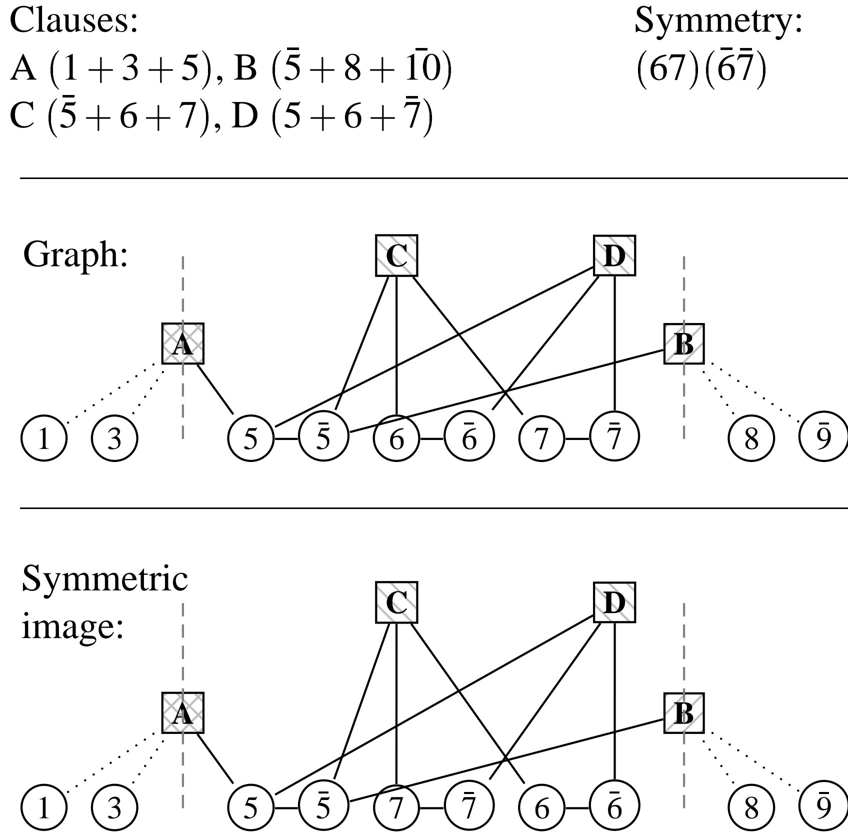

Figure 6: Window-based opportunistic symmetry detection for a CNF instance with ten variables and four clauses. Vertical dashed lines capture the window to which search for symmetries is limited. Each clause that includes literals from both within and beyond the window are represented by vertices of unique colors (dashed boxes). Symmetries are only allowed to permute vertices within the current window, therefore vertices and edges beyond the current window are not included in the graph for window-based symmetry-detection. This reduces the size of graph automorphism problems.

correspond to hyperedges) and heuristically finding a min-cut linear arrangement of those vertices using recursive balanced bisection [2]. We then consider cuts along the resulting variable ordering, and those cuts are relatively small by construction. Note that cuts in Lemma 6.1.4 correspond to pairs of cuts in a given variable ordering as shown by vertical dashed lines in Figure 6. Furthermore, in window-based symmetry-detection only clausal vertices can be recolored, therefore min-cut linear arrangement naturally minimizes the number of recolored vertices.

We concatenate lists of permutation generators produced for different windows, consider the group generated by all those and use GAP [51] to produce an irredundant list of generators of this "global" group. Symmetry-breaking clauses are constructed from those generators. Observe that when applying symmetry detection to a given window, we can only detect symmetries that permute variables in that window only. Therefore, potentially more symmetries can be found if windows are allowed to overlap. On the other hand, if overlaps are allowed some symmetries may be detected in multiple windows. Thus, producing symmetry-breaking clauses independently from each window and concatenating them may cause considerable redundancy. This is why we call GAP if windows are allowed to overlap. The trade-off between run time,

