25.1 A Physics-Inspired Oscillator-Based Mixed-Signal Optimization Engine for Solving 50-Variable 218-Clause 3-SAT Problems with 100% Solvability and 31.7µs Solution Time

Evangelos Dikopoulos, Ying-Tuan Hsu*, Luke Wormald*, Wei Tang, Zhengya Zhang, Michael P. Flynn

University of Michigan, Ann Arbor, MI

*Equally Credited Authors (ECAs)

The Boolean satisfiability (SAT) problem is a fundamental NP-complete problem, and efficiently solving it would revolutionize fields like optimization, artificial intelligence, cryptography, software, and hardware verification. Physics-inspired computers offer significant advantages, including continuous-time (CT) operation, massive parallelism, and increased energy efficiency. Solvers that map the optimization objective to a dynamical system of spins [1] have been shown to outperform classical discrete optimization solvers. Recent work maps 3-SAT to systems of coupled spins but suffers from long solution times or low solvability [2,3], and is limited to problems with only 20 variables [3]. [4] decomposes 3-SAT problems to an all-to-all connected analog Ising machine, but the proposed iterative compute scheme results in ms-level solution times. A digital solver based on an array of processing elements (PEs) [5] has reported competitive performance, but it does not account for the substantial preprocessing time and energy required to embed the problem into the available hardware - embedding itself is a complex optimization problem. Furthermore, the system in [5] can connect only up to 32 clauses to a given variable, limiting its ability to solve real-world satisfiability problems where some spins are highly connected.

This work advances the field of Combinatorial Optimization Problem (COP) solvers by demonstrating a massively parallel oscillator-based direct 3-SAT engine that leverages physics-inspired heuristics in a mixed-signal compute fabric for unprecedented solutions times and energy efficiency without requiring any preprocessing. The prototype solves 20-variable 3-SAT problems more than 4 times faster than state-of-the-art solvers [3,5], is 22 times faster for larger (≥ 50 variables) problems than the fully connected state-of-the-art solver [6], and does not require problem embedding. We introduce innovations in the architecture and at the circuit level (Fig. 25.1.1): (i) A new continuous time dynamical system with multi-bit bidirectional spin-clause interconnect naturally escapes local minima and explores the solution space. (ii) A robust and highly scalable mixed analog/digital crossbar-based feedback system enables unrestricted all-to-all 3-SAT connectivity. Current summation-based feedback enables scalable true CT asynchronous computation. (iii) A relaxation-oscillator (RXO) compute node with Dynamic CT Injection (DaCTI) enables state-of-the-art solution time and reduces solution energy by 8 times (for 50 variables) compared to a constantly oscillating node. DaCTI also alleviates the need for a custom sampler (e.g. unlike [4]), and reduces PVT sensitivity. (iv) A low-power feedback interface alleviates the effect of RC delay in the current summation nodes.

The proposed solver directly maps optimization problems to hardware to significantly accelerate computations by taking advantage of CT operation, parallelism, multi-bit operation, and highly efficient mapping. A 3-SAT problem seeks a truth value assignment to literals (SAT variables or their negations) in a Boolean formula $\bigwedge_{i=1}^{m}(l_{i1} \vee l_{i2} \vee l_{i3})$ that makes the formula true. In this work, the problem variables correspond to oscillator compute nodes, and the graph directly programs the mixed-signal feedback system that implements clauses. Figure 25.1.2 illustrates this architecture, where the 3-SAT graph directly maps onto an array of relaxation oscillator spins. The clause cells, implemented with NOR3 gates and current sources, bidirectionally interconnect with the spins through twin analog/digital crossbars. The feedback system continuously adjusts spin phases and states through CT interactions until it converges to a solution.

A key design highlight of our architecture is the multi-bit bidirectional interconnect between spins and clauses. We leverage programmable mixed-signal feedback provided by twin analog/digital crossbars. Each CT spin has a binary output value (i.e., the SR latch state) and an internal phase (i.e., the capacitor voltage). In the feedforward path, the digital crossbar propagates the CT binary states of the spins to the clause inputs. The outputs of the NOR3 gates in the clause cells (i.e., SAT or UNSAT) enable current sources feeding crossbar feedback current lines (Fig. 25.1.3). For each spin, the currents from associated clauses sum on the current summation lines (i.e. ISUM). The feedback currents rotate the spin's phase proportionally to the number of its broken clauses. This architecture improves the average solution time by more than 4 times compared to [3], where the feedback generates hard binary spin flips.

