

SKI-SAT: A CMOS-Compatible Hardware for Solving SAT Problems

Ahmet Yusuf Salim¹, Bart Selman, *Member, IEEE*, Henry Kautz², *Member, IEEE*,
Zeljko Ignjatovic¹, *Member, IEEE*, and Selçuk Köse¹, *Member, IEEE*

Abstract—Nature-inspired computation is receiving increasing attention. Various Ising machine (IM) implementations have recently been proven to be effective in solving numerous combinatorial optimization problems including maximum cut, low density parity check (LDPC) decoding, and Boolean satisfiability (SAT) problems. In this paper, a novel method is presented to solve SAT or MAX-SAT problems with a CMOS circuit implementation. The technique solves a SAT problem by mapping the SAT variables onto quantized capacitor voltages generated by an array of nodes that interact through a network of coupling units. The nodal interaction is achieved through coupling currents produced by the coupling units, which charge or discharge capacitor voltages, implementing a gradient descent along the SAT problem’s cost function to minimize the number of unsatisfied clauses. The system also incorporates a unique low-complexity perturbation scheme to avoid settling in local minima, greatly enhancing the performance of the system. The simulation results demonstrate that the proposed SKI-SAT is a high-performance and low-energy alternative that surpasses existing software-based SAT solvers by significant margins, achieving more than 10 times faster solution and over 300 times less power.

Index Terms—Combinatorial optimization problems, Ising machine, Boolean satisfiability, SAT, CMOS, nature-based computing.

I. INTRODUCTION

DIGITAL computers revolutionized the world by automating complex tasks, enabling instant communication, and providing unmatched performance for data processing and artificial intelligence applications. However, recent years have seen an incredible surge in energy-intensive applications for cryptocurrency and neural network training. In digital systems information is represented and processed in one of the two discrete states while analog computers process and store information as analog values such as voltages or currents with

inherently higher resolution enabling time and power-efficient solutions for certain problems. The drawback associated with analog computers is their lack of versatility, but they can outperform digital computers in certain specific functions. Although widespread use of analog computers for general purpose computing is unlikely, analog computers can be a favorable option for specific applications.

In recent years, a subset of analog computation models, known as nature-inspired computation, has garnered attention for providing new avenues to solve NP-hard and NP-complete problems. Since analytical solutions don’t typically exist for these problems, digital computers are often required to rely on rather inefficient trial-and-error methods. An Ising machine is an emerging nature-inspired computation model that maps a problem into an energy space. The system naturally creates a gradient descent to the Hamiltonian function based on its energy, allowing Ising machines to progress toward the optimal solution with the help of random perturbations. The system is, however, likely to settle into a sub-optimal solution without a perturbation scheme [1]. The efficiency of Ising machines in solving quadratic unconstrained binary optimization (QUBO) problems has been demonstrated to be almost a million times faster than conventional simulated annealing [1] owing to the trivial mapping of QUBO to Ising formula. However, functions containing higher-order polynomials cannot be directly mapped to the quadratic Ising machines and therefore require preprocessing. The work in [2] demonstrates a methodology that first transforms a 3-SAT problem into a QUBO format, which is then decomposed into sub-problems and mapped onto a ring oscillator-based 49-spin Ising machine. However, the auxiliary spin overhead created during the conversion of third-order terms in SAT cost function to quadratic terms in QUBO severely limits the size of 3-SAT problems that can be solved on a specific size QUBO solver (e.g., a 49-spin QUBO hardware in [2] is able to solve a 3-SAT problem with up to 20 literals and 91 clauses only by decomposing into sub-Hamiltonians). Therefore, supporting larger problems with more variables becomes increasingly challenging and time consuming. Hizzani et al. [3] report on a memristor based system that employs a Hopfield Neural Network, which eliminates the need for order reduction for third-order polynomials present in a 3-SAT problem. This makes the design intrinsically compatible with polynomial unconstrained binary optimization (PUBO) problems, particularly those involving SAT. How-

Received 1 November 2024; revised 23 January 2025 and 28 February 2025; accepted 18 March 2025. This work was supported in part by the Defense Advanced Research Projects Agency (DARPA) Quantum-Inspired Classical Computing (QuICC) Program under the Air Force Research Laboratory (AFRL) under Contract FA8650-23-C-1034. This article was recommended by Associate Editor A. Ascoli. (*Corresponding author: Ahmet Yusuf Salim.*)

Ahmet Yusuf Salim, Zeljko Ignjatovic, and Selçuk Köse are with the Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY 14627 USA (e-mail: asalim@ur.rochester.edu; zeljko.ignjatovic@rochester.edu; selcuk.kose@rochester.edu).

Bart Selman is with the Department of Computer Science, Cornell University, Ithaca, NY 14850 USA (e-mail: selman@cs.cornell.edu).

Henry Kautz is with the Department of Computer Science, University of Virginia, Charlottesville, VA 22903 USA (e-mail: henry.kautz@virginia.edu).
Digital Object Identifier 10.1109/TCSI.2025.3554145

ever, the hardware implementation of this work imposes significant physical design constraints due to the quadratically growing word lines. This indicates that scalability for larger problems remains a challenging issue. Furthermore, Elmitwalli et al. [4] demonstrate how a QuBRIM system [1], originally designed to solve MaxCUT problems, can be altered to address higher-order combinatorial optimization problems, including SAT and the Traveling Salesman Problem, particularly for LDPC decoding, without the need for order reduction through multi-body interactions. However, the proposed system's limitation lies in the nodal interactions being constrained by the linearity of the DAC units and dynamic range for programmable resistance, which reduces the accuracy of the solution to a certain degree.

The proposed **Salim-Köse-Ignjatovic** SAT hardware solver (SKI-SAT) offers a scheme that is not restricted by polynomial terms, as it inherently supports polynomial terms of any order. Additionally, the coupling precision limitations of the previous studies are no longer a concern for SKI-SAT, thanks to the binary nature of its coupling. Finally, SKI-SAT is easily scalable and cost-efficient thanks to the compact footprint, which is compatible with standard CMOS processes and can accommodate hundreds or even thousands of clauses.

II. BACKGROUND

Solving a SAT problem can be described as answering the question of whether a boolean (i.e., logic) function evaluates true for at least one combination of the input variables. Despite the simple definition, SAT solvers have widespread applications in computer science, including electronic design automation, hardware verification [5], and cryptanalysis [6], making them a focal point of research. SAT is formally an NP-complete problem, which means that it is highly likely that its general solution requires worst-case exponential time. Modern SAT solvers use branching heuristics and dynamic programming to guide their search through the exponential space of possible solutions. Although such techniques are often successful in practice, there remain many problems that resist solution. Continuous advancements in algorithms and hardware have been improving SAT solver performance each year. However, the inherent computational complexity of these problems often leads to lengthy solution times, constrained by the limits of exponential-time algorithms in a search space as broad as 2^N combinations for an N variable function. In recent years, various Ising machine implementations have emerged to solve QUBO problems, such as MaxCUT problems. However, not all NP optimization problems can be effectively expressed in QUBO form; many include cubic or even higher-order terms. These problems with higher-order terms require additional hardware, potentially reducing the effectiveness of an Ising machine when attempting to map such problems to its architecture.

A SAT problem is typically presented in conjunctive normal form (CNF), where the CNF of a Boolean function contains a conjunction of clauses. When each clause consists of logical disjunction of exactly k literals, a SAT problem is typically termed as k -SAT problem. For example, in a 3-SAT problem (i.e., $k = 3$), the literals in each of the N_C clauses may take

any one of the N variables or their complements, as shown in Eq. (1).

$$\ell_{m,k} \in \{X_1, X_2, \dots, X_N, \overline{X_1}, \overline{X_2}, \dots, \overline{X_N}\},$$

$$\text{where } 1 \leq m \leq N_C, \quad 1 \leq k \leq 3 \quad (1)$$

3-variable clauses are created by a disjunction of any three literals as shown in Eq. (2)

$$C_i = (\ell_{i,1} \vee \ell_{i,2} \vee \ell_{i,3}) \quad (2)$$

The CNF form F of a 3-SAT problem can then be written as shown in Eq. (3)

$$F(X_1, X_2, \dots, X_N) = C_1 \wedge C_2 \wedge \dots \wedge C_{N_C} \quad (3)$$

A generalization of the SAT problem described above is the Maximum Satisfiability problem (or MAX-SAT problem). Unlike a SAT problem which asks whether there is at least one assignment of all variables X_i , $1 \leq i \leq N$ that renders all clauses C_j , $1 \leq j \leq N_C$ true, a MAX-SAT problem asks for an assignment of variables X_i that maximizes the number of clauses that are made true with that assignment.

