

OPEN

# Analog Coupled Oscillator Based Weighted Ising Machine

Jeffrey Chou, Suraj Bramhavar, Siddhartha Ghosh  & William Herzog

We report on an analog computing system with coupled non-linear oscillators which is capable of solving complex combinatorial optimization problems using the weighted Ising model. The circuit is composed of a fully-connected 4-node LC oscillator network with low-cost electronic components and compatible with traditional integrated circuit technologies. We present the theoretical modeling, experimental characterization, and statistical analysis of our system, demonstrating single-run ground state accuracies of 98% on randomized MAX-CUT problem sets with binary weights and 84% with 5-bit weight resolutions. Solutions are obtained within 5 oscillator cycles, and the time-to-solution has been demonstrated to scale directly with oscillator frequency. We present scaling analysis which suggests that large coupled oscillator networks may be used to solve computationally intensive problems faster and more efficiently than conventional algorithms. The proof-of-concept system presented here provides the foundation for realizing such larger scale systems *using existing hardware technologies* and could pave the way towards an entirely novel computing paradigm.

Solving certain classes of combinatorial optimization (CO) problems has proven to be notoriously difficult using standard von Neumann computing architectures. A canonical example is the traveling salesman problem, for which exact algorithms scale very poorly with problem size. Applications of CO problems span many disciplines, including business operations, scheduling, traffic routing, finance, big data, drug discovery, machine learning, and many other systems requiring the minimization of a complex energy landscape with multivariate inputs<sup>1</sup>. As the decades-long progress of digital CMOS technologies begins to plateau, there is a growing desire to find alternative computing methodologies which can address the challenges of these traditionally difficult problems. One strategy to solve CO problems lies in mapping them to the Ising Hamiltonian for spin glass systems and finding the ground state solution<sup>2,3</sup>. The Ising model dates back many decades but was re-popularized by D-Wave Systems in an attempt to exploit quantum mechanical phenomena to speed up computation. Novel quantum annealing machines<sup>4–7</sup> and digital CMOS annealing accelerators<sup>8</sup> are garnering significant attention in the hopes of solving such problems faster than conventional algorithms performing heuristic optimizations. The Ising Hamiltonian can also be mapped to the comparable field of Hopfield neural networks, which have been previously explored to solve CO problems<sup>1,9–11</sup>.


Recently, alternative classical methods to solve the Ising model have emerged using optoelectronic parametric oscillators<sup>12–14</sup>, memristor cross-bar arrays<sup>11,15,16</sup>, electronic oscillators<sup>17,18</sup>, and GPU based algorithms<sup>19,20</sup>. An analysis of one such coupled oscillator system revealed the potential for a significant speedup over digital computing algorithms at large node sizes<sup>21</sup>. Even algorithmic approaches which emulate nonlinear dynamical systems but are run on conventional computing hardware have been shown to match and even surpass the performance of state-of-the-art algorithms, motivating the desire to build these systems in physical hardware<sup>22,23</sup>. Scaling up the optoelectronic oscillator Ising machine remains challenging owing to its time-multiplexed architecture and highly complex and costly optoelectronic setup<sup>12,13</sup>. All of the proposed systems bear close resemblance to those described in the field of stochastic computing, where so-called ‘p-bits’, or probabilistic representations of digital bits have been shown in simulation to be capable of solving invertible logic and combinatorial optimization problems<sup>3,24,25</sup>. Analog implementations of these circuits typically rely on magnetic tunnel junctions, which could conceivably be implemented using standard CMOS technology, but a scalable physical implementation of this technology has yet to be realized. In a similar vein, the all-electronic oscillator concept initially proposed by Wang and Roychowdhury introduces the tantalizing prospect of creating a similar system using readily available electronic components interconnected in a parallel fashion and is particularly well suited for chip-scale integration and scaling using present day technologies<sup>17</sup>. In this case, the ‘p-bit’ is represented by the phase of an oscillator,

Massachusetts Institute of Technology Lincoln Laboratory, Lexington, Massachusetts, USA. Jeffrey Chou and Suraj Bramhavar contributed equally. Correspondence and requests for materials should be addressed to W.H. (email: [whertzog@ll.mit.edu](mailto:whertzog@ll.mit.edu))

