

DEMOTIC: A Differentiable Sampler for Multi-Level Digital Circuits

Arash Ardakani
University of California, Berkeley
arash.ardakani@berkeley.edu

Minwoo Kang
University of California, Berkeley
minwoo_kang@berkeley.edu

Kevin He
University of California, Berkeley
kevinjhe@berkeley.edu

Qijing Huang
NVIDIA
jennyhuang@nvidia.com

Vighnesh Iyer
University of California, Berkeley
vighnesh.iyer@berkeley.edu

Suhong Moon
University of California, Berkeley
suhong.moon@berkeley.edu

John Wawrzynek
University of California, Berkeley
johnw@berkeley.edu

ABSTRACT

Efficient sampling of satisfying formulas for circuit satisfiability (CircuitSAT), a well-known NP-complete problem, is essential in modern front-end applications for thorough testing and verification of digital circuits. Generating such samples is a hard computational problem due to the inherent complexity of digital circuits, size of the search space, and resource constraints involved in the process. Addressing these challenges has prompted the development of specialized algorithms that heavily rely on heuristics. However, these heuristic-based approaches frequently encounter scalability issues when tasked with sampling from a larger number of solutions, primarily due to their sequential nature. Different from such heuristic algorithms, we propose a novel differentiable sampler for multi-level digital circuits, called DEMOTIC, that utilizes gradient descent (GD) to solve the CircuitSAT problem and obtain a wide range of valid and distinct solutions. DEMOTIC leverages the circuit structure of the problem instance to learn valid solutions using GD by re-framing the CircuitSAT problem as a supervised multi-output regression task. This differentiable approach allows bit-wise operations to be performed independently on each element of a tensor, enabling parallel execution of learning operations, and accordingly, GPU-accelerated sampling with significant runtime improvements compared to state-of-the-art heuristic samplers. We demonstrate the superior runtime performance of DEMOTIC in the sampling task across various CircuitSAT instances from the ISCAS-85 benchmark suite. Specifically, DEMOTIC outperforms the state-of-the-art sampler by more than two orders of magnitude in most cases.

KEYWORDS

Circuit Satisfiability, Gradient Descent, Multi-level Circuits, Verification, and Testing.

ACM Reference Format:

Arash Ardakani, Minwoo Kang, Kevin He, Qijing Huang, Vighnesh Iyer, Suhong Moon, and John Wawrzynek. 2025. DEMOTIC: A Differentiable Sampler for Multi-Level Digital Circuits. In *30th Asia and South Pacific Design Automation Conference (ASPDAC '25)*, January 20–23, 2025, Tokyo, Japan. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3658617.3697760>

1 INTRODUCTION

Circuit satisfiability (CircuitSAT) solving is an integral part of testing and verification of digital circuits in modern front-end applications such as logic rewriting, false path analysis, property checking, logic synthesis and equivalence checking [6, 22, 23, 30, 33]. CircuitSAT samplers play a crucial role in generating diverse samples from the solution space, aiding in validation, analysis, and optimization of digital circuit designs [9]. Exhaustive exploration and the creation of diverse solutions are essential to guaranteeing that the design fulfills its functional requirements and operates correctly across different scenarios. Consequently, numerous sampling techniques have been developed to detect edge cases and outliers, ensure representativeness, bolster robustness, and promote the broad applicability of findings [10].

High-throughput sampling is fundamental in the realm of CircuitSAT, playing a vital role in various critical tasks such as constrained random verification (CRV) [17]. Its primary function lies in boosting efficiency and scalability by enabling swift exploration of vast solution spaces, which is especially crucial when dealing with complex digital circuits containing numerous inputs, outputs, and intermediate signals. Moreover, it expands coverage across solution spaces, assisting in the detection of uncommon solutions and intricate edge cases inherent in CircuitSAT problems. Furthermore, high-throughput sampling enhances statistical reliability by producing larger sample sizes, thereby reducing sampling variability in CircuitSAT formula analysis.

A common approach for solving the CircuitSAT problem and obtaining diverse solutions involves transforming the CircuitSAT problem into a Boolean Satisfiability (SAT) problem, then employing robust and advanced solvers to solve it effectively [15]. More precisely, the CircuitSAT solving process begins with the formulation of the logical constraints of a digital logic into a Boolean formula, typically represented in conjunctive normal form (CNF).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
ASPDAC '25, January 20–23, 2025, Tokyo, Japan
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0635-6/25/01
<https://doi.org/10.1145/3658617.3697760>

This formula encodes the functionality of the underlying circuits into a set of Boolean clauses according to the interconnections between the circuit’s inputs, outputs and internal signals, ensuring that the resulting CNF accurately represents the original circuit’s behavior [32].

The conversion process can sometimes be complex and computationally intensive, especially for large circuits. In some cases, the resulting CNF might not always be *compact* due to different factors such as the size of the circuit and the number of intermediate variables. For instance, circuits containing many gates or components may result in a large CNF since each gate or component can potentially introduce additional variables and clauses. The intermediate variables can also introduce additional variables and clauses, contributing in the size of the resulting CNF. This increase in complexity of CNF can impact the efficiency of SAT solvers.

SAT solvers employ a variety of techniques to search for a satisfying assignment to the variables in the CNF. Modern SAT solvers [4, 11, 24] often utilize conflict-driven clause learning (CDCL) algorithm [20, 21] which heavily relies on heuristics such as conflict-driven backtracking and clause learning. These heuristics help guide the search process of CDCL for a satisfying assignment efficiently. Due to the sequential nature of these heuristics and their reliance on branching and backtracking, current state-of-the-art (SOTA) SAT solvers are executed on CPUs. Accordingly, SOTA SAT samplers, which integrate a SAT solver as part of their algorithms, depend on a sequential process and are also optimized for execution on CPUs.

