

Late Breaking Results: New Computational Results and Hardware Prototypes for Oscillator-based Ising Machines

Tianshi Wang
University of California, Berkeley
Berkeley, CA, USA

Leon Wu
University of California, Berkeley
Berkeley, CA, USA

Jaijeet Roychowdhury
University of California, Berkeley
Berkeley, CA, USA

Abstract

In this paper, we report new results on a novel Ising machine technology for solving combinatorial optimization problems using networks of coupled self-sustaining oscillators. Specifically, we present several working hardware prototypes using CMOS electronic oscillators, built on breadboards/perfboards and PCBs, implementing Ising machines consisting of up to 240 spins with programmable couplings. We also report that, just by simulating the differential equations of such Ising machines of larger sizes, good solutions can be achieved easily on benchmark optimization problems, demonstrating the effectiveness of oscillator-based Ising machines.

1 Introduction

The Ising model is a mathematical model originally used to study ferromagnetism. It describes spins coupled in a graph that try to minimize a collective energy, *aka*, the Ising Hamiltonian.

$$\min H \triangleq - \sum_{1 \leq i < j \leq n} J_{ij} s_i s_j - \sum_{i=1}^n h_i s_i, \quad \text{s.t. } s_i \in \{-1, +1\}, \quad (1)$$

where n is the number of spins; $\{J_{ij}\}$ and $\{h_i\}$ are real coefficients.

Finding optimal spin configurations that minimize the Ising Hamiltonian, *aka*, the Ising problem, is in general difficult [1], even with purposely built digital accelerators [2]. A physical implementation of coupled spins that directly perform the minimization in an analog way, namely an Ising machine, therefore becomes very attractive for potential speed and power advantages. As many difficult real-world optimization problems are equivalent to the Ising problem [3], Ising machines have been attracting considerable research attention in recent years, with incarnations mostly based on novel devices, such as optical cavities [4], nanomagnets [5], and quantum circuits [6]. A recent work — oscillator-based Ising machine (OIM) [7, 8] — shows that almost all types of nonlinear self-sustaining oscillators are suitable to represent Ising spins physically. As many tried-and-tested types of such oscillators already exist, this scheme offers the advantages of scalability to large numbers of spins, high-speed and low-power operation, and straightforward design and fabrication using standard circuit technologies. In the recent months, we have achieved several results from testing the feasibility of the OIM idea:

1. We have built several hardware prototypes, starting from coupling 8 CMOS LC oscillators on a breadboard, moving on to soldering 32 of them on perfboards, then to PCB designs of size 64 and 240 with programmable couplings. We have tested all prototypes on many instances of Ising problems of their corresponding sizes; every prototype can achieve global optima for these problems. We plan to make them **open hardware projects**, with schematics, design files, tutorials and code released to the public, so that researchers (and hobbyists) can reproduce the results and improve upon the design.
2. We have simulated larger-sized OIMs, trying them on **all 54 problems in the G-set** [9] (available at [10]) — a widely used benchmark set for MAX-CUT problems (which have a direct mapping to Ising problems). Much to our surprise, without changing any parameters (coupling strength, noise level, *etc.*) across different problems, the results match 21/54 and improve upon 17/54 previously published optimal solutions.

The remainder of this paper contains a brief overview of OIM’s mechanism and more details on the new results listed above.

2 Oscillator-based Ising Machines

OIM’s operation relies on a special type of injection locking, known as Subharmonic Injection Locking (SHIL). Under SHIL, when an oscillator is perturbed by an input at twice its natural frequency (often called a SYNC signal), it can develop bistable phase-locked states, separated by a 180° phase difference.




Figure 1: An Ising model and its OIM schematic.

Multiple such oscillators, whose phase is “binarized” by a common SYNC through SHIL, can be networked such that they synchronize to binary phase configurations that minimize Ising Hamiltonians, as illustrated in Fig. 1. It can be shown [7] that when the coupling coefficients (often represented by conductances connecting electronic oscillators¹) are proportional to J_{ij} s and h_i s in (1), the oscillator network naturally minimizes a global Lyapunov function that can be made equivalent to the Ising Hamiltonian, thus physically implementing an Ising machine.