Instance	Satis-fiable?	#variables and #clauses	Plain CHAFF sec	Time -out %	Symmetries					Speed-up: total / search only
					Finding sec	# of	#generators		CHAFF sec	
							cycles			
WINDOW-BASED SYMMETRY FINDING (1000 variables per window)										
2dlx_ca_mc*	UNS	3250;24.6K	6.54	0%	3.17	1	-	0	6.54	0.67; 1.00
2pipe	UNS	892; 6695	2.08	0%	10.47	128	10	7	1.30	0.18; 1.63
2pipe_1_ooo	UNS	834; 7026	2.55	0%	9.02	8	10	3	1.80	0.24; 1.41
2pipe_2_ooo	UNS	925; 8213	3.43	0%	11.09	32	10	5	2.80	0.25; 1.23
3pipe	UNS	2468;27.5K	36.44	0%	3.63	4	10	2	36.20	0.91; 1.01
4pipe	UNS	5237;80.2K	337.61	0%	9.32	2	10	1	334.0	0.98; 1.01
5pipe	UNS	9471;195K	325.92	0%	29.42	2	10	1	325	0.92; 1.00

Table 6: Results for window-based symmetry-finding. Labeling is identical to that of Table 4. Typically all or a large fraction of all symmetries are discovered, compared to data in Table 4.

incomplete symmetry-finding and redundancy among windows depends on their overlap. Similarly, the window size affects the trade-off between run time and incomplete symmetry-finding. We observe good empirical performance with windows of size 1000. Results in Table 6 show that our window-based technique found all or a significant portion of all symmetries for the micro-processor verification benchmarks [54] in a fraction of the run time spent by complete symmetry-finding. If a randomized variable ordering is used, one could combine local permutation generators found for different orderings.

6.2 Improving SAT formulations

One way to reduce the run time of symmetry-finding is to learn how to detect (or predict) symmetries from domain-specific knowledge. Given the well-understood structure and symmetries of the `hole`, `chnl` and `fpga` benchmarks, we evaluated this approach on (randomized) `grout` benchmarks. We noticed that permuted variables in many cases correspond to neighboring tracks, e.g., if two connections are routed in parallel through several grid cells, there is considerable freedom (symmetry) in track assignment. To break this symmetry, we added clauses that preserve the relative order of tracks taken by every pair of connections routed through the same two edges of a grid cell. In other words, if one connection is routed through track 2 when entering the cell, and another connection is routed through track 3 when entering the cell, then the connections are allowed to leave the cell through tracks 2 and 3 resp., 1 and 2 resp. or 1 and 3 resp. Such constraints speed-up CHAFF: each `grout` instance is now solved in *0.50-0.80 seconds versus 19-45 seconds*. More dramatic speed-ups are achieved for `grout` instances built with larger routing grids. Even if we apply symmetry-detection to modified instances, it completes much faster than on original instances because no symmetries are found. It may also be possible to add domain-specific symmetry-breaking clauses to SAT instances from [54] and improve CHAFF runtime according to results in Table 4.

7 Conclusions

Our work addresses solving difficult instances of Boolean (CNF) satisfiability that exhibit structural symmetries. While the utility of our approach on easy instances is not clear at this moment, the difficulty of domain-specific classes of CNF-SAT instances is often known, and adequate SAT algorithms can be chosen. Otherwise, several SAT solvers can be executed in parallel until one of them finishes. On a single processor, this may buy exponential speed-ups at the cost of a constant-factor slow-down. Therefore, our focus on difficult instances is well justified. Additionally, our experiments identify a number of difficult instances whose difficulty is apparently due to symmetries and redundant search caused by them.

We describe an automated flow that finds symmetries in given CNF instances and uses them to speed up SAT search. This flow includes symmetry detection, pre-processing of given CNF instances and an application of an existing state-of-art SAT solver. When compared to the SAT solver alone, applied to given CNF instances without pre-processing, our flow dramatically speeds up the solution of two well-known provably-difficult benchmark families — pigeonhole problems and Urquhart benchmarks. Notably, methods proposed in a previous work [18] cannot detect any non-trivial symmetries in Urquhart (`Urq`) benchmarks.

