

Low Power State Assignment Using ILP Techniques

Assim Sagahyoon^{#1}, Fadi A. Aloul^{#2}, Alexander Sudnitson^{*3}

[#]*Department of Computer Science & Engineering, American University of Sharjah, UAE*

¹asagahyoon@aus.edu

²faloul@aus.edu

^{*}*Department of Computer Engineering
Tallinn University of Technology, Estonia*

³alsu@cc.ttu.ee

Abstract - State assignment for finite state machines is a critical optimization problem in the synthesis of sequential circuits. In this paper we address the state assignment problem from a low power perspective. We experiment with Boolean Satisfiability and Integer Linear Programming techniques to solve the assignment problem where the primary goal is the reduction of switching activity during state transitions. We also detect and evaluate the use of symmetries in speeding up the search process. These techniques have been applied to the MCNC benchmark circuits and yielded promising results.

Keywords - State Assignment, Power, Integer Linear Programming, Boolean Satisfiability

I. INTRODUCTION

With the ever increasing integration scale, power consumption has emerged as a major design constraint for integrated circuits. During logic synthesis of sequential circuits, i.e. finite state machines (FSMs), assigning binary codes to each state in the circuit is a critical step in the low power design of the circuit. The state encoding problem for low power has been explored by a number of researchers. Recent work has attempted to address this issue using various techniques. In [7], using a probabilistic description of the circuit, a state assignment algorithm that minimizes the Boolean distance between the codes of the state with high transition probability was proposed. The primary objective was to reduce the average switching activity of the input and output state variables by minimizing the number of bit changes during state transition. Symbolic methods for re-encoding a circuit to reduce the dissipated power are discussed in [17]. The authors experimented with two encoding strategies. One based on recursive weighted no-bipartite matching, and one on recursive mincut bi-partitioning. The authors conclude that even though the results are promising, their synthesis procedure needs refinements in both circuit transformation and the power estimation phases. Encoding of states for a given user-specified input sequence was addressed in [18]. Simulation was used to determine the relative frequency for all state transitions. By defining and using a register switching rate as a cost, the authors then used simulated annealing to solve the state assignment problem. In [12], given an FSM description and the input probabilities, the total state transition probabilities for each

edge in a state transition graph of the circuit are computed first. This is done by modeling the FSM as a Markov chain. A cost function that is the summation of the products of distinct code words with the Hamming distance between any two adjacent states is used to measure the effectiveness of the assignments. Similar to [7], the goal is to assign codewords with minimum Hamming distances to states with higher transition probabilities. A spanning tree based encoding approach was tested in [25]. The assignment problem was formulated as a hypercube embedding problem, where the embedding process is directed by a maximum spanning tree of the attraction graph of the FSM. In [6], heuristic techniques were used to visit the state transition graph of the circuit and assign a priority to the symbolic states. Next, an encoding technique that follows the priority established in the first part is used to assign binary codes to the states. As assignment method that utilized dynamic loop information extracted from FSM profiling data was presented in [16]. The authors then experimented with three different loop-based state assignment algorithms, depth-first search, loop-based depth first and per state encoding. An m -block partitioning technique for the state assignment to reduce the number of feedback cycles and keep low switching activities among state variables is discussed in [26]. The objective was both, to improve power consumption and testability of the circuit. Re-engineering (targeting low power) an FSM by constructing a functionally equivalent but topologically different design based on an optimization objective was presented in [31]. The authors argue that this will allow the exploration of a larger solution space making it possible to obtain solutions better than the optimal ones for the original circuit. They used genetic algorithms and heuristics to re-engineer existing low power state encoding procedures with favorable results.

The last decade have seen a remarkable growth in the use of Boolean Satisfiability (SAT) models and algorithms for solving various problems in Electronic Design Automation (EDA). This is mainly due to the fact that SAT algorithms have seen tremendous improvements in the last few years, allowing larger problem instances to be solved in different applications domains [5, 8, 22, 27]. Such applications include formal verification [9], FPGA routing [23], global routing [4], logic synthesis [21], and power optimization [28].

