# BRIM: Bistable Resistively-Coupled Ising Machine

Richard Afoakwa, Yiqiao Zhang, Uday Kumar Reddy Vengalam, Zeljko Ignjatovic, and Michael Huang Department of Electrical and Computer Engineering, University of Rochester {richard.afoakwa, yiqiao.zhang, uvengala, zeljko.ignjatovic, michael.huang}@rochester.edu

Abstract—

Physical Ising machines rely on nature to guide a dynamical system towards an optimal state which can be read out as a heuristical solution to a combinatorial optimization problem. Such designs that use nature as a computing mechanism can lead to higher performance and/or lower operation costs. Quantum annealers are a prominent example of such efforts. However, existing Ising machines are generally bulky and energy intensive. Such disadvantages may be acceptable if these designs provide some significant intrinsic advantages at a much larger scale in the future, which remains to be seen. But for now, integrated electronic designs of Ising machines allow more immediate applications. We propose one such design that uses bistable nodes, coupled with programmable and variable strengths. The design is fully CMOS compatible for on-chip applications and demonstrates competitive solution quality and significantly superior execution time and energy.

Index Terms—Ising machine, optimization, CMOS accelerators, nature-based computing, quantum annealing

#### I. Introduction

The power of computing machinery has improved by orders of magnitude over the past decades. At the same time, the need for computation has been spurred by the improvement and continues to require better mechanisms to solve a wide array of modern problems. For a long time, the industry focused on improving general-purpose systems. In recent years, special-purpose designs have been increasingly adopted for their efficacy in certain type of tasks such as encryption and network operations [1], [2] More recently, machine learning tasks have become a new focus and many specialized architectures are proposed to accelerate these operations [3], [4]. Much of this work is to construct a more efficient architecture where the control overhead as well as the cost of operation becomes much lower than traditional designs.

In a related but different track of work, researchers are trying to map an entire algorithm to physical processes such that the resulting state represents an answer to the mapped algorithm. Quantum computers marketed by D-Wave Systems are prominent examples. Different from circuit model quantum computers [5], [6], D-Wave machines perform quantum annealing [7]. The idea is to map a combinatorial optimization

Copyright and Reprint Permission: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923. For reprint or republication permission, email to IEEE Copyrights Manager at pubs-permissions@ieee.org. All rights reserved. Copyright ©2021 by IEEE.

<sup>1</sup>Recent theoretical works have claimed increasingly strong equivalence between the two modes of quantum computing [8], [9].

problem to a system of qubits such that the system's energy maps to the metric of minimization. Then, when the system is controlled to settle down to the ground state, the state of qubits can be read out, which corresponds to the solution of the mapped problem.

It is as yet not definitive whether D-Wave's systems can reach some sort of quantum speedup. But one thing is clear: machines like these can indeed find some good solutions to an optimization problem, and in a very short amount of time too. Indeed, a number of alternative designs have emerged recently all showing good quality solutions for non-trivial sizes (sometimes discovering better results than the best known answer from all prior attempts) in milli- or micro-second latencies [10]–[12]. These systems all share the property that a problem can be mapped to the machine's setup and then the machine's state evolves according to the physics of the system. This evolution has the effect of optimizing a particular formula called the Ising model (more on that later). Reading out the state of such a system at the end of the evolution thus has the effect of obtaining a solution (usually a very good one) to the problem mapped.

For example, in some systems, the Hamiltonian is closely related to the Ising formula. Naturally, the system seeks to enter a low-energy state. In other systems, a Lyapunov function of the system can be shown to be related to the Ising formula. In general, these systems can be thought of as optimizing an objective function (in the form of the Ising formula) due to physics. Hence, they are generally referred to as *Ising machines*. Clearly, unlike in a von Neumann machine, there is no explicit algorithm to follow. Instead, nature is effectively carrying out the computation. Ising machines have been implemented in a variety of ways with very different (and often complex) physics principles involved. It is unclear (to us at least) whether any particular form has a fundamental advantage that will manifest in a very large scale.

Note that these systems can not guarantee reaching the ground state in practice.<sup>2</sup> Nonetheless, some systems find a good answer with high speed and a good energy efficiency, as we shall see later with concrete examples. In this paper, we propose a novel CMOS-compatible Ising machine which uses circuit elements' physical properties to achieve nature-based computation. This design is completely different from other efforts of using CMOS circuit to build machines that

<sup>2</sup>Theoretical guarantee in some ideal setup may exist. For instance, adiabatic quantum computing theory says that when the annealing schedule is *sufficiently* slow and in the absence of noise (zero kelvin) the system is guaranteed to stay in the ground state [13].

*simulate* an annealer. We perform a detailed analysis of the design and show that it is a compelling design and superior in many respects to existing Ising machines and accelerators of simulated annealing.

#### II. BACKGROUND AND RELATED WORK

We first explain the background of Ising machines and discuss the state of the art in implementations.

#### A. Ising model

The Ising model is used to describe the Hamiltonian of a system of spins.<sup>3</sup> The model is a general one that describes a system with many nodes (*e.g.*, atoms), each with a spin ( $\sigma_i$ ) which takes one of two values (+1, -1). The energy of the system is a function of pair-wise coupling ( $J_{ij}$ ) of the spins and each spin's reaction ( $h_i$ ) to some external magnetic field ( $\mu$ ). The resulting Hamiltonian is as follows:

$$H = -\sum_{(i < j)} J_{ij} \sigma_i \sigma_j - \mu \sum_i h_i \sigma_i \tag{1}$$

If we ignore the external field, the Hamiltonian simplifies to

$$H = -\sum_{(i < j)} J_{ij} \sigma_i \sigma_j \tag{2}$$

This simplified version is more useful for the purpose of our discussion. Henceforth, when we refer to the Ising model or formula, we mean Eq. 2.