III. THEORETICAL ANALYSIS, SYSTEM DESIGN, AND CIRCUIT LEVEL IMPLEMENTATION

A. Theoretical Analysis

To construct the SKI-SAT solver machine, which refers to a CMOS circuit topology that minimizes the number of unsatisfied clauses in F from Eq. (3), first an appropriate penalty or cost function is derived. For this purpose, a negation of the CNF form is considered, as shown in Eq. (4), where the satisfiability of the CNF form F is converted to the dissatisfiability of a logical Boolean function \overline{F} .

$$\overline{F}(X_1, X_2, \dots, X_N) = \overline{C_1} \vee \overline{C_2} \vee \dots \vee \overline{C_{N_C}} \quad (4)$$

The negated CNF form \overline{F} is then converted into a penalty or cost function H by replacing the logical OR and AND operations in \overline{F} with addition and multiplication operations, respectively. Additionally, complemented variables $\overline{X_k}$ are replaced with $(1 - X_k)$. By way of example, Eq. (5) shows the conversion of \overline{F} containing 3 clauses and 6 variables into the corresponding penalty function H .

$$F = (X_1 \vee \overline{X_2} \vee X_5) \wedge (\overline{X_3} \vee \overline{X_4} \vee X_5) \wedge (\overline{X_6} \vee X_4 \vee X_2)$$

$$\xrightarrow{\text{negation}}$$

$$\overline{F} = (\overline{X_1} \wedge X_2 \wedge \overline{X_5}) \vee (X_3 \wedge X_4 \wedge \overline{X_5}) \vee (X_6 \wedge \overline{X_4} \wedge \overline{X_2})$$

$$\xrightarrow{\text{convert to cost}}$$

$$H = \overline{X_1}X_2\overline{X_5} + X_3X_4\overline{X_5} + X_6\overline{X_4}\overline{X_2}$$

$$= (1 - X_1)X_2(1 - X_5) + X_3X_4(1 - X_5)$$

$$+ X_6(1 - X_4)(1 - X_2) \quad (5)$$

The next step in constructing the SKI-SAT circuit topology that minimizes penalty function H is to derive a gradient vector of H within the Hamming space $\{0, 1\}^N$ spanned by variables X_k , $1 \leq k \leq N$, as shown in Eq. (6).

$$\nabla H = \left[\frac{\partial H}{\partial X_1} \quad \frac{\partial H}{\partial X_2} \quad \dots \quad \frac{\partial H}{\partial X_N} \right]^T \quad (6)$$

Following example in Eq. (5), the gradient of H becomes:

$$\nabla H = \begin{bmatrix} -X_2(1 - X_5) \\ (1 - X_1)(1 - X_5) - X_6(1 - X_4) \\ X_4(1 - X_5) \\ X_3(1 - X_5) - X_6(1 - X_2) \\ -(1 - X_1)X_2 - X_3X_4 \\ (1 - X_4)(1 - X_2) \end{bmatrix} \quad (7)$$

Alternatively, partial derivatives of the cost function may be expressed in the following form:

$$\nabla H = \begin{bmatrix} -X_2\overline{X_5} \\ \overline{X_1}X_5 - X_6\overline{X_4} \\ X_4\overline{X_5} \\ X_3\overline{X_5} - X_6\overline{X_2} \\ -\overline{X_1}X_2 - X_3X_4 \\ \overline{X_4}X_2 \end{bmatrix} \quad (8)$$

Since all variables in Eq. (8) are logical (i.e., taking values 0 or 1), the multiplicative terms can be replaced with AND logic gates as shown in the following Eq. (9):

$$\nabla H = \begin{bmatrix} -X_2 \wedge \overline{X_5} \\ \overline{X_1} \wedge X_5 - X_6 \wedge \overline{X_4} \\ X_4 \wedge \overline{X_5} \\ X_3 \wedge \overline{X_5} - X_6 \wedge \overline{X_2} \\ -\overline{X_1} \wedge X_2 - X_3 \wedge X_4 \\ \overline{X_4} \wedge X_2 \end{bmatrix} \quad (9)$$

Or alternatively, with NOR gates as shown in Eq. (10).

$$\nabla H = \begin{bmatrix} -\overline{X_2} \vee X_5 \\ \overline{X_1} \vee X_5 - \overline{X_6} \vee X_4 \\ \overline{X_4} \vee X_5 \\ \overline{X_3} \vee X_5 - \overline{X_6} \vee X_2 \\ -\overline{X_1} \vee \overline{X_2} - \overline{X_3} \vee \overline{X_4} \\ \overline{X_4} \vee X_2 \end{bmatrix} \quad (10)$$

An N -dimensional real vector \vec{v} is defined, which resides in a unit hypercube Ω_N , meaning that for each element v_i of \vec{v} , $v_i \in [0, 1]$. In addition, a surjective function $f: \Omega_N \rightarrow \{0, 1\}^N$ is defined, such that the real vector \vec{v} is mapped onto a Hamming vector \mathbf{X} , following the mapping rule in Eq. (11).

$$\mathbf{X} = f(\vec{v}) \quad \text{s.t.} \quad X_i = \begin{cases} 0, & \text{if } v_i < 0.5 \\ 1, & \text{if } v_i \geq 0.5 \end{cases}, \quad 1 \leq i \leq N \quad (11)$$

By following the construction rule $\frac{d}{dt}\vec{v} = -\alpha\nabla H$, a set of differential equations (12) that govern the operation of SKI-SAT circuit is defined. Eq. (12) describes the gradient decent capability of SKI-SAT in minimizing the penalty function H .

$$\frac{dv_i}{dt} = -\alpha \frac{\partial H}{\partial X_i}, \quad 1 \leq i \leq N \quad (12)$$

Following the example of SAT problem with 3 clauses and 6 variables defined in Eq. (5), the corresponding set of differential equations is shown in Eq. (13).

$$\frac{dv_1}{dt} = -\alpha \frac{\partial H}{\partial X_1} = +\alpha \overline{X_2} \vee X_5$$

$$\begin{aligned} \frac{dv_2}{dt} &= -\alpha \frac{\partial H}{\partial X_2} = -\alpha \overline{X_1} \vee X_5 + \alpha \overline{X_6} \vee X_4 \\ \frac{dv_3}{dt} &= -\alpha \frac{\partial H}{\partial X_3} = -\alpha \overline{X_4} \vee X_5 \\ \frac{dv_4}{dt} &= -\alpha \frac{\partial H}{\partial X_4} = -\alpha \overline{X_3} \vee X_5 + \alpha \overline{X_6} \vee X_2 \\ \frac{dv_5}{dt} &= -\alpha \frac{\partial H}{\partial X_5} = +\alpha X_1 \vee \overline{X_2} + \alpha \overline{X_3} \vee \overline{X_4} \\ \frac{dv_6}{dt} &= -\alpha \frac{\partial H}{\partial X_6} = -\alpha \overline{X_4} \vee X_2 \end{aligned} \quad (13)$$

The gradient descent nature of the system described by Eqs. (12) and (13), and its ability to minimize the cost function H (i.e., enforcing the number of dissatisfied clauses C_i in F to be non-increasing over time) can be demonstrated by calculating the rate of change of H as

$$\frac{dH}{dt} = \sum_{i=1}^N \frac{\partial H}{\partial X_i} \frac{dX_i}{dt} = -\frac{1}{\alpha} \sum_{i=1}^N \frac{dv_i}{dt} \frac{dX_i}{dt} \leq 0 \quad (14)$$

Since the rate of change of real variables v_i and their corresponding logic variables X_i are of the same polarity, the summation term to the right of Eq. (14) is non-negative. The resulting rate of change of H is therefore non-positive (i.e., cost function H is non-increasing over time).

The constant α in Eq. (12) could be chosen to be equal to I_{ref}/C , where I_{ref} is the reference/constant current and C is a capacitance value. Rearranging Eq. (12) produces a set of equations that describe the current values into a set of capacitors C_i as

$$I_{C_i} = C \frac{dv_i}{dt} = -I_{ref} \frac{\partial H}{\partial X_i}, \quad 1 \leq i \leq N \quad (15)$$

where voltage v_i on a capacitor C_i represents its state, which can range from 0V to V_{dd} or a normalized range from 0 to 1. The corresponding logical variable X_i is produced at the output of a comparator circuit which compares v_i against a threshold V_{th} which is typically $V_{dd}/2$ - or normalized to 0.5, as previously described in Eq. (11). As an example, a SKI-SAT circuit solving a 3-SAT problem with 6 variables in 3 clauses defined in Eq. (5) can be described as

$$\begin{aligned} I_{C_1} &= -I_{ref} \frac{\partial H}{\partial X_1} = +I_{ref} \overline{X_2} \vee X_5 \\ I_{C_2} &= -I_{ref} \frac{\partial H}{\partial X_2} = -I_{ref} \overline{X_1} \vee X_5 + I_{ref} \overline{X_6} \vee X_4 \\ I_{C_3} &= -I_{ref} \frac{\partial H}{\partial X_3} = -I_{ref} \overline{X_4} \vee X_5 \\ I_{C_4} &= -I_{ref} \frac{\partial H}{\partial X_4} = -I_{ref} \overline{X_3} \vee X_5 + I_{ref} \overline{X_6} \vee X_2 \\ I_{C_5} &= -I_{ref} \frac{\partial H}{\partial X_5} = +I_{ref} X_1 \vee \overline{X_2} + I_{ref} \overline{X_3} \vee \overline{X_4} \\ I_{C_6} &= -I_{ref} \frac{\partial H}{\partial X_6} = -I_{ref} \overline{X_4} \vee X_2 \end{aligned} \quad (16)$$

B. System Design and Circuit Level Implementation

In order to construct the system governed by the described differential equations in Section III-A, certain circuit topologies and their objectives are presented in this section. The

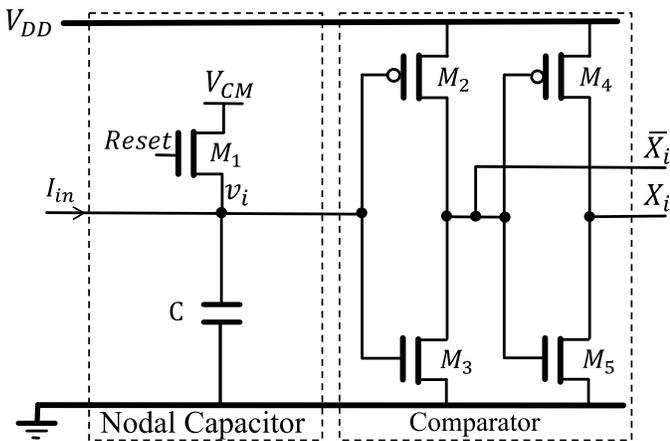


Fig. 1. Example CMOS circuit for a computational node in SKI-SAT. The reset voltage V_{CM} is typically chosen to be equal or close to threshold V_{th} of the comparator.

