Cut-Based Inductive Invariant Computation

Michael Case[1,2] **Alan Mishchenko**[1] **Robert Brayton**[1]

[1] Department of EECS, University of California, Berkeley, CA

[2] IBM Systems and Technology Group, Austin, TX

{casem, alanmi, brayton}@eecs.berkeley.edu

Abstract

This paper presents a new way of computing inductive invariants in sequential designs. The invariants are useful for strengthening inductive proofs in difficult unbounded model checking instances. Candidate invariants are derived from a set of m-feasible cuts in the logic network and proved by induction. Thus, the proposed computation is very scalable, and it is possible to flexibly trade computational effort for the expressiveness of the proved invariants. Experimental results on several benchmark families show that the proposed strengthening proves many hard properties that are unsolved by other model checkers. The implementation is publicly available in the synthesis and verification system ABC.

1 Introduction

Model checking [11][28] safety and liveness properties involves proving that a safety property holds on all reachable states [2]. Many safety properties can be verified by proving the property on an inductive superset of reachable states. If the superset can be represented compactly, then such a method is easier and more scalable than deriving the exact set of reachable states. Finding such an inductive superset is called *inductive strengthening*.

This paper introduces a new way of deriving and proving an additional inductive property, or *invariant*, that (1) can be effective for inductive strengthening and (2) leads to a flexible and scalable computation that can trade computational effort for increased expressive power of the invariant derived. As a byproduct, the same invariant can be used as a source of external don't cares for circuit restructuring in sequential logic synthesis.

The proposed invariant consists of a set of clauses derived using *m*-input cuts of nodes in the sequential circuit. A cut is a boundary separating the node from the primary inputs and register outputs. Therefore the invariant is expressed in terms of groups of adjacent nodes in the network.

This computation illustrates the synergy between logic synthesis and verification [8]. In the past, external don't cares for logic synthesis were obtained by computing the set of unreachable states characterized by a function of the register outputs. However, even if these can be computed, scalability motivates the use of windowing where the computation is temporarily restricted to a node being optimized and a group of surrounding nodes. To use the external don't cares they must be projected onto the inputs of the window. The useful unreachable states have nontrivial projections onto the inputs of the windows, which form cuts in the network. The new idea is to skip computing the unreachable states and directly compute its projections onto various cuts. Although motivated by logic synthesis, we find these invariants to be useful for inductive strengthening in unbounded verification as well.

Induction [15][5][12] is a practical model checking method, applicable to large designs whose size and logic complexity often cause other methods (such as BDD-based reachability, interpolation, localization, etc) to fail. A property is *inductive* if it satisfies two conditions: (*base case*) it holds in the design's initial state(s), and (*inductive case*) if it holds in a particular state, then it holds in all states reachable from that state in one transition. Induction is scalable because both the base and inductive cases can be formulated as incremental instances of Boolean satisfiability (SAT) [12], which can be solved efficiently using modern SAT solvers [13].

In our method, the properties are clauses plus a target property, and the groups of variables participating in the clauses are derived using efficient *m*-cut computation, which is adopted from LUT-based technology mapping [24]. It avoids exhaustive cut enumeration [27] and computes only a small subset of useful cuts using priority heuristics similar to those in [14].

The initial set of candidate clauses is detected using two types of random simulation, combinational and sequential. Minterms at a cut that appear under combinational but never under sequential simulation are recorded. A candidate clause is the complement of such a minterm. The set of candidates is iteratively refined using SAT-based induction, and if the greatest fixed-point also satisfies the inductive base case then the conjunction of all clauses in the fixed point set yields the proposed inductive invariant, an over-approximation of the reachable states.

To make this computation efficient, a flexible framework has been developed for trading the number and expressiveness of the candidate invariant clauses for computation time. The clauses are proved in batches, each of which successively refines the already computed approximation of the reachable states. The process is stopped when the target property becomes inductive, or when the number of clauses successfully proved is sufficient for the calling application.

Scalability is achieved by using heuristics for candidate clause generation and filtering. One heuristic limits clauses to those derived for cuts a few logic levels from the register outputs. Inductive proofs for such shallow clauses can be processed efficiently by partitioning the design and solving partitions in parallel without compromising the completeness of the result. A similar approach was used in [25] to partition inductive proofs for register correspondence.