A physical system with such a Hamiltonian naturally tends towards low-energy states and thus serves as a convenient machine to solve a problem with a formulation equivalent to the Ising formula – provided we can configure parameters  $(e.g.,\ J_{ij})$  to match that of the problem.

# B. Optimization problems and mapping issues

A group of optimization problems naturally map to an Ising machine. Perhaps the most straightforward problem to map is (weighted) Max-Cut. Given a graph, G=(V,E), a cut is a partition of vertices into two sets of, say,  $V^+$  and  $V^-$ , where  $V^-=V-V^+$ .

The Max-Cut problem tries to find a cut such that the combined weight of the edges spanning the two sets of vertices is maximum. In other words, the maximum cut is

$$\underset{V^{+} \in \mathbb{P}(V)}{\operatorname{arg max}} \left( \sum_{(i,j) \in E;} W_{ij} \right)$$

$$\underset{i \in V^{+}:}{\sum} \underset{i \in V^{-}}{W_{ij}}$$
(3)

where  $W_{ij}$  is the weight of edge (i, j). (We will refer to the resulting  $\sum W_{ij}$  as the cut *value* in this paper.)

It is easy to see the resemblance between Eq. 2 and 3. In fact, if we set the coupling weight  $(J_{ij})$  to be the negative of edge weight  $(-W_{ij})$  then the Ising formula is simply twice the negative cut value plus a problem-specific constant  $(\sum W_{ij})$  as follows (for notational simplicity, for  $i \geq j$  we set  $W_{ij}$  to 0):

<sup>3</sup>Though commonly called the Ising model, the model itself existed before Ernst Ising (read "Easing") solved analytically a one-dimensional system.

$$H = -\sum_{\sigma_i = -\sigma_j} J_{ij} \sigma_i \sigma_j = \sum_{\sigma_i = -\sigma_j} W_{ij} \sigma_i \sigma_j + \sum_{\sigma_i = \sigma_j} W_{ij} \sigma_i \sigma_j$$

$$= -\sum_{\sigma_i = -\sigma_j} W_{ij} + \sum_{\sigma_i = \sigma_j} W_{ij} = -2\sum_{\sigma_i = -\sigma_j} W_{ij} + \sum_{\sigma_i = \sigma_j} W_{ij}$$
(4)

Hence if the machine finds the ground state of the Hamiltonian, we have the maximum cut. Finding out the maximum cut of an arbitrary graph is an NP-hard problem. Practical algorithms only try to find a good answer. Similarly, existing Ising machines (including our design) are all Ising *sampling* machines that typically provides a good sample of a lowenergy state, with no guarantee of optimality.

Because of the trivial mapping of the Max-Cut problem to the Ising formula, designers of Ising machines, often focus on this optimization problem. However, other optimization problems can also be mapped to an Ising machine. Indeed, every problem in the original NP-complete set [14] can be expressed by an Ising formulation specifically designed for that problem [15]. Note that Ising formulation may require more nodes than that of the original formulation and usually requires additional time to compute coupling coefficients in the Ising formulation from the original formulation. This transformation is largely straightforward and the need for it is problem-dependent and thus shared by all Ising machines.

Another transformation, however, may be necessary depending on the machine's coupling topology. While we will get into the details as we discuss the machines, it is worth emphasising up front the significant impact of the issue. If a machine has only local connections, then spins mapped to nodes not directly connected have to rely on additional, auxiliary spins. An alternative description is that if a machine has a limited connection topology, then the graph of a problem needs to be transformed (e.g., using minor embedding [16]) into a new graph that observes the limitation imposed by the machine. Consequently, a graph of size N may contain many more nodes (e.g.,  $\frac{N^2}{2}$ ) after the transformation. Fig. 1 illustrates these transformations in the process of solving a problem on an Ising machine.

## C. Quantum mechanical and optical Ising machines

There are many natural systems that can be described by the Ising model. Take two existing systems with relatively large footprints for example. D-Wave's quantum annealers use superconducting qubits as the basic building block. These bits are then coupled together with couplers forming a connection topology known as the Chimera graph. This is an important architectural constraint that limits the typology of the problem that can be mapped to the machine. As we will see later, despite supporting nominally more than 2000 spins, many of our benchmarks can not be mapped to the machine. Another disadvantage of the system is the cryogenic operating condition (15mK) needed for the quantum annealer. This requirement consumes a significant portion of the 25KW power of the machine [17].

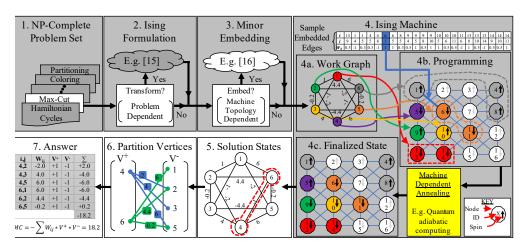


Fig. 1. High-level overview of a generic Ising machine workflow solving (weighted) Max-Cut problem (box-1). Ising formulation (box-2) is a required property for solving diverse problem sets. Minor-embedding (box-3) is a topology-dependent design factor and *not* required for all Ising machines. Box-4a and -4b shows how an example embedded graph fits unto the machine. The yellow box represents the machine dynamics.

Coherent Ising machines (CIMs) are another recent example of Ising sampling machines [10], [18]–[21]. In a CIM, an optical device called OPO (optical parametric oscillator) is used to generate and manipulate the signal to represent one spin. Unlike in a D-Wave Ising machine, the coupling between spins in CIM is relatively straightforward in principle. As a result, CIM implementations have always supported all-to-all coupling. The authors also emphasized that the 2000-node CIM is therefore far more capable than D-Wave 2000Q which can only map all-to-all problems of size 64 [22].