SAT solvers have traditionally been used to solve *decision* problems. Given a set of Boolean variables and constraints expressed in products-of-sum form (also known as *conjunctive normal form* (CNF)), the goal is to identify a variable assignment that will satisfy all constraints in the problem or prove that no such assignment exists. Recently, SAT solvers have been extended to handle Pseudo-Boolean (PB) constraints [4, 10, 14, 15, 29]. PB constraints are more expressive and can replace an exponential number of CNF constraints. Another key advantage of PB constraints is the ability to express *optimization* problems which are traditionally handled as integer linear programming (ILP) problems. Hence, SAT solvers can now handle both decision and optimization problems and compete with the best available generic ILP solvers such as IBM-ILOG CPLEX [19].

In this paper, we propose using SAT-based and generic ILP solvers to handle the state assignment problem. We formulate the problem as an optimization instance and use SAT-based ILP solvers to find a solution. The problem is also solved using generic ILP solvers, e.g. CPLEX [19], for performance comparison. Interestingly, presented empirical results show that SAT-based ILP solvers outperform generic commercial ILP solvers when solving such instances.

Previous work has shown that breaking symmetries in SAT 0-1 ILP formulas effectively prunes the search space and can lead to significant runtime speedups [1]. The idea is that breaking symmetries prevents symmetric images of search paths from being searched, thus pruning the search tree [2, 11]. Since industrial instances, in general, have some structure, they are likely to have symmetries.

In this paper, we statically detect and break the symmetries in the generated SAT 0-1 ILP instances and evaluate the advantage of that on SAT-based and generic ILP solvers. Our results indicate the presence of symmetries in the tested instances. Furthermore, as suggested by previous work [30], SAT-based ILP solvers perform better on instances with broken symmetries while generic ILP solvers show mixed performance.

The rest of the paper is organized as follows: Section 2 provides background information on the state assignment problem and Boolean satisfiability; Section 3 explains how to formulate the state assignment problem as a SAT 0-1 ILP problem. Section 4 gives a summary of the experimental results and the paper is concluded in Section 5.

II. BACKGROUND

In this section, we review the state assignment problem and we introduce some key concepts related to Boolean satisfiability. This will serve in clarifying our proposed approach to solve the state assignment problem.

A. The State Transition Graph

The state assignment problem entails the codification of states in an FSM, and is a known NP-complete problem [24]. In [7], a brief attempt was made to formulate the state encod-

ing as an integer linear programming (ILP) problem. However, at the time there were no commercial or state-of-the-art solvers available and the authors only briefly discussed a framework to formulate the problem.

In this work, prior to formulating the problem as an ILP instance, the state transition graph (STG) for each benchmark circuit is developed. That is, given the FSM description and the input probabilities; we compute the transition probabilities for each edge in the STG, by modeling the FSM as a Markov chain. Thereafter, all unreachable states and self-loops are eliminated from the graph. The STG is then transformed into an undirected graph converting all multiple edges into a single undirected edge with weights that equals to the sum of the directed edge probabilities [7, 6]. Thus the weights on the edges are proportional to the total probability of transition between every two connected states.

B. Boolean Satisfiability

Boolean Satisfiability (SAT) is often used as the underlying model in the field of computer aided designs of integrated circuits. A number of SAT solvers have been proposed and implemented [5, 8, 22, 27]. These solvers employ powerful algorithms that are sufficiently efficient to deal with large-scale SAT problems that typically arise in the design automation domain. Most of these algorithms claim competitive results in runtime efficiency and robustness.

In SAT, given a formula f , the objective is to identify an assignment to a set of Boolean variables that will satisfy a set of constraints. If an assignment is found, it is known as a satisfying assignment, and the formula is called satisfiable. Otherwise if an assignment doesn't exist, the formula is called unsatisfiable. The constraints are typically expressed in conjunctive normal form (CNF). In CNF, the formula consists of the conjunction (AND) of m clauses $\omega_1, \dots, \omega_m$ each of which consists of the disjunction (OR) of k literals. A literal l is an occurrence of a Boolean variable or its complement. Hence, in order to satisfy a formula, each of its clauses must have at least one literal evaluated to true.

As an example, a CNF instance $f(a, b, c) = (a + \bar{b}) \cdot (a + b + c)$ consists of 3 variables, 2 clauses, and 5 literals. The assignment $\{a = 0, b = 1, c = 0\}$ leads to a conflict, whereas the assignment $\{a = 0, b = 0, c = 1\}$ satisfies f .