Each spin's phase rotates proportionally to the number of associated broken clauses, enabling a parallel search for the states that maximize satisfied clauses. The internal analog phase adds stochasticity, while the guaranteed rotation of spins with broken clauses prevents getting stuck in local minima. After spins reset to zero, the system only settles when a full 3-SAT solution is found. Measurements show no benefit with random initial conditions over fixed ones, eliminating the need for a random number generator.

A key design consideration is the use of a relaxation oscillator with dynamic CT injection (Fig. 25.1.4), which offers significant benefits for dynamical system solvers, particularly when compared to ring oscillators [4]: (i) In-oscillator current and voltage DACs allow multi-bit precise feedback; (ii) Dynamic CT operation allows for significant energy savings; (iii) Current source charging significantly reduces sensitivity to PVT variations; (iv) The digital output of the relaxation oscillator is robust and does not burden the oscillator.

The proposed RXO-based compute node with dynamic CT operation provides significant energy savings. While a standard relaxation oscillator is more PVT robust [7] than a ring oscillator [4], the continuous charging still wastes energy by repeatedly changing the spin state. In the proposed RXO, the phase rotates only in the presence of a feedback signal, saving energy. Measurements show that DaCTI improves the solver's solution time by 1.9× for 20 variables and 6× for 50 variables. Furthermore, DaCTI removes the need for a multi-phase sampling system [4] since the spins freeze to the final solution.

The proposed CT current-domain feedback architecture (DaCTI) implements a non-linear quantization of the feedback current to regularize the feedback and isolate the oscillator capacitor from the feedback system. We leverage 4-level in-oscillator flash-ADCs to quantize each spin's feedback current, from the analog crossbar. Limiting the feedback signal to four levels with this ADC relieves the system's bias toward changing the state of spins that are highly connected and consequently more likely to receive feedback. The ADC drives an in-oscillator 4-level IDAC that integrates regularized current on the RXO capacitors and a 4-level Voltage DAC (VDAC) that changes the comparator threshold. The combined use of the IDAC and VDAC enables solution times that are orders of magnitude lower compared to IDAC alone.

Another significant benefit is that this architecture does not load the spin's capacitor and isolates the oscillator from the current summation line large (280fF) parasitic capacitance. A low-power trans-impedance amplifier (TIA) terminates each vertical current summation line in the analog crossbar with a low impedance, minimizing the effect of the parasitic capacitance. We optimize the analog feedback path delay by pre-charging the clause current sources' parasitic capacitance, reducing delay by 7×. In the feedforward path of the system, spin outputs are rebuffered at each intersection of the digital crossbar.

The 28nm CMOS prototype has a core area of 0.58mm^2. An on-chip digital controller orchestrates the compute cycles and facilitates extensive performance measurements. Loading a problem to the system does not require preprocessing. The host programs the 3-SAT graph to crossbar connection points through a scan chain. The prototype's performance was rigorously evaluated with the well-known SATLIB suite (Fig. 25.1.5). All 1000 problems in the 20-variable and 50-variable libraries were evaluated 100 times each. For 20-variable problems, the prototype achieves a mean solution time of 1.6µs, an improvement of more than 4× over [3,5], and an Energy Delay Product (EDP) that surpasses both (Fig. 25.1.6). The mean solution time and energy for 50-variable problems are 31.7µs and 268.9nJ respectively, surpassing the previous state-of-the-art solution for large 3-SAT instances without preprocessing [6] (60-variable problems) by more than 22× and 4×, respectively. All tested instances are considered hard problems. The proposed satisfiability solver (Fig. 25.1.7) introduces a physics-inspired continuous-time architecture that is scalable, highly efficient, and enables solution times and energy efficiency beyond the state-of-the-art without requiring any preprocessing.

Acknowledgement:
This work was supported by DARPA QuICC.