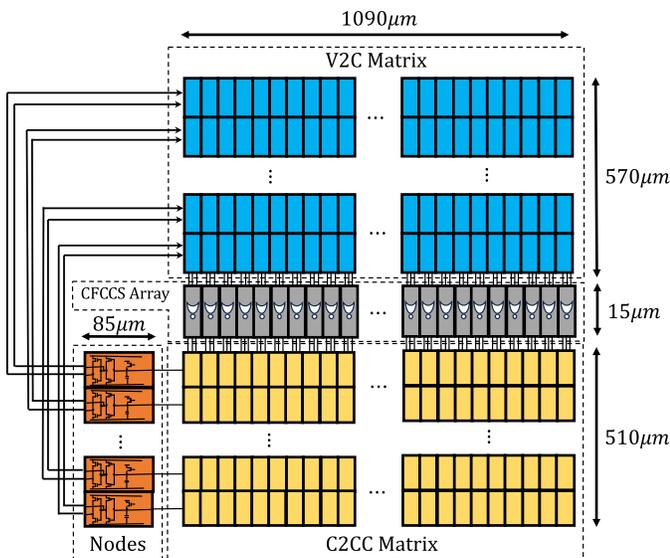


Fig. 2. SKI-SAT top-level architecture with dimensions given for 50 literals and 218 clauses.

SKI-SAT circuit topology consists of N nodes for an N -variable function. Each node comprises a capacitor with voltage v_i across its terminals, an input to supply charging or discharging current to the nodal capacitor, and a comparator. The comparator compares the capacitor voltage v_i against a threshold voltage V_{th} , which could be set in the middle between the power rails. This comparison produces a binary output variable X_i , which can be $0V$ or V_{dd} , along with its complement \bar{X}_i . Additionally, a reset switch is placed between V_{CM} and the nodal capacitor, allowing the nodal variables to be initialized to the middle rail before the system begins operation. An example schematic of the SKI-SAT node is shown in Fig. 1 where the comparator is implemented as a simple inverter.

Moreover, the two outputs (X_i and \bar{X}_i) from this computational node are connected to a programmable Variables-to-Clauses (V2C) array containing $N \times N_C$ units, as shown in the top-level circuit architecture of Fig. 2. A detailed discussion about the physical design is found in Section III-C.

The V2C array provides digital outputs to clause formation and coupling control signal generation array (CFCCS) of size $1 \times N_C$ to form clauses and generate control signals associated with these clauses. Each unit cell in the V2C array contains memory elements to store the polarity bit as well as an address applied to the input of a digital address decoder. If the variable X_i or its complement \bar{X}_i is present in the C_j clause of the CNF form F , cell at the (i, j) location in the V2C array is activated to connect the variable to the clause C_j . Depending on the polarity bit stored in the memory cell of the V2C unit, either variable X_i or its complement \bar{X}_i is selected to be connected to the CFCCS array. The selection between a variable X_i or its complement \bar{X}_i could be achieved through a digital multiplexer (MUX) which is controlled by the polarity bit. The V2C unit further includes a set of output buffers that connect the variable selected through the polarity selection MUX to one of the k readout lines. k is the number of literals in the clause C_j that run through all rows of the j^{th} column and are connected to the inputs of j^{th} CFCCS unit. The CFCCS array is preferably located at the periphery of the V2C array as illustrated in Fig. 2.

The output buffer in the V2C unit can be implemented with a total of k 3-state inverter-based buffers with their inputs connected to the output of the polarity selection MUX and with output-enable (OE) ports connected to outputs of the address decoder. The digital address decoder located within the V2C cell converts the address stored in the memory of the V2C unit into at least k control signals connected to the OE ports of the output buffer such that at most one of the 3-state inverter-based buffers is enabled at the time. Conversely, if a literal is not present in a clause, none of the OE ports are activated and the variable present at the input of the V2C cell remains disconnected from a clause. Fig. 3 shows an example circuit of the V2C unit corresponding to a clause with three literals. The unit contains 3 memory elements e.g., latches with access switches. One of the memory elements is dedicated to storing the polarity selection bit D_s , as well as the 2-bit address bits D_0 and D_1 for readout line selection, as shown in the schematic. A 2-to-4 address decoder is used to either connect the variable or its complement to one of the three readout lines L_1 , L_2 , and L_3 , or allow the unit and its corresponding variable to remain unconnected to any clause in the CFCCS array.

The CFCCS array contains N_C units and each unit has a total of k inputs. These inputs are connected to the k readout lines of the corresponding j^{th} column in the V2C array. Each unit in the CFCCS array contains a total of k NOR logic gates, with one NOR gate for each literal in clause C_j . Every NOR gate has $k-1$ inputs, which are connected to the inputs of the CFCCS unit. A $k-1$ input NOR gate corresponding to the m^{th} literal in clause C_j takes all literals at its inputs except the m^{th} literal. In other words, the inputs to a $k-1$ input NOR gate corresponding to the m^{th} literal are connected to all the readout lines from the j^{th} column in the V2C array, except the readout line corresponding to the m^{th} literal.

Each NOR gate in the j^{th} column produces a digital output $Z_{j,m}$ (where $1 \leq m \leq k$) which is connected to Clause-to-Coupling-Current (C2CC) array of size $N \times N_C$.

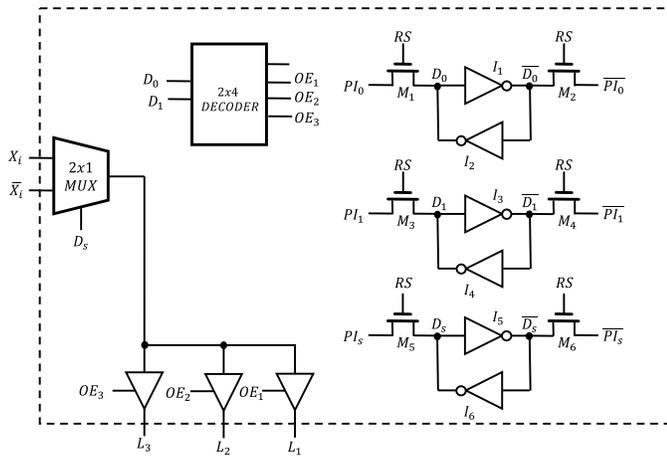


Fig. 3. Example unit element of the Variables-to-Clauses $N \times N_C$ array for a 3-SAT implementation.

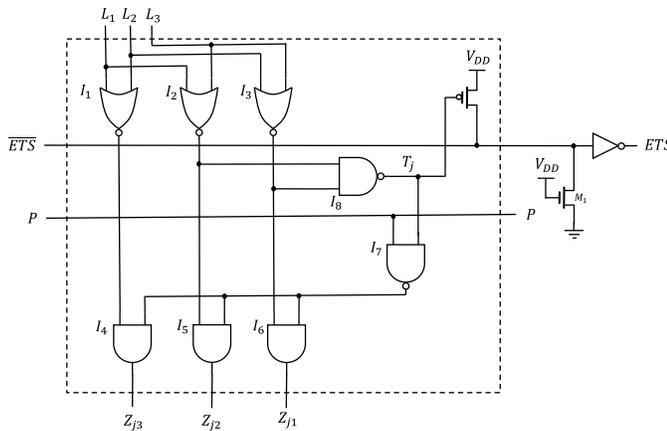


Fig. 4. The CFCCS unit with clause perturbation logic and perturbation signal P , satisfiability signal T_j and early termination signal ETS .

Each unit in the CFCCS array contains additional logic that receives outputs from the NOR gates to calculate satisfiability signal T_j of clause C_j . The logic producing satisfiability signal T_j can be configured so that T_j is equal to logic one if C_j is satisfied (or evaluates as TRUE) or $T_j = 0$, otherwise. For example, in the case of a clause with 3 literals ($k = 3$), the satisfiability signal T_j of clause C_j may be formed by applying a 2-input NAND gate on the outputs from any of the two NOR gates in the j^{th} CFCCS unit as shown in Fig. 4.

The SKI-SAT circuit described is further supplied with a perturbation circuit that perturbs the clause formation and coupling control signal generation logic in the CFCCS array. The CFCCS units are also supplied with additional logic gates (gates I_4 to I_7) that monitor the satisfiability signal T_j and a perturbation control signal P shared among all CFCCS units. An example CFCCS unit employing clause-perturbation logic is shown in Fig. 4.