Received: 25 June 2019

Accepted: 27 August 2019

Published online: 15 October 2019



**Figure 1.** (a) Schematic of a fully connected 4-node system. (b) Degenerate solutions for 4-node problem. (c) Simulation of the oscillator phases with initial conditions set to incorrect state. (d) Probability distribution of solution states after 1000 trials with randomized initial conditions.

which can probabilistically settle to one of two values. Previous reports have experimentally demonstrated a Chimera-graph architecture with a total of 240 nodes and described its performance on 20 random problems<sup>26</sup>. However, a rigorous and statistically significant analysis of the accuracy of this experimental system has not been demonstrated. In the reported physical implementation, resistive coupling is used to apply the connecting weights between their oscillators, where the absolute resistance values determines the coupling weights. However, specific details of the oscillator and interconnect architectures employed in these demonstrations are currently publicly unavailable, thus precluding our ability to determine the operation and limitations of this approach.

In this paper, we build upon this initial work and demonstrate a 4 node, fully-connected, differential LC (inductor-capacitor) oscillator based analog circuit with standard electronic components which accurately maps to the Ising model. One key difference between the circuit presented here and that demonstrated previously<sup>17,26</sup> involves the interconnection architecture and oscillator coupling scheme. Whereas previous implementations used a simple resistor to apply the connection weights directly, the circuit demonstrated in this work implements gain ratios in a cross-bar architecture. Specific advantages gleaned from using the cross-bar architecture include the ability to program weights with larger bit-depth and easy scalability in CMOS for large and fully-connected systems<sup>27</sup>. We provide a thorough analysis of system performance, describe in detail the circuit architecture including a new modality used to program variable interconnection strengths, and demonstrate the system capabilities in solving a variety of MAX-CUT problems with both binary and multi-bit weight values. To the best of our knowledge, this is the first demonstration of an all-electronic oscillator-based Ising machine with multi-bit weights. The thorough statistical analysis presented herein provides valuable insight into the viability of these systems as computing platforms when scaled to larger node counts.


## Theory

Much of the theoretical framework for using coupled oscillators as Ising solvers has been described previously, but will be summarized here for convenience<sup>28</sup>. The Ising Hamiltonian is given by the following equation:

$$H = -\sum_{i,j} J_{ij} s_i s_j - \sum_i h_i s_i \quad (1)$$

where  $V$  represents the number of nodes in a particular problem set,  $J_{ij}$  represents the weight values interconnecting the nodes, and  $s = [s_1 \dots s_V]$  represents the solution space where  $s_i$  can take the value of either +1 (spin  $\uparrow$ ) or -1 (spin  $\downarrow$ ). One common benchmark optimization problem, known as the MAX-CUT problem, is defined by taking an undirected graph and finding a bisection of that graph which maximizes the cut set. This problem can be mapped directly to the Ising Hamiltonian above, and the solution to the problem is represented by the state  $s$  which minimizes  $H$ . A particular problem of interest is typically defined by a graph  $G(V, E)$ , where  $V$  represents the number of vertices and  $E$  represents the number of edges. An example of a 4-node system is shown in Fig. 1(a). If  $J = 1$  for all connections, and we neglect the Zeeman term ( $h_i = 0$ ), the minimum solution is defined by 6 degenerate  $2 \times 2$  solution sets, as shown in Fig. 1(b). It has been shown previously that these graphs can be represented by a network of coupled nonlinear oscillators whose phase dynamics are described by the Kuramoto model<sup>29</sup>, and that this model maps directly to the Ising Hamiltonian if the phases of these oscillators take values of either  $0^\circ$  or  $180^\circ$ <sup>17</sup>. One mechanism used to polarize the phases is to introduce an injection-locking signal at twice the natural frequency of the oscillators<sup>17,30</sup>. This 'super-harmonic' injection locking signal, and its application in the context of LC-oscillator systems is mathematically similar to the case of the degenerate optical parametric oscillator used in previous optical Ising machines<sup>12,13</sup>, where optical pump pulses at twice the optical oscillation frequency are used to create binary phase values.