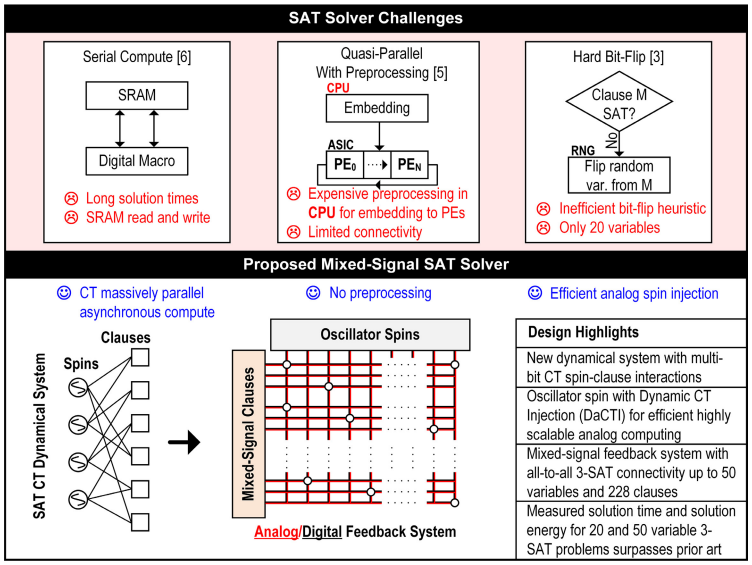

Figure 25.1.1: Overview of recent Boolean satisfiability (SAT) solver architectures and their limitations. Introduction to the proposed mixed-signal SAT solver and design highlights.

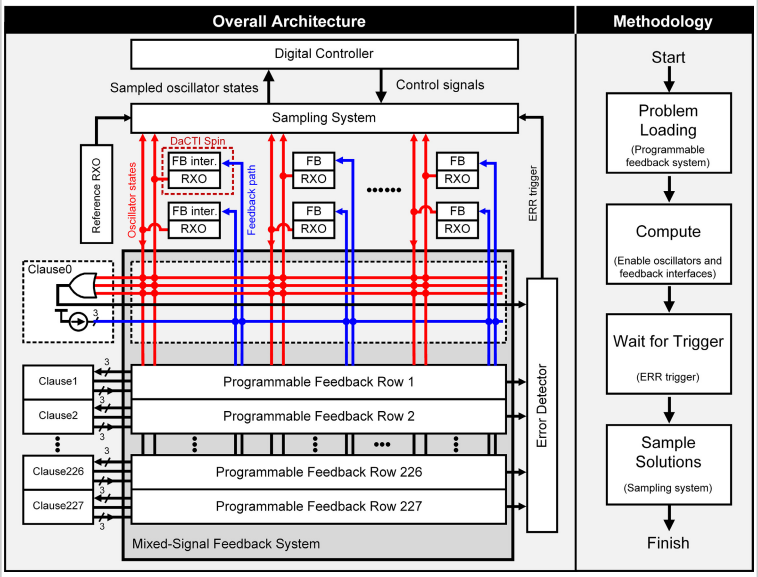

Figure 25.1.2: Proposed solver architecture with key design blocks and methodology.

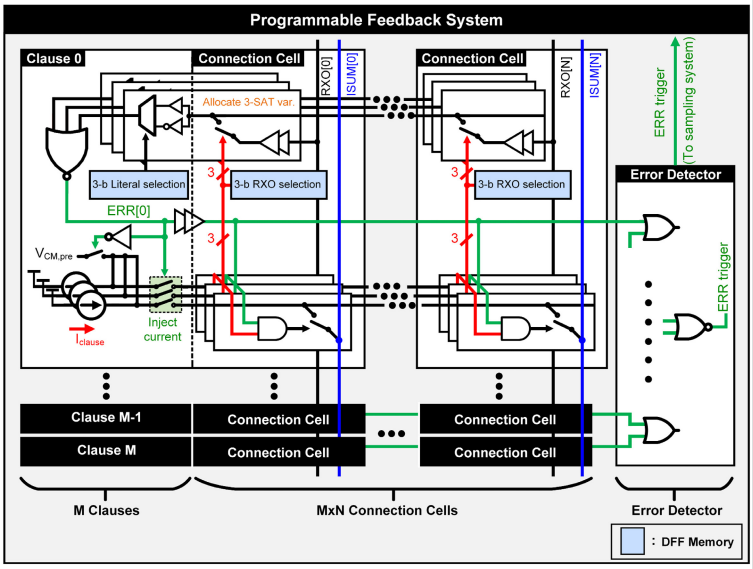

Figure 25.1.3: Crossbar-based spin feedback architecture. Current summation line (ISUM) architecture with TIA and pre-charge improves delay.