CIM is not without its disadvantages. To support 2000 spins, kilometers of fibers are needed. Temperature stability of the system is thus an acute engineering challenge. Efforts to scale beyond the currently achieved size (of about 2000) have not been successful as the system runs into stability problems. Also worth noting is that the coupling between nodes is – at least in the current incarnation – implemented via computation external to the optical cavity. There is a rather intensive computational demand (100s of GFLOPS) [23]. Every pulse's amplitude and phase are detected and its interaction with all other pulses calculated on an auxiliary computer (FPGA). The computation is then used to modulate new pulses that are injected back into the cavity. Strictly speaking, the current implementation is a nature-simulation hybrid Ising machine. Thus, beyond the challenge of constructing the cavity, CIM also requires a significant supporting structure that involves fast conversions between optical and electrical signals.

These room-sized Ising machines are certainly worthwhile creations for the sake of science. In particular, investigations are needed to see whether the theoretical underpinning for these machines is relevant in practice. As we shall see later, both models have significant room for improvements.

#### D. Electronic oscillator-based Ising machines

A network of coupled oscillators is another physical implementation of an Ising machine. After sufficient time, the

coupled oscillators will synchronize forming stable relative phase relationship. While many factors (e.g., amplitude, stochastic noise) will influence the phase of each oscillator, the following formula is a simplified steady-state description of phase relationship for N oscillators:

$$\frac{d}{dt}\phi_i(t) = \sum_{i=1}^{N} J_{ij} sin\Big(\phi_j(t) - \phi_i(t)\Big)$$
 (5)

Note that this simplified model ignores certain elements (*e.g.*, diffusion due to noise) and is thus an approximation of a more complicated reality. Given such a differential equation describing a dynamic system, it can be shown that a Lyapunov function in the following form exists [11]:

$$H(\Phi(t)) = -\sum_{i < j} J_{ij} cos(\phi_j(t) - \phi_i(t))$$
 (6)

This means that the system will generally evolve along a trajectory that minimizes the Lyapunov function [25]. As a result, the system's stable states represent good solutions that minimize the right hand side of Eq. 6. On a closer inspection, we see the resemblance of Eq. 6 and the Ising model (Eq. 2). Specifically, when all phases  $(\phi_i)$  are either 0 or  $\pi$ , the two formulae are the same.<sup>5</sup> A number of oscillator-based Ising machines have been recently proposed [11], [26], [27].

<sup>4</sup>The observation of such synchronization dates back to at least the 17th century when Huygens observed synchronization of two pendulums [24]. Synchronization phenomenon is the subject of research efforts in a wide variety of fields. Large-scale synchronization of firefly flashings and rhythmic applause in a large crowd of audiences are but two examples in the general underlying principles beyond mechanical objects.

 $^5 \mathrm{In}$  fact, the formulation of Eq. 6 is similar to the classic XY spin model (again ignoring external field): each spin can point to any direction along an "XY" plane and thus can be represented by a phase  $(\phi_i)$ . Ising model is thus a special case of the XY model. In other words, a system of coupled oscillators form an "XY machine" (not an Ising machine). An XY state can be quantized into an Ising state  $(\phi_i=0,\pi)$  in a number of different ways. For the sake of this paper, let us simply imagine direct quantization which rounds the phase to the nearest multiple of  $\pi$ .

All these examples use LC tank oscillators. While this is a common practice for analog circuit designers and relatively straightforward for discrete-element prototypes, the use of LC tanks introduce non-trivial practical challenges in integrated circuit (IC) designs. The lack of high quality inductors and the usually high area costs of incorporating them are common challenges for integrated RF circuitry. These desktop Ising machines are a significant improvement (at least in size) over room-sized Ising machines. But, for genuine wide-spread applications, we believe a clean-slate IC-focused design is a valuable direction to pursue. Needless to say, we believe there will be significant cross-pollination of different approaches and future practice may very well be a confluence of multiple styles of Ising machines.

#### E. Accelerated simulated annealing

Finally, a set of chips have been designed to accelerate simulated annealing [28] or a variant of the classic algorithm. These chips are often described as having tens of thousands of spins [29], [30]. In these designs, the spins are *virtual* in that they are bits in memory and manipulated by an algorithm (simulated annealing). These machines are specially built to accelerate that algorithm. Hence we refer to such a machine as an *Accelerated Simulated Annealer* or ASA for short.

These ASAs differ fundamentally from *physical* Ising machines. In a physical Ising machine, nature guides the spins to a preferable state according to physical laws. Thus, the machine can achieve ultimate speed and energy efficiency in principle – though it is entirely possible that a particular physics exploited is slow or energy-intensive to control; or it may be expensive to enable the physics, such as in creating the cryogenic environment required for quantum annealing.

## F. Taxonomy

Existing Ising machines can be categorized based on whether they use physical or virtual spins. Note that this aspect is more of a continuum than a binary distinction. In the case of CIM, for instance, the interaction of the spins happens physically in the fiber. But the appropriate amplitude of feedback is controlled using external calculation. Fig. 2 shows this high-level classification, where the horizontal dimension summarizes the contrast between physical and virtual systems.

A second important differentiator of Ising machines is the connection topology. As discussed before, with a local connection, there is a limitation on the kind of problems that can be mapped to the machine. When a problem does not map directly, a transformation is needed to convert the problem, usually by requiring many auxiliary nodes. The nominal number of spins a machine provides is therefore a very poor representation of the machine's capability. We will provide a quantitative analysis on this point in Sec. IV. At this point, we qualitatively place existing Ising machines based on their connection topology on the vertical axis in Fig. 2.