GPU acceleration has demonstrated substantial throughput performance advantages across diverse applications especially in machine learning [19]. In general, algorithms that exhibit parallelism and can be decomposed into independent tasks are generally suitable for GPU acceleration. This appears to align with the requirements for generating various satisfying solutions to the CircuitSAT problem, if there was a sampling method performing regular and data-parallel computations. In addition to parallelism, introducing a sampling method that operates directly on circuits by leveraging the spatial and temporal dimensions of digital computations—without converting to CNF—is crucial, as the CNF conversion process may not consistently produce the most optimal or compact representation. To this end, we introduce a novel differentiable sampler for multi-level digital circuits, called DEMOTIC¹, that utilizes gradient descent (GD) for learning diverse solutions to the CircuitSAT problem. We re-frame the CircuitSAT problem as a multi-output regression task, where each logic gate is modeled with a probabilistic representation. We subsequently formulate a loss function by incorporating specified constraints into the circuit. This approach allows us to maintain the integrity of the circuit structure while transforming the sampling process into a learning process. This process enables the generation of independent solutions to the CircuitSAT problem in a parallel fashion, enabling acceleration with GPUs. In summary, we make the following contributions in this paper.

- We introduce a novel differentiable sampler for multi-level digital circuits, called DEMOTIC.

- We use a probabilistic representation to model logic gates and convert the CircuitSAT problem into a multi-output regression task.
- Our proposed sampling method maintains the original structure of the underlying digital circuit without the need for the conversion into any Boolean formulation.
- DEMOTIC enables performing the learning process in parallel, leading to GPU-accelerated sampling of satisfying formulas for the CircuitSAT problem.
- We demonstrate the performance of DEMOTIC across different CircuitSAT instances from the ISCAS-85 benchmark suite [13].

2 PRELIMINARIES

2.1 CircuitSAT Sampling

CircuitSAT sampling refers to the task of sampling solutions from the solution space of a given CircuitSAT problem. In a Boolean circuit, variables can only take on binary values of either 0 or 1. A digital circuit is composed of various logic gates such as AND, OR, and NOT gates, which manipulate these Boolean variables. The output of such a circuit is produced based on how these logic gates operate.

In CircuitSAT, the objective is to determine whether a given circuit, representing a Boolean expression, has an assignment of binary values to its variables that results in the circuit output valuation to 1. The sampling aspect introduces a probabilistic dimension to this problem. Instead of finding a single solution for satisfiable problems, CircuitSAT sampling aims to generate multiple solutions or samples from the set of all possible solutions. Sampling solutions from CircuitSAT instances is an integral part of the design verification process, with significant applications in CRV [17].

CRV is a verification methodology employed in the design and testing of modern digital circuits. It involves generating random input stimuli for the design under test (DUT) while adhering to a set of predefined constraints. This method helps explore a wide range of possible input scenarios, increasing the likelihood of identifying design flaws and ensuring robustness. Inputs to the DUT are generated randomly within specified constraints. This randomness helps cover a broad spectrum of test scenarios, including corner cases that might not be easily detected with directed tests. Constraints are rules or conditions that the random inputs must satisfy. These can be functional constraints based on the design specifications or operational constraints based on practical considerations. Constraints ensure that the generated random inputs are valid and meaningful for the DUT.

For instance, consider a 4-bit multiplier with two unsigned 4-bit binary inputs and an 8-bit product as its primary output. To restrict our bug search to input pairs where the product is less than 128, we need to set the most significant bit of the product to 0 and find random satisfying solutions. It is worth mentioning that while exhaustive search can be used for design verification with complete coverage, it is often impractical for complex designs due to scalability and efficiency limitations. CRV provides a more feasible, efficient, and effective approach by leveraging constraints and randomization to achieve high coverage and uncover critical issues within a manageable timeframe.

¹The code of DEMOTIC is available at <https://github.com/arashardakani/Demotic>.

One common approach for CircuitSAT sampling is the use of SAT solvers with sampling capabilities. These solvers are designed to not only determine the satisfiability of a Boolean formula but also to sample solutions from the solution space. There are various techniques for efficient SAT solving, including backtracking algorithms like Davis-Putnam-Logemann-Loveland (DPLL) algorithm [8], stochastic local search methods like WalkSAT [28], and CDCL algorithms [20, 21]. Over the past years, various algorithms and techniques have been developed for CircuitSAT/SAT sampling including randomized algorithms, Markov chain Monte Carlo (MCMC) methods, and sampling-based heuristics [9, 12, 16, 18, 29]. These approaches typically involve iteratively exploring the solution space, selecting candidate solutions based on certain criteria, and stochastically accepting or rejecting them.

The process of CircuitSAT sampling using SAT solvers involves translating the structure and logic of the given circuit into an equivalent Boolean formula, typically represented in CNF. The CNF consists of a conjunction of multiple clauses (i.e., AND of multiple clauses), where each clause is a disjunction of literals (i.e., OR of literals). Literals are referred to as Boolean variables or their complements. During the conversion process, the number of inputs, outputs and intermediate signals directly contributes to the number of variables in the CNF. The functionality of logic gates and their interconnections determine the number of clauses in the CNF.

The size and complexity of the CNF formula vary significantly depending on factors such as the number of gates, the depth of the circuit, and the number of inputs, outputs and intermediate signals. In general, the CNF formula tends to contain significantly more bit-wise operations than its corresponding circuit. The introduced complexity to the SAT instance as the result of the conversion exponentially increases the time required to find a solution using SAT solvers since the SAT problem is NP-complete. This issue becomes worse when the complexity of digital circuits increases, making SAT sampling a non-trivial task, especially for large or complex circuits.

2.2 Multi-Output Regression Task

A multi-output regression task is a statistical technique used to predict multiple target variables simultaneously from a set of input variables [5]. In this task, the primary objective is to develop a model that accurately captures the relationships between input and output variables. The model can be constructed using various techniques such as linear regression and neural networks. The model is then trained using a dataset where both input-output pairs are known. During training, the model’s parameters are adjusted by minimizing the distance between the predicted outputs and the desired target variables. A common metric to measure such a distance is mean squared error (MSE) or ℓ_2 -loss.

3 METHODOLOGY

In this section, we describe our differentiable solver/sampler for multi-level digital circuits. While the common approach in solving CircuitSAT typically involves converting the underlying circuit into CNF and employing a SAT solver to find the satisfying solution, we take a completely different approach. Instead, we re-frame the CircuitSAT problem as a multi-output regression task, transforming it into a learning problem. Digital circuits are inherently discrete

Table 1: Probability model of logic gates.