3 Hardware Prototypes

The schematics of several OIM hardware prototypes are summarized in Fig. 2. They all use CMOS LC oscillators made with cross-coupled inverters (from TI SN74HC04N ICs), fixed inductors, trimmer capacitors and a 5V single supply. OIM8 and OIM32 use 33μH inductors with capacitors tuned to around 30pF, for a natural frequency of 5MHz. We could manually plug in resistors and potentiometers on the breadboard to try different problems on OIM8, then read back results using oscilloscopes. For OIM32, rotary potentiometers were soldered on perfboards as couplings. Next to each potentiometer, we designed male pin connectors such that the polarity of each connection can be controlled by shorting different pins using female jumper caps (color coded green and pink for positive and negative couplings). Furthermore, we soldered TI SN74HC86N Exclusive-OR (XOR) gate ICs to convert the oscillator phases to voltage levels, which then power on-board LEDs for visualization and readout (by two 16-channel logic analyzers). We observed that for small-sized (8 and 32) Ising problems, global optima can be achieved easily using these prototypes.

OIM64 and OIM240 use digital potentiometers from AD5206 ICs (6-channel potentiometers with 8-bit accuracy). Because these ICs are designed primarily for audio processing and do not have multi-MHz bandwidth, we reduced oscillator frequency to 1MHz. Both prototypes consist of multiple PCBs. OIM64 connects 64 oscillators in a 8x8 2D toroidal grid, with 192 couplings, each made of one channel of AD5206 and a SPDT switch for setting its polarity. Even though it was not easy to “program” OIM64 due to the use of physical switches, we tried it on 10 randomly generated toroidal Ising grid instances, achieving the global optimum for each one.

In OIM240, we improved the design to use the position of the potentiometer wiper to switch polarity, thus eliminating the use of switches and making the coupling truly programmable. On each PCB, we implemented 12 oscillators with a denser connectivity; 20 such PCBs were plugged into a motherboard through edge connectors, and interconnected in a 4x5 toroidal grid, implementing a total of 240 oscillators with 1200 couplings. The motherboard also distributes CLK, data lines, address lines for programming the 200 AD5206 ICs and for reading oscillator states, all controlled by an Arduino module on the motherboard that communicates with a PC through USB. When operating OIM240, we flip on the supply digitally, wait 1ms for oscillators to synchronize, then read back the solution. Even with all the overhead from serial reading, solutions can be read back every 3.5ms. OIM240’s operation consumes ~5W of power for all the oscillators and peripheral circuitry, excluding only the LEDs.

We tested OIM240 with many randomly generated Ising problems (with each of the 1200 couplings randomly chosen from 0, -1, +1). A typical histogram for the energy levels of the measured solutions is shown in Fig. 2 (c). Note that a random (trivial) solution has an energy around 0, whereas the best polynomial-time algorithm (based on SDP) guarantees to achieve 87.8% of the global optimum. In comparison, results from OIM240 center around a very low energy, and achieve the global optimum multiple

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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. DAC '19, June 2–6, 2019, Las Vegas, NV, USA
© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6725-7/19/06...\$15.00
https://doi.org/10.1145/3316781.3322473

¹A negative coefficient can be implemented conveniently using positive conductances to cross-couple two differential oscillators.