# III. ARCHITECTURE OF THE PROPOSED ISING MACHINE

In this section, we start with a simplified system to provide some intuition about how common electronics can also make

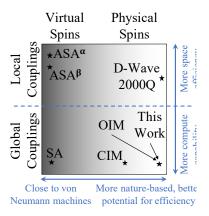


Fig. 2. High-level taxonomy of Ising machine design space showing a continuum of capabilities.  $ASA^{\alpha}$  and  $ASA^{\beta}$  refer to machines in works [29] and [30] respectively.

a physical Ising machine (Sec. III-A); then describe in more detail the system architecture (Sec. III-B); and finally present theoretical analysis of the system's optimum-seeking behavior (Sec. III-C).

#### A. Overview & intuition

As already discussed before, existing physical Ising machine have different strengths and weaknesses. The room-sized machines are vehicles for continued scientific exploration of the underlying principles. It is particularly useful to show the difference between ideal theoretical capabilities and what can be achieved in practice. For instance, according to quantum adiabatic theory, the system's Hamiltonian needs to be changed sufficiently slowly to guarantee that the system stays in ground state. In practice, D-Wave's quantum annealer appears to offer a fixed annealing schedule. Operated as such, the quantum annealer predictably provides sub-optimal solutions, as we will show later.

The question then becomes: can we build better (smaller, less power-intensive) physical Ising machines. And the answer is: yes, with electronics. In digital designs, electronic devices are often thought of as no more than the building blocks of functional units. But their behavior is also subject to physical laws that can be leveraged to perform nature-based computation. As it turns out, practical physical Ising machines can be built out of common devices such as capacitors and resistors. We start with one such simple design to show the working principle. Of course, this design is not yet a high-performance system. But as we will show later, with a proper architecture and careful design of key elements, a physical Ising machine built out of electronics is much more compelling than existing proposals. Additionally, it can be fabricated entirely in a CMOS process.

**Intuition:** In the Ising model, when two nodes (say, i, and j) are strongly and positively coupled (i.e.,  $J_{ij}$  is large and

<sup>&</sup>lt;sup>6</sup>We are unable to find ways of adjusting this schedule.

positive), their spins are likely to be parallel  $(\sigma_i = \sigma_j)$ . In this way, the term  $-J_{ij}\sigma_i\sigma_j$  will contribute to lowering the energy. Conversely, a strong negative coupling  $(J_{ij}$  is large and negative) will likely lead to anti-parallel spins  $(\sigma_i = -\sigma_j)$ . Finally, weak coupling  $(|J_{ij}|$  is small) suggests that the two spins are more likely to be independent.

This behavior can be easily mimicked with resistively coupled capacitors. Consider representing a node with a capacitor, where the polarity of the voltage across it represents the spin of the node. More specifically, in Fig. 3, if a node has a spin of "-1", the voltage at the upper terminal (top plate of the capacitor, labeled "+") is lower than that of the lower terminal (bottom plate of the capacitor). We can then connect nodes with different resistors. A strong coupling means high conductance (or low resistance), so that voltages of two nodes can more easily equilibrate. So, we set  $R_{ij} \propto 1/J_{ij}$ . The sign of coupling can also be achieved by connecting either the same or opposite polarity in the corresponding capacitors. Fig. 3 shows a simple 4-node system mapped from a logical graph of a Max-Cut problem with the labeled edge weights. <sup>7</sup> It is not difficult to see that the solution should separate the nodes into  $\{1,4\}$  and  $\{2,3\}$ . Let's see how the machine functions.

The graph translates to couplings in a straightforward manner:  $R_{ij} = R/J_{ij} = -R/W_{ij}$ , where the sign indicates polarity of coupling. For instance, nodes 1 and 4  $(W_{1,4} = -1)$  are parallel/positively coupled  $(R_{1,4} = -R/W_{1,4} = R)$ . So a resistor R connects the upper terminals of nodes 1 and 4 and another connects the two lower terminals. Nodes 1 and 3  $(W_{1,3} = +0.5)$  are antiparallel/negatively coupled  $(R_{1,3} = -R/W_{1,3} = -2R)$ , so a 2R resistor connects the upper terminal of Node 1 and the lower terminal of Node 3, and another 2R resistor connects their remaining terminals.

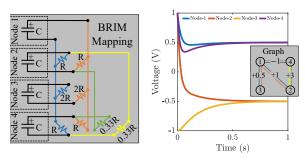


Fig. 3. Sample Max-Cut problem mapping to a 4-node resistive Ising machine (left) and corresponding solution, cut value of +4.5, (right).

Once initialized with random polarities, these coupled capacitors can indeed seek some equilibrium as shown in Fig. 3 (right). In this example, the polarity of the capacitors at equilibrium indeed gives the best solution to the Max-Cut problem.

While this oversimplified design confirms the intuition that it *can* find a solution, it is far from a robust design. For

example, depending on the initial state and system scale, the voltages at equilibrium can be  $0\,V$  or just too low for reliable readout. The equilibrium is also temporary because leakage will make all nodes decay to  $0\,V$  eventually, rather than staying at the desired voltage levels. Nevertheless, the resistively-coupled capacitor network is at the core of our proposed Ising machine. To induce and maintain the nodes at equilibrium, we can introduce a local feedback unit to make the node voltages bistable. For brevity, we will refer to such a Bistable, Resistively-coupled Ising Machine as BRIM.

## B. Architecture of an integrated design

We now discuss the architecture of a more complete system designed for integrated circuits. The system is illustrated in Fig. 4 and consists of: ① bistable nodes (N1 to N4) and their digital interface (DFF 1-4); ② coupling units connecting the nodes  $(CU_{ij})$ ; ③ programming units to configure coupling weights (including: memory, multiplexers, digital-to-analog converters, and the column control systems); and ④ annealing control. We discuss each in turn as follows:

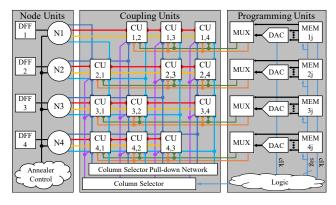


Fig. 4. Block diagram of BRIM system components. Nodes are  $N_i$  and coupling units  $CU_{ij}$ . The multi-colored interconnecting links are wires.