Gate	Input Probability	Output Probability	Derivative w.r.t Input
NOT	P_1	$P_y = \bar{P}_1 = 1 - P_1$	$\frac{\partial P_y}{\partial P_1} = -1$
AND	P_1, P_2	$P_y = P_1 P_2$	$\frac{\partial P_y}{\partial P_1} = P_2, \frac{\partial P_y}{\partial P_2} = P_1$
OR	P_1, P_2	$P_y = 1 - \bar{P}_1 \bar{P}_2$	$\frac{\partial P_y}{\partial P_1} = \bar{P}_2, \frac{\partial P_y}{\partial P_2} = \bar{P}_1$
NAND	P_1, P_2	$P_y = 1 - P_1 P_2$	$\frac{\partial P_y}{\partial P_1} = -P_2, \frac{\partial P_y}{\partial P_2} = -P_1$
NOR	P_1, P_2	$P_y = \bar{P}_1 \bar{P}_2$	$\frac{\partial P_y}{\partial P_1} = -\bar{P}_2, \frac{\partial P_y}{\partial P_2} = -\bar{P}_1$
XOR	P_1, P_2	$P_y = \bar{P}_1 P_2 + P_1 \bar{P}_2$	$\frac{\partial P_y}{\partial P_1} = 1 - 2P_2, \frac{\partial P_y}{\partial P_2} = 1 - 2P_1$
XNOR	P_1, P_2	$P_y = P_1 P_2 + \bar{P}_1 \bar{P}_2$	$\frac{\partial P_y}{\partial P_1} = 2P_2 - 1, \frac{\partial P_y}{\partial P_2} = 2P_1 - 1$

and non-differentiable. Therefore, we first need to relax the CircuitSAT problem into a continuous form while accurately capturing the structure and behavior of the circuit. To accomplish this, we leverage the probability model of digital gates, as shown in Table 1. This probability model is commonly used in different domains such as stochastic computing [3] and dynamic power estimation of digital circuits [14]. We use these probabilities to model each gate in the circuit. The result of such modeling is a differentiable formulation of the underlying circuit that accurately describes its functionality while preserving its spatial structure. Of course, the outcome of this model remains identical to the original circuit in its discrete form for any binary input valuations.

Given the differentiable model of the circuit obtained by replacing its discrete logic gates with their corresponding probability model, our objective now is to generate a set of inputs that satisfy a desired constraint. This constraint could pertain to any desired valuation of intermediate signals or outputs. To generate satisfying solutions to the CircuitSAT problem, we represent the input variables to the circuit as $\mathbf{V} \in \mathbb{R}^{b \times n}$, where n represents the number of variables and b denotes the batch size. We define the matrix \mathbf{V} as the parameters of an embedding layer in our circuit model, which will be updated during the learning process. It is worth mentioning that the number of variables in our sampling method is significantly fewer than that of SAT samplers, remaining the same as the number of inputs in the circuit. This discrepancy arises because SAT samplers deal with the CNF of the circuit, where each gate or component introduces additional variables. The embedding layer converts the real-value input variables of the circuit into probabilities in the range from 0 to 1 using the sigmoid function $\sigma(\cdot)$, expressed as:

$$\mathbf{P} = \sigma(\mathbf{V}) = \frac{1}{1 + e^{-\mathbf{V}}}, \quad (1)$$

where $\mathbf{P} \in [0, 1]^{b \times n}$ represents the input probabilities to the underlying circuit. The circuit functionality is then computed as:

$$\mathbf{Y} = \mathcal{F}(\mathbf{P}), \quad (2)$$

where $\mathcal{F} : [0, 1]^{b \times n} \rightarrow [0, 1]^{b \times m}$ denotes the probabilistic model of the circuit. The matrix $\mathbf{Y} \in [0, 1]^{b \times m}$ denotes the m outputs across b data batches. The ℓ_2 -loss function \mathcal{L} can be constructed by measuring the distance between \mathbf{Y} and the target output valuation

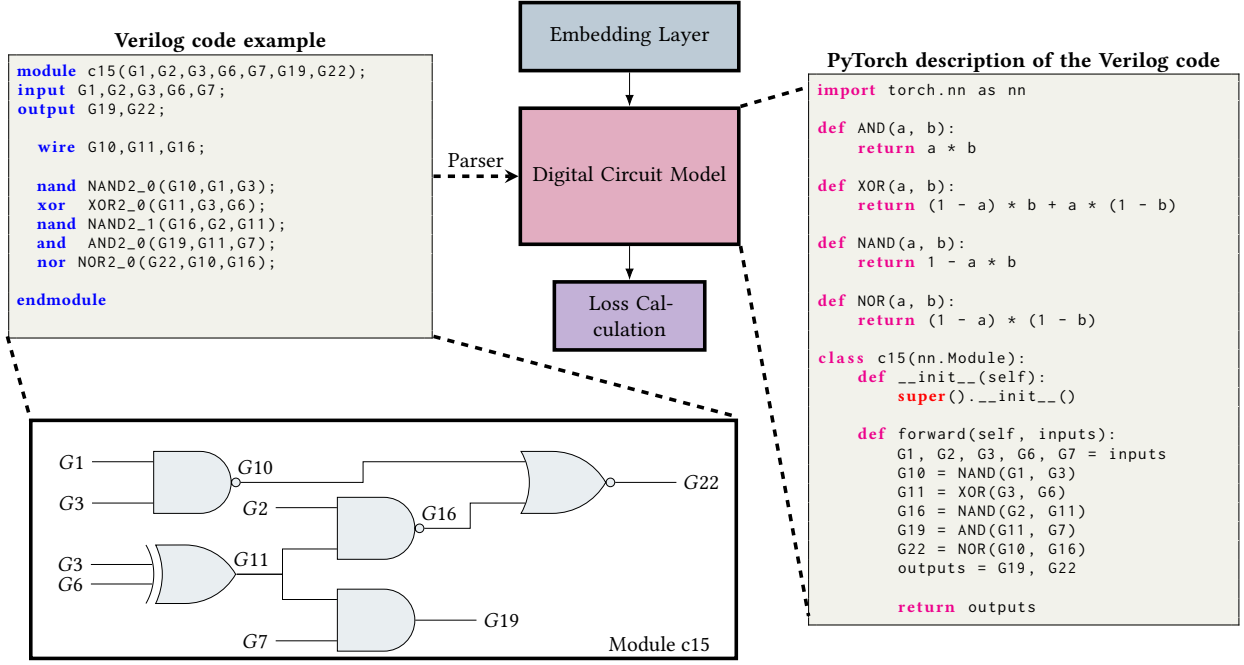


Figure 1: An overview of DEMOTIC is shown. DEMOTIC takes a Verilog instance describing a combinational circuit and parse it into its corresponding probabilistic model described in PyTorch. The embedding layer converts the learnable real-value inputs into probabilities. The ℓ_2 -loss function is calculated in each training iteration and the input variables are updated using GD.