In the LC-oscillator system, the phase evolution of each oscillator in the network ( $\theta_i(t)$ ) can be represented by the following differential equation:



**Figure 2.** (a) Circuit diagram of the LC oscillator circuit. The input coupling is inserted at the gates of transistors M3 and M4. (b) Multiply and accumulate cross-bar array composed of digital potentiometers. (c) Photograph of full breadboard system.

$$\frac{d\theta_i}{dt} = -\sum_{j=1}^V J_{ij} \sin(\theta_i(t) - \theta_j(t)) - A(t) \sin(2\theta_i(t)) \quad (2)$$

where  $V$  is the total number of oscillators (or vertices) and  $A$  represents the amplitude of the injection-locking signal applied to each oscillator. The time  $t$  in the equation is defined in dimensionless units in relation to the oscillation period. The addition of Gaussian phase fluctuations, representing noise in the system, converts equation Eq. (2) into a network of stochastic differential equations (SDEs), whose solutions can be approximated iteratively using the Euler-Maruyama method. Annealing is accomplished by gradually increasing the injection locking term  $A$  according to the formula:

$$A(t) = A_0 \left(1 - e^{-\frac{t}{\tau}}\right) \quad (3)$$


with  $\tau$  set to 5 oscillation cycles. One solution simulated for the example problem described above is shown in Fig. 1(c), where the initial condition was intentionally selected to be an incorrect solution state ( $s = [\uparrow\uparrow\uparrow\uparrow]$ ). Even in this scenario, the oscillator phases evolve to settle at one of the 6 correct solutions ( $s = [\downarrow\downarrow\uparrow\uparrow]$ ). The system settles to the ground state within 3 oscillation cycles. This simulation was then run 1000 times with randomized initial conditions, and the system always settles to a correct solution state with fairly uniform probability (shown in Fig. 1(d)). The average settling time for the 1000 trials was 2.3 oscillation cycles.

## Results

**Circuit implementation.** The oscillator circuit employs a differential injection-locked frequency divider topology, as shown in Fig. 2(a)<sup>31</sup>. This architecture is selected to offer the option of synchronizing the oscillator with an incident super-harmonic injection locking signal. Transistors (Supertex TN0702) M1 and M2 form a cross-coupled pair, which serves as the negative resistance component, necessary for unity loop gain. The coupling signal from the other oscillators is applied differentially through transistors M3 and M4. This specific coupling circuitry is typically used for quadrature LC oscillators and employs an injection locking based coupling scheme, which has been previously mapped to the generalized Adler's equation<sup>32</sup>, and can further be mapped to the Kuramoto model<sup>28</sup>. Current source I2 provides the bias current for the coupling signal. The output voltage of the oscillator is tapped at nodes VoL and VoR, directly out from the oscillating LC tank ( $L = 100 \mu\text{H}$ ,  $C = 0.1 \mu\text{F}$ ). The LC tank circuit has a resonant frequency of 50 kHz, and is composed of L1, C1, and L2, C2. Current source I1 provides a biasing current for the oscillator circuit and can also be used for injection locking to help polarize the phases to  $0^\circ$  and  $180^\circ$ .

A differential analog multiply and accumulate circuit in a cross-bar array is used apply the necessary signals into input nodes ViL and ViR, as shown in Fig. 2(b). The coupling coefficient polarity is controlled by the polarity of the output of the differential summing amplifier (Texas Instruments THS4140) to the nodes ViR and ViL. Digital potentiometers (Analog Devices AD5272), R12–R34, are employed to control the individual gains of each of the input oscillator signals. A key advantage of this interconnect circuit versus a more straight-forward resistive approach<sup>17</sup> stems from the ability to scale the number of fully-connected nodes without adding resistive loading to each oscillator by using a simple buffer circuit. The actual gain term is determined by the ratio of the feedback resistor,  $R_{FB} = 1 \text{ k}\Omega$ , and the digital potentiometers. An Arduino microcontroller board was used to apply the digital I2C communication signals to the digital potentiometers. Precise tuning of the bias voltages for the oscillator and coupling circuit is critical to ensure accurate solution performance of the system.