1) Nodes: The bistable nodes  $N_i$  are shown at the left side of Fig. 4. In Fig. 5 we show a more detailed illustration of one node. Recall that in the simplified circuit (Fig. 3) the capacitor's voltage can be too low (even zero) as compared to electronic noise level to reliably indicate the node's spin. A feedback circuit is therefore needed to stabilize the voltages at the desired levels  $(e.g., \pm V_{dd})$ . Two conditions are required: ① the capacitor should be charged according to its polarity when the voltage is between  $-V_{dd}$  and  $+V_{dd}$ , and discharged when the voltage exceeds this range; and ② at low voltages, the feedback circuit should supply a low current in order not to overwhelm signals coupled from other nodes.

Combining these considerations, we can design a feedback circuit with the current-voltage (IV) curve as shown in Fig. 5. Because of the slanted "Z" shape of the IV curve, we call such sub-system a ZIV diode. As seen in the IV curve, for capacitor voltages between  $-V_{dd}$  and  $V_{dd}$ , the ZIV diode acts as an active element charging its voltage closer to  $\pm V_{dd}$ . Conversely,

 $<sup>^{7}</sup>$ Recall the coupling and edge weight relation,  $J_{ij}=-W_{ij}$ , in the Ising formula.

for voltages outside the range, the ZIV diode acts essentially as a linear resistor to discharge the capacitor.

Each node is supplied with two pairs of switches to set the capacitor voltage to a known spin state. This is useful for specific initialization or to change the system state as described in the "Annealing Control" (Sec. III-B4). Finally, the spin of the node can be read out by a digital latch  $(e.g., a\ D\ flip-flop)$  connected to the ZIV diode's V+ terminal. This converts the spin to a bit  $(1\ or\ 0)$ .

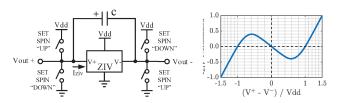


Fig. 5. Simplified schematic of one BRIM node (left) and ZIV diode's IV characteristic (right).

2) **Coupling units:** With an array of BRIM nodes, we have the flexibility to couple them in any desired topology. As we will discuss quantitatively later, an all-to-all coupling is computationally more useful and is thus our choice.

In our system, the coupling is (uni)directional. There are two separate coupling units (CU) connecting node  $N_i$  to  $N_j$ :  $CU_{ij}$  and  $CU_{ji}$ . This directional coupling is achieved by connecting the node capacitor through a voltage buffer before sending to the CU. The coupling coefficient of both directions are of course the same  $(J_{ij} = J_{ji})$ . In principle, a bidirectional coupling has similar effects and is a simpler design point. But, empirically, we found directed coupling tends to produce better solution quality at the expense of area. A more detailed analysis as to the reason for this phenomenon is left as future work.

Because of the directed coupling scheme, each node has separate input and output terminals (two each for fully-differential coupling), as shown in different colors in Fig. 4. Similarly, each CU also has four input/output terminals, with one pair of transistors connecting the same polarity input/output nodes (parallel coupling) and another pair to establish antiparallel coupling, as shown in Fig. 6.

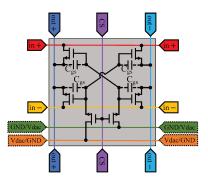


Fig. 6. Schematic of BRIM coupling unit. Color schemes are matched to system-level block diagram in Fig. 4. CS is column selector.

The two pairs of transistors in the CU are biased in trioderegion to act as variable/programmable resistors. Depending on the polarity of the coupling coefficient  $(J_{ij})$ , only one pair of the transistors are turned on and biased to an appropriate non-zero  $V_{GS}$ . For example, to establish a negative (antiparallel) coupling, the gates of the upper-right and lower-left transistors are biased to a non-zero voltage  $(V_{dac})$ , while the other pair of transistors are biased at 0V (GND).

- 3) **Programming units:** Both the initial nodal values and the coupling resistance are programmable. Programming of a resistor is achieved through a transistor with adjustable gate voltage. To the right of the coupling unit array in Fig. 4 is the programming array. This array consists of digital memory for storing the weights which drive an array of digital-to-analog converters (DACs). A small amount of such DACs are sufficient to program all the coupling units in a time-interleaved fashion per column. In such configuration, we need corresponding column selectors and pull-down logic as shown below the coupling units in Fig. 4.
- 4) Annealing control: With the basic network of coupled bistable nodes, the system can reach the local optimum determined by the initial state. Two commonly used mechanisms for annealers are incorporated to allow the system to escape local optima. First, the coupling strength is globally increased over time. In this way, at the very beginning, the machine is only weakly coupled, rendering the energy landscape relatively flat. This helps the system explore the landscape in a coarse granularity. We choose an exponential annealing schedule because (a) it is a common practice, and (b) it can be conveniently achieved by discharging an appropriately sized capacitor as the global annealing scheduler. The voltage from this capacitor is then used to control the variable gain buffer in each node to achieve the change in coupling strength.

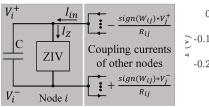
Second, we also adopt a similar strategy as that used in simulated annealing. By performing a "spin flip" of select nodes (*i.e.*, to change the spin to its opposite value using the switches in 5), we can enter a neighboring state in the global phase space. This allows the system to escape the current basin of attraction and explore new regions. In simulated annealing, the probability of such bit flips is a function of both energy difference due to the bit flip and the current temperature. For implementation convenience, we only use the temperature to decide the probability/frequency of spin flips. The temperature follows the same exponential annealing schedule discussed above. In other words, the frequency of spin flips decays exponentially.