matrix $\mathbf{T} \in \{0, 1\}^{b \times m}$ as follows:

$$\mathcal{L} = \sum_{b,m} \|\mathbf{Y} - \mathbf{T}\|_2^2. \quad (3)$$

The above loss function can be minimized, and the input variables (i.e., \mathbf{V}) can be updated using GD in an iterative manner. Upon convergence, the b solutions to the CircuitSAT problem are obtained by converting the soft input values (i.e., \mathbf{V}) into hard values (i.e., $\bar{\mathbf{V}} \in \{0, 1\}^{b \times n}$).

Fig. 1 illustrates the overview of DEMOTIC. DEMOTIC is equipped with a parser to convert the circuit described in either bit-blasted Verilog or Berkeley Logic Interchange Format (BLIF) into its corresponding probabilistic model. Consequently, DEMOTIC can describe combinational circuits and generate satisfying solutions for any arbitrary constraint on the circuit. Such a sampling paradigm can also benefit from GPU acceleration due to the parallel independent computations across the data batches, enabling a high-throughput sampling procedure.

To better understand our methodology, let us consider a quantitative example using the module “c15” shown in Fig. 1. We set the output node $G19$ to 1 as an output constraint, while the output node $G22$ can take any value of either 0 or 1. Therefore, the goal in this example is to find a pair of inputs such that the output node $G19$ is equal to 1. In this example, the input nodes contributing to our output constraint are $G3$, $G6$, and $G7$. These inputs are learned iteratively using gradient descent. The remaining input nodes, $G1$, $G2$, and $G3$, will not be updated and can take any arbitrary binary values. During each training iteration, each input node is updated

by computing the derivative of the loss function with respect to each input node.

To illustrate the process, we generate two samples. In the first step, we randomly assign two values to each input node as follows:

$$\mathbf{v}_{G3} = \begin{bmatrix} 0.1 \\ -0.2 \end{bmatrix}, \mathbf{v}_{G6} = \begin{bmatrix} 0.5 \\ -0.4 \end{bmatrix}, \mathbf{v}_{G7} = \begin{bmatrix} -0.7 \\ -0.8 \end{bmatrix}, \quad (4)$$

where the concatenation of the above vectors forms the matrix \mathbf{V} . Next, the input probabilities to the circuit are calculated using the sigmoid function:

$$\mathbf{p}_{G3} = \begin{bmatrix} 0.5250 \\ 0.4502 \end{bmatrix}, \mathbf{p}_{G6} = \begin{bmatrix} 0.6225 \\ 0.4013 \end{bmatrix}, \mathbf{p}_{G7} = \begin{bmatrix} 0.3318 \\ 0.3100 \end{bmatrix}. \quad (5)$$

Using the probability model of each gate shown in Table 1, the probabilities of the intermediate node $G11$ and the output node $G19$ are calculated as follows:

$$\mathbf{p}_{G11} = \begin{bmatrix} 0.4939 \\ 0.4902 \end{bmatrix}, \mathbf{p}_{G19} = \begin{bmatrix} 0.1639 \\ 0.1520 \end{bmatrix}. \quad (6)$$

Given the target value of 1 for the output node $G19$, the loss is calculated as:

$$\mathcal{L} = (\mathbf{p}_{G19} - 1)^2 = \begin{bmatrix} (0.1639 - 1)^2 \\ (0.1520 - 1)^2 \end{bmatrix} = \begin{bmatrix} 0.6991 \\ 0.7192 \end{bmatrix}. \quad (7)$$

The above computations are commonly referred to as forward computations. To update the value of the loss of the input variables, we need to calculate the derivative of the loss with respect to each input variable, which is referred to as backward computations. This involves using the derivatives of each gate (as shown in Table 1) and

applying the chain rule. The process is derived as follows:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \mathbf{v}_{G3}} &= \frac{\partial \mathcal{L}}{\partial \mathbf{p}_{G19}} \frac{\partial \mathbf{p}_{G19}}{\partial \mathbf{p}_{G11}} \frac{\partial \mathbf{p}_{G11}}{\partial \mathbf{p}_{G3}} \frac{\partial \mathbf{p}_{G3}}{\partial \mathbf{v}_{G3}} = 2\mathbf{p}_{G19} \cdot \mathbf{p}_{G7} \cdot (1 - 2\mathbf{p}_{G6}) \\ &\quad \cdot \sigma(\mathbf{v}_{G3}) \cdot (1 - \sigma(\mathbf{v}_{G3})) = \begin{bmatrix} 0.0339 \\ -0.0257 \end{bmatrix}, \\ \frac{\partial \mathcal{L}}{\partial \mathbf{v}_{G6}} &= \frac{\partial \mathcal{L}}{\partial \mathbf{p}_{G19}} \frac{\partial \mathbf{p}_{G19}}{\partial \mathbf{p}_{G11}} \frac{\partial \mathbf{p}_{G11}}{\partial \mathbf{p}_{G6}} \frac{\partial \mathbf{p}_{G6}}{\partial \mathbf{v}_{G6}} = 2\mathbf{p}_{G19} \cdot \mathbf{p}_{G7} \cdot (1 - 2\mathbf{p}_{G3}) \\ &\quad \cdot \sigma(\mathbf{v}_{G6}) \cdot (1 - \sigma(\mathbf{v}_{G6})) = \begin{bmatrix} 0.0065 \\ -0.0126 \end{bmatrix}, \\ \frac{\partial \mathcal{L}}{\partial \mathbf{v}_{G7}} &= \frac{\partial \mathcal{L}}{\partial \mathbf{p}_{G19}} \frac{\partial \mathbf{p}_{G19}}{\partial \mathbf{p}_{G7}} \frac{\partial \mathbf{p}_{G7}}{\partial \mathbf{v}_{G7}} = 2\mathbf{p}_{G19} \cdot \mathbf{p}_{G11} \\ &\quad \cdot \sigma(\mathbf{v}_{G7}) \cdot (1 - \sigma(\mathbf{v}_{G7})) = \begin{bmatrix} -0.1831 \\ -0.1778 \end{bmatrix}, \end{aligned} \quad (8)$$

where “ \cdot ” denotes element-wise multiplication.