The analog coupling coefficients from the Ising Hamiltonian ( $J_{ij}$ ) are mapped linearly to the ratio of the gains between the various oscillators. The digital potentiometers employed here have 1024 tap points, with a maximum



**Figure 3.** (a) System level diagram of the oscillator control system. Blue lines indicate digital signals while red lines indicate analog signals. (b) Schematic timing diagram of the oscillator phases during system operation.

resistance of 20 kΩ. The conversion from the analytical coupling coefficients to the digital potentiometers' programmed values ( $D_{ij}$ ), which range from 1–1024, is shown the following equation:

$$D_{ij} = \frac{\beta\alpha}{J_{ij}} + \beta(1 - \alpha) \quad (4)$$


where  $\beta = R_{min} * 1024 / 20 \text{ k}\Omega$ , and  $\alpha$  is the mapping scalar. To prevent high currents at the summing amplifier at high coupling coefficient values, we set  $R_{min} = 760 \Omega$ . The mapping from the coupling coefficients,  $J_{ij}$ , to the resistance values is scaled based on the  $\alpha$  term to enable maximum dynamic range. A value of  $\alpha = 2.5$  is used in this paper.

**Experimental setup.** A schematic of the experimental setup for the 4 node system is shown in Fig. 3(a). A National Instruments Labview program transmits the user-defined digital potentiometer values to an Arduino Uno board. The Arduino board then transmits the digital I2C resistors values to the six unique weights (J12–J34). The analog oscillator voltage signals are monitored by the Labview program with a data acquisition board, at a maximum sampling rate of 500 kHz.


A representative phase-time trace of a single run while the weights are being programmed is shown in Fig. 3(b). The digital potentiometers are programmed sequentially in time, due to the serial nature of the digital scheme. The time period of the weight programming is dependent on the loop speed of the Arduino board and the complexity of the code – which has a maximum loop frequency of 117 kHz<sup>33</sup>. After the last weight (J34) is programmed, the oscillators arrive at the new solution state. Phases of the oscillators are each measured relative to oscillator 1, whose phase is given a default binary value of “0”. The binary phase outputs are then determined with a simple threshold, where phases differences  $>90^\circ$  with respect to oscillator 1 are assigned “1”, and phases  $<90^\circ$  are assigned “0”. We note that, unlike the simulation, no annealing of the coupling strength, is required to arrive at the solution. This is likely due to the low node count of the system. With more nodes, simulations of this oscillator network using coupled differential equations have shown that ramping up either the coupling strength or an external super-harmonic injection locking signal will be required in order for the system to preferentially settle to the correct solution. The super-harmonic injection locking signal serves a dual purpose in that it replaces the rounding feature described above, forcing the phases of each oscillator to their respective binary values (either  $0^\circ$  for “0” or  $180^\circ$  for “1”). As pointed out previously, this signal is required to extract sensible answers using large node count systems<sup>17</sup>.

A high sampling rate phase measurement of the phase transition of the oscillators is shown in Fig. 4. The phases were obtained via performing a moving window Fourier transform on the time-domain amplitude signal. The red dashed weight change line is when the last potentiometer (J34) has been digitally switched. From this measurement, we observe a 100 μs solution time, which translates to 5 cycles of the 50 kHz oscillators, as shown in Fig. 4(a). A potentially attractive pathway to scale the solution time of these systems involves simply increasing the oscillator frequency, as the settling time is expected to scale with oscillator period. In an attempt to validate this hypothesis, Fig. 4(b) shows the same experiment but with oscillator frequencies reduced by one order of magnitude, to 5 kHz, by utilizing 10 μF capacitors. We experimentally observe the time-to-solution also increases by approximately one order of magnitude, which suggests that the time-to-solution does indeed scale with the oscillator frequencies. The  $90^\circ$  threshold algorithm translates the final oscillator phase states to a binary solution of [0 1 1 0], if we use the order of the oscillator number. This solution is one of the 6 correct degenerate solutions for a fully-connected MAX-CUT problem with 4 nodes and equal connection weights.

Time trace data of the oscillators under various binary weight configurations is shown in Fig. 5 to demonstrate the functioning system. The voltage amplitudes of each oscillator are shifted by 10 V to help distinguish the individual oscillators. All oscillation voltage amplitudes in this experiment range from 0 to 10 V. With the appropriate binary weights shown in Fig. 5(a–d), each oscillator can be made to be  $180^\circ$  out of phase with all other oscillators.




**Figure 4.** Measured phase versus time data during the last weight change, J34, with (a) 50 kHz and (b) 5 kHz oscillators. A solution time of approximately 5 cycles is measured. The time-to-solution directly scales with oscillator frequency.



**Figure 5.** (a–e) Measured oscillator amplitude data at various binary weight configurations. The lines connecting the oscillators correspond to a weight of 1, while an empty space corresponds to a weight of 0.