Despite the problem being NP-Complete, there have been dramatic improvements in SAT solver technology over the past decade. This has led to the development of several powerful SAT solvers that are capable of handling problems consisting of thousands of variables and millions of constraints [5, 8, 22, 27].

Recently, SAT solvers [4, 10, 14, 15, 29] have been extended to handle pseudo-Boolean (PB) constraints which are linear inequalities with integer coefficients that can be expressed in the normalized form [4] of:

$$a_1x_1 + a_2x_2 + \dots + a_nx_n \geq b \quad (1)$$

where $a, b \in Z$ and x_i are Boolean variables. PB constraints can, in some cases, replace an exponential number of CNF constraints. They have been found to be very efficient in expressing “counting constraints” [4]. Furthermore, PB extends SAT solvers to handle *optimization* problems as opposed to only *decision* problems. Subject to a given set of CNF and PB constraints, one can request the minimization (or maximization) of an objective function which consists of a linear combination of the problem’s variables. Note that each CNF constraint can be viewed as a PB constraint. For example the CNF constraint $(\bar{a} + b)$ can be viewed as the PB constraint $\bar{a} + b \geq 1$. PB constraints represent 0-1 ILP inequalities.

Detecting and breaking symmetries in SAT instances have been shown to help prune the search space explored by a SAT solver. The basic framework for utilizing symmetries was proposed in [11], and later extended by [2, 3] to account for phase-shift symmetries and consider only generators of the group of symmetries. In [1], the authors extended previous work to detect and break symmetries in SAT PB, i.e. 0-1 ILP, instances.

Symmetries in a SAT 0-1 ILP instance are first *detected* by reduction to graph automorphism and then *broken* by adding symmetry breaking predicates (SBPs) to the formulation. In the graph automorphism step, the instance is represented by a graph and the automorphism problem for that graph is solved using graph automorphism software packages, such as saucy [13]. SBPs, representing the generators of group of symmetries, are then added to the SAT instances in CNF clause format.

III. PROBLEM FORMULATION AND IMPLEMENTATION

In this section we show how to formulate the state assignment problem as a 0-1 ILP instance. We used an approach similar to the one discussed in [7]. The goal is to find the state assignment that leads to the minimum number of weighted transitions. To illustrate the approach, assume an FSM with x states and a defined weight between every two connected states. The number of bits, referred to as B , needed to encode each state is $\log_2 x$. The objective is to find a unique state assignment to each state while minimizing the weighted hamming distance between the adjacent states.

Two sets of variables are defined for the problem:

- A Boolean variable b_i^l that represents bit l of state i . A total of $x \log_2 x$ variables are defined. A value of 1 (0) for each variable indicates that the corresponding bit is a 1 (0) in the original problem.
- A Boolean variable $e_{i,j}^l$ that represents the XOR operation between the l bits of states i and j . A total of $\binom{x}{2} \log_2 x$ variables are defined. A value of 1 (0) for each variable indicates that the corresponding bit assignments are different (similar) in the original problem.

The following set of constraints are generated:

- The XOR relation between the e and b variables for all state bits must be defined using the following constraint:




Fig. 1. An illustrative example showing an FSM with 4 states and 4 edges. The weight of edge is shown.

$$\sum_{l=1}^B (e_{i,j}^l = b_i^l \oplus b_j^l) \quad \forall i, j, i \neq j \quad (2)$$

Each XOR relation of the form $(x = y \oplus z)$ is expressed using 4 CNF constraints as follows:

$$\begin{aligned} &(\bar{x} \vee y \vee z) \wedge (x \vee \bar{y} \vee z) \wedge \\ &(x \vee y \vee \bar{z}) \wedge (\bar{x} \vee \bar{y} \vee \bar{z}) \end{aligned} \quad (3)$$

This relation yields a total of $4B \binom{x}{2}$ 3-CNF constraints.

- Each state must have a unique state assignment. This is represented using the following PB constraint:

$$\sum_{l=1}^B e_{i,j}^l \geq 1 \quad \forall i, j, i \neq j \quad (4)$$

This relation yields a total of $B \binom{x}{2}$ PB constraints.

The optimization goal is to minimize the weighted hamming distance between the states. This is expressed using the following:

$$\text{Min} \left(\sum_{k=1}^n w_{i,j}^k \sum_{l=1}^B e_{i,j}^l \right) \quad (5)$$

where n is the number of edges in the FSM and $w_{i,j}^k$ is the weight of the edge between states i and j .