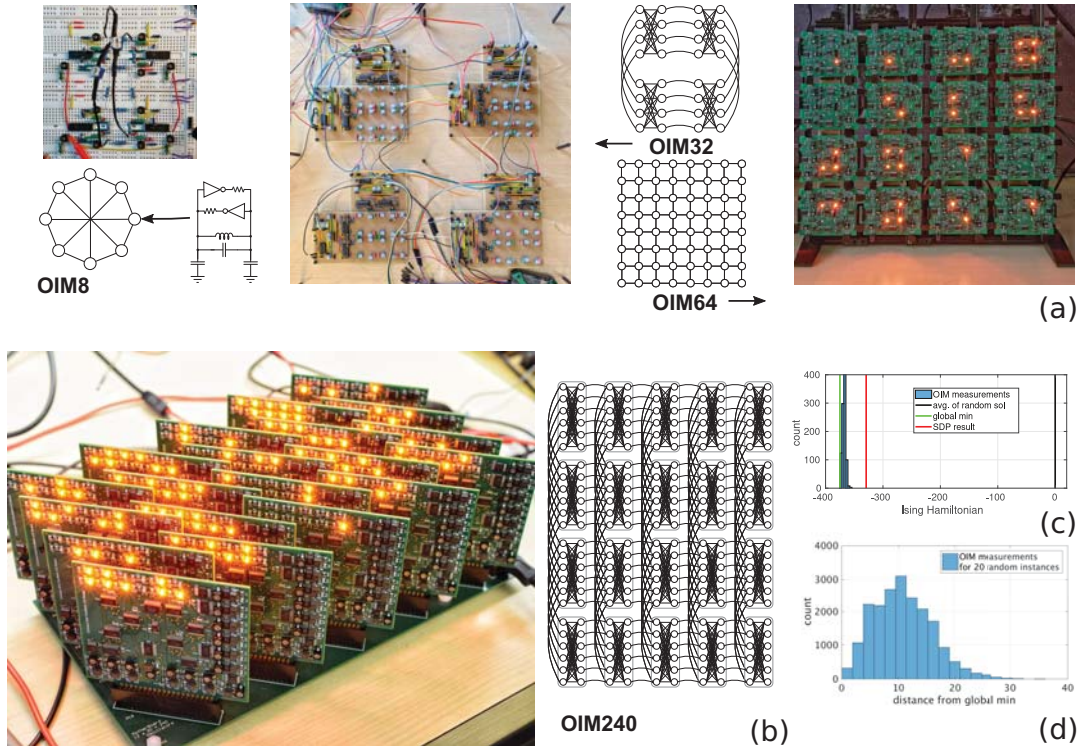


Figure 2: OIM prototypes: (a) photos and schematics of OIM8, OIM32, OIM64; (b) OIM240; (c) energy levels of 1000 measured solutions from OIM240, on a random instance of size-240 Ising problem; (d) energy levels for 20 such random instances (showing the distances to their respective global minimum).

times. We performed the same measurements for 20 different random Ising problems, with the distances of solutions from their respective global optima² shown in Fig. 2 (d). The fact that OIM240 is finding highly non-trivial solutions indicates that it indeed physically implements a working Ising machine.

4 Results on MAX-CUT Benchmark Problems

We simulated OIMs for solving the Ising version of the MAX-CUT problems from the G-set, with sizes ranging from 800 to 3000, with partial results shown in Tab. 1. Full results for all 54 problems will be open-source released with the code (a simple C++ script for integrating stochastic differential equations). Each problem was simulated with 100 random instances. In Tab. 1, we compare the results with the best-known results listed in [10] (mostly from the heuristic algorithm Scatter Search (SS) [11]). We also list results from a recent study applying simulated annealing to MAX-CUT [12], the only one we could find that contains results for all the G-set problems. Although we could not rerun the best-known results reported in [10], we made the comparison more fair by running our simulation with an older desktop that are comparable with the environment used by others.

Benchmark	SS	Time (s)	SA	Time (s)	OIM	Time (s)
G1	11624	139	11621	295	11624	52.6
G11	562	172	564	189	564	6.7
C21	930	233	927	195	931	14.6
C31	3288	1336	3309	214	3301	59.1
G41	2386	1017	2405	208	2401	37.8
G51	3846	513	3841	234	3846	18.4

Table 1: Results (cut sizes) from OIM and several heuristics run on MAX-CUT benchmarks in the G-set. Time reported in this table is for a single run.