To achieve perturbations in SKI-SAT, which allows for escaping local minima in the penalty function, each time a pulse in P is detected by the CFCCS units, all CFCCS units whose satisfiability signals T_j indicate that clause C_j is TRUE (or satisfied) generate logic zero on all of their

output lines $Z_{j,m}$. As a result, all satisfied clauses during the pulse duration in P do not contribute to the coupling currents. Therefore, these clauses do not affect the changes in nodal states v_i . In the absence of a pulse in P , all CFCCS units regardless of the value of their satisfiability signal T_j contribute to the coupling currents. The pulse stream P is generated so that the start of each pulse (i.e., the rising edge of the pulse) is random in time, while the duration of the pulse is fixed. In addition, the pulse stream P is generated such that the average pulse rate (i.e., number of pulses per second) diminishes over the annealing time. Further details about the perturbation mechanism deployed in SKI-SAT as well as a comparison to other perturbation methods are provided in Section IV-C.

The CFCCS units further incorporate logic gates to form an early termination signal (ETS) which serves as a trigger signal to sample the machine's current state of variables and store it as a solution found by the machine, as illustrated in Fig. 4. To achieve this, each CFCCS unit contains a PMOS transistor driven by the clause satisfiability signal T_j which together with the PMOS transistors from all CFCCS units and the shared NMOS transistor loading the \overline{ETS} readout line form an N_C -input NAND gate. The \overline{ETS} signal is then inverted to create ETS .

Outputs from all AND gates (a total of k outputs) in the j^{th} CFCCS unit are connected to corresponding j^{th} column of the Clause-to-Coupling-Current (C2CC) array containing $N \times N_C$ coupling units. For example, if clause C_j has three literals ($k = 3$), the j^{th} CFCCS unit produces three output signals $Z_{j,m}$ where $1 \leq m \leq 3$ that are sent to the coupling units in the j^{th} column of the C2CC array. A coupling unit at (i, j) location within the C2CC array corresponds to i^{th} node and j^{th} clause. The coupling units are provided with an output supplying either a positive or negative reference current I_{ref} or zero current depending on the polarity bit stored in the unit as well as the control signal supplied by the corresponding CFCCS units. Outputs from all coupling units in one row of the C2CC array are connected together for current summing and connected to the input of the node corresponding to that row, as illustrated in Fig. 2. In addition, each coupling unit in the C2CC array contains memory elements to store polarity bit as well as an address for digital multiplexer. For example, if the variable X_i is present in the C_j clause as the m^{th} literal, coupling unit at the (i, j) location in the C2CC array is configured through its MUX to receive control signal on $Z_{j,m}$ and its polarity bit is set to logic 1. In this configuration, each time the control signal $Z_{j,m}$ is logic 1, the coupling unit at location (i, j) provides a positive reference current I_{ref} at its output. Otherwise, when $Z_{j,m} = 0$ is received by the coupling unit (i, j) , its output current is set to zero. Likewise, in the case when a complemented variable \overline{X}_i is the m^{th} literal in clause C_j , the polarity bit stored in the coupling unit (i, j) is set to zero and each time $Z_{j,m} = 1$ is received, the coupling unit provides a negative reference current ($-I_{ref}$) at its output. Otherwise, when $Z_{j,m} = 0$ is received, the coupling unit produces no current at its output. The outputs from all coupling units in the i^{th} row of the C2CC array are connected together to the input of the i^{th} node.

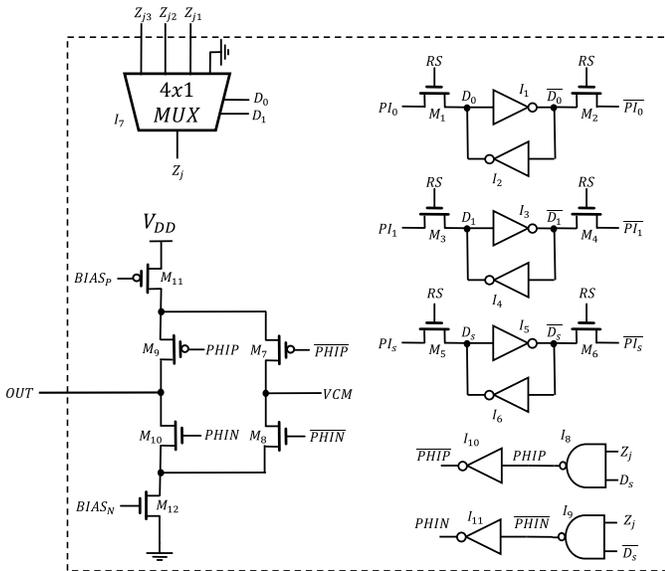


Fig. 5. Clause to the coupling current (C2CC) unit for a 3-literal SAT solver implementation.

An example schematic of the C2CC unit for 3-SAT implementation is shown in Fig. 5. In this circuit, three memory elements retain information about the polarity of D_s and address bits for the selection of one of the inputs $Z_{j,m}$ supplied by the corresponding CFCCS unit in the j^{th} column. The address bits are utilized for selection through the multiplexer I_7 . Depending on the polarity of D_s , either the NAND gate I_8 or I_9 will be activated to pull down either \overline{PHIP} or $PHIP$, respectively, if Z_j is set to high. When the polarity is positive (indicating D_s is high), transistors M_9 and M_8 are turned on, while M_7 and M_{10} are turned off. This setup connects the current source M_{11} to the output, leading to a positive current at the output of the coupling unit. Conversely, if the polarity is negative, M_9 and M_8 will turn off, and M_7 and M_{10} will turn on, thereby connecting the current sink M_{12} producing a negative current at the output. In the case where Z_j is set to low, the switches M_7 and M_8 are engaged and connect the current source M_{11} and current sink M_{12} to virtual ground V_{CM} , while switches M_9 and M_{10} are turned off leaving the output floating. It should be noted that switches M_7 and M_8 are not required for the operation of the coupling unit (i.e., they provide an auxiliary feature), however, employing these switches improves the linearity and settling speed of the current sources.

Although specific circuit parameters are not provided in this work, the following design guidelines have been considered. The V2C and CFCCS units consist entirely of standard cell library instances that are chosen with appropriate drive strength to drive their respective loads with acceptable delays. For the C2CC units and nodes, careful consideration is given to ensure that kT/C noise remains negligible and does not affect variable assignments throughout the annealing period. Furthermore, the current sources in the C2CC units are designed to remain in saturation, with a target of limiting current variation due to mismatch to a maximum of 5% standard deviation.

C. Physical Design of the System

The physical layout of a circuit requires meticulous attention to detail to minimize parasitic passive elements, such as stray capacitance and resistances, to achieve the desired performance goals. Additionally, the design aims to create a compact and smart footprint for individual cells detailed in Section III-B, ensuring they can be easily connected to neighboring cells without requiring additional routing and spacing. This approach facilitates the scalability of the overall architecture. With these considerations, the following units are designed: a V2C unit measuring $5 \mu\text{m}$ by $11.4 \mu\text{m}$, a CFCCS unit measuring $5 \mu\text{m}$ by $15 \mu\text{m}$, and a C2CC unit with dimensions of $5 \mu\text{m}$ by $10.2 \mu\text{m}$. Together, the V2C, CFCCS, and C2CC units form a column per clause, with a width of $5 \mu\text{m}$ and a height equal to the number of variables multiplied by the combined heights of the V2C and C2CC units, plus the height of one CFCCS unit. This configuration forms a matrix structure, with number of clause determining the width along the X-axis, while the height of the matrix is governed by the number of variables. For example, a SKI-SAT solver with 50 nodes and 218 clauses would have a width of $N_C \cdot 5 \mu\text{m}$, resulting in 1.09 mm. The height would be $N \cdot (11.4 + 10.2) \mu\text{m} + 15 \mu\text{m}$, equaling 1.095 mm. Additionally, a SKI-SAT computational node is designed with dimensions of $85 \mu\text{m} \times 10.2 \mu\text{m}$, with the nodes aligned to match the height of the C2CC units. The physical architecture and dimensions detailed here are visualized by the Fig. 2. Furthermore, Eq. (17) can be used to approximate the area requirements of a SKI-SAT solver containing N variables and N_C clauses:

$$A(N, N_C) = N \cdot A_{\text{node}} + (N \times N_C) \cdot (A_{V2C} + A_{C2CC}) + N_C \cdot A_{CFCCS} \quad (17)$$

For instance, a 50 variable and 218 clauses would occupy an area of 1.2369 mm^2 . The perturbation signal generating circuit proposed in [7] takes up an additional 0.0013 mm^2 bringing the total active area for a complete system to 1.2382 mm^2 excluding the auxiliary units for programming and communicating with the circuit.

Furthermore, a physical implementation of SKI-SAT in CMOS will introduce certain nonidealities. Two primary challenges that could limit the system's performance as it scales to hundreds or even thousands of variables are expected. The first is the propagation delay within the closed loop of the signal traveling from the nodes to the V2C, then to the CFCCS and C2CC, and finally back to the nodes. This delay will grow with a greater number of nodes and clauses. In order to address the growing propagation delay, the overall speed of the system can be slowed down proportionally by reducing coupling current I_{ref} or increasing nodal capacitance such that the propagation delay continues to remain negligible. So with the penalty of increasing the annealing period (i.e. the solution time), the delay problem should be alleviated. The second problem we expect to deal with at larger scale is the mismatch between current mirrors in the C2CC units and threshold variation of nodal comparators. The latter can be readily addressed by

auto-zeroing methods. The former will require a calibration scheme that reduces the mismatch variations to an acceptable level that does not impede the solution’s success rate.