We offer constructions of realistic satisfiable and unsatisfiable SAT instances, arising in routing applications, that are unusually difficult for their size. Unlike most existing SAT benchmarks, our benchmark families enable studies of the asymptotic performance of SAT solvers.

Since symmetry-finding is a bottleneck, we speed it up using opportunistic approaches. In one, we only look for symmetries that permute small groups of variables. Those groups are determined by sliding a fixed-sized window along a given variable ordering. The second approach attempts to improve the construction of SAT instances by detecting symmetries in domain-specific terms so that new clauses can be added during construction. We find astronomically many symmetries in randomized `Urq` and `grout` benchmarks. This refutes a conventional-wisdom argument claiming that significant randomization necessarily destroys symmetries. We explain symmetries in `grout` benchmarks and break them using domain-specific knowledge.

Our proposed flow does not require source code modifications in SAT solvers and should work with most back-track SAT solvers. We successfully validated our flow with CHAFF [41], BerkMin [10] and GRASP [45] (GRASP results are not included in this paper). Experiments performed with publicly available versions of WalkSAT [46] indicate that symmetry-breaking clauses do not improve runtimes and even make them worse. This was observed by others and is the focus of on-going work by Preswitch, Kautz and Selman.

We stress that the proposed flow may not be useful on SAT benchmarks that (i) are easy, or (ii) do not have symmetries. Many difficult SAT instances do not have symmetries [16]. On the other hand, many DIMACS benchmarks [21] have large numbers of symmetries, but are easy and can be solved faster than their symmetries can be found by existing methods.

Our on-going research seeks (i) faster symmetry detection, e.g., via incomplete algorithms, (ii) finding [some] semantic symmetries that are not necessarily syntactic, (iii) more efficient constructions of symmetry-breaking clauses, and (iv) the use of partial/conditional symmetries. The latter were already shown useful in BDD-based model checking [23], SAT-solvers based on backtracking [13, 31] and more general constraint-satisfaction solvers [6].

Acknowledgements

This work is funded by the DARPA/MARCO Gigascale Silicon Research Center, an Agere Systems/SRC Research fellowship and a fellowship from the ACM/IEEE Design Automation Conference.

References

[1] D. Achlioptas, C. P. Gomes, H. A. Kautz, and B. Selman, "Generating Satisfiable Problem Instances", *AAAI 2000*, pp. 256-261.

[2] F. Aloul, I. Markov and K. Sakallah, "Faster SAT and Smaller BDDs via Common Structure", *Intl. Conf. on Computer-Aided Design (ICCAD) 2001*, pp. 443-448.

[3] L. Babai and E. M. Luks, "Canonical Labeling of Graphs", *Symp. on the Theory of Computing (STOC) '83*, pp. 171-183.

[4] L. Babai, R. Beals and P. Takácsi-Nagy, "Symmetry and Complexity", *Symp. Theory of Comp. (STOC) '92*, pp. 438-449.

[5] L. Babai, "Automorphism Groups, Isomorphism, Reconstruction", Chapter 27, pp. 1447-1541, In (R. L. Graham, M Grötschel and L. Lovász, eds, Handbook of Combinatorics, vol. 2, MIT Press, 1995).

[6] R. Backofen and S. Will, "Excluding Symmetries in Constraint-Based Search", *Intl. Conf. on Principles and Practice of Constraint Programming (CP'99)*, Lecture Notes in Computer Science, vol. 1713, pp. 73-87, Springer-Verlag, 1999.