The rest of the paper is organized as follows. Section 2 describes further background and relations with previous work. Section 3 describes the algorithms used for inductive strengthening. Section 4 discusses application to logic synthesis. Section 5 reports experimental results. Section 6 concludes the paper and outlines future work.

2 Background and Related Research

A *Boolean network* is a directed acyclic graph (DAG) with nodes corresponding to logic gates and directed edges corresponding to wires connecting the gates. The terms Boolean network, design and circuit are used interchangeably in this paper.

A node n has zero or more *fanins,* i.e. nodes that are driving n, and zero or more *fanouts*, i.e. nodes driven by n. The *primary inputs* (PIs) are nodes without fanins in the current network.

A *fanin* (*fanout*) *cone* of node n is a subset of all nodes of the network, reachable through the fanin (fanout) edges from the given node. A *topological order* of nodes in the network is any order in which any node appears later in the order than any of its fanins.

If the network is *sequential*, the memory elements are assumed to be D-flip-flops. The terms memory elements, flip-flops, and registers are used interchangeably in this paper. The registers are assumed to have a fixed binary initial state. If a register has an unknown or a don't-care initial state, it can be transformed to have 0-initial state by adding a new PI and a MUX controlled by a special register that produces 0 in the first frame and 1 afterwards. So without loss of generality, we consider only registers with a 0 initial state. *The set of reachable states* includes the initial state and all the states reachable from it by any input sequence.

An *And-Inverter Graph* (AIG) is a Boolean network composed of two-input ANDs and inverters represented as complemented attributes on the edges.

A *cut C* of node n, called *root*, is a set of nodes of the network, called *leaves*, such that each path from a PI to n passes through at least one leaf. A cut is *m-feasible* if its size does not exceed m and is *dominated* if it contains a cut of the same root.

A *literal* of a Boolean variable is the variable or its complement. Given a set of variables x (e.g. the set of leaves of cut C), a *minterm* is a product of literals, one for each variable in the set. The complement of a minterm is a *clause*. A product of clauses is a *Conjunctive Normal Form* (CNF). *Boolean satisfiability* (SAT) is the problem of determining whether a variable assignment exists that will cause a CNF to evaluate to 1. A *SAT-based method* is a method that reduces a given problem to SAT and solves it using a SAT solver.

Simulation is a way of computing node values in a circuit under given input values. *Random simulation* uses random or biased random input values. Simulation assigns values at the inputs and evaluates the internal nodes in a topological order. Simulation is typically performed *bitwise*, where 32 or 64 input patterns are evaluated using a single machine operation. Simulation information of a node is stored in a bit-string composed of many machine words. It is computed by a sequence of bitwise operations using the simulation information of the node's fanins.

A *combinational invariant* is a relation among arbitrary signals in the network that holds in all states. A *sequential invariant* is a relation that holds in all reachable states, but possibly fails in one or more unreachable states. A sequential invariant can be seen as a characterization of a set of states for which it holds, a set that includes the set of reachable states and possibly some unreachable states. An example sequential invariant is equality among two registers outputs that does not hold combinationally (in all states) but holds sequentially (in all reachable states).

A *candidate sequential invariant* is an invariant that has not yet been proved (e.g. by induction or interpolation) but is suspected to hold (e.g. after several rounds of simulation). Such invariants express properties that should be proved, e.g. mutual exclusion of the values at two primary outputs. *Model checking* focuses on proving user-specified properties.

Induction is often used to prove sequential invariants. Its use for sequential designs was pioneered in [15] and further developed in [5][26][17]. A sequential invariant is *inductive* when: (*base case*) it holds in the initial state, and (*inductive case*) if it holds in a state, then it holds in all states reachable from that state in one transition. An invariant provable by induction is known as an inductive invariant.

In an efficient implementation of induction, the base and the inductive cases are formulated as SAT instances and solved by a SAT solver. The solution is incremental because while each property in the invariant is checked independently, the same solver instance can be used in all the checks. This can be done efficiently using an incremental interface of a modern SAT solver [13]. More details on SAT-based induction can be found in [25].

The set of all reachable states is an inductive sequential invariant. However, not every sequential invariant is inductive. For example, consider an unreachable state s that has a transition into it from state t. The complement of the minterm composed of register variables representing s is a sequential invariant because it holds in all reachable states. This invariant is not inductive because it holds in t but not in s. The state failing the inductive case (in this case t) is called an *induction leak*.