At this point, each variable is updated using the gradient descent update rule. This involves subtracting the derivative of the loss, scaled by the learning rate, from the corresponding input variables. Given a learning rate of $\gamma = 10$, the new values of the input variables at the end of this iteration are obtained as follows:

$$\begin{aligned} \mathbf{v}_{G3} &= \mathbf{v}_{G3} - \gamma \frac{\partial \mathcal{L}}{\partial \mathbf{v}_{G3}} = \begin{bmatrix} -0.2389 \\ 0.0569 \end{bmatrix}, \mathbf{v}_{G6} = \begin{bmatrix} 0.4349 \\ -0.2741 \end{bmatrix}, \\ \mathbf{v}_{G7} &= \begin{bmatrix} 1.1311 \\ 0.9783 \end{bmatrix}. \end{aligned} \quad (9)$$

This process can be repeated multiple times until convergence. However, even after one iteration in this specific example, we obtain two valid and distinct solutions by rounding the input variables to their nearest discrete values after applying the sigmoid function. In this example, the two input pairs of $(v_{G3} = -0.2389, v_{G6} = 0.4349, v_{G7} = 1.1311)$ and $(v_{G3} = 0.0569, v_{G6} = -0.2741, v_{G7} = 0.9783)$ are rounded to $(\tilde{v}_{G3} = 0, \tilde{v}_{G6} = 1, \tilde{v}_{G7} = 1)$ and $(\tilde{v}_{G3} = 1, \tilde{v}_{G6} = 0, \tilde{v}_{G7} = 1)$, respectively. As demonstrated through this example, the forward and backward computations of the two samples are independent of each other. This allows for the parallel execution of the learning process across multiple samples (i.e., batches), enabling GPU acceleration.

4 EXPERIMENTAL RESULTS

In this section, we showcase how our differentiable method tackles sampling in CircuitSAT problems. To achieve this, we have created a prototype for DEMOTIC using PyTorch. PyTorch is an open-source machine learning framework that blends Torch’s efficient GPU-accelerated backend libraries with a user-friendly Python interface. For a comprehensive evaluation, we use the ISCAS-85 benchmark suite, comprising 11 combinational circuits [13]. The results of DEMOTIC were obtained from running on a system equipped with an Intel Xeon E5 – 2698 with 2.2GHz clock rate and 8 NVIDIA V100 GPUs with 32GB of memory each. We report the runtime performance of DEMOTIC in term of throughput, measured as the number of valid and distinct solutions per second, while using a single NVIDIA V100 GPU. To obtain the experimental results of DEMOTIC for the CircuitSAT problems of the ISCAS-85 benchmark suite, we used GD as the optimizer. We set the learning rate to 15, the batch size to 500,000, and the number of iterations to 10.

4.1 Runtime Performance

We use all 11 combinational circuits from the ISCAS-85 benchmark suite, encompassing designs ranging from relatively simple to moderately complex. These circuits serve as standardized test cases for evaluating algorithm performance in tasks such as logic synthesis, technology mapping, simulation, and testing. We convert these circuits into CircuitSAT sampling problems by randomly assigning specific binary values to some of their output nodes. The objective is to identify a set of inputs that would yield the desired values for those fixed outputs. The size of the solution space for such problems is proportional to the number of inputs in these circuits. Table 2 summarizes the sampling performance of DEMOTIC in terms of throughput for all the combinational circuits in the ISCAS-85 benchmark suite. Throughput is measured as the number of unique solutions generated per second. We report the experimental results corresponding to the best throughput obtained from each sampler in Table 2.

For comparison purposes, we evaluate DEMOTIC’s performance against state-of-the-art SAT sampler baselines, namely UNIGEN3 [29], CMSGEN [12], and DIFFSAMPLER [2]. UNIGEN3 and CMSGEN are highly optimized C++ implementations, whereas DIFFSAMPLER is a GPU-accelerated SAT sampler implemented in Python using the high-performance numerical computing library JAX. To this end, we first need to convert the CircuitSAT problems into their CNF formulas under the same aforementioned output constraints. We employ the Tseytin transformation, which takes a combinational logic circuit as input and produces its corresponding CNF [31]. The size of the solution space for the resulting SAT problems is proportional to the number of variables in their CNF representation. Table 2 presents the performance of the baseline samplers for the obtained SAT instances. UNIGEN3 and CMSGEN were executed on server-grade Intel Xeon Gold 6254 CPU with a clock rate of 3.1GHz and 790GB of RAM. Similar to DEMOTIC, the results of DIFFSAMPLER were obtained from running on a system equipped with an Intel Xeon E5 – 2698 with 2.2GHz clock rate and 8 NVIDIA V100 GPUs with 32GB of memory each.

The experimental results presented in Table 2 showcase the superior performance of DEMOTIC in the sampling task, surpassing state-of-the-art samplers by over two orders of magnitude in most cases. This is because the conversion to CNF introduces additional variables and operations in the form of clauses depending on the complexity of the underlying circuit, as shown in Table 2, undermining the performance of baseline samplers across all the CircuitSAT instances except for “c17”. Due to the limited number of inputs in the CircuitSAT instance for “c17”, only 18 unique solutions exist when constraining the circuit’s second output to 1. This restriction reduces DEMOTIC’s performance in this scenario, as the GPU becomes under-utilized. Consequently, CMSGEN performs more efficiently in this case.

Fig. 2 illustrates the scaling patterns of runtime performance relative to the number of unique solutions generated by each sampler. The analysis reveals two key findings: Firstly, DEMOTIC overall demonstrates superior efficiency compared to UNIGEN3, CMSGEN and DIFFSAMPLER, especially when sampling larger numbers of solutions. Secondly, our method exhibits more efficient scalability, as

Table 2: The runtime performance of DEMOTIC, UNI GEN3, CMS GEN and DIFF SAMPLER is evaluated in terms of unique solution throughput. Throughput is measured under the case where each method is tasked with generating a minimum of 1000 distinct solutions (except for “c17”) within a timeout (TO) of 2 hours.