With these supports, our BRIM is used similarly to other Ising machines: first, program the weights; then, select the annealing time; and finally, read out the state of the nodes after the completion of the annealing schedule. Note that the system can be used in a number of different ways: the annealing time can be adjusted; the spin flip frequency can be tuned; the machine can also be used together with a software-based search algorithm (*e.g.*, simulated annealing), perhaps by searching a subspace.

### C. Theoretical analysis

It is often desirable to show some theoretical foundation for nature-based computing machines. We present a Lyapunov stability analysis of BRIM to show its ability for optimum-seeking. Fig. 7 shows a simplified node i in BRIM where other nodes are coupled into i from the right. Writing the current equation, we get Eq. 7, where  $J_{ij}$  is the effective conductance of the coupling taking into account the polarity of the coupling:

$$\frac{dv_i}{dt} = \frac{1}{C}(I_{in} - I_{ZIV}) = \frac{1}{C} \left[ \sum_{j \neq i} J_{ij} v_j - g_{ZIV}(v_i) \right]$$
(7)



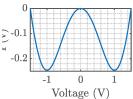


Fig. 7. Right: Single BRIM node unit coupled to other nodes. Right: Potential well of the bistable node.

BRIM is a continuous time system where the state is summarized by  $v(t) = [v_1(t), v_2(t), \dots, v_N(t)]^T$ . (For brevity, in the following,  $v_i(t)$  will be abbreviated as  $v_i$ .) Recall that in Lyapunov analysis, if a function H(v) can be found such that  $\frac{\mathrm{d}H(v)}{\mathrm{d}t} = 0$  at point  $v^e$  and  $\frac{\mathrm{d}H(v)}{\mathrm{d}t} < 0$  in the region around  $v^e$ , then  $v^e$  is the equilibrium point [25]. In our case, this point is the solution to the equation set  $\frac{\mathrm{d}v_i}{\mathrm{d}t} = 0$ ; i = 1..N. In other words, once the system enters a region, it will inevitably evolve towards lowering H(v) (because its time derivative is negative) until it reaches the equilibrium point  $v^e$ .

negative) until it reaches the equilibrium point  $v^e$ . To ensure  $\frac{\mathrm{d} H(v)}{\mathrm{d} t}$  is non-positive, we can construct it to be in a negative square form. This can be achieved by imposing the following construction rule.

$$\frac{\partial H(v)}{\partial v_i} = -\alpha \frac{\mathrm{d}v_i}{\mathrm{d}t}; \ \alpha > 0 \tag{8}$$

Following the chain rule, it is not difficult to see that:

$$\frac{\mathrm{d}H(v)}{\mathrm{d}t} = \sum_{i} \left( \frac{\partial H(v)}{\partial v_i} \frac{\mathrm{d}v_i}{\mathrm{d}t} \right) = -\alpha \sum_{i} \left( \frac{\mathrm{d}v_i}{\mathrm{d}t} \right)^2 \quad (9)$$

One choice of a Lyapunov function satisfying the conditions in Eq. 8 is shown in Eq. 10, where  $P(v_i)$  is obtained from  $g_{ZIV}(v_i)$  by integration over  $v_i$ .

$$H(v) = \frac{\alpha}{C} \left( -\sum_{i < j} J_{ij} v_j v_i + \sum_i P(v_i) \right)$$
 (10)

It is important to notice that the "Z" shape of  $g_{ZIV}(v)$  will give P(v) a double-well profile (as shown in Fig. 7-right) with two stable equilibrium points at voltages corresponding to two nontrivial zero-crossings of the IV curve  $(e.g., v_i = \pm 1V)$  and a saddle point at  $v_i = 0V$ . Given the double-well profile of P(v), the state of a stable solution will consist of voltages

at (or at least very close to) these equilibria  $(e.g., \pm 1V)$ . Thus, the second term of Eq. 10 will be (close to) a constant. Consequently, minimizing H(v) is equal to minimizing the first term  $-\sum_{i < j} J_{ij} v_i v_j$ , which is the Ising formula.

Note that this analysis does not guarantee that the system will converge to a *global* minimum as it depends on the energy landscape and initial conditions. This is similar to other annealers: None has strong guarantees for reaching ground state in a typical usage scenario (as opposed to ideal conditions) or in an efficient manner. For example, it is shown that simulated annealing can reach ground state in a system with a finite phase space. However, the time it takes to do so may be longer than enumerating the space [31].

#### IV. EXPERIMENTAL ANALYSIS

In this section, we provide some experimental analysis of BRIM by

- 1) Comparing Ising machines at a high level (Sec. IV-C);
- 2) Discussing architectural design choices (Sec. IV-D); and
- 3) Showcasing a discrete-component prototype (Sec. IV-E).

First, we describe the machines (Sec. IV-A) the benchmarks (Sec. IV-B) used in the analyses.

#### A. Ising machines

We compare BRIM to 4 other machines using both physical and virtual spins. We use results reported in literature when direct measurement or modeling is unavailable.

- 1) **D-Wave:** we use 2000Q which is the latest quantum annealer [32]. We launch jobs using the API provided by D-Wave [33]. For each graph problem, we collect 50 samples.<sup>8</sup>
- 2) **CIM:** is an optical Coherent Ising Machine [10]. There is no known public access to the actual hardware and no model available for simulation. Thankfully, there are reported results for two commonly used benchmarks that allow us to make meaningful comparisons.
- 3) **OIM:** is an electronic Oscillator-based Ising Machine [11]. For simulation, we use the Kuramoto model-based code provided in [11].
- 4) ASA: refers to a number of related designs of Accelerated Simulated Annealers [29], [30]. These accelerators use virtual spins and are straightforward to model based on the description in literature. We focus on one with 30,000 nominal spins. In this design, the coupling follows a near-neighbor pattern dubbed the King's graph. All nodes are grouped into 4 groups. Every annealing step  $(0.22 \, \mu s)$ , nodes in one of the groups will process in parallel: they read off neighbor's spin and the associated weights to compute whether keeping the same spin or inverting its current spin provides a lower energy in the neighborhood. In addition, random bit flips similar to those in standard simulated annealing algorithms are also adopted.