Several ways of strengthening induction are known:

- Extending simple induction to k-step induction ($k > 1$) [15].
- Using unique-state constraints [12].
- Using equivalences expressed over register outputs (register correspondence) or over arbitrary signals in the network (signal correspondence) [15][26][17][25].
- Applying signal correspondence after timeframe expansion, hoping this will capture additional equivalences among signals across different timeframes in the original design.
- Using implications of signals in the network [5][9][10].
- Using p-th invariants, that is, invariants that hold starting from frame p from the initial state ($p > 1$) [16].
- Incrementally computing inductive clauses in terms of register variables using counter-examples to induction [7].

We propose using inductive strengthening based on generating an invariant in the form of a set of m-literal clauses. The method is a generalization of [9], [10], and [7], as shown in Section 3.4.

3 Computing inductive invariants

This section presents a new algorithm for computing inductive invariants in a sequential network. The algorithm is presented for AIGs but it is equally applicable to general logic networks.

The overall pseudo-code of the algorithm is shown in Figure 3. Details are given in the subsections listed in the parentheses.

The computation starts by enumerating for each node a subset of m-cuts using procedure **aigEnumerateCuts** (Section 3.1).

Next, two rounds of simulation are performed. For each cut of size m, all 2^m value assignments of the cut leaves give a set of minterms that will be tested with simulation. Simulation information is used to determine if each minterm is likely to appear only under unreachable states, and for minterms that likely only hold in unreachable states, the complement gives a candidate invariant clauses. A number of candidate clauses are collected and filtered using simulation information in the procedure **aigComputeCandidates** (Section 3.2).

A set of clauses representing the candidate invariant is checked by the base case and then by an iterative refinement procedure **performInductiveCase,** similar to that of van Eijk [15]. When this procedure terminates, the conjunction of the set of remaining

clauses, if it is non-empty, represents an inductive invariant. If strengthening is not sufficient (determined by procedure **checkSufficient** whose definition is application-specific), another round of invariant computation is performed where the candidates considered are those not contained in the already proved set. As a result, new invariants that are proved provide ever tighter approximations of the state space (Section 3.3).

Figure 3.1 gives the top-level pseudocode for our proposed procedure, and Figure 3.2 illustrates the discovery of candidate sequential invariants on an example circuit.

```
set of clauses computeInvariants( aig N, parameters P )
{
    // compute m-cuts for all nodes
    set of cuts Cuts = EnumerateCuts( N, P );

    // perform two rounds of simulation
    simulation patterns Comb = SimulateComb( N, P );
    simulation patterns Seq = SimulateSeq( N, P );

    // iterate while the set of clauses is not sufficient
    set of clauses S = ∅;
    while ( !checkSufficient( S ) ) {

        // compute candidate clauses
        clauses C = ComputeCandidates( N, P, Cuts, Comb, Seq );

        // refine the candidates using the base case
        C = PerformBaseCase( C, N, P );

        // refine the candidates using van Eijk's loop
        do {
            C = PerformInductiveCase( C, N, P );
        } while ( CheckChanges( C ) );

        // add newly proved invariant to the set
        S = S ∪ C;
    }
    return S;
}
```

Figure 3.1. Pseudo-code for computing inductive invariants.

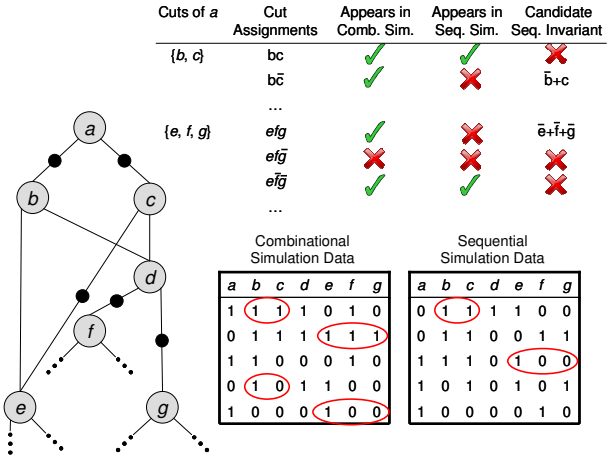

Figure 3.2. Example candidate sequential invariant derivation. On the sample AIG, all nodes are AND gates and marks denote complemented edges.