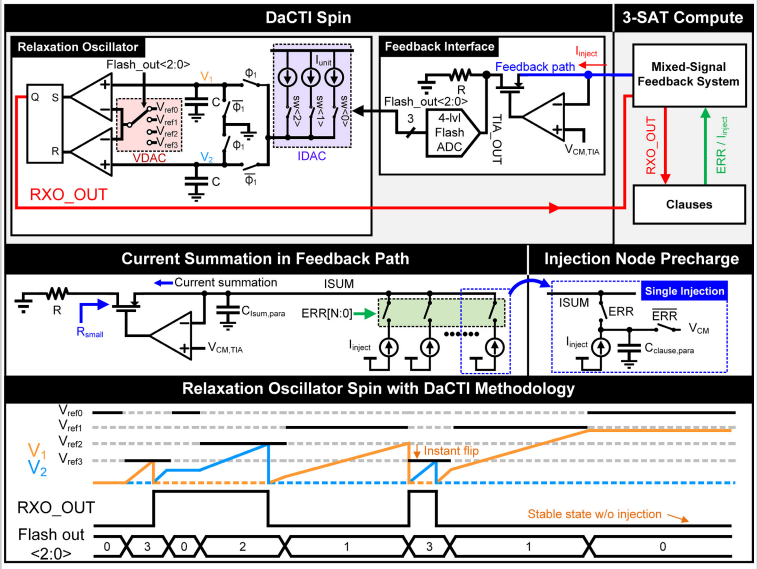

Figure 25.1.4: Relaxation-oscillator-based spin with Dynamic CT Injection (DaCTI). Current summation architecture and DaCTI timing diagram.

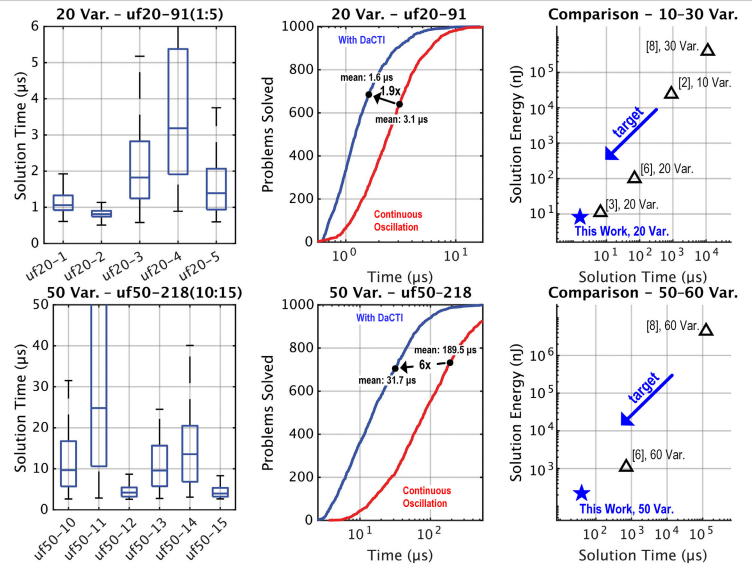

Figure 25.1.5: Measured results for 3-SAT problems with 20 variables/91 clauses and 50 variables/218 clauses from the SATLIB benchmark. Solution time distribution of instances and cumulative results for all benchmark problems. Prior art comparison.

	CICC 2022 [2]		ISSCC 2023 [8]		ISSCC 2024 [5]		ISSCC 2023 [6]		VLSI 2024 [3]	This Work	
Technique	Energy-based Dynamics		Neural Network		PE Array		In-Memory Clauses		Stochastic Bit Flip	Dynamic CT Injection	
Architecture	Analog, CT		Digital		Digital, DT		Digital, DT		Analog, DT/CT	Analog, CT	
# of Variables/ Clauses[A]	10/42	20/84	30/126	60/252	20/91	50/218	20/86	60/258	20/91	20/91	50/218
Solvability[B] (%)	92	74	74	31.5	100	98	NA	72	100	100	100
# of Problems Tested	200[4]		200[4]		100[3]		1000[4]		1000[3]	1000[3]	
All-To-All 3-SAT Connectivity	YES		YES		NO[2], limited connectivity		YES		YES	YES	
Pre-processing	NO		NO		REQUIRED[1], time/ energy not measured		NO		NO	NO	
Solution Time[C] (µs)	900	NA	11e3[5]	125e3[5]	7.0[1]	18.7[1]	70	713	6.6	1.6	31.7
Solution Energy[D] (nJ)	24.3e3	NA	398e3[5]	4425e3[5]	2.1[1]	20.8[1]	100	1098	11	7.8	268.9
EDP[E] (µs x nJ)	21.9e6	NA	4480e6	553e9	14.7[1]	389[1]	7000	783e3	72.6	12.6	8515
Area (mm²)	4		0.4		1.1		0.93		0.37	0.58	
Process	65nm		65nm		65nm		65nm		65nm	28nm	