 $<sup>^8</sup>$ In terms of timing, we do not specify any constraints, and adopt D-Wave default values. Specifically:  $20 \,\mu s$ ,  $198 \,\mu s$ ,  $21 \,\mu s$ , and  $11.7 \,m s$  respectively for annealing, data readout, inter-sample delay, and qubits programming [34].

As explained in Sec. II (box 3 of Fig. 1), machines without all-to-all coupling (D-Wave and ASA) require preprocessing of input graph through minor embedding [16]. This process currently takes significant time on a conventional computer (more details in Sec. IV-D) and can often fail for larger graphs. ASA uses the King's graph which is even more limited than the Chimera graph in D-Wave, the minor embedding process takes even longer time to complete and results in more physical nodes needed. In our high-level analysis, we ignore the significant time needed for this preprocessing.

BRIM is modeled using Eq. 7. We perform simulations using MATLAB's nonstiff, single step, 5th-order differential solver (ode45). Finally, we also performed the classic Simulated Annealing (SA) [35] algorithm as a reference using MATLAB. The execution time is measured on server cluster nodes of Intel Xeon Platinum 8268 CPUs at 2.9GHz with 371GB of RAM.

# B. Benchmarks

To compare these systems, we use a set of popular graphs called "Gset" (and their derivatives) with diverse node sizes and edge densities. These graphs are just weighted, undirected graphs and not associated with any specific problem. They naturally correspond to an Ising formula. Thus, we are comparing different machines optimizing the same set of Ising formulae.

Because of the direct mapping between an Ising formula and the Max-Cut problem, and that many algorithms have been developed for optimizing Max-Cut, it is convenient for researchers, especially Ising machine designers, to *interpret* these graphs as specifying a Max-Cut problem, which we follow in this paper. Note that this does not mean we can only solve a Max-Cut problem as already discussed in Sec. II.

The graphs we use can be divided into the following groups:

- 1) Regular graphs: We use the original Gset graphs from Stanford [36]. These graphs have between 800 and 20,000 nodes. The edges as well as the weights of such edges, were generated probabilistically, sometimes between +1 and -1, and sometimes all +1. We only use those graphs with less than 2000 nodes in our experiments.
- 2) Small graphs: Although supporting nominally 2048 spins, D-Wave's machine can not map even the smallest graph in Gset. We therefore generate graphs with smaller node sizes (*e.g.*, 120) and/or edge densities so that they can be successfully map onto D-Wave. For this purpose, we used rudy, a machine-independent graph generator [37], which is the same generator used to produce the "Gset" graph suite.
- 3) Tiny graphs: Finally, we also generated fully-connected graphs with random edge weights, and node sizes ranging from 16 to 32 (in increments of 4). Each node size has 20 sample graphs, for a total of 100 graphs. For these graphs, we are able to enumerate all possible spin combinations to determine actual maximum cut.

#### C. High-level comparison

It is important to keep in mind that Ising machines are far from mature. Early designs and prototypes are necessarily experimental in nature and thus lack the polish of, say, a conventional architecture. Much of the performance difference may be due to the art of prototyping rather than the fundamental science of the mechanism being exploited. This is perhaps especially the case for D-Wave, as in our comparison, it is the only actual hardware that we have access to. (CIM and OIM both have hardware prototypes but are not accessible to us.)

We start with a crude, high-level comparison of different Ising machines. There are several practical factors that make this comparison crude and incomplete. First, there is no single problem that can be measured on all machines. This is primarily because CIM only reported raw data on a very specific set of benchmarks and we are unaware of any reliable model of the physics that is publicly available. Second, the workload of optimization usually allows tradeoff between speed and quality of the solution. Ideally, we will fix one metric (say, execution time) and compare the other (quality of solution). But in some cases, such control is unavailable to us. Third, the execution result depends on initial conditions. So any single run is subject to random chances. The common practice of using these machines is doing multiple runs and taking the best solution, which we follow. But this value should still be regarded as a random variable.

1) Room-sized machines: With these caveats in mind, Table I shows the estimated power, chip area (when applicable), and the execution time and solution quality of a few workloads. The cut value itself changes a lot with the workload. Therefore, for solution quality, we use the distance from the (presumed) ground-state solution (the lower the distance the better). This is a much more stable metric. We use the best reported cut value anywhere as the (presumed) ground-state.

TABLE I
HIGH-LEVEL COMPARISON OF DIFFERENT ISING MACHINES. *Time* IS
ANNEALING TIME. *Dist.* IS SOLUTION DISTANCE FROM BEST. SMALL
GRAPHS AVERAGE SOLUTION: 374. TINY GRAPHS AVERAGE SOLUTION:
23. G22 AND G39 BEST SOLUTIONS: 13359 [38]; 2408 [39].