With this graph structure of only one node connected to the remaining three, it intuitively follows that the energy minimization is achieved when the central node is out of phase with the others. Under the fully connected system in Fig. 5(e), there exist 6 degenerate solutions, in which any 2 oscillators are  $180^\circ$  out of phase with any other 2 oscillators. Here we only show one of the six solutions.




**Figure 6.** Experimentally measured probability of solution energy for (a) 1 bit and (b) 5 bit weight resolutions after more than 2000 trials with random weights. Solution energy of E0 represents the ground state solution. Error bars represent the standard deviation. (c) The measured ground state probability versus weight bit resolution. (d) Estimated solution time in number of oscillator periods based on a 99% confidence. The corresponding shaded regions represent the standard deviation.


**Statistical analysis.** In order to validate the accuracy and reliability of the system, more than 2000 automated computer controlled experiments with random weights were performed. Figure 6(a,b) shows the probability distribution of the arrival of each solution energy with a 1 bit and 5 bit weight values, respectively. The probability of achieving the ground state for 1 bit and 5 bit weights are 98% and 84%, respectively. The ground state probability was determined by 10 trials of each of the 2000 randomly generated problems. Clearly, the probability of achieving the ground state decreases with increasing weight bit resolutions. We observe that the solutions for all of the experiments follow a Boltzmann distribution with the peak of the curve at the ground state.

To fully characterize the ground state solution probability versus weight bit resolution, the results of running over 2000 random problems for weight bit resolutions from 1 to 5 is shown in Fig. 6(c). The ground state accuracy decreases monotonically as the bit resolution is increased. The total number of oscillator periods,  $N_T$ , required to reach ground state with 99% certainty is defined as  $N_T = N_c \log(1 - 0.99) / \log(1 - P(b))$ . Here,  $P(b)$  is the measured ground state probability from a single run for bit resolution  $b$ , and  $N_c$  is the number of oscillator periods for a single run. We obtain the  $P(b)$  from Fig. 6(a) and plot the results in Fig. 6(d). Thus to achieve the ground state with 5 bits of weight resolution, an average of 13 periods would be required.

To characterize the cause of the non-ground state solutions for the non-binary weights, Fig. 7 plots the measured ground state probability versus the difference between the ground state Ising energy and the next highest energy state ( $\Delta E = E_1 - E_0$ ) for randomly generated problems with 5 bit weights. Intuitively, this parameter characterizes the depth of the ground state potential well, with easier problems exhibiting large  $\Delta E$ , and more difficult problems exhibiting small  $\Delta E$ . The data shows that as  $\Delta E$  decreases, the probability of finding the ground state on a single trial also decreases. For example, when  $\Delta E < 2$ , the probability of achieving the ground state falls below 0.9. Noise in the system causes difficulty in distinguishing energy-near solutions from the ground state solution. The minimum  $\Delta E$  value inversely scales with weight bit resolution, which is the cause of the decreased accuracy for higher bit problems. Alternatively, binary weight systems typically have very high ground state probabilities due to the large  $\Delta E$  values. For example, in the 4 node binary weight system, the minimum energy difference between E1 and E0, is  $\Delta E = 2$ , which gives rise to very high accuracy solutions. The complete absence of non-unitary probabilities for  $\Delta E > 2$  implies that the system does not become trapped in local minimum, so long



**Figure 7.** Experimentally measured ground state probability versus solution energy difference between the ground state Ising energy and the next highest energy state ( $\Delta E = E_1 - E_0$ ) for randomly generated 5 bit weights.




**Figure 8.** (a) Simulated ground state probability on a circular graph versus number of oscillators. (b) Scaling simulation of the coupled oscillator system with 50 kHz and 1 MHz oscillator frequencies. For comparison, data extracted from previous Ising machine demonstrations for an identical problem<sup>13</sup> (orange dots) is plotted alongside the simulated data, illustrating a potential advantage of this coupled oscillator system.

as the  $\Delta E$  is large enough. To compare the experimental results to simulations, we overlay a similar plot generated by simulating a set of randomly generated 5-bit weight problems with the SDE simulation technique described above. The simulated sample set includes 1000 different problems, and each point represents the mean ground state probability value of all the problems with a given  $\Delta E$  value. The behavior of the SDE simulated system matches reasonably well to the behavior measured experimentally.