IV. SIMULATION RESULTS AND DISCUSSION

This section presents the simulations performed on the proposed system and compares its performance to related work. The performance of the proposed design is evaluated both at the circuit and behavioral levels: the former ensures the circuit’s integrity, while the latter offers insights into the statistical performance of the system at larger scales.

A. Circuit-Level Simulation

The circuits described in the previous section are implemented in the Cadence Virtuoso environment using the TSMC 65nm process design kit. To reduce netlist size and simulation complexity, ideal components are used for the memory units (i.e., the memory latches in V2C unit of Fig. 3 and C2CC unit of Fig. 5 are replaced with DC voltage sources). Since memory does not affect system performance and is merely required for programmability (i.e., it is accessed only once to program the SAT problem onto the hardware), using ideal components does not compromise the reliability of the simulations. A uniform random 3-SAT Boolean function consisting of 50 literals and 218 clauses is chosen from the SATLIB [8] to evaluate the performance of the SKI-SAT circuit solver. SATLIB [8], a widely adopted benchmark, is utilized in this paper due to its incorporation of hard randomly generated formulas; such formulas arise when the ratio of clauses to variables is at a certain critical ratio [9]. In the following step, a circuit with 50 nodes, 50 by 218 units of V2C and C2CC, and 218 units of CFCCS is instantiated and programmed using a SKILL code to implement the function.

The circuit is initialized by resetting the nodal capacitors at each node to the common mode voltage, which is half the power supply – 0.6V for this process. The top strip in Fig. 6 illustrates the potential difference across the nodal capacitors, starting from the middle rail and progressively diverging toward either the ground or power supply levels, with the exception of a few variables, which remains in the vicinity of the middle rail while being slightly higher or lower than the comparator’s approximate threshold of 0.6V.

This behavior is expected and commonly observed in simulations, as some variables may cease charging or discharging once the system reaches a stable equilibrium. The middle strip of Fig. 6 depicts the quantized outputs of the nodal capacitors, which correspond to the digital values of the literals. Finally, the bottom strip shows the number of false clauses gradually decreasing, eventually reaching zero, indicating that the function is satisfied with the final combination of ‘1’s and ‘0’s present at the nodal outputs.

Although SKI-SAT is designed to decrease the number of unsatisfied clauses over time through its gradient decent nature, perturbations are needed to escape local minima in search for a variable assignment that maximizes the number of satisfied clauses. The particular perturbation scheme used in SKI-SAT introduces a “don’t care” state, during which the

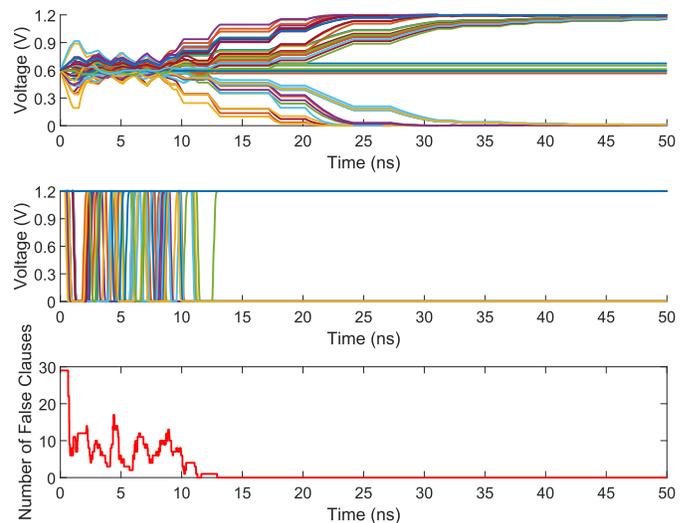


Fig. 6. A circuit-level simulation of SKI-SAT finding a solution for uf50-218/05. The top strip illustrates the evolution of nodal capacitor voltages, with most converging to either power rail. The middle strip displays the corresponding digitized literal values derived from these voltages. The bottom strip depicts the cost function plummeting to global minima indicating that the current variable assignment satisfies the function.

currents from satisfied clauses are disconnected from nodal capacitors. Temporarily disconnecting the satisfied clauses from contributing to the charging currents effectively allows the system to enforce an assignment of variables that would satisfy the remaining unsatisfied clauses with disregard to the already satisfied clauses, which might push the system to a higher energy state. The bottom strip in Fig. 6 shows that the number of unsatisfied clauses does not decrease monotonically, but occasionally increases to a higher energy level during the “don’t care” phases. This behavior is critical to maintain a robust solver that avoids getting stuck at local minima, which is a common issue for all stochastic solvers.

For this particular function, the system reaches the global minimum within 15 ns, consuming an average current of 24.23 mA during the 50 ns annealing period from a 1.2 V supply, resulting in 29.07 mW power demand, excluding the memory elements and perturbation unit. When the pseudo-random number generator circuit described in [7] is used to generate the perturbation signal P , the total active power requirement increases to 30.21 mW. The breakdown of contributions from the major blocks is illustrated in Fig. 7. The majority of energy consumption occurs before the machine settles into a local or global minimum. During this period, variables rapidly change values to satisfy more clauses. The machine reaches the global minimum around 13 ns, marked by the most significant switching activity observed in Fig. 6. After this, the current consumption of the nodes and V2C becomes insignificant, as literals no longer change values, eliminating the need for current drawn by the inverters within the nodes and V2C units. Meanwhile, the C2CC continuously supplies current to the nodal capacitors to maintain the found solution. The ETS shown in Fig. 4 can flag the solution is found and halt the system before the annealing period is over which further enhances the power efficiency by removing the need for replenishing the nodal capacitors.

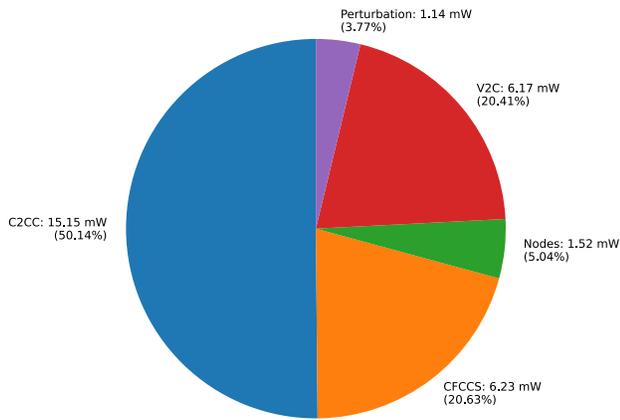


Fig. 7. Break down of power ratings of individual blocks while solving the uf50-218/05.

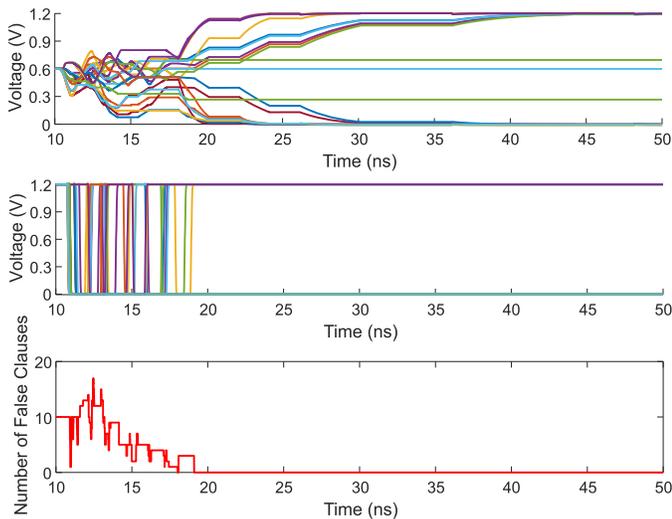


Fig. 8. SATLIB instance uf20-91/014 settling into a correct assignment. During the first 10 nanoseconds of the simulation, the system exhibits rapid switching activity as it searches for the global minimum. Random perturbations, which occasionally interrupt the gradient descent, prevent the cost function from decreasing monotonically.

A similar experiment is run on the SATLIB [8] instance uf20-91/014 for the purposes of investigating the probabilistic characteristics of the SKI-SAT on a commonly used benchmark item. Fig. 8 shows the 20 variables diverging and eventually settling into one of the global minima, thereby finding a satisfying assignment. For brevity, the initialization phase during the first 10 ns is omitted. Stochastic solvers inherently possess a probabilistic chance of solving a given function, and SKI-SAT is not an exception. To analyze the statistical properties of such solvers, it is necessary to perform a large number of runs. Circuit simulations, however, are notoriously computationally expensive and call for realistic simplifications. With this in mind, circuit simulations are only carried out on SAT problems with a smaller number of variables and a smaller number of runs. The obtained statistical results are used to verify the validity of the SKI-SAT behavioral model developed in MATLAB, which is then used to estimate SKI-SAT performance for large scale SAT problems.