3.1 Cut computation

For two sets of size m or less cuts A and B, the operation $A \lozenge B$ is defined as:

$$A \lozenge B = \{\, u \cup v \mid u \in A, v \in B, |u \cup v| \le m \,\}.$$

Let $\Phi(n)$ denote the set of m-feasible cuts of node n. If n is an AND node, let n_1 and n_2 denote its fanins. $\Phi(n)$ is computed using the sets of cuts of its fanins:

$$\Phi(n) = \begin{Bmatrix} \{\{n\}\} & : n \in \text{PI} \\ \{\{n\}\} \cup \Phi(n_1) \lozenge \Phi(n_2) & : \text{otherwise} \end{Bmatrix}.$$

Performing cut computation for the nodes in a topological order guarantees that the fanin cuts, $\Phi(n_1)$ and $\Phi(n_2)$, are available when the node cuts, $\Phi(n)$, are computed. The set of computed cuts is filtered by removing dominated cuts. This reduces runtime and memory without sacrifiying the expressiveness of cuts computed.

The above complete cut enumeration [25] is practical for small m ($m < 6$) because the number of cuts is approximately linear in the size of the circuit. For larger m, the above procedure can be supplemented with a method to compute a subset of all m-cuts meeting some criteria. These cuts are called *priority cuts* [24]. The criterion used to prioritize the cuts for invariant computation is to prefer cuts with a larger average number of fanouts of the leaves of a cut. A similar criterion was used in [14].

In our implementation, m is parameterizable in order to give the user control over the expressiveness and the number of the m-feasible cuts. For all benchmarks discussed in Section 5, we find m=4 to be sufficient.

In the example of Figure 3.2, an example AIG is shown along with two sample cuts for the node *a*.

3.2 Collecting candidates

To help form candidate invariants from cuts, two rounds of simulation are performed: combinational and sequential. *Combinational simulation* assumes random values at the primary inputs and register outputs, which are treated as additional primary inputs. *Sequential simulation* assumes random values at the primary inputs while the register outputs are set to the initial state. This sequential simulation iterates over the circuit several times, setting the register outputs to the register inputs computed on the previous step, thereby accumulating simulation data for many reachable states. The combinational and sequential simulation differ in the assignments that are made to the registers; combinational simulation produces values under any state while sequential produces values under reachable states.

Candidate clauses are collected by considering the m-cuts of all nodes in the AIG. Each node has two types of simulated minterms. A cut is analyzed to determine what values appear at the cut inputs. Suppose assignment $\tilde{x}_0\tilde{x}_1 \ldots \tilde{x}_{m-1}$ appears N times at the cut inputs under combinational simulation but does not appear under sequential simulation. This indicates that this assignment may be produced at least N states and the assignment is likely not produced in any reachable states. Thus, the complement of this assignment, the clause $\bar{\tilde{x}}_0 \vee \bar{\tilde{x}}_1 \vee \ldots \vee \bar{\tilde{x}}_{m-1}$, excludes many unreachable states and is likely true for all reachable states. All such clauses are accumulated and used as candidates.

An example of this method is illustrated in Figure 3.2. For the cuts $\{b, c\}$ and $\{e, f, g\}$ the assignments $b\bar{c}$ and efg were seen to occur in combinational simulation and hence are not vacuous. The same assignments were not seen in the sequential simulation and likely cannot be produced on the reachable states. Complementing these assignments gives two candidate clauses.

It should be noted that neither the combinational or sequential simulation includes all possible states. Because the combinational simulation is not exhaustive, some minterms may incorrectly be classified as vacuous and excluded from the set of candidate

clauses. This affects the number of clauses we can prove, but it does not affect the correctness of the overall method. Moreover, since such minterms do not readily appear under combinational simulation, they are not likely to substantially refine the characterization of the state space. Likewise, the sequential simulation is also not exhaustive. This causes minterms to be promoted to candidate clauses while they may not hold for all reachable states. This is not a problem because the candidate clauses will be refined with induction.