CircuitSAT Instance	# Inputs	# Outputs	# Logic Gates	# Variables (CNF)	# Clauses (CNF)	DEMOTIC	UNI GEN3	CMS GEN	DIFF SAMPLER
c17	5	2	6	25	19	850	15	2, 928	36
c432	36	7	160	539	516	2, 054, 518	1.5	10, 070	105
c499	41	32	202	683	717	1, 123, 605	1.5	5, 704	28
c880	60	26	383	1198	1115	510, 760	0.2	4, 379	15
c1355	41	32	546	1683	1613	648, 736	0.2	3, 109	0.9
c1908	33	25	880	2436	2381	367, 720	TO	2, 213	TO
c2670	233	140	1269	3642	3274	323, 617	TO	1, 385	TO
c3540	50	22	1669	4680	4611	65, 156	TO	1, 073	TO
c5315	178	123	2307	6994	6698	180, 085	TO	655	TO
c6288	32	32	2416	7280	7219	40, 325	TO	14	TO
c7552	207	108	3513	9971	9661	64, 483	TO	430	TO

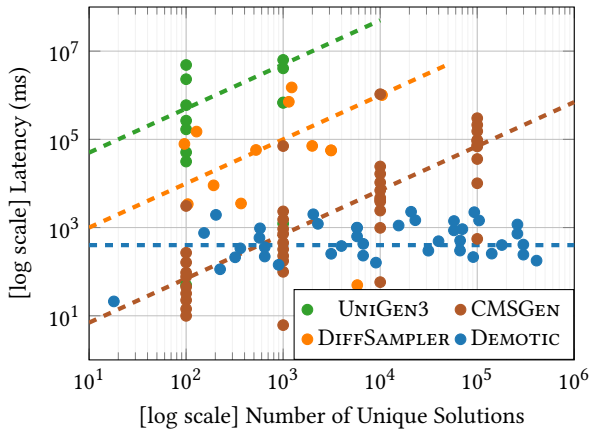


Figure 2: Log-Log plot of runtime in millisecond against the count of unique satisfying solutions found within that run time across all the CircuitSAT problems from the ISCAS-85 benchmark. Dotted lines denote the trend of the runtime performance for each sampler.

evidenced by the linear scaling of the time required for generating larger numbers of solutions.

4.2 Learning Dynamics

In this section, we provide a detailed analysis of the learning dynamics of DEMOTIC over time. For all experiments, we set the batch size to 500, 000 and the learning rate to 15 unless stated otherwise. We excluded the module “c17” from our experiments due to its limited number of inputs.

Figure 3 illustrates the learning progress of DEMOTIC in terms of the number of unique solutions discovered across 10 iterations. The learning curves show that as the number of iterations increases, the number of unique solutions learned by DEMOTIC also increases. While there is no theoretical guarantee that gradient descent will reach the global minimum in non-convex landscapes, including our continuous formulation of CircuitSAT problems, our experiments demonstrate its effectiveness in finding solutions that perform well

in the continuous form. Even if these solutions aren’t the absolute global minimum in the continuous form, they still satisfy the CircuitSAT constraints in the discrete form.

The convergence rate for each CircuitSAT problem varies depending on the complexity and structure of the underlying circuit, as well as the chosen hyper-parameters, such as the learning rate. More complex circuits and sub-optimal hyper-parameter settings typically result in slower convergence rates. Conversely, simpler circuits and well-tuned hyper-parameters tend to lead to faster convergence. Choosing an appropriate learning rate, as the most important hyper-parameter in our experiments, is crucial for effective model training. Very low learning rates can make slow convergence of the learning process, requiring many iterations to reach an optimal solution, which increases computational cost and time. On the other hand, very high learning rates can make the model oscillating around the minimum and leading to poor convergence. For instance, Fig. 4 illustrates the learning progress of DEMOTIC across different learning rates ranging from 1 to 20 for the CircuitSAT problem “c2670”. In this example, the learning rate of 15 provides the best convergence among the tested rates. In contrast, the learning rate of 1 results in the slowest convergence, while the learning rate of 20 leads to slower convergence than the learning rate of 15.

While increasing the number of iterations can lead to learning more unique solutions, it does not necessarily result in higher throughput, as shown in Fig. 5. In fact, the majority of solutions are learned by the end of the first iteration. Specifically, the number of solutions at the end of the first iteration is higher than the number of new unique solutions learned at the end of each subsequent iteration. Given that the latency of each iteration is roughly the same, the throughput of generating unique solutions decreases as the number of iterations increases, as depicted in Fig. 5.

This observation suggests that running DEMOTIC for only one iteration may be sufficient to obtain the desired number of distinct solutions by adjusting the batch size. However, this conclusion holds only when there is no GPU memory constraint for the underlying circuit. GPU acceleration of CircuitSAT sampling incurs GPU

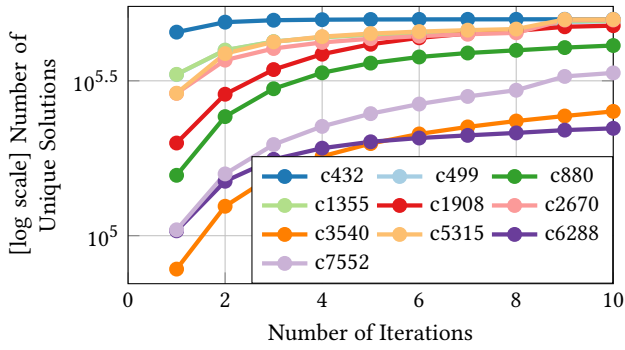


Figure 3: Log learning plot of DEMOTIC showing the number of unique satisfying solutions across different iterations for the CircuitSAT problems from the ISCAS-85 benchmark.




Figure 4: Log learning plot of DEMOTIC showing the number of unique satisfying solutions across different iterations and learning rates for the CircuitSAT problem of “c2670”.

memory usage depending on the size of the CircuitSAT problem and the batch size. Fig. 6 shows the GPU memory usage of the CircuitSAT problems, measured by “nvidia-smi”, across different batch sizes. This figure illustrates the significant growth in GPU memory usage for larger batch sizes. In scenarios where generating a large number of unique samples is targeted but there are constraints on GPU memory usage, the inevitable solution is to run the learning process for more iterations, albeit at the cost of lower throughput.