[7] P. Beame, R. Karp, T. Pitassi and M. Saks, "The efficiency of Resolution and Davis-Putnam Procedure", to appear in *SIAM Journal on Computing*.
`http://www.cs.washington.edu/homes/beame/papers/resj.ps`

[8] B. Benhamou and L. Sais, "Tractability through symmetries in propositional calculus", *Journal of Autom. Reasoning*, vol. 12, (no.1), Feb. 1994. pp. 89-102.

[9] A. Bernasconi, V. Ciriani, F. Luccio, L. Pagli, "Fast Three-Level Logic Minimization Based on Autosymmetry", *Design Automation Conf.*, 2002, pp. 425-430.

[10] E. Goldberg and Y. Novikov, "BerkMin: A Fast and Robust SAT Solver", *Design Automation and Test in Europe (DATE) 2002*, pp. 142-149.

[11] D. Bosnacki, D. Dams and L. Holenderski, "A Heuristic for Symmetry Reductions with Scalarsets", *Intl. Symposium on Formal Methods for Increasing Software Productivity (FME) 2001*, Lecture Notes in Computer Science, Springer-Verlag, 2001.

[12] L. Brisoux, E. Gregoire, L. Sais, "Improving backtrack search for SAT by means of redundancy", *Intl. Symp. Foundations of Intelligent Systems (ISMIS) '99*. Warsaw, Poland; Springer-Verlag 1999, p. 301-309.

[13] C. A. Brown, L. Finkelstein, and P. W. Purdom. "Backtrack searching in the presence of symmetry". (T. Mora, editor), *Intl. Conf. on Applied Algebra, Algebraic Algorithms and Error Correcting codes, 6th intl. conf.*, pp. 99-110. Springer-Verlag, 1988.

[14] V. Chvatal and E. Szemeredi, "Many hard examples for resolution", *Journal of ACM*, 35(4), pp. 759–768, 1988.

[15] E.M. Clarke et al., (Edited by: Hu, A.J.; Vardi, M.Y.) "Symmetry Reductions in Model Checking", in *Intl. Conf. Computer Aided Verification (CAV) '98*, pp. 159-171.

[16] S. A. Cook and D. G. Mitchell, "Finding Hard Instances of the Satisfiability Problem: A Survey", In *Satisfiability Problem: Theory and Applications*, DIMACS Series in Discr. Math. and Theor. Comp. Sci, 25, pp. 1–17. Amer. Math. Soc., 1997.

[17] J. Crawford, "A theoretical analysis of reasoning by symmetry in first-order logic", *The AAAI Workshop on Tractable Reasoning held at the Tenth National Conference on Artificial Intelligence (AAAI-92)*, San Jose, CA.

[18] J. Crawford, M. Ginsberg, E. Luks and A. Roy, "Symmetry-breaking predicates for search problems", *5th Intl Conf. Principles of Knowledge Representation and Reasoning (KR) '96*, Cambridge, MA, pp. 148-159.

[19] M. Davis and H Putnam. "A Computing Procedure For Quantification Theory", *Journal of the ACM*, vol. 7(3), pp. 201–215, 1960.

[20] M. Davis, G. Logemann, and D. Loveland, "A Machine Program for Theorem Proving", *Journal of the ACM*, (5)7, pp. 394–397, 1962.

[21] DIMACS Boolean Satisfiability Challenge Benchmarks:
`ftp://dimacs.rutgers.edu/pub/challenge/sat/benchmarks/cnf`

[22] C.A.J. van Eijk, E.T.A.F. Jacobs, B. Mesman and A. H. Timmer, "Identification and Exploration of Symmetries in DSP Algorithms", in *Design Automation and Test in Europe (DATE) '99*, March 1999, pp. 602-608.

[23] E. A. Emerson and R. J. Trefler, "From Asymmetry to Full Symmetry: New Techniques for Symmetry Reduction in Model Checking", *Conf. on Correct Hardware Design and Verification Methods (CHARME) '99*, Lecture Notes on Computer Science, Springer 1999.

[24] M. Hall Jr., "The Theory of Groups", McMillan, 1959.