Except for small circuits and small cut sizes, the number of candidate clauses can be large. For example, on a circuit with 1K registers and 15K AIG nodes, there may be 50K candidate clauses computed using the set of all 4-cuts. In such cases, the invariants can be filtered by the following heuristic: if a candidate clause is falsified by a large number of vectors in the combinational simulation then it is likely to characterize a large number of unreachable states. The number of combinational simulation vectors that can falsify a clause determine its score, and our implementation has a user-controlled parameter which limits the number of the highest-scoring clauses considered. This heuristic plays an important role in selecting useful candidates. In our experiments it was sufficient to limit the candidates to the 5000 highest-scoring.

The set of candidate clauses can lead to a stronger inductive invariant if it is supplemented with the candidate clauses expressing one-hotness conditions. These conditions are two-literal clauses involving register outputs and can be easily computed using sequential simulation information. Most of these additional clauses cannot be collected as candidates using cuts because cuts include literals in the vicinity of a particular node, while one-hotness, if applicable, can relate registers that are far apart. We found that adding the candidate one-hotness conditions often improves the performance of the algorithm. One reason for this is that many industrial designs use one-hot encoding for at least some of the registers.

3.3 Proving candidates

The well-known van Eijk procedure [15] is used to process the candidates and prove some of them. First, those candidates that do not satisfy the base case are removed. Second, the inductive case is performed by asserting the clauses in the first frame and checking them in the next frame. The counter-examples are used to refine the remaining candidate invariants. The failing clauses are removed and refinement is iterated until a fixed point is reached. If non-empty, the conjunction of the clauses in this fixed point represents an inductive invariant.

To derive a sufficiently tight invariant, the van Eijk procedure can be sequentially applied to several different sets of candidate invariants. An invariant proved in a run is assumed in the next run. Since the proved clauses form an invariant, there is no need to re-prove them; only new clauses need to be proved. This results in accumulating clauses, which increasingly refine the invariant. New candidate clauses are collected only if they refine the current invariant. If cuts of the given size do not yield additional clauses, the cut size can be increased to find new candidates to continue refining the invariant. This strengthening enhances van Eijk's procedure and allows tighter invariants to be found efficiently.

The "sufficiency" of the resulting invariant depends on the application. In model checking, it is sufficient if the invariant implies the target property. In logic synthesis, it is sufficient if it contains "enough" flexibility to do substantial logic restructuring.

In model checking, the procedure can stop as soon as the proved invariant implies the target property. For this, the target property is added to the set of candidate clauses. If the property remains in the fixed point, it is proved. Otherwise, a new set of clauses is considered that provides a tighter approximation of the reached state set and has a better chance to prove the target property.

3.4 Comparison with previous work

For a description of other SAT-based approaches to model checking, refer to [28] and for an overview of recent work in induction strengthening refer to [10][7].

The proposed approach can be seen as a generalization of three previous approaches [9][10][7]. The following is a comparison:

- Computation of *m*-cuts scales better than that of Boolean implications between signal pairs because priority cuts [24] only take linear-time in circuit size to compute while computing implications takes quadratic-time [10].
- The *m*-literal clauses have more expressive power than the Boolean implications of [9][10] which are essentially are two-literal clauses.
- Our flexible framework for inductively proving groups of *m*-literal clauses is similar to [9], with novel heuristics to prioritize clauses according to their expressive power.
- The *m*-clauses are computed in terms of internal variables rather than register outputs as done in [7], which increases the expressive power of the invariants.
- The *m*-clause candidates are computed by simulation rather than from counter-examples as done in [7], which is less time-consuming and avoids the risk of not having inductive sub-clauses.
- The inductive proof for *m*-clauses, with the cuts limited to a few levels from the register outputs, can use partitioning similar to [25] which increases the possibility that the proposed approach works for designs of any size.
- Adding signal-correspondence and one-hotness invariants, which was not used in [9][10][7] gives additional strength to the proposed approach.

4 Application to Logic Synthesis

The inductive invariants proved by this method compactly represent unreachable state information useful as flexibility in circuit restructuring during logic synthesis with don't-cares [21].

The following are advantages of this approach compared to using other types of sequential flexibility:

- **Complete set of unreachable states**

Except for small circuits, the reachable state set is hard or impossible to obtain. BDD-based methods for computing this set mostly fail on circuits with more than a 50-100 registers. Another disadvantage is that, if the unreachable state information represented with BDDs is used in sequential synthesis, sequential equivalence checking (SEC) is very hard because it doubles the number of registers. In contrast, when the proposed invariants are used, sequential verification tends to be easier because the inductive nature of the invariants tends to increase inductiveness of the associated SEC problems.