## Discussion

To understand the expected performance as the system size is scaled, the SDE solver discussed previously was extended beyond the case of  $V=4$  and used to calculate the probability that the system can find the ground state as a function of  $V$ . We note that, for this scaling study, the annealing schedule described in Eq. (3) was used with  $\tau=5$  cycles. This value was intentionally chosen to be small to allow for a pessimistic assumption of the solution time scaling trends. The problem set for this test was a set of circular graphs, known as Mobius ladder graphs, with each node connected to its two nearest neighbors and one node directly opposite (shown in Fig. 8(a)). The ground-state solutions for this problem set were determined through brute force calculation<sup>34</sup>, which allows for the solution from the SDE solver to be compared against a global energy minimum. Figure 8(a) plots the probability  $P$  that the system reaches the ground state solution after 500 independent runs for each problem set, showing a natural roll-off from  $\sim 100\%$  ground state probability when  $V=8$  to roughly 3% probability when  $V=300$ . Translating these results to calculate the overall time required to find a solution for a given problem ( $t_s$ ), we first find  $N_t$  by defining  $N_t$  as the number of oscillation periods required for a single run of the simulated system to reach a minimum value and multiplying by the number of trials required to reach the ground state based on the



**Figure 9.** Solution time scaling simulation for randomized cubic graphs using coupled oscillator approach (blue) vs. mean-field algorithm run on conventional GPU (orange). Dashed lines represent log-linear fits to the respective data. In the case of the coupled oscillator simulation, the fit was determined using the last 8 data points.

expected probability  $P$  (as described previously). Finally, multiplying  $N_i$  by the period of one oscillation yields the final solution time  $t_s$ . Figure 8(b) shows these computed values assuming a system built with 50 kHz oscillators (to match the current experimental circuit), along with a linear fit to a logarithmic trend line. The results are compared against the identical problem set and a similar analysis from a previous Ising machine implementation<sup>13</sup> (orange points). For illustrative purposes, a similar trend line is plotted for a system assuming an oscillator frequency of 1 MHz. Direct scaling with oscillator frequency was demonstrated in Fig. 4, and the additional line is shown here to highlight the potential speedup that can be achieved against existing implementations with relatively modest hardware specifications. It is also important to note that this scaling analysis focuses entirely on ground state solutions. As shown in Fig. 6(b), when the ground state solution is not reached, the system often settles to a value very close to the absolute minimum, which can be sufficient for many optimization problems.

Moving beyond simple Mobius ladder graph problems, we extend the analysis to randomized cubic graphs, where each node is connected to three others with the connections determined at random. For each node size, 100 different problem instances were generated and the exact ground state was found using the BiqMac MAX-CUT solver<sup>35</sup>. Each problem set was then run 100 times using the coupled oscillator simulator, and the ground state probability ( $P$ ) and solution time ( $t_s$ ) were determined as described previously. Figure 9 shows the mean  $t_s$  calculated for all 100 problem sets at each node size, assuming a 1 MHz oscillator. The data is plotted against an identical experiment on cubic graphs solved using a GPU-accelerated mean-field algorithm, which has shown ~20x speedup when compared against existing Ising machine implementations<sup>19</sup>. It is clear from the simulation that a crossover point is reached at  $V \sim 150$ , with the coupled oscillator approach trending much more favorably as the system size is scaled. As a point of reference, the BiqMac solver running on a modern CPU took approximately 10–20 seconds to solve random cubic problem instances with  $V = 200$ . This analysis lends credibility to the promise of using coupled oscillator systems to solve non-trivial computational problems.

## Conclusion

In summary we demonstrate the simulation, design, experiment, and characterization of a parallel, all-to-all connected coupled oscillator system which correctly maps to the weighted Ising model. The system is capable of successfully solving random MAX-CUT problems with 98% success probability with binary weights. The proof-of-concept breadboard experimental demonstration enabled a study of accuracy as a function of the interconnect weight bit-resolution. We show that the performance maps accurately to SDE simulations, and use these simulations to predict the system behavior at larger node sizes, which show favorable scaling when compared against existing optoelectronic implementations and GPU-based algorithms for similar problem sets. A full comparison on more densely connected problem sets will be the subject of future work.