[25] Th. W. Hungerford, "Algebra", *Graduate Texts in Mathematics*, vol. 73, Springer, 1973.

[26] C. N. Ip and D. L. Dill, "Better verification through symmetry", *Formal Methods in System Design*, 9(1/2), pp.41-75, 1996.

[27] V. Kravets and K. Sakallah, “Generalized Symmetries of Boolean Functions”, *Intl. Conf. on Computer-Aided Design 2000*, pp. 526-532.

[28] V. Kravets and K. Sakallah, “Constructive Library-Aware Synthesis Using Symmetries”, *Intl. Conf. Design Automation and Test in Europe (DATE) 2001*, pp. 208-213.

[29] B. Krishnamurthy, “Short Proofs For Tricky Formulas”, *Acta Informatica*, vol. 22, pp.327–337, 1985.

[30] J. Köbler, U. Schöning and J. Torán, “Graph Isomorphism Is Low For PP”, *Computational Complexity*, vol. 2, no. 4, 1992, pp. 301-330, `http://citeseer.nj.nec.com/obler92graph.html`.

[31] C. M. Li, B. Jurkowiak and P. W. Purdom, “Integrating Symmetry Breaking Into A DLL Procedure”, *Intl. Symp. on Boolean Satisfiability (SAT)*, Cincinnatti, 2002, pp. 149-155.

[32] A. Lozano and V. Raghavan, ”On the complexity of moving vertices in a graph”, Tech Report LSI-98-30-R, Universitat Politéchnica de Catalunya, 1998, `http://citeseer.nj.nec.com/25551.html`.

[33] E. Luks, “Isomorphism of Graphs of Bounded Valence Can Be Tested in Polynomial Time”, Proc. *IEEE Symp. on Foundations of Comp. Sci.* (1980), pp. 42-49.

[34] E. Luks, “Hypergraph isomorphism and structural equivalence of boolean functions”, *Symp. on Theory of Computing (STOC) ‘99*, pp. 652-658.

[35] E. Luks and A. Roy, “Symmetry Breaking in Constraint Satisfaction”, *Intl. Conf. of Artificial Intelligence and Mathematics*, Ft. Lauderdale, Florida, Jan 2-4, 2002.

[36] G. S. Manku, R. Hojati and R. Brayton, “Structural symmetry and model checking”, *Intl. Conf. Computer-Aided Verification (CAV) ’98*, pp. 159-171.

[37] I. McDonald and B. Smith, “Partial Symmetry Breaking”, Technical Report, APES-49-2002, May 2002. `http://www.dcs.st-and.ac.uk/~apes/apesreports.html`

[38] B. D. McKay, “Practical Graph Isomorphism”, *Congressus Numerantium*, 30 (1981), pp. 45-87.

[39] B. D. McKay, “Nauty user’s guide” (version 1.5), Technical report TR-CS-90-02, Australian National University, Computer Science Department, ANU, 1990. `http://cs.anu.edu.au/~bdm/nauty/`

[40] T. Miyazaki, “The Complexity of McKay’s Canonical Labeling Algorithm”, In *L. Finkelstein and W. M. Kantor, eds, Groups and Computation II*, Workshop on Groups and Computation, DIMACS Series on Discrete Mathematics and Theor. Computer Science, 1996.

[41] M. Moskewicz, C. Madigan, Y. Zhao, L. Zhang and S. Malik, “Chaff: Engineering an Efficient SAT Solver”, *Design Automation Conf.*, 2001, pp. 530-535.

[42] G. Nam, F. Aloul, K. Sakallah, and R. Rutenbar, "A Comparative Study of Two Boolean Formulations of FPGA Detailed Routing Constraints", *Intl. Conf. on Physical Design (ISPD 2001)*, pp. 222-227.

[43] M. Prasad, P. Chong and K. Keutzer, "Why is ATPG easy?", *Design Automation Conf (DAC) '99*, pp. 22-28.