- **Equivalences in terms of internal signals**

Signal equivalences in terms of internal signals (signal correspondences) have been shown to be a powerful vehicle for capturing sequential flexibility. Sequential synthesis based on this flexibility can lead to substantial reductions in area and register count [25]. However, the best use of this flexibility for circuit restructuring, is to collapse the equivalent nodes into a single node and remove the others. This reduces the circuit but does not allow for a more fine-grain circuit restructuring

afforded by the *m*-cut invariants. This is why signal equivalence should be computed and used as a preprocessing step before using the proposed inductive invariant.

- **Implications in terms of internal signals**

Signal implications among internal signals provide additional expressive power, compared to signal equivalences and can be useful in logic synthesis [9]. Detection of implications can be done similarly to the proposed invariants, using simulation information. However, *m*-literal clauses are more expressive compared to implications (2-literal clauses). In addition, collecting implications is harder and may require a procedure quadratic in the number of nodes, while collecting *m*-literal clauses is linear when priority cuts are used.

5 Experimental Results

The proposed algorithms are implemented in ABC [1] as command *indcut*. The SAT solver is a modified version of MiniSat-C_v1.14.1 [13]. The workstation used has two dual-core AMD Opteron 2218 CPUs with 16GB RAM, and runs x86_64 GNU/Linux. Only one core was used in the experiments.

Experiments were performed using two suites of model checking benchmarks: a set of PicoJava II benchmarks [19] and the TIP benchmarks [12]. Other benchmark suites from the model checking competition [4] were also evaluated: (a) the TIP benchmarks, (b) the AMBA benchmarks (all *unsat*), and (c) the L2S benchmarks (9 *unsat* cases). The unsat cases from the latter two suites could be solved easily using signal correspondence (ABC command *ssw*) [25] after combinational synthesis (ABC command *dcompress2*). Since the proposed algorithm is developed as a method to be applied when other methods fail, we do not report its performance on the AMBA and L2S suites.

Before using the proposed algorithm to solve the properties, the benchmarks were first heavily synthesized: 1) sequentially constant or structurally redundant latches were removed, 2) combinational synthesis, 3) removal of redundant latches via latch correlation analysis, 4) combinational synthesis, 5) sequential signal correspondence, 6) combinational synthesis, 7) inductive cut computation + sequential don't-care based resynthesis. This set of synthesis operations was necessary to reduce the design complexity and aid our later proof.

ABC command *indcut* was used in all reported experiments with the following default set of parameters: induction depth (K = 1), cut size (M = 4), the limit on the number of candidate clauses collected (C = 5000), the maximum level of the nodes whose cuts are considered (L = 8), the number of times invariant computation was iterated (B = 1).

5.1 PicoJava benchmarks

The complete set of PicoJava benchmarks includes 20 *unsat* problems. After the preprocessing steps outlined above, 9 out 20 problems were already solved.

The remaining benchmarks were preprocessed and then solved by command *indcut*. Design statistics both before and after preprocessing are shown in the columns "Original Design" and "Preprocessed Design," respectively. "Clauses Proved" gives the total number of clauses proved by *induct*. No more than 5000 clauses can be proved because this was the imposed limit on the number of candidates. Finally, the column "Runtime" gives the time in seconds needed to prove the property with *induct*.

Table 5.1 shows that on average 75% of the 5000 candidate invariants are proved by *induct*. Although the set of proved clauses is incomplete, it was sufficient to imply the target property for all of the considered problems.

5.2 TIP benchmarks

These benchmarks are among the smallest and the most well-studied model checking benchmarks [12]. The original set of 158 testcases includes both *sat* and *unsat* problems. First, this set was filtered by removing all problems provable by signal correspondence with induction depth K = 4 (ABC command *ssw -F* 4) or disproved by BMC of depth 100 (ABC command *bmc –F* 100). This led to a subset containing 51 "hard" TIP problems.

Applying *indcut* with default settings these 51 problems solved 41 of them, with runtime for each benchmark not exceeding 1 second. Interestingly, some of the benchmarks solved by *indcut* could also be proved by signal correspondence with very large induction depth. Thus, *cmu_periodic_N* could be proved by *ssw -F* 96 (*K* = 96) in 30 sec, while *indcut* solved it in 0.2 sec.