[1]Time and energy for pre-processing is not measured. [2]Can connect a given variable to only 32 clauses. [A]The Clauses to Variables Ratio (CVR) is a difficulty metric for SAT problems. [B]Number of problems that can be solved in a given time. [C]Mean time to find a solution. [D]Mean energy to find a solution. [E]Energy Delay Product (EDP), i.e. Solution Time x Solution Energy. [3]Used the SATLIB Benchmark (https://www.cs.ubc.ca/~hoos/SATLIB/benchm.html). [4]Benchmark not available. [5]Median solution time for 99% satisfiability.

Figure 25.1.6: Performance summary and comparison with prior SAT solvers.

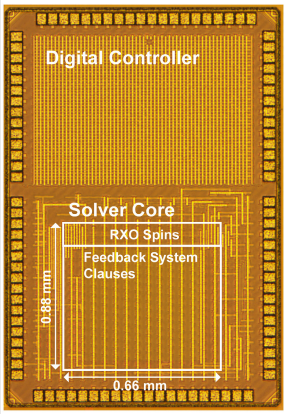

Prototype	CT Mixed-Signal SAT Solver
Technology	28 nm
Area	0.58 mm²
Supply	0.9 V
Variables (Max.)	50
Clauses (Max.)	228
Mean Solution Time	1.6 µs - 20 Var.
	31.7 µs - 50 Var.
Mean Solution Energy	7.8 nJ - 20 Var.
	268.9 nJ - 50 Var.
Solvability	100% - 20,50 Var.

Figure 25.1.7: Prototype micrograph and summary.

References:
[1] B. Molnar et al., "A continuous-time MaxSAT solver with high analog performance," Nature Communications 9, Article number: 4864, 2018.
[2] M. Chang et al., "An Analog Clock-free Compute Fabric base on Continuous-Time Dynamical System for Solving Combinatorial Optimization Problems," IEEE Custom Integrated Circuits Conference (CICC), Newport Beach, CA, USA, 2022, pp. 1-2.
[3] Q. Zhang et al., "A Stochastic Analog SAT Solver in 65nm CMOS Achieving 6.6µs Average Solution Time with 100% Solvability for Hard 3-SAT Problems," IEEE Symposium on VLSI Technology and Circuits (VLSI Technology and Circuits), Honolulu, HI, USA, 2024, pp. 1-2.
[4] H. Cılasun et al., "3SAT on an all-to-all-connected CMOS Ising solver chip," Nature Scientific Reports 14, Article number: 10757, 2024.
[5] C. Shim, J. Bae and B. Kim, "VIP-Sat: A Boolean Satisfiability Solver Featuring 5×12 Variable In-Memory Processing Elements with 98% Solvability for 50-Variables 218-Clauses 3-SAT Problems," IEEE International Solid-State Circuits Conference (ISSCC), San Francisco, CA, USA, 2024, pp. 486-488.
[6] S. Xie et al., "Snap-SAT: A One-Shot Energy-Performance-Aware All-Digital Compute-in-Memory Solver for Large-Scale Hard Boolean Satisfiability Problems," IEEE International Solid-State Circuits Conference (ISSCC), San Francisco, CA, USA, 2023, pp. 420-422.
[7] Evangelos Dikopoulos et al., "A Relaxation Oscillator-Based Probabilistic Combinatorial Optimization Engine for Soft Decoding of LDPC Codes," European Solid-State Electronics Research Conference (ESSERC), Bruges, Belgium, 2024, pp. 717-720.
[8] D. Kim, N. M. Rahman and S. Mukhopadhyay, "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," IEEE International Solid-State Circuits Conference (ISSCC), San Francisco, CA, USA, 2023, pp. 28-30.

979-8-3315-4101-9/25/$31.00