[44] M. Prasad, E. Goldberg and R. Brayton, "Using Problem Symmetry In Search Based Satisfiabiliy Problems", *Design Automation and Test in Europe (DATE) 2002*, pp. 134-142.

[45] J. P. M. Silva and K. A. Sakallah, "GRASP: A New Search Algorithm for Satisfiability", *IEEE Trans. On Computers*, vol. 48, no. 5, May 1999, pp. 506-521.

[46] B. Selman, H. A. Kautz and B. Cohen, "Noise Strategies for Improving Local Search", *National Conference on Artificial Intelligence (AAAI' 94)*, pp. 337-343.

[47] B. Selman, D. Mitchell and H. Levesque, "Generating Hard Satisfiability Problems", *Artificial Intelligence*, vol. 81, no. 1-2 (1996) pp. 17–29.

[48] A. Seress, "An introduction to computational group theory", *Notices Amer. Math. Soc.*, vol. 44 (1997), no. 6, 671-679.

[49] B. M. Smith, K. E. Petrie and I. P. Gent, "Models and Symmetry breaking for 'Peaceable Armies of Queens", Technical Report, APES-50-2002, May 2002. `http://www.dcs.st-and.ac.uk/~apes/apesreports.html`

[50] L. H. Soicher, "GRAPE: A System For Computing With Graphs and Groups", *in "Groups and Computation"* (L. Finkelstein and W.M. Kantor, eds), *DIMACS Ser. in Discr. Math. & Theor. Comp. Sci.* 11, pp. 287-291. `www-groups.dcs.st-andrews.ac.uk/~gap/Share/grape.html`

[51] E. L. Spitznagel, "Review of Mathematical Software, GAP", *Notices Amer. Math. Soc.*, 41 (7), (1994), pp. 780–782. `http://www.gap-system.org/`

[52] A. Urquhart, "Hard Examples for Resolution", *Journal of ACM*, 24(1), pp.209-219, 1987.

[53] A. Urquhart, "The Symmetry Rule in Propositional Logic", 1996.

[54] M. N. Velev, and R.E. Bryant, "Effective Use of Boolean Satisfiability Procedures in the Formal Verification of Superscalar and VLIW Microprocessors", *DAC 2001*, pp. 226-231. `http://www.ece.cmu.edu/~mvelev/#BENCHMARKS`

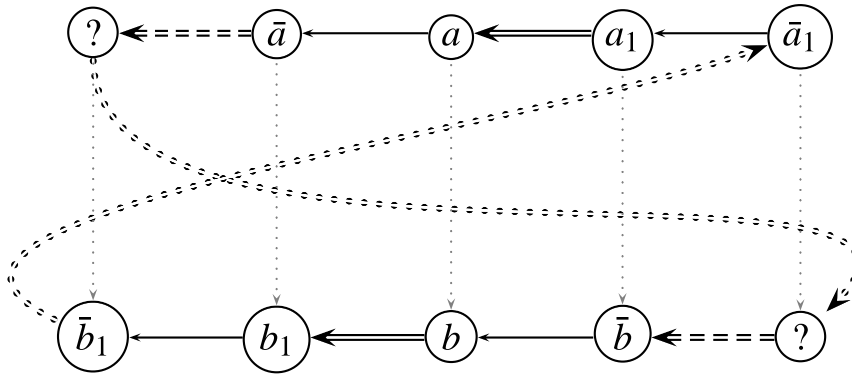

Figure 7: An illustration of the proof of Theorem 2.3.2. Only vertices of color 2 are shown. All labels are *literals* and not necessarily *variables*. Solid and dashed double-arrows show edges between those vertices and the direction of logic implications in the chains. Double edges correspond to binary clauses and single edges are Boolean consistency edges. Dotted arrows show a permutation of vertices implied by $a \mapsto b$ and $\bar{a} \not\mapsto \bar{b}$. Double dashed arrows show implications that complete the circular chain. Note that the original graph is undirected, and the choice of direction for horizontal edges is a part of the proof.