The 10 remaining benchmarks not solved by *indcut* are: *cmu_dme*1*_B, cmu_dme*2*_B, irst_dme*4*_B, irst_dme*5*_B, irst_dme*6*_B, nusmv_dme*1-16*_B, nusmv_dme*2-16*_B, texas_two_proc*_6*_E, vis_coherence*_3*_E, vis_coherence*_4*_E*. These benchmarks could not be solved by *indcut* even when we modified the default set of parameters. In all cases, a subset of clauses was proved inductively, but the resulting invariant was not sufficient to imply the target property, while other candidate clauses implying it were not inductive. We believe that none of the model checkers submitted to the model checking competition [4] were able to solve these 10 benchmarks.

6 Conclusions and Future Work

This paper proposes a new method for inductively strengthening the model checking of safety properties. The method supplements existing methods and is useful for proving hard *unsat* problems.

In combination with other synthesis and verification algorithms implemented in ABC, the proposed method solved 334 of the 344 benchmarks from the model checking competition [4]. The remaining 10, plus another 27 of the 344 problems solved by the proposed method, were not solved by any of the entrants in the 15 minutes allowed for each example. The hardest PicoJava example took less than 7 seconds.

In summary, the contributions of this paper are:

- A new efficient method for expressing candidate invariants using *m*-clauses formulated for the nodes in the circuit.
- A scalable hierarchical approach to proving the candidate invariants, which trades off computational effort for the number and expressiveness of invariants generated.
- Experiments using several benchmark suites to show that the proposed method can solve many difficult problems.

Future work will include:

- Further experiments and fine tuning using benchmarks contributed by industrial collaborators.
- Integrating the induction strengthening engine into robust equivalence and model checkers.
- Using the computed invariant clause sets as don't-cares for circuit restructuring in logic synthesis.
- Performing direct comparison with industrial model checkers.

Acknowledgements

This work was supported in part by SRC contracts 1361.001 and 1444.001, NSF grant CCF-0702668, and the California Micro Program with industrial sponsors Actel, Altera, Calypto, Intel, Magma, Synopsys, Synplicity, and Xilinx.

References

[1] Berkeley Logic Synthesis and Verification Group. *ABC: A System for Sequential Synthesis and Verification*. Release 70930. http://www-cad.eecs.berkeley.edu/~alanmi/abc

[2] A. Biere, C. Artho, V. Schuppan, "Liveness checking as safety checking". Proc. Intl. Workshop on Formal Methods for Industrial Critical Systems (FMICS'02), ENTCS, Vol. 66(2).

[3] A. Biere. *AIGER format and toolbox*. http://fmv.jku.at/aiger/

[4] A. Biere and T. Jussila. *Hardware model checking competition at CAV'06*. http://fmv.jku.at/hwmcc/

[5] P. Bjesse and K. Claessen. "SAT-based verification without state space traversal". *Proc. FMCAD'00.* LNCS, Vol. 1954, pp. 372-389.

[6] P. Bjesse and J. Kukula, "Automatic generalized phase abstraction for formal verification", *Proc. ICCAD'06*, pp. 1076-1082.

[7] A. R. Bradley and Z. Manna, "Checking safety by inductive generalization of counterexamples to induction", *Proc. FMCAD '07*.

[8] R. Brayton, "The synergy between logic synthesis and equivalence checking", *Keynote at FMCAD'07*. http://www.cs.utexas.edu/users/hunt/FMCAD/2007/presentations/fmcad07_brayton.ppt

[9] M. L. Case, A. Mishchenko, and R. K. Brayton, "Inductively finding a reachable state space over-approximation", *Proc. IWLS '06*, pp. 172-179. http://www.eecs.berkeley.edu/~alanmi/publications/2006/iwls06_inv.pdf

[10] M. L. Case, A. Mishchenko, and R. K. Brayton, "Automated extraction of inductive invariants to aid model checking", *Proc. FMCAD '07*, pp. 165-172. http://www.eecs.berkeley.edu/~alanmi/publications/2007/fmcad07_ind.pdf

[11] E. Clarke, O. Grumberg, and D. Peled. *Model checking*. MIT Press, 1999.

[12] N. Een and N. Sörensson, "Temporal induction by incremental SAT solving", *Proc. BMC'03*, ENTCS, Vol. 89(4).