Parameters	D-Wave	CIM	OIM	ASA	BRIM	
Power (W)	25 K	210	-	$\approx 1$	$\approx 250  m$	
Area $(mm^2)$	-	-	-	$4.3 \times 5.5$	$\approx 5$	
Effective Spins	126	2000	-	160	2000	
G22 Dist.	-	46	56	-	46	
Time	-	5 ms	6 ms	-	$0.25\mu s$	
G39 Dist.	-	47	66	-	46	
Time	-	5 ms	6 ms	-	$0.25\mu s$	
Tiny Graphs Dist.	6	-	0	0	0	
Time	$20 \mu s$	-	$20 \mu s$	$20\mu s$	$20\mu s$	
Small Graphs Dist.	12	-	0	7	0	
Time	$20 \mu s$	-	$20 \mu s$	$20\mu s$	$20  \mu s$	
Avg. Gset Dist.	-	-	9.7	-	2.6	
Time	-	-	0.8ms	-	$2.2\mu s$	

First, we look at CIM using G22 and G39 because these are tested on CIM and reported (Fig. 3 of reference [10]). These graphs cannot be mapped on D-Wave. For other machines, we try to match CIM's solution quality and show execution

time. We see that BRIM can obtain similar quality solutions 4 orders of magnitude faster with a power consumption also about 4 orders of magnitude better. Here, we model a 2000-node BRIM built with 45 nm CMOS technology with all-to-all connections.<sup>9</sup>

The next two rows show small- and tiny-graph comparisons intended to contrast D-Wave to others. Here, because D-Wave's current interface allows only fixed annealing time, we use the same annealing time for all machines. Recall that these are the best answers in 50 runs for each. We see that in these cases BRIM always achieves the best result among all machines, while D-Wave does not.

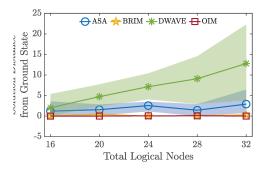


Fig. 8. Performance (solution quality) for Ising machines under test. All systems have  $20\,\mu s$  annealing time. Solid lines represent averages, and shaded regions represent range.

In the case of tiny graphs, we can enumerate and verify that BRIM, OIM, and ASA always reach the ground state. Yet, even for these tiny graphs, D-Wave does not reach the ground state, with the *best* solution having a distance of 6. A more detailed view of the solution quality can be seen in Fig. 8. ASA and BRIM can both achieve ground state often. But BRIM can do so with a higher probability (66%) than ASA (25%). These are concrete examples where the theoretical ability to reach the ground state for an Ising machine is just that: a theoretical ability.

To sum, we see that the room-sized machines (CIM and D-Wave) do not show any tangible advantage in solving optimization problems. D-Wave can only map the smaller problems due to the connection limits. Even on these smaller problems, its solution quality trails behind others. Recall this limited connection means additional compute time to perform minor embedding, which at the moment takes  $5\,s$ , on average. CIM is less power-hungry and can map bigger problems due to its all-to-all connections. Nevertheless, there is no tangible advantage in any figure of merit. Again, these machines may (or may not) prove useful for scientific exploration and may (or may not) show some qualitative superiority at some other scale or at a future time.

 $^9 \text{For}$  a coarse grain, order-of-magnitude estimate; ① Area from preliminary Cadence layout shows CUs are  $1\,\mu m \times 1\,\mu m$  each, a node is  $9.5\mu m \times 9\,\mu m$ , and a DAC is  $2\,\mu m \times 1\,\mu m$ . ② Power for a 6-node system (Fig. 12) shows 620  $\mu W$ . The Gsets (Table I) are all sparse (with 16.4 edges/node ratio), we simply extrapolated to  $250\,mW$  for 2000-node BRIM.

The most important take-away point is that just because the machine leverages nature to perform computation does not necessarily make it efficient. Much engineering diligence is needed to convert some theoretical possibilities to tangible practical benefits.

2) Electronic designs: Next, we look at OIM. While the current OIM prototype is a desktop machine, in principle it can be scaled down in size and up in frequency. The primary unknown is to what extent the nodes can scale down and still operate like ideal LC-tanks. For this study, we assume a 100MHz frequency. Admittedly, this is nothing more than a rough guess. We see that OIM is perhaps comparable to BRIM, though subjectively, we feel that the practical challenges are far more daunting than in BRIM.

Finally, we compare ASA with BRIM. ASA is essentially a specialized computer doing algorithmic search for better cut values. Clearly, ASA takes advantage of the tremendous cumulative improvements of CMOS technology to provide relatively fast and efficient computation. However, it is still a modified von-Neumann machine (with its only program hardwired). In contrast, BRIM is a physical Ising machine where nature is performing the computation. Indeed, here, we see that to get similar quality solutions on small graphs, BRIM is still producing better results while consuming a tenth of the power, with a 5x lower area cost, and providing 12x more spins. It is tempting to think that the area comparison is unfair because ASA offers a large number of nominal spins. We will show next that the nominal number of spins is an extremely poor, if not useless, metric.

3) **Recap:** We see that: ① for the scale of problems we are discussing, room-sized machines are not advantageous; ② ASA, OIM, and BRIM are three possible candidates for chipscale applications with perhaps different strengths. We look at these models in more detail next.

## D. Detailed analyses

1) The impact of topology: We hope by now it is starting to be clear that BRIM is a compelling design. Next, we want to get into details to understand some of the intrinsic strengths that made it a compelling design. We first discuss an important issue in Ising machine architecture. As already mentioned, when an Ising machine uses a more local coupling network, it has a more limited ability to map problems. Many problems would need far more physical nodes than logical nodes in the problem. As shown in Fig. 9, even with an edge density of only 6%, it takes about 1600 nodes (in D-Wave's Chimera configuration) to map a problem of size 126. Worse still, many nominally unused spins cannot be used due to their location. In fact, no size 127 problems we tried can be mapped to the 2048 nominal nodes. Moreover, it takes up to 10 seconds to perform this minor embedding for Chimera (and even longer for King's graph). This is much longer than the annealing itself and more than needed to get a better solution using conventional simulated annealing. Finally, the auxiliary spins degrade the machine's ability to find good solutions. As we saw in Table I, D-Wave and ASA have trouble reaching the ground state for small graphs. All in all, at the moment, we see no evidence that using local connections (at least these specific ones) is helpful in any practical way.