Appendix

Below we complete the proof of Theorem 2.3.3 by showing that Boolean consistency is preserved. Recall that we model two-literal clauses by edges that connect vertices of color 2. Such edges may potentially map to Boolean consistency edges, and we must prove that impossible.

The informal argument below proceeds by contradiction and amounts to a proof by induction, but better *explains* correctness to a human reader. Assume, without the loss of generality, that there exists a literal a that maps to a literal b, but $\bar{a}$ does not map to $\bar{b}$. Then the edge $a\bar{a}$ must map into an edge bb_1, where b_1 is a literal ($b_1 \neq \bar{a}$) sharing a binary clause with b. In other words, $\bar{a} \mapsto b_1$. Similarly, a literal a_1, distinct from $\bar{a}$ and sharing a binary clause with a, must map into $\bar{b}$ (the case when literals a_1 and b_1 correspond to the same variable can be analyzed explicitly, but generally falls into cases (ii) and (iii) below). Two immediate consequences involve vertices $\bar{b}_1$ and $\bar{a}_1$. Since $\bar{a}_1$ is adjacent to a_1, it must map into a vertex adjacent to $\bar{b}$ and distinct from the vertices that we have encountered so far. Similarly, a vertex adjacent to $\bar{a}$ and distinct from the vertices that we have encountered so far must be mapped to $\bar{b}_1$. These two new vertices are shown with question marks in Figure 7. Observe that any binary clause can be thought of as a logic implication between its values. For example, $(a+b)$ can be viewed as $(a=0) \Rightarrow (b=1)$ or $(b=0) \Rightarrow (a=1)$, depending on the direction we need.

Through the process outlined above, a spurious symmetry leads to two chains of implications, such that one is mapped to the other by the symmetry. The fact that the symmetry is spurious allows one to continue the chain by adding two new implications at a time, in either direction. However, since we are dealing with a finite graph, this process must end at some point by involving variables used before, whence we arrive at a circular chain of implications as shown in Figure 7. Indeed, suppose we continue the two chains to the left, considering two

more vertices in the graph. One of them ($\bar{b}_1$ in Figure 7) must represent the negation of a literal we considered at the previous step (b_1), and the other (the vertex in the upper left corner with a question mark) does not have to. We are interested in two cases for this unknown vertex: (i) represents the negation of a previously encountered literal, or (ii) this vertex does not represent the negation of a previously encountered literal, but was encountered before. If the vertex was not encountered before, we simply continue expanding the two chains to the left until we find ourselves in case (i) or case (ii).

If case (i) is observed when extending the chains in one direction, e.g., to the left in Figure 7, the two chains become connected. Moreover, because existing chains consist of pairs of complementary literals, one chain can only connect to "the opposite" end of the other chain. After that, we can still continue the united chain to the right. If case (i) is observed on the right, a circular chain of implications found. For example, if the two vertices with question marks in Figure 7 correspond to literals z and $\bar{z}$, we get $(a_1+a)(\bar{a}+z)(\bar{z}+\bar{b})(b+b_1)(\bar{b}_1+\bar{a}_1)$.

Case (ii) requires more sophisticated analysis because not every chain of binary clauses is a chain of implications according to Definition 2.3.1. For example, if the two vertices with question marks in Figure 7 both correspond to the literal z, then simply connecting the two chains by z yields $(a_1+a)(\bar{a}+z)(z+\bar{b})(b+b_1)(\bar{b}_1+\bar{a}_1)$. The literal z enters the two incident binary clauses *with the same polarity* and therefore does not propagate implications in either direction. Therefore, *instead of* connecting the two chains, we continue both of them by adding the Boolean consistency edge $z\bar{z}$. While this does not, by itself, create a circular chain of implications, it reduces the supply of vertices that were not encountered yet. By finiteness, case (i) should happen sooner or later, and a circular chain of implications will be found. □