[13] N. Een and N. Sörensson, "An extensible SAT-solver". *SAT '03*. http://www.cs.chalmers.se/Cs/Research/FormalMethods/MiniSat

[14] N. Een, "Cut sweeping", *Cadence Technical Report 2007*. http://minisat.se/downloads/CutSweeping.ps.gz

[15] C. A. J. van Eijk. "Sequential equivalence checking based on structural similarities", *IEEE TCAD*, Vol. 19(7), July 2000, pp. 814-819.

[16] F. Lu and K.-T. Cheng. "Sequential equivalence checking based on *k*-th invariants and circuit SAT solving". *Proc. HLDVT'05*.

[17] F. Lu and T. Cheng. "IChecker: An efficient checker for inductive invariants". *Proc. HLDVT '06*, pp. 176-180.

[18] A. Kuehlmann, V. Paruthi, F. Krohm, and M. K. Ganai, "Robust boolean reasoning for equivalence checking and functional property verification", *IEEE Trans. CAD*, Vol. 21(12), 2002, pp. 1377-1394.

[19] K. L. McMillan and N. Amla, "Automatic abstraction without counterexamples." *Proc. TACAS '03*, LNCS, Vol. 2619, Springer, pp. 2-17.

[20] K. L. McMillan. "Interpolation and SAT-Based model checking". *Proc. CAV'03*, pp. 1-13.

[21] A. Mishchenko and R. Brayton, "SAT-based complete don't-care computation for network optimization", *Proc. DATE '05*, pp. 418-423.

[22] A. Mishchenko, S. Chatterjee, and R. Brayton, "DAG-aware AIG rewriting: A fresh look at combinational logic synthesis", *Proc. DAC '06*, pp. 532-536. http://www.eecs.berkeley.edu/~alanmi/publications/2006/dac06_rwr.pdf

[23] A. Mishchenko, S. Chatterjee, R. Brayton, and N. Een, "Improvements to combinational equivalence checking", *Proc. ICCAD '06*, pp. 836-843 http://www.eecs.berkeley.edu/~alanmi/publications/2006/iccad06_cec.pdf

[24] A. Mishchenko, S. Cho, S. Chatterjee, and R. Brayton, "Combinational and sequential mapping with priority cuts", *Proc. ICCAD '07,* pp. 354-361. http://www.eecs.berkeley.edu/~alanmi/publications/ 2007/iccad07_map.pdf

[25] A. Mishchenko, M. Case, R. Brayton, and S. Jang, "Scalable and scalably-verifiable sequential synthesis", *Submitted DAC'08*. http://www.eecs.berkeley.edu/~alanmi/publications/2008/dac08_vss.pdf

[26] H. Mony, J. Baumgartner, V. Paruthi, and R. Kanzelman. "Exploiting suspected redundancy without proving it". *Proc. DAC'05*, pp. 463-466.

[27] P. Pan and C.-C. Lin, "A new retiming-based technology mapping algorithm for LUT-based FPGAs," *Proc. FPGA '98*, pp. 35-42.

[28] M. Prasad, A. Biere, and A. Gupta. "A survey of recent advances in SAT-based formal verification", *Intl. Journal on Software Tools for Technology Transfer (STTT)*, Springer 2005, Vol. 7 (2), pp. 156-173. http://fmv.jku.at/papers/PrasadBiereGupta-STTT-7-2-2005.pdf

[29] E. Sentovich et al. "SIS: A system for sequential circuit synthesis". *Tech. Rep. UCB/ERI, M92/41*, ERL, Dept. of EECS, UC Berkeley, 1992.

Table 5.1. Experimental results for proving unsatisfiability on the PicoJava benchmarks [19].

Example	Original Design			Preprocessed Design			Indcut Performance	
	PI	Reg	AIG	PI	Reg	AIG	Clauses Proved	Runtime, sec
pj006	1277	703	17542	1277	332	16160	4732	5.11
pj007	396	314	7224	396	108	6040	1183	2.60
pj008	446	338	7555	446	139	6142	4659	2.90
pj009	336	269	6844	336	76	5566	2555	3.25
pj010	366	295	7493	366	89	6705	3261	2.53
pj015	1322	775	18964	1322	370	18607	4338	6.78
pj016	1190	671	17000	1190	303	16149	4675	5.11
pj019	476	383	10467	476	64	7884	4311	4.48