A significant advantage of the coupled oscillator system presented here lies in the fact that it can be scaled using existing low-cost hardware and lends itself especially well to an integrated circuit implementation. One major challenge, and an important area of future study, involves the interconnect architecture required to densely connect very large numbers of oscillators ( $> 1e3$ ). While time-multiplexed approaches have been implemented in recent instances<sup>12,13</sup>, frequency- or code-division multiplexed systems could also provide interesting pathways to efficiently allocate connection resources. Another interesting area for study involves the use of measurement and feedback approaches to increase the ground state solution probability for larger problem sizes<sup>36,37</sup>. As transistor scaling begins to limit the progress in conventional digital computing architectures, the demand for alternative strategies continues to grow, and oscillator-based Ising machines could represent a novel and important platform to push the bounds of certain classes of computational problems.

## References

1. Smith, K. A. Neural Networks for Combinatorial Optimization: A Review of More Than a Decade of Research. *Inf. J. Comput.* **11**, 15–34 (1999).
2. Lucas, A. Ising formulations of many NP problems. *Front. Phys.* **2** (2014).
3. Yoshimura, C. *et al.* Uncertain behaviours of integrated circuits improve computational performance. *Sci. Rep.* **5**, 16213 (2015).
4. Johnson, M. W. *et al.* Quantum annealing with manufactured spins. *Nature* **473**, 194–198 (2011).
5. Nigg, S. E., Lörch, N. & Tiwari, R. P. Robust quantum optimizer with full connectivity. *Sci. Adv.* **3**, e1602273 (2017).
6. Puri, S., Andersen, C. K., Grimsmo, A. L. & Blais, A. Quantum annealing with all-to-all connected nonlinear oscillators. *Nat. Commun.* **8**, 15785 (2017).
7. Vinci, W. *et al.* Hearing the Shape of the Ising Model with a Programmable Superconducting-Flux Annealer. *Sci. Rep.* **4**, 5703 (2014).
8. Yamaoka, M. *et al.* A 20 k-Spin Ising Chip to Solve Combinatorial Optimization Problems With CMOS Annealing. *IEEE J. Solid-State Circuits* **51**, 303–309 (2016).
9. Fang, Y., Yashin, V. V., Levitan, S. P. & Balazs, A. C. Pattern recognition with “materials that compute”. *Sci. Adv.* **2**, e1601114 (2016).
10. Maffezzoni, P., Bahr, B., Zhang, Z. & Daniel, L. Oscillator Array Models for Associative Memory and Pattern Recognition. *IEEE Trans. Circuits Syst. Regul. Pap.* **62**, 1591–1598 (2015).
11. Guo, X. *et al.* Modeling and Experimental Demonstration of a Hopfield Network Analog-to-Digital Converter with Hybrid CMOS/Memristor Circuits. *Front. Neurosci.* **9** (2015).
12. Inagaki, T. *et al.* A coherent Ising machine for 2000-node optimization problems. *Science* **354**, 603–606 (2016).
13. McMahon, P. L. *et al.* A fully programmable 100-spin coherent Ising machine with all-to-all connections. *Science* **354**, 614–617 (2016).
14. Hamerly, R. *et al.* Experimental investigation of performance differences between coherent Ising machines and a quantum annealer. *Sci. Adv.* **5**, eaau0823 (2019).
15. Shin, J. H., Jeong, Y. J., Zidan, M. A., Wang, Q. & Lu, W. D. Hardware Acceleration of Simulated Annealing of Spin Glass by RRAM Crossbar Array. In *2018 IEEE International Electron Devices Meeting (IEDM)* 3.3.1–3.3.4, <https://doi.org/10.1109/IEDM.2018.8614698> (2018).
16. Cai, F. *et al.* Harnessing Intrinsic Noise in Memristor Hopfield Neural Networks for Combinatorial Optimization. *ArXiv190311194 Cs* (2019).
17. Wang, T. & Roychowdhury, J. OIM: Oscillator-based Ising Machines for Solving Combinatorial Optimisation Problems. *ArXiv190307163 Cs* (2019).
18. Parihar, A., Shukla, N., Jerry, M., Datta, S. & Roychowdhury, A. Computing with dynamical systems based on insulator-metal-transition oscillators. *Nanophotonics* **6**, 601–611 (2017).