A. An Illustrative Example

In this subsection we use the FSM shown in Figure 1 to provide the reader with an example that clearly illustrates the various steps of the formulation. The shown FSM has a total of 4 states. Therefore, 2 bits are needed to represent each state. A maximum of $\binom{4}{2}$, i.e. 6, state combinations can exist. A total of 8 b and 12 e variables are declared. The XOR relation constraints are:

$$\begin{aligned} e_{A,B}^1 &= b_A^1 \oplus b_B^1 & e_{B,C}^2 &= b_B^2 \oplus b_C^2 \\ e_{A,B}^2 &= b_A^2 \oplus b_B^2 & e_{C,D}^1 &= b_C^1 \oplus b_D^1 \\ e_{B,C}^1 &= b_B^1 \oplus b_C^1 & e_{C,D}^2 &= b_C^2 \oplus b_D^2 \\ e_{B,D}^1 &= b_B^1 \oplus b_D^1 & & \\ e_{B,D}^2 &= b_B^2 \oplus b_D^2 & & \end{aligned} \quad (6)$$

Since only 4 edges exist, 8 e variables were used only. The XOR relations generate a total of 32 3-CNF constraints.

TABLE I. EXPERIMENTAL RESULTS USING THE SAT-BASED 0-1 ILP SOLVER MINISAT+ AND THE GENERIC ILP SOLVER CPLEX 10.0. TIME IS IN SECONDS. VALUE REPRESENTS THE MINIMUM OPTIMIZATION COST FOUND BY THE SOLVER.

Circuit		0 Extra Bits				1 Extra Bits				2 Extra Bits			
Name	# States	CPLEX		MiniSAT+		CPLEX		MiniSAT+		CPLEX		MiniSAT+	
		Time	Value	Time	Value	Time	Value	Time	Value	Time	Value	Time	Value
dk15	4	0.02	1060433	0	1060433	0.03	1060433	0.01	1060433	0.08	1060433	0.03	1060433
s8	5	0.13	161635	0.05	161635	0.49	161635	0.82	161635	0.99	161635	1.78	161635
s27	6	0.65	894772	3.9	894772	1.45	894772	14.82	894772	8.63	894772	14.48	894772
dk14	7	2.14	1273863	17.58	1273863	29.34	1273863	122.91	1273863	60.05	1273863	369.3	1273863
dk27	7	1.24	1309519	2.38	1309519	10.8	1297614	7.28	1297614	60.2	1297614	82.43	1297614
dk17	8	2.9	1063784	7.78	1063784	50.98	1063784	40.76	1063784	62.62	1063784	38.23	1063784
ex6	8	3.98	1071092	7.38	1071092	28.39	1071092	76.59	1071092	67.43	1064262	182.16	1064262
ex3	10	0.09	0	>10000	0	0.11	0	>10000	0	0.06	0	>10000	0
opus	10	22.59	835510	11.28	835510	355.8	835510	121.43	835510	2831	835510	482.6	835510
s386	13	569.2	967095	42.57	967095	>10000	910845	1216	910845	4583	899996	561.8	899996
ex4	14	126.2	478257	0.93	478257	30.36	478257	0.48	478257	24.52	478257	0.32	478257
dk512	15	>10000	1223202	1229.2	1218738	>10000	1200880	>10000	1200880	>10000	1200880	>10000	1200880
mark1	15	>10000	1057205	863.5	1057205	>10000	1047797	749.6	1047797	>10000	1040749	1399	1040749
kirkm.	16	>10000	761685	25.71	761685	>10000	761608	149.4	761608	>10000	761540	167.9	761540
ex1	19	>10000	715921	>10000	723029	>10000	689671	>10000	712127	>10000	686277	>10000	1482254
ex2	19	1028.6	1000000	0.04	1000000	45.75	1000000	0.04	1000000	79.77	1000000	0.03	1000000
tma	20	>10000	250624	>10000	248733	>10000	249697	5121	248241	>10000	251161	>10000	248241
Total		51758	14124597	32212	14125350	60554	13997458	37622	14018458	57779	13970733	43301	14763790