5 RELATED WORK

Several SAT formula sampling techniques have been explored in the literature. UNIGEN3, for instance, offers approximate uniformity guarantees [26], while CMSGEN and QUICKSAMPLER [9] prioritize sampling efficiency. Previous research has also investigated the use of data-parallel hardware for SAT solving, primarily focusing on parallelizing CDCL or other heuristic-based SAT solving algorithms [7, 25]. Attempts have been made to frame a SAT instance as a constrained numerical optimization problem, as seen in recent work like MATSAT [27] and NEUROSAT [1]. Nevertheless, these approaches have fallen short in showcasing the efficacy of GPU-accelerated formula sampling on standard benchmarks, which are larger and more diverse than the small, random instances typically examined in earlier research. A new differentiable sampling method named DIFFSAMPLER was recently introduced in [2]. This method

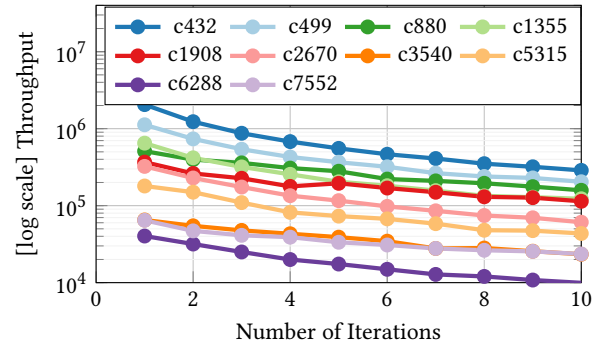


Figure 5: Log plot of the throughput of DEMOTIC, measured by the number of unique satisfying solutions per second across different iterations for the CircuitSAT problems from the ISCAS-85 benchmark.

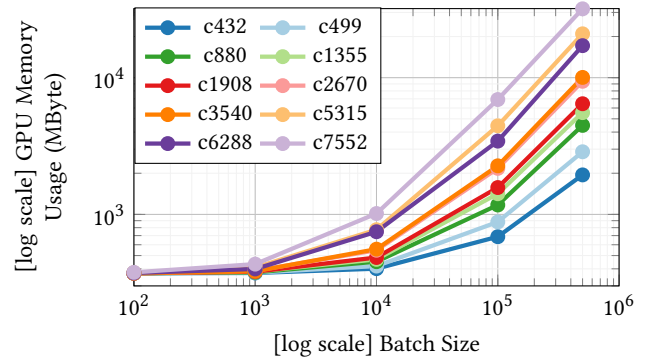


Figure 6: Log-log plot of GPU memory usage of DEMOTIC in megabytes, measured by “nvidia-smi” across different batch sizes for the CircuitSAT problems from the ISCAS-85 benchmark.

allows for GPU-accelerated SAT sampling on standard benchmarks and achieved a comparable runtime performance with respect to UNIGEN3 and CMSGEN.

6 CONCLUSION

CircuitSAT problems are typically transformed into Boolean Satisfiability (SAT) problems, where the sampling task is performed using SAT samplers, albeit with computational complexities, especially for large circuits. To reduce the computational complexity of the CircuitSAT sampling task and leverage GPU acceleration, this paper introduced a novel differentiable sampler called DEMOTIC. DEMOTIC re-frames the CircuitSAT problem as a multi-output regression task, utilizing gradient descent for learning diverse solutions. By maintaining the circuit’s structure without Boolean conversion, DEMOTIC enables parallel learning and GPU-accelerated sampling, offering significant advancements in CircuitSAT sampling methodology. We have demonstrated the exceptional performance of DEMOTIC in the sampling task across various CircuitSAT instances, where it outperformed state-of-the-art samplers by more than two orders of magnitude in most cases.