Fig. 9 demonstrates five successful evaluations and five unsuccessful evaluations out of ten iterations with differ-

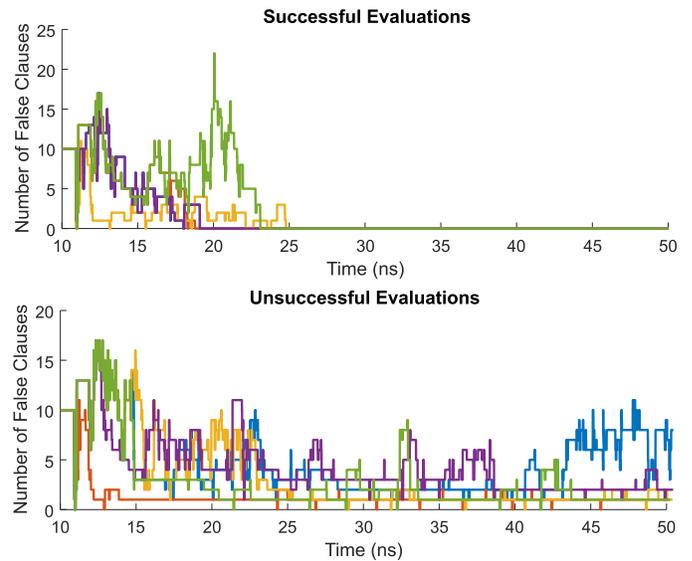


Fig. 9. Penalty function over time for a 20-variable 3-SAT instance (uf20-91/014) as solved by SKI-SAT solver for 10 iterations resulting in 50% success rate. *Top trace*: 5 iterations resulting in a successful evaluation - i.e., global minimum corresponding to zero unsatisfied clauses is found. *Bottom trace*: 5 unsuccessful iterations where the machine remains in the vicinity of local minima.

ent perturbation sequences for the uf20-91/014, resulting in a success rate of 0.5 for SKI-SAT for this particular function.

Furthermore, simulations show that the absence of perturbations causes the solver to become stuck in a local minimum due to the greedy gradient descent even when the transient noise is included in simulations. This observation highlights the importance of random perturbations to system performance. In addition, including the transient noise in simulations do not change the success rate, demonstrating the proposed circuit's resilience against inherently present electronic noise. Further, certain instances from SATLIB are quite trivial such that the correct solution is found rapidly without any perturbation. However, such instances are rare and unlikely to exist for instances with larger energy landscapes that contain far more literals and clauses.

B. High-Level Model

The SKI-SAT behavioral model is implemented in MATLAB as a fixed-step solver with discrete states. The fixed time step of $\Delta t = 20$ ps is chosen for this experiment. This value is empirically determined to balance numerical stability and computational efficiency. It ensures that the cost function during gradient descent (i.e., the number of unsatisfied clauses) remains stable and monotonically non-increasing, addressing potential numerical instability associated with the Forward Euler method [10]. Additionally, the chosen time step is not excessively small, avoiding prohibitively high simulation costs. In each time step, discrete voltage values on nodal capacitors are incremented or decremented by an integer multiple of ΔV . Assuming the nodal capacitance value of 200 fF and a reference current of 10 μA (i.e., current produced by the current source M_{11} and current sink M_{12} of the C2CC unit

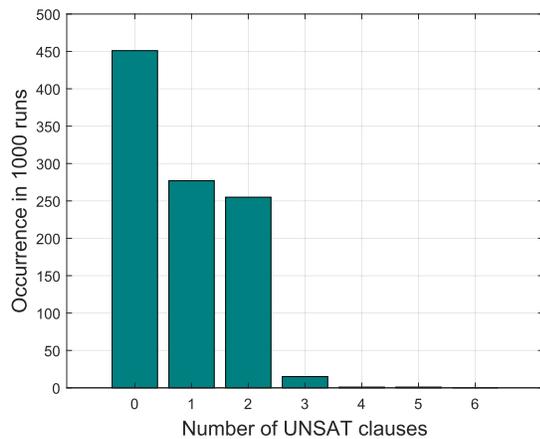


Fig. 10. Solution histogram for 1,000 iterations of SKI-SAT MATLAB model solving uf20-91/014 with 45.1% success rate.

in Fig. 5), the voltage increment ΔV is determined as

$$\Delta V = \frac{I_{\text{ref}} \times \Delta t}{C} = 1 \text{ mV} \quad (18)$$

The MATLAB script starts with fetching the Boolean function and initializes two sets of variables for the analog voltage values across the nodal capacitors and digitized outputs of the literals. In this model, the capacitor voltages are initialized to mid-rail value of $V_{dd}/2$ with additional random component to account for kT/C noise present in the circuit. Additionally, a random perturbation sequence is generated, clocked at a period of 320 ps, meaning the perturbation may change state every 16 simulation steps, establishing a realistic clocking speed. The code keeps track of how many clauses remain unsatisfied at each step and records the first step where all clauses are satisfied while simulating a system where the capacitor voltages are updated based on the satisfaction of clauses at every step. The model utilizes discrete voltage increments and decrements to simulate analog behavior. The literal values are updated by comparing the voltage values against the threshold. The system is in greedy mode when there is no perturbation and tries to diminish the number of unsatisfied clauses. On the other hand, the perturbation stops the greedy behavior and lets some of the satisfied clauses break by aborting the inputs from satisfied clauses. The perturbation increases the likelihood of the system eventually descending into global minima. Running the uf20-91/014 for 1,000 repeats with the described behavioral model reveals the success rate of 45.1%, as shown in Fig. 10, which is fairly consistent with the 50% success rate obtained through circuit simulations.

C. Perturbations in SKI-SAT

A key feature of the SKI-SAT is related to how it escapes local minima through the utilization of random perturbations. The primary goal of perturbation is to allow the system to occasionally choose a random path through the energy landscape, which does not align with the steepest gradient descent direction and does not necessarily lead to a reduction in the number of unsatisfied clauses. The SKI-SAT employs an highly-effective and resource-efficient solution to achieve this through the use of a single and global random

TABLE I
TRUTH TABLE OF CFCCS UNIT SHOWING INPUTS P, L_1, L_2, L_3 , AND CORRESPONDING OUTPUTS $T_j, Z_{j1}, Z_{j2},$ AND Z_{j3}

Inputs				Outputs			
P	L_1	L_2	L_3	T_j	Z_{j1}	Z_{j2}	Z_{j3}
0	0	0	0	0	1	1	1
0	0	0	1	1	0	0	1
0	0	1	0	1	0	1	0
0	0	1	1	1	0	0	0
0	1	0	0	1	1	0	0
0	1	0	1	1	0	0	0
0	1	1	0	1	0	0	0
0	1	1	1	1	0	0	0
1	0	0	0	0	1	1	1
1	0	0	1	1	0	0	0
1	0	1	0	1	0	0	0
1	0	1	1	1	0	0	0
1	1	0	0	1	0	0	0
1	1	0	1	1	0	0	0
1	1	1	0	1	0	0	0
1	1	1	1	1	0	0	0

(or pseudo-random) signal P . As shown in Fig. 4, the CFCCS units generate coupling signals $Z_{j,m}$ either to change variable assignments to achieve satisfiability or, if only one variable satisfies the clause, to maintain its current state. When a pulse is detected at P , the coupling signals $Z_{j,m}$ for satisfied clauses cease, allowing only unsatisfied clauses to generate coupling signals. During this phase, some of the satisfied clauses are expected to become unsatisfied while previously unsatisfied clauses will turn into satisfied, which will ultimately assist the system in finding the global minimum. The truth table for the CFCCS unit that summarizes the function of the perturbation signal P is depicted in Table I. The perturbation signal P is configured such that the perturbation pulse length is constant (e.g., 320 ps) and the pulse density (i.e., probability of entering perturbation mode) decays linearly over time.

In order to demonstrate the significance of the perturbations to system performance, various perturbation scenarios have been considered. A SATLIB instance, specifically uf50-218/0100, is simulated using the MATLAB model 1,000 times with three different approaches. First, when no perturbations are introduced forcing the machine to rely solely on its gradient descent in search for a global minimum, the success rate was quite low, with only 2 successful runs out of 1,000. Second, an alternative perturbation strategy is implemented, where $8mV_{RMS}$ Gaussian-distributed noise is randomly injected into the nodal capacitors. The RMS value of $8mV_{RMS}$ for injected noise is determined as optimal to maximize the success rate in this experiment. This approach improves the outcome, yielding 5 successful cases with zero unsatisfied clauses at the end of the annealing periods. Finally, with SKI-SAT's own perturbation mechanism, a success rate of 29.3% is achieved. The efficiency of the SKI-SAT perturbation as compared to the alternative method and the no-perturbation case is illustrated in Fig. 11.