The unique state assignment condition is expressed using the following 4 PB constraints:

$$\begin{aligned}
 e_{A,B}^1 + e_{A,B}^2 &\geq 1 & e_{C,D}^1 + e_{C,D}^2 &\geq 1 \\
 e_{B,C}^1 + e_{B,C}^2 &\geq 1 & e_{B,D}^1 + e_{B,D}^2 &\geq 1
 \end{aligned} \tag{7}$$

Finally, the optimization goal is expressed using:

$$\text{Min} \left(\begin{aligned} &3 \cdot e_{A,B}^1 + 3 \cdot e_{A,B}^2 + 6 \cdot e_{B,C}^1 + 6 \cdot e_{B,C}^2 \\ &12 \cdot e_{C,D}^1 + 12 \cdot e_{C,D}^2 + 9 \cdot e_{B,D}^1 + 9 \cdot e_{B,D}^2 \end{aligned} \right) \tag{8}$$

The optimization instance is passed to the ILP solver which returns the assignment: $((b_A^1, b_A^2), (b_B^1, b_B^2), (b_C^1, b_C^2), (b_D^1, b_D^2)) = ((0, 0), (0, 1), (1, 0), (1, 1))$. The assignment yields the minimum possible optimization cost of 36 that the given FSM can experience.

IV. EXPERIMENTAL RESULTS

In this section, we report and discuss the experimental results obtained for the state assignment technique. The results for the MCNC benchmark circuits [20] are presented in Table I. We used the SAT-based 0-1 ILP solver MiniSAT+ [15], in addition to the generic commercial ILP solver, IBM-ILOG CPLEX 10.0 [19]. All experiments were conducted on a Intel Xeon 3Ghz workstation running Linux and equipped with 4 GB of RAM. We used the default settings for MiniSAT+ and CPLEX. A time-out limit of 10,000 seconds was

set for all experiments. Since some SAT solvers don't accept fractions for edge weights and to improve the results accuracy, weights were multiplied by 1M.

Table I lists the experimental results for MiniSAT+ and CPLEX. The first two columns show the name and size of the circuit; *Time* is the runtime in seconds needed to solve the problem; *Value* is the optimal cost achieved by the solver as explained in equation (5). The table also includes three categories labeled 0-, 1-, and 2- extra bits. In the first category, i.e. 0 extra bits, the minimum number of bits were used to represent the states, i.e. $\log_2 x$ where x is the number of states. Note that we are not constrained by the using the minimum number of bits to encode the states, and hence we decided to experiment by increasing the number of bits and try different solutions. The increase in the number of bits needs to be reasonable otherwise this increase will lead to a complex combinational component of the circuit. The second and third categories, i.e. 1 and 2 extra bits, used $(\log_2 x + 1)$ and $(\log_2 x + 2)$ bits, respectively, to represent each state.

Overall, the results obtained show the superiority of MiniSAT+ over CPLEX in most instances. Although CPLEX outperformed MiniSAT+ in some smaller circuits, for complex larger circuits MiniSAT+ produced better results. For example, when considering the *kirkman* circuit, CPLEX timed-out after 10,000 seconds while MiniSAT+ found the optimal solution in 26 seconds.

TABLE II. EXPERIMENTAL RESULTS USING THE SAT-BASED 0-1 ILP SOLVER MINISAT+ AND THE GENERIC ILP SOLVER CPLEX 10.0. ALL INSTANCES WERE PRE-PROCESSED WITH SHATTERPB AND UPDATED WITH SYMMETRY BREAKING PREDICATES (SBPs). TIME (IN SECONDS) REPRESENTS THE SOLVER RUNTIME. VALUE REPRESENTS THE MINIMUM OPTIMIZATION COST FOUND BY THE SOLVER.