Note that we were simulating a fixed duration equivalent to 1000 cycles of oscillation, which will be much faster on a physical hardware than the simulation time we show here. But even simulating OIM's differential equations yields a good solution quality. Unlike other algorithms, OIM simulation does not know about the energy function or relative energy changes, which are implicit in the dynamics of differential equations, yet it proves effective and fast.




Figure 3: Hamiltonian values achieved by several variants of OIM.

²We ran simulated annealing for a long time (1min) and for multiple times, then treated the best results as global optima.

Furthermore, we removed SYNC from OIM, reran the 54 benchmarks and compared the results in Fig. 3. Without SYNC, the system becomes a simple coupled oscillator system with phases taking a continuum of values (as opposed to binary values), which we then threshold as Ising solutions. From Fig. 3, the solutions become much worse, indicating SYNC and SHIL are essential to OIM's operation. We have also studied OIM's performance with variability in the natural frequency of oscillators (a spread following a Gaussian distribution of 1% or 5% standard derivation). From Fig. 3, we observe that even with non-trivial variability, the solution quality is not affected by much.

5 Conclusion and Outlook

While much future work is still needed to establish OIM's advantages or competence against other Ising machine proposals, the new preliminary results we show in this paper demonstrate that OIM is indeed a feasible and attractive Ising machine technology. Even with PCB implementations using decade-old 7404 inverters, 1MHz oscillators and 5V supply, OIM can return solutions in milliseconds with only 5W of power consumption. Future CMOS IC implementations at a larger scale with GHz nano-oscillators under lower supply voltages hold exciting promises in outperforming conventional hardware for Ising problems; our results shown here will serve as a solid foundation for these future developments.

References

- [1] Francisco Barahona. On the computational complexity of Ising spin glass models. *Journal of Physics A: Mathematical and General*, 15(10):3241, 1982.
- [2] M. Yamaoka, C. Yoshimura, M. Hayashi, T. Okuyama, H. Aoki and H. Mizuno. A 20k-spin Ising Chip to Solve Combinatorial Optimization Problems with CMOS Annealing. *IEEE Journal of Solid-State Circuits*, 51(1):303–309, 2016.
- [3] A. Lucas. Ising formulations of many NP problems. *arXiv preprint arXiv:1302.5843*, 2013.
- [4] T. Inagaki, Y. Haribara, K. Igarashi, T. Sonobe, S. Tamate, T. Honjo, A. Marandi, P. L. McMahon, T. Umeki, K. Enbutsu and others. A Coherent Ising machine for 2000-node Optimization Problems. *Science*, 354(6312):603–606, 2016.
- [5] K. Y. Camsari, R. Faria, B. M. Sutton and S. Datta. Stochastic p-Bits for Invertible Logic. *Physical Review X*, 7(3):031014, 2017.
- [6] Z. Bian, F. Chudak, R. Israel, B. Lackey, W. G. Macready and A. Roy. Discrete optimization using quantum annealing on sparse Ising models. *Frontiers in Physics*, 2:56, 2014.
- [7] T. Wang and J. Roychowdhury. Oscillator-based Ising Machine. *arXiv preprint arXiv:1709.08102*, 2017.
- [8] T. Wang and J. Roychowdhury. OIM: Oscillator-based Ising Machines for Solving Combinatorial Optimisation Problems. *arXiv preprint arXiv:1903.07163*, 2019.
- [9] C. Helmberg, F. Rendl. A spectral bundle method for semidefinite programming. *SIAM Journal on Optimization*, 10(3):673–696, 2000.
- [10] The G-set benchmarks for MAX-CUT. Website: <http://www.optiscom.es/maxcut>.
- [11] R. Marti, A. Duarte, M. Laguna. Advanced scatter search for the max-cut problem. *INFORMS Journal on Computing*, 21(1):26–38, 2009.
- [12] T. Myklebust. Solving maximum cut problems by simulated annealing. *arXiv preprint arXiv:1505.03068*, 2015.