D. Performance Comparison

A number of selected benchmark instances from SATLIB are used for comparison between a SAT solving IM [2], a hardware implementation of AmoebaSAT [11], WalkSAT [12], and SKI-SAT. AmoebaSAT [13] is a bio-inspired algorithm

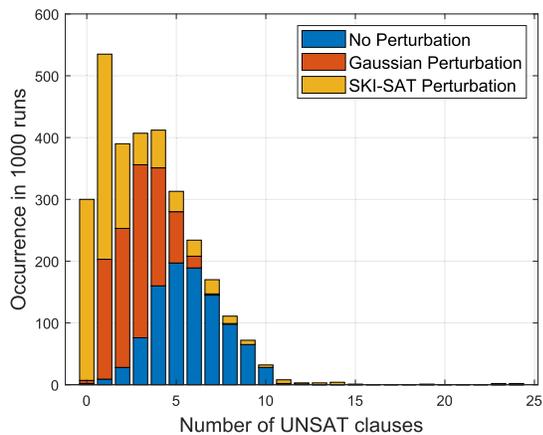


Fig. 11. SKI-SAT perturbation is effective in improving the success rate.

that has been demonstrated to solve SAT problems using nanoelectromechanical devices. Hara et al. [11] report the clock frequency and average number of cycles required for solution over 100 repeats of an FPGA implementation of AmoebaSAT. The provided information allows the calculation of the total time to solution (TtS) by multiplying the clock period by the average number of cycles. Furthermore, the latest version of the WalkSAT solver, found in the GitLab repository [14], is executed with the default settings on an Apple Silicon M1 CPU with 8 GB of RAM. The WalkSAT algorithm sets a predefined number of steps, known as the cutoff number, after which it either finds a solution or exits the iteration and restarts with a new initial assignment. Sweeping the cutoff parameter shows that the default cutoff setting provides a good balance. The instances tested for this paper neither suffer from a low cutoff rate, which leads to a low success rate, nor from an excessively large limit that causes the solver to get stuck in a cycle, taking millions of flips to escape without success. However, it is not guaranteed that WalkSAT will find a solution in every trial. Therefore, it is necessary to calculate the TtS in a way that accounts for unsuccessful cases. The number of repetitions required to solve a function with a 99% probability given a certain success rate (SR) is determined as

$$N_{99\%} = \left\lceil \frac{\log_{10}(0.01)}{\log_{10}(1-SR)} \right\rceil \quad (19)$$

The average time per assignment multiplied by the number of runs needed to solve an instance with 99% probability is used to estimate TtS for WalkSAT as shown in [15]. Lastly, the SKI-SAT behavioral model in MATLAB is employed to solve the same instances. The simulations assume a 300 ns annealing period, a 15 μ A reference current, and a 320 ps perturbation clock period, with the probability of entering perturbation mode decreasing from 90% to 50% during the annealing period. This model is used to project solution times for SKI-SAT, taking into account the success rate in the same manner as for WalkSAT. Table II summarizes the benchmarking results for the four methods.

The SKI-SAT outperforms the advanced WalkSAT, achieving solution times that are more than 10 times faster. In addition, a crucial figure of merit to compare different

TABLE II
COMPARISON OF TtS FOR HW AMOEBASAT, WALKSAT, AND SKI-SAT FOR DIFFERENT BENCHMARKS, IN UNITS OF μ S

Benchmark	IM [2]	AmoebaSAT [11]	WalkSAT [14]	SKI-SAT
uf20-91/011	12830	-	8	0.6
uf20-91/012	6300	-	6	0.3
uf20-91/013	27580	-	19	1.8
uf20-91/014	28070	-	16	1.2
uf20-91/015	4090	-	9	0.9
uf50-218/0100	-	4.13	187	4.2
uf50-218/0410	-	4.36	141	4.2
uf50-218/0767	-	8.30	365	10.8
uf100-430/0285	-	356	5872	344.7
uf150-645/0100	-	1832	6026	137.7
uf225-960/028	-	3078	1508	27.6

solvers is the energy to solution (EtS). EtS is determined by multiplying the time to solution (TtS) with the average power demand (P) to determine the energy required to solve a function as

$$EtS = TtS \times P \quad (20)$$

While WalkSAT solvers executed on von Neumann computing platforms consume energy on the order of Watts, the SKI-SAT circuit described in Section IV-A consumes about 14.27 mW (including perturbation sequence generation logic [7]) while solving the uf20-91/014 3-SAT instance. Apple's M1 CPU consumes about 7.5 W of power while executing the WalkSAT solver for the same instance, indicating hundreds of times power efficiency improvement. In order to estimate the power consumption, the powermetrics feature [16] is used to sample the average power drawn by the CPU with certain intervals. The idle CPU power is then subtracted from the CPU power observed while WalkSAT was running. It should be noted the EtS is slightly better when the CPU is operated in low power mode, however, this case yields a curtailed performance in terms of TtS. Consequently, the EtS figure of merit for SKI-SAT is up to thousands of times better than that of WalkSAT, offering a high-performance and low-power hardware alternative.

To further corroborate performance of SKI-SAT, particularly from a statistical perspective, and enable a broader comparison with existing published work, SKI-SAT is tested on 100 instances of uf20, uf50, uf100, and uf150 from SATLIB [8] with every instance run 1000 times. Fig. 12 shows the median TtS is 1.2 μ s for 20 variable instances and increases to 68.7 μ s for 150 variables, shown with corresponding inter-quartile ranges (IQR). Although the observed TtS growth in this result appears polynomial rather than exponential, this is contributed to the limited sample size (i.e., 100 instances per batch) and the small problem sizes used in the experiment. An exponential growth trend is expected instead, consistent with all known software or hardware based solvers for SAT problems with the constraint density N_C/N around critical value, as proved by Karp et al [17] and empirically verified in [9] and [18]. A similar trend is observed in EtS which starts from single-digit nJ for 20-variable instances and climbs to μ Js for 150-variable instances.

The presented SKI-SAT performance and scalable architecture merits a detailed comparison with existing work in the literature. Table III offers extensive studies with custom integrated circuit implementations that summarize the hardware

TABLE III
COMPARISON OF SAT SOLVERS

	CICC 2022 [19]	ISCAS 2024 [3]	ISSCC 2023 [20]	ISSCC 2023 [21]	ISSCC 2024 [22]	Nature 2024 [2]	Nature 2024 [23]	This Work
Technology	65nm CMOS	28nm CMOS	65nm CMOS	65nm CMOS	65nm CMOS	65nm CMOS	32nm CMOS	65nm CMOS
Supply Voltage	1–1.2V	Not reported	1.2V	0.7–1.2V	1–1.4V	Not reported	Not reported	1.2V
Power	27 mW	Not reported	32.5 mW	Not reported	Not reported	Not reported	13.35 mW	30.21 mW
Signal Domain	Mixed	Mixed	Mixed	Digital	Digital	Mixed	Digital	Mixed
Active Area (mm²)	4 ^b	Not reported	0.4	0.93	1.116	0.28	0.0214 ^d	1.2369 ^d
Maximum # of Variables	50	Not reported	60	128	Not reported	Not reported	100	50
Maximum # of Clauses	212	Not reported	252	1024	Not reported	Not reported	430	218
Experimented Problem (# of variables/# of clauses)	Hard 3-SAT (10/42)	Hard 3-SAT (50/218&150/645)	Hard 3-SAT (30/126&60/252)	Hard 3-SAT (60/258)	Hard 3-SAT (20/91&50/218)	Hard 3-SAT (20/91)	Hard 3-SAT (20/91&100/430)	Hard 3-SAT (50/218&150/645)
TtS	18 ms ^a	14 & 400 μ s ^c	11.25 & 125 ms	2.852 ms	7 & 37.4 μ s	12.83 ms	1.8 & 51.8 μ s ^c	4.2 & 68.7 μ s ^c
EtS	Not reported	80 & 5500 nJ ^{ac}	Not reported	4392 nJ	2.1 & 41.6 nJ	128 μ J	24 & 3430 nJ ^c	54.6 & 8293 nJ ^c

^a Estimated from plots.

^b Estimated from die micrograph.

^c Simulation result.

^d Estimated area.

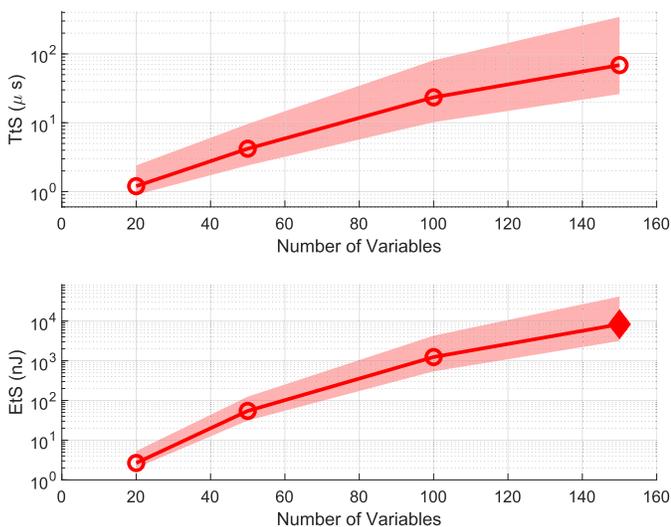


Fig. 12. Median TtS values shown in the top plot are calculated across batches of 100 instances from the SATLIB benchmark sets [8], which consist of problems with 20, 50, 100, and 150 variables. Each of the SAT instance in a batch with TtS corresponding to the median value is then simulated in Cadence environment to estimate power and corresponding EtS as shown in the bottom plot. Due to limited computational resources preventing circuit simulation of 150-variable problems, the EtS value for 150 variables was extrapolated using a second-order polynomial fit based on the known data points at 20, 50, and 100 variables. The shaded area in both plots represent the interquartile range (25–75%), while the markers denote the median for each problem size.

and performance specifications. The TtS and EtS values are derived from the plots and data provided in the referenced papers, as these metrics were not always explicitly reported. These values are considered key figures of merit for this study. For instance, while SKI-SAT consumes more power per iteration than [3] for 50-variable instances, its significantly shorter TtS results in a much lower EtS, reflecting the energy required to reach a solution with 99% probability.

V. CONCLUSION

A new power-efficient hardware-based SAT solver, SKI-SAT, is proposed in this paper that is capable of solving problems with more than quadratic terms and optimized for seamless CMOS implementation. The work demonstrates a

highly scalable architecture that inherently supports third-order polynomials found in SAT problem's cost functions. The circuit implementation validates the architecture as an effective solution for SAT problems, requiring over 300 times less power as compared to software-based solvers such as WalkSAT. Also as compared to WalkSAT, the behavioral model shows that SKI-SAT achieves, at least 10-fold reduction in solution time for the selected benchmark instances. Finally, the key performance metric, energy to solution (EtS), is orders of magnitude better than that of conventional SAT-solving algorithms. In conclusion, the proposed SKI-SAT solver demonstrates significant improvements in solution time and power efficiency while supporting seamless CMOS implementation.