Circuit		0 Extra Bits w/ SBP				1 Extra Bits w/ SBP				2 Extra Bits w/ SBP			
Name	# States	CPLEX		MiniSAT+		CPLEX		MiniSAT+		CPLEX		MiniSAT+	
		Time	Value	Time	Value	Time	Value	Time	Value	Time	Value	Time	Value
dk15	4	0.02	1060433	0	1060433	0.06	1060433	0.01	1060433	0.22	1060433	0.01	1060433
s8	5	0.21	161635	0.03	161635	0.62	161635	0.6	161635	1.75	161635	1.03	161635
s27	6	0.6	894772	2.03	894772	2.77	894772	4.27	894772	10.46	894772	8	894772
dk14	7	2.53	1273863	7.25	1273863	15.21	1273863	14.4	1273863	76.99	1273863	88.25	1273863
dk27	7	2.07	1309519	1.45	1309519	12.88	1297614	2.34	1297614	42.68	1297614	4.91	1297614
dk17	8	5.58	1063784	4.43	1063784	16.85	1063784	9.76	1063784	61.19	1063784	32.91	1063784
ex6	8	4.9	1071092	3.6	1071092	22.58	1071092	9.92	1071092	357.4	1064262	41.2	1064262
ex3	10	0.45	0	>10000	0	0.76	0	>10000	0	0.84	0	>10000	0
opus	10	50.09	835510	3.35	835510	145.4	835510	7.37	835510	1294	835510	12.28	835510
s386	13	952.6	967095	11.32	967095	4122	910845	61.67	910845	7043	899996	42.18	899996
ex4	14	139	478257	0.94	478257	44.64	478257	0.7	478257	79.29	478257	1.44	478257
dk512	15	>10000	1224690	144	1218738	>10000	1200880	1015.5	1200880	>10000	1200880	2109	1200880
mark1	15	7438	1057205	120.4	1057205	>10000	1047797	353.45	1047797	>10000	1041977	320.3	1040749
kirkm.	16	>10000	761685	5.14	761685	>10000	761608	19.53	761608	>10000	761540	27.32	761540
ex1	19	>10000	703609	>10000	681296	>10000	699856	>10000	680282	>10000	694910	>10000	726522
ex2	19	57.07	1000000	0.07	1000000	96.75	1000000	0.06	1000000	83.72	1000000	0.05	1000000
tma	20	>10000	253692	8524	248733	>10000	252887	7463	248241	>10000	249029	>10000	249469
Total		48653	14116841	28828	14083617	54480	14010833	28963	13986613	59052	13978462	32689	14009286

When testing the advantage of adding extra bits to encode the states, results show that extra bits lead to better optimization cost values, but at the expense of increasing search runtimes and added circuit complexity.

In an effort to prune the search space and speed up the ILP solver, the tested instances were pre-processed and checked for symmetries, using the ShatterPB tool [1]. As shown in Table III, ShatterPB was able to find symmetries in all instances in a short amount of time. Table II shows the results after detecting the symmetries and adding symmetry breaking predicates (SBPs) to the instances. Upon comparing the results of Tables 1 and 3, it is evident that MiniSAT+ significantly benefited from the addition of the SBPs. For example, in the *tma* circuit, MiniSAT+ was able to solve the instance with symmetries in 8524 seconds while it timed-out when solving the original instance. In terms of the effect of symmetry breaking on the performance of CPLEX, the results are mixed. For the category with 0 extra bits, CPLEX performed better on 2 instances (*mark1* and *ex2*), but worse on 6 instances. This observation agrees with earlier work [30] that suggests that the generic ILP solver CPLEX is actually slowed down by the addition of SBPs. Since CPLEX is a commercial tool and the algorithms used by it are not publicly known, it is difficult to pinpoint a reason for this disparity.

V. CONCLUSIONS

Starting from a probabilistic description of an FSM, this work has attempted to find a state assignment solution that minimizes the switching activity of the state variables. The main contributions of this work are as follows: (1) We showed how to formulate the assignment problem as a SAT 0-1 ILP problem. (2) We experimented and compared the performance of advanced Boolean satisfiability and generic ILP solvers when solving the SAT 0-1 ILP state assignment problem. Results indicate that for larger topologies, the SAT-based 0-1 ILP solver, MiniSAT+, outperforms the commercial generic ILP solver, CPLEX. (3) We showed empirically that increasing the number of state variables does not lead to noticeable improvement in reaching an optimal solution. (4) We tested the state assignment instances for the existence of symmetries and were able to detect symmetries in all instances. The detected symmetries were broken and used to improve the search runtime. Presented results indicate that detecting and breaking symmetries in such instances significantly improves the 0-1 ILP SAT solver runtime. As far as we know, this is the first work that shows that state assignment problem contains symmetries that can be utilized to speedup search. The effects of the code assigned using the technique presented here on the switching activity of the combinational part of the circuits will be studied in future work.