19. King, A. D., Bernoudy, W., King, J., Berkley, A. J. & Lanting, T. Emulating the coherent Ising machine with a mean-field algorithm. *ArXiv180608422 Quant-Ph* (2018).
20. Tiunov, E. S., Ulanov, A. E. & Lvovsky, A. I. Annealing by simulating the coherent Ising machine. *Opt. Express* **27**, 10288 (2019).
21. Haribara, Y., Utsunomiya, S. & Yamamoto, Y. A Coherent Ising Machine for MAX-CUT Problems: Performance Evaluation against Semidefinite Programming and Simulated Annealing. In *Principles and Methods of Quantum Information Technologies* (eds Yamamoto, Y. & Semba, K.) 251–262, [https://doi.org/10.1007/978-4-431-55756-2\\_12](https://doi.org/10.1007/978-4-431-55756-2_12) (Springer Japan, 2016).
22. Di Ventra, M. & Traversa, F. L. Perspective: Memcomputing: Leveraging memory and physics to compute efficiently. *J. Appl. Phys.* **123**, 180901 (2018).
23. Di Ventra, M. & Pershin, Y. V. The parallel approach. *Nat. Phys.* **9**, 200–202 (2013).
24. Sutton, B., Camsari, K. Y., Behin-Aein, B. & Datta, S. Intrinsic optimization using stochastic nanomagnets. *Sci. Rep.* **7**, 44370 (2017).
25. Camsari, K. Y., Faria, R., Sutton, B. M. & Datta, S. Stochastic  $\mathbb{S}^p$ -Bits for Invertible Logic. *Phys. Rev. X* **7**, 031014 (2017).
26. Wang, T., Wu, L. & Roychowdhury, J. Late Breaking Results: New Computational Results and Hardware Prototypes for Oscillator-based Ising Machines. *ArXiv190410211 Cs* (2019).
27. Li, C. *et al.* Analogue signal and image processing with large memristor crossbars. *Nat. Electron.* **1**, 52–59 (2018).
28. Wang, T. & Roychowdhury, J. Oscillator-based Ising Machine. *ArXiv170908102 Phys* (2017).
29. Acebrón, J. A., Bonilla, L. L., Pérez Vicente, C. J., Ritort, F. & Spigler, R. The Kuramoto model: A simple paradigm for synchronization phenomena. *Rev. Mod. Phys.* **77**, 137–185 (2005).
30. Wang, T. & Roychowdhury, J. PHLOGON: PHase-based LOGic using Oscillatory Nano-systems. In *Unconventional Computation and Natural Computation* (eds Ibarra, O. H., Kari, L. & Kopecki, S.) 353–366 (Springer International Publishing, 2014).
31. Rategh, H. R. & Lee, T. H. Superharmonic injection-locked frequency dividers. *IEEE J. Solid-State Circuits* **34**, 813–821 (1999).
32. Mirzaei, A., Heidari, M. E., Bagheri, R., Chehrizi, S. & Abidi, A. A. The Quadrature LC Oscillator: A Complete Portrait Based on Injection Locking. *IEEE J. Solid-State Circuits* **42**, 1916–1932 (2007).
33. Go Speed Racer... Arduino Speed Test - learn.sparkfun.com. Available at, <https://learn.sparkfun.com/blog/1687> (Accessed: 18th May 2019).
34. Godsil, C. & Royle, G. F. *Algebraic Graph Theory*. (Springer-Verlag, 2001).
35. Rendl, F., Rinaldi, G. & Wiegale, A. Solving Max-Cut to optimality by intersecting semidefinite and polyhedral relaxations. *Math. Program.* **121**, 307 (2008).
36. Leleu, T., Yamamoto, Y., McMahon, P. L. & Aihara, K. Destabilization of Local Minima in Analog Spin Systems by Correction of Amplitude Heterogeneity. *Phys. Rev. Lett.* **122**, 040607 (2019).
37. Albash, T., Martin-Mayor, V. & Hen, I. Analog errors in Ising machines. *Quantum Sci. Technol.* **4**, 02LT03 (2019).

## Acknowledgements

This material is based upon work supported by the United States Air Force under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force.

## Author Contributions

J.C. and S.B. wrote the manuscript text. J.C. conceived, built, and tested the breadboard circuit system. S.G. aided in building the breadboard circuit. S.B. performed the scaling calculations and plots. W.H. advised the program. All authors reviewed the manuscript.

## Additional Information

**Competing Interests:** The authors declare the following patent application: U.S. Patent Application No. 62/826,080.

**Publisher's note:** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2019