REFERENCES

- [1] Saeed Amizadeh, Sergiy Matushevych, and Markus Weimer. 2018. Learning to solve circuit-sat: An unsupervised differentiable approach. In *International Conference on Learning Representations*.
- [2] Arash Ardakani, Minwoo Kang, Kevin He, Vighnesh Iyer, Suhong Moon, and John Wawrzynek. 2024. Late Breaking Results: Differential and Massively Parallel Sampling of SAT Formulas. In *Proceedings of the 61st ACM/IEEE Design Automation Conference (DAC)*.
- [3] Arash Ardakani, François Leduc-Primeau, Naoya Onizawa, Takahiro Hanyu, and Warren J. Gross. 2017. VLSI Implementation of Deep Neural Network Using Integral Stochastic Computing. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 25, 10 (2017), 2688–2699. <https://doi.org/10.1109/TVLSI.2017.2654298>
- [4] Gilles Audemard and Laurent Simon. 2018. On the Glucose SAT Solver. *International Journal on Artificial Intelligence Tools* 27, 01 (2018), 1840001. <https://doi.org/10.1142/S0218213018400018>
- [5] Hanen Borchani, Gherardo Varando, Concha Bielza, and Pedro Larranaga. 2015. A survey on multi-output regression. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 5, 5 (2015), 216–233.
- [6] Aaron R. Bradley. 2011. SAT-based model checking without unrolling. In *Proceedings of the 12th International Conference on Verification, Model Checking, and Abstract Interpretation (Austin, TX, USA) (VMCAI'11)*. Springer-Verlag, Berlin, Heidelberg, 70–87.
- [7] CS Costa. 2013. *Parallelization of sat algorithms on gpus*. Technical Report. Technical report, INESC-ID, Technical University of Lisbon.
- [8] Martin Davis, George Logemann, and Donald Loveland. 1962. A machine program for theorem-proving. *Commun. ACM* 5, 7 (jul 1962), 394–397. <https://doi.org/10.1145/368273.368557>
- [9] Rafael Dutra, Kevin Laeuffer, Jonathan Bachrach, and Koushik Sen. 2018. Efficient sampling of SAT solutions for testing. In *Proc. of the International Conference on Software Engineering*.
- [10] Rafael Tupynambá Dutra. 2019. *Efficient sampling of SAT and SMT solutions for testing and verification*. University of California, Berkeley.
- [11] Niklas Eén and Niklas Sörensson. 2003. An Extensible SAT-solver.. In *SAT (Lecture Notes in Computer Science, Vol. 2919)*, Enrico Giunchiglia and Armando Tacchella (Eds.), Springer, 502–518.
- [12] Priyanka Golia, Mate Soos, Sourav Chakraborty, and Kuldeep S. Meel. 2021. Designing Samplers is Easy: The Boon of Testers. In *Proc. of Formal Methods in Computer-Aided Design (FMCAD)*.
- [13] M.C. Hansen, H. Yalcin, and J.P. Hayes. 1999. Unveiling the ISCAS-85 benchmarks: a case study in reverse engineering. *IEEE Design & Test of Computers* 16, 3 (1999), 72–80. <https://doi.org/10.1109/54.785838>
- [14] David Harris and N Weste. 2010. *Cmos vlsi design*. ed: Pearson Education, Inc (2010).
- [15] Chih-Jen Hsu, Wei-Hsun Lin, Chi-An Wu, and Kei-Yong Khoo. 2014. ICCAD-2014 CAD contest in simultaneous CNF encoder optimization with SAT solver setting selection and benchmark suite. In *Proceedings of the 2014 IEEE/ACM International Conference on Computer-Aided Design (San Jose, California) (ICCAD '14)*. IEEE Press, 357–360.
- [16] Russell Impagliazzo, Valentine Kabanets, Antonina Kolokolova, Pierre McKenzie, and Shadab Romani. 2017. Does Looking Inside a Circuit Help?. In *42nd International Symposium on Mathematical Foundations of Computer Science (MFCS 2017) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 83)*, Kim G. Larsen, Hans L. Bodlaender, and Jean-Francois Raskin (Eds.), Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 1:1–1:13. <https://doi.org/10.4230/LIPIcs.MFCS.2017.1>
- [17] Nathan Kitchen and Andreas Kuehlmann. 2007. Stimulus generation for constrained random simulation. In *2007 IEEE/ACM International Conference on Computer-Aided Design*. 258–265. <https://doi.org/10.1109/ICCAD.2007.4397275>
- [18] Nathan Kitchen and Andreas Kuehlmann. 2009. A Markov chain Monte Carlo sampler for mixed Boolean/integer constraints. In *Computer Aided Verification: 21st International Conference, CAV 2009, Grenoble, France, June 26–July 2, 2009. Proceedings 21*. Springer, 446–461.
- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems*, F. Pereira, C.J. Burges, L. Bottou, and K.Q. Weinberger (Eds.), Vol. 25. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf
- [20] Joao Marques-Silva, Ines Lynce, and Sharad Malik. 2021. *Chapter 4: Conflict-driven clause learning SAT solvers*. IOS Press BV, 133–182. <https://doi.org/10.3233/FALA200987> Publisher Copyright: © 2021 The authors and IOS Press. All rights reserved..
- [21] J.P. Marques Silva and K.A. Sakallah. 1996. GRASP-A new search algorithm for satisfiability. In *Proceedings of International Conference on Computer Aided Design*. 220–227. <https://doi.org/10.1109/ICCAD.1996.569607>
- [22] Alan Mishchenko and Robert K. Brayton. 2005. SAT-Based Complete Don't-Care Computation for Network Optimization. In *Proceedings of the Conference on Design, Automation and Test in Europe - Volume 1 (DATE '05)*. IEEE Computer Society, USA, 412–417. <https://doi.org/10.1109/DATE.2005.264>
- [23] Alan Mishchenko, Satrajit Chatterjee, Robert Brayton, and Niklas Eén. 2006. Improvements to combinational equivalence checking. In *Proceedings of the 2006 IEEE/ACM International Conference on Computer-Aided Design (San Jose, California) (ICCAD '06)*. Association for Computing Machinery, New York, NY, USA, 836–843. <https://doi.org/10.1145/1233501.1233679>
- [24] M.W. Moskewicz, C.F. Madigan, Y. Zhao, L. Zhang, and S. Malik. 2001. Chaff: engineering an efficient SAT solver. In *Proceedings of the 38th Design Automation Conference (IEEE Cat. No.01CH37232)*. 530–535. <https://doi.org/10.1145/378239.379017>
- [25] Muhammad Osama, Anton Wijs, and Armin Biere. 2021. SAT solving with GPU accelerated inprocessing. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 133–151.
- [26] Yash Pote and Kuldeep S. Meel. 2022. On Scalable Testing of Samplers. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [27] Taisuke Sato and Ryosuke Kojima. 2021. MatSat: a matrix-based differentiable SAT solver. *arXiv preprint arXiv:2108.06481* (2021).
- [28] Bart Selman, Henry A Kautz, Bram Cohen, et al. 1993. Local search strategies for satisfiability testing. *Cliques, coloring, and satisfiability 26* (1993), 521–532.
- [29] Mate Soos, Stephan Gocht, and Kuldeep S. Meel. 2020. Tinted, Detached, and Lazy CNF-XOR solving and its Applications to Counting and Sampling. In *Proceedings of International Conference on Computer-Aided Verification (CAV)*.
- [30] Shihheng Tsai and Chung-Yang (Ric) Huang. 2009. A false-path aware formal static timing analyzer considering simultaneous input transitions. In *Proceedings of the 46th Annual Design Automation Conference (San Francisco, California) (DAC '09)*. Association for Computing Machinery, New York, NY, USA, 25–30. <https://doi.org/10.1145/1629911.1629921>
- [31] Grigori S Tseitin. 1983. On the complexity of derivation in propositional calculus. *Automation of reasoning: 2: Classical papers on computational logic 1967–1970* (1983), 466–483.
- [32] Miroslav N. Velev. 2004. Efficient translation of boolean formulas to CNF in formal verification of microprocessors. In *Proceedings of the 2004 Asia and South Pacific Design Automation Conference (Yokohama, Japan) (ASP-DAC '04)*. IEEE Press, 310–315.
- [33] He-Teng Zhang, Jie-Hong R. Jiang, and Alan Mishchenko. 2021. A Circuit-Based SAT Solver for Logic Synthesis. In *2021 IEEE/ACM International Conference On Computer Aided Design (ICCAD) (Munich, Germany)*. IEEE Press, 1–6. <https://doi.org/10.1109/ICCAD51958.2021.9643505>