ACKNOWLEDGMENT

This research was partially supported by the European Union through the European Regional Development Fund.

REFERENCES

[1] F. Aloul, A. Ramani, I. L. Markov, and K. Sakallah, "Symmetry-Breaking for Pseudo-Boolean Formulas," in *ACM Journal of Experimental Algorithmics (JEA)*, vol. 12, article 1.3, 2007.

[2] F. Aloul, A. Ramani, I. L. Markov, and K. Sakallah, "Solving difficult instances of boolean satisfiability in the presence of symmetry," in *IEEE Transactions on Computer Aided Design (TCAD)*, 22(9), 1117-1137, 2003.

[3] F. Aloul, K. Sakallah, and I. Markov, "Efficient Symmetry Breaking for Boolean Satisfiability," in *IEEE Transactions on Computers*, 55(5), 549-558, May 2006.

[4] F. Aloul, A. Ramani, I. Markov, and K. Sakallah, "Generic ILP versus Specialized 0-1 ILP: An Update", in *Proc. of the International Conference on Computer-Aided Design (ICCAD)*, 450-457, 2002.

[5] G. Audemard, L. Simon, "Predicting Learnt Clauses Quality in Modern SAT Solver," in *Proc. of the Twenty-first International Joint Conference on Artificial Intelligence (IJCAI)*, July 2009.

[6] P. Bacchetta, L. Daldoss, D. Sciuto, and C. Silvano, "Low-Power State Assignment techniques for Finite State Machines," in *Proc. of the IEEE Int'l Symposium on Circuits and Systems (ISCAS)*, 641-644, 2000.

[7] L. Benini and G. De Micheli, "State Assignment for Low Power Dissipation," in *IEEE Journal of Solid-State Circuits*, 30(3), 258-268, 1995.

[8] A. Biere, "PicoSAT Essentials," in *Journal on Satisfiability, Boolean Modeling and Computation (JSAT)*, vol. 4, 75-97, 2008.

[9] A. Biere, A. Cimatti, E. Clarke, M. Fujita, and Y. Zhu, "Symbolic Model Checking using SAT procedures instead of BDDs", in *Proc. of the Design Automation Conference (DAC)*, 317-320, 1999.

[10] D. Chai and A. Kuehlmann, "A Fast Pseudo-Boolean Constraint Solver", in *Proc. of the Design Automation Conference (DAC)*, 830-835, 2003.

[11] J. Crawford, M. Ginsberg, E. Luks, A. Roy, "Symmetry-breaking predicates for search problems," in *Proc. of the 5th International Conference on Principles of Knowledge Representation and Reasoning*, 148-159, 1996.

[12] L. Daldoss, D. Sciuto, and C. Silvano, "State Encoding for Low Power Embedded Controllers," in *Proc. of the IEEE International Symposium on Circuits and Systems (ISCAS)*, 421-424, 1998.

[13] P. Darga, M. Lifiton, K. Sakallah, and I. L. Markov, "Exploiting structure in symmetry generation for cnf," in *Proc. of the Design Automation Conference (DAC)*, 530- 534, 2004.

[14] H. Dixon and M. Ginsberg, "Inference Methods for a Pseudo-Boolean Satisfiability Solver," in *Proc. of the National Conference on Artificial Intelligence (AAAI)*, 635-640, 2002.

[15] N. Een and N. Sorensson, "An Extensible SAT-solver," in *Proc. of the International Conference on Theory and Applications of Satisfiability Testing (SAT)*, 502-508, 2003.

[16] R. Eggermont, S. Cotoñana, and C. Lageweg, "Profiling-Based State Assignment for Low Power Dissipation," in *Proc. of ProRISC*, 2004.

[17] G. Hachtel, M. Hermida, A. Pardo, M. Poncino, and F. Somenzi, "Re-Encoding Sequential Circuits to Reduce Power Dissipation," in *Proc. of the Int'l Conference on Computer Aided Design (ICCAD)*, 70-73, 1994.

[18] M. Koegst, G. Franke, and K. Feske, "State Assignment for FSM Low Power Design," in *Proc. of European Design Automation Conference*, 28-33, 1996.