REFERENCES

- [1] Y. Zhang, U. K. R. Vengalam, A. Sharma, M. Huang, and Z. Ignjatovic, "QuBRIM: A CMOS compatible resistively-coupled Ising machine with quantized nodal interactions," in *Proc. IEEE/ACM Int. Conf. Comput. Aided Design (ICCAD)*, Oct. 2022, pp. 1–8.
- [2] H. Cilasun et al., "3SAT on an all-to-all-connected CMOS Ising solver chip," *Sci. Rep.*, vol. 14, no. 1, May 2024, Art. no. 10757.
- [3] M. Hizzani et al., "Memristor-based hardware and algorithms for higher-order Hopfield optimization solver outperforming quadratic Ising machines," in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, May 2024, pp. 1–5.
- [4] E. Elmitwalli, Z. Ignjatovic, and S. Kose, "Utilizing multi-body interactions in a CMOS-based Ising machine for LDPC decoding," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 71, no. 1, pp. 40–50, Jan. 2024.
- [5] Y. Vazel, G. Weissenbacher, and S. Malik, "Boolean satisfiability solvers and their applications in model checking," *Proc. IEEE*, vol. 103, no. 11, pp. 2021–2035, Nov. 2015.
- [6] F. Legendre, G. Dequen, and M. Krajecki, "Encoding hash functions as a SAT problem," in *Proc. IEEE 24th Int. Conf. Tools Artif. Intell.*, vol. 1, Nov. 2012, pp. 916–921.
- [7] J. Wu, A. Y. Salim, E. Elmitwalli, S. Köse, and Z. Ignjatovic, "A pseudo-random number generator for multi-sequence generation with programmable statistics," 2024, *arXiv:2501.00193*.
- [8] H. Hoos and T. Stützle, "SATLIB: An online resource for research on SAT," in *Proc. SAT*, Jan. 2000, pp. 283–292.
- [9] B. Selman and S. Kirkpatrick, "Critical behavior in the computational cost of satisfiability testing," *Artif. Intell.*, vol. 81, nos. 1–2, pp. 273–295, Mar. 1996.
- [10] E. Hairer and G. Wanner, *Solving Ordinary Differential Equations II: Stiff and Differential-Algebraic Problems* (Springer Series in Computational Mathematics), vol. 14, 2nd ed., Berlin, Germany: Springer-Verlag, 1996.

- [11] K. Hara, N. Takeuchi, M. Aono, and Y. Hara-Azumi, "Amoeba-inspired stochastic hardware SAT solver," in *Proc. 20th Int. Symp. Quality Electron. Design (ISQED)*, Mar. 2019, pp. 151–156.
- [12] B. Selman, H. Kautz, and B. Cohen, "Noise strategies for improving local search," in *Proc. AAAI*, Seattle, WA, USA, Oct. 1994, pp. 337–343.
- [13] M. Aono et al., "Amoeba-inspired nanoarchitectonic computing: Solving intractable computational problems using nanoscale photoexcitation transfer dynamics," *Langmuir*, vol. 29, no. 24, pp. 7557–7564, Jun. 2013.
- [14] H. Kautz. *Walksat*. GitLab Repository. Accessed: Sep. 25, 2024. [Online]. Available: <https://gitlab.com/HenryKautz/Walksat>
- [15] R. Hamerly et al., "Experimental investigation of performance differences between coherent Ising machines and a quantum annealer," *Sci. Adv.*, vol. 5, no. 5, May 2019, Art. no. eaau0823.
- [16] Apple Inc. (2024). *Powermetrics Manual Page*. Accessed: Oct. 30, 2024. [Online]. Available: <https://www.unix.com/man-page/osx/1/powermetrics/>
- [17] R. M. Karp, "Reducibility among combinatorial problems," in *Complexity of Computer Computations*. Cham, Switzerland: Springer, 1972, pp. 85–103.
- [18] B. Selman, D. G. Mitchell, and H. J. Levesque, "Generating hard satisfiability problems," *Artif. Intell.*, vol. 81, nos. 1–2, pp. 17–29, Mar. 1996.
- [19] M. Chang, X. Yin, Z. Toroczka, X. Hu, and A. Raychowdhury, "An analog clock-free compute fabric base on continuous-time dynamical system for solving combinatorial optimization problems," in *Proc. IEEE Custom Integr. Circuits Conf. (CICC)*, Apr. 2022, pp. 1–2.
- [20] D. Kim, N. M. Rahman, and S. Mukhopadhyay, "29.1 a 32.5mW mixed-signal processing-in-memory-based k-SAT solver in 65nm CMOS with 74.0% solvability for 30-variable 126-clause 3-SAT problems," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2023, pp. 28–30.
- [21] S. Xie et al., "29.2 snap-SAT: A one-shot energy-performance-aware all-digital compute-in-memory solver for large-scale hard Boolean satisfiability problems," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2023, pp. 420–422.
- [22] C. Shim, J. Bae, and B. Kim, "30.3 VIP-SAT: A Boolean satisfiability solver featuring 5×12 variable in-memory processing elements with 98% solvability for 50-variable 218-clause 3-SAT problems," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, vol. 67, 2024, pp. 486–488.
- [23] T. Bhattacharya et al., "Computing high-degree polynomial gradients in memory," *Nature Commun.*, vol. 15, no. 1, p. 8211, Sep. 2024, doi: [10.1038/s41467-024-52488-y](https://doi.org/10.1038/s41467-024-52488-y).



Ahmet Yusuf Salim received the B.S. degree in electronics and communication engineering from Istanbul Technical University, Istanbul, Türkiye, in 2022. He is currently pursuing the Ph.D. degree in electrical and computer engineering with the University of Rochester, Rochester, NY, USA. He worked with ATEK Microwave, Istanbul, and held an internship position with Advanced Micro Devices (AMD), Rochester, NY, USA. His research interests include analog and RF integrated circuit design, Ising machines, and combinatorial optimization problems.



Bart Selman (Member, IEEE) received the Ph.D. degree from the University of Toronto. He is currently a Professor of computer science with Cornell University. He previously was with AT&T Bell Laboratories. He has (co-)authored over 100 publications, including six best paper awards. His papers have appeared in venues spanning Nature, Science, Proceedings of the National Academy of Sciences, and a variety of conferences and journals in AI and computer science. His research interests include computational sustainability, efficient reasoning procedures, planning, knowledge representation, and connections between computer science and statistical physics. He is a fellow of American Association for Artificial Intelligence and a fellow of American Association for the Advancement of Science. He has received the Cornell Stephen Miles Excellence in Teaching Award, the Cornell Outstanding Educator Award, the NSF Career Award, and the Alfred P. Sloan Research Fellowship.



Henry Kautz (Member, IEEE) received the A.B. degree from Cornell University and the Ph.D. degree from the University of Rochester. From 2018 to 2022, he was the Division Director for Information and Intelligent Systems (IIS) of the National Science Foundation. He was the Founding Director of the Goergen Institute for Data Science, University of Rochester. He is currently a Professor of computer science with the University of Virginia, Charlottesville. His research interests include practical algorithms for solving worst-case intractable problems in logical and probabilistic reasoning, pervasive healthcare applications of AI, and social media analytics. He received the 2018 ACM-AAAI Allen Newell Award for career contributions that have breadth within computer science and that bridge computer science and other disciplines.



Zeljko Ignjatovic (Member, IEEE) received the B.S. degree in electrical engineering from the University of Novi Sad, Serbia, in 1999, and the M.S. and Ph.D. degrees in electrical and computer engineering from the University of Rochester, Rochester, NY, USA, in 2001 and 2004, respectively. He is currently an Associate Professor of electrical and computer engineering with the University of Rochester. His research interests include analog and mixed-signal circuit design, development of analog-to-digital converters, image sensors, CMOS-based Ising machines, radar and ultrasound imaging techniques, and related signal processing methods.



Selçuk Köse (Member, IEEE) received the B.S. degree in electrical and electronics engineering from Bilkent University, Ankara, Türkiye, in 2006, and the M.S. and Ph.D. degrees in electrical engineering from the University of Rochester, Rochester, NY, USA, in 2008 and 2012, respectively. He worked with TÜBİTAK, Ankara, Türkiye; Intel Corporation, Santa Clara, CA, USA; and Freescale Semiconductor, Tempe, AZ, USA. He is currently a Professor with the Department of Electrical and Computer Engineering, University of Rochester. Prior to joining the University of Rochester, he was an Associate Professor with the University of South Florida, Tampa, FL, USA. His current research interests include integrated voltage regulation, 3-D integration, hardware security, graphene nanoribbon transistors, and cryogenic electronics.

Dr. Köse was a recipient of the NSF CAREER Award, the Cisco Research Award, the USF College of Engineering Outstanding Junior Researcher Award, and the USF Outstanding Faculty Award. He is/was an Associate Editor of IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—I: REGULAR PAPERS, *Nature Computer Science* (Springer), and *Microelectronics Journal* (Elsevier). He has served on the technical program and organization committees of various conferences.