[19] ILOG CPLEX, <http://www.ilog.com/products/cplex>

[20] MCNC Benchmarks, http://www.cbl.ncsu.edu/CBL_Docs/Bench.htm

[21] S. Memik and F. Fallah, "Accelerated Boolean Satisfiability-Based Scheduling of Control/Data Flow Graphs for High-Level Synthesis," in *Proc. of the International Conference on Computer Design (ICCD)*, 2002.

TABLE III. SYMMETRY BREAKING RESULTS. S, G, AND T REPRESENTS THE NUMBER OF SYMMETRIES, NUMBER OF GENERATORS, AND THE TIME (IN SECONDS) NEEDED BY SHATTERPB TO FIND THE SYMMETRIES, RESPECTIVELY.

Circuit		0 Extra Bits			1 Extra Bits			2 Extra Bits		
Name	# States	S	G	T	S	G	T	S	G	T
dk15	4	2	1	0.01	6	2	0	24	3	0
s8	5	6	2	0	24	3	0.01	120	4	0
s27	6	6	2	0	24	3	0	120	4	0.01
dk14	7	6	2	0.01	24	3	0.01	120	4	0.01
dk27	7	6	2	0.01	24	3	0	120	4	0.01
dk17	8	6	2	0.01	24	3	0.02	120	4	0.03
ex6	8	6	2	0.01	24	3	0.02	120	4	0.03
ex3	10	8.7E7	12	0.22	4.4E8	13	0.33	2.6E9	14	0.5
opus	10	24	3	0.04	120	4	0.07	720	5	0.1
s386	13	24	3	0.11	120	4	0.19	720	5	0.29
ex4	14	24	3	0.14	120	4	0.24	720	5	0.38
dk512	15	24	3	0.19	120	4	0.33	720	5	0.5
mark1	15	48	4	0.24	240	5	0.41	1440	6	0.63
kirk.	16	24	3	0.25	120	4	0.41	720	5	0.65
ex1	19	120	4	0.8	120	4	1.06	24	3	1.04
ex2	19	5040	6	1.04	120	4	1.02	24	3	1.05
tma	20	120	4	1	24	3	1.03	6	2	1.07

[22] M. Moskewicz, C. Madigan, Y. Zhao, L. Zhang, and S. Malik, "Chaff: Engineering an Efficient SAT Solver", in *Proc. of the Design Automation Conference (DAC)*, 530-535, 2001.

[23] G. Nam, F. A. Aloul, K. A. Sakallah, and R. Rutenbar, "A Comparative Study of Two Boolean Formulations of FPGA Detailed Routing Constraints", in *IEEE Transactions on Computers*, 53(6), 688-696, June 2004.

[24] J. Nelson, K. Tumer, and J. Ghosh, "Designing Genetic Algorithms for the State Assignment Problem," in *IEEE Transactions on Systems Man and Cybernetics*, 25(4), 687-694, 1999.

[25] W. Noth and R. Kolla, "Spanning Tree Based State Encoding for Low Power Dissipation," in *Proc. of the Design Automation and Test in Europe Conference (DATE)*, 168-174, 1999.

[26] S. Park, S. Cho, S. Yang, and M. Ciesielski, "A New State assignment Technique for Testing and Low Power," in *Proc. of the Design Automation Conference (DAC)*, 510-513, 2004.

[27] K. Pipatsrisawat and A. Darwiche, "RSat 2.0: SAT Solver Description," *Technical report D153. Automated Reasoning Group, Computer Science Department, University of California, Los Angeles*, 2007.

[28] A. Sagahyoon and F. Aloul, "Using SAT-Based Techniques in Power Estimation," in *Microelectronics Journal*, vol. 38, issues 6-7, 706-715, 2007.

[29] H. Sheini and K. Sakallah, "Pueblo: A Modern Pseudo-Boolean SAT Solver," in *Proc. of the Design, Automation, and Test Conference in Europe (DATE)*, vol. 2, 684-685, 2005.

[30] A. Ramani, F. Aloul, I. Markov, and K. Sakallah, "Breaking Instance-Independent Symmetries in Exact Graph Coloring," in *Journal of Artificial Intelligence Research (JAIR)*, Volume 26, 289-322, August 2006.

[31] L. Yuan, G. Qu, T. Villa, and A. Sangiovanni-Vincentelli, "FSM Re-Engineering and Its Application in Low Power State Encoding," in *Proc. of the Asia and South Pacific Design Automation Conference (ASPDAC)*, 254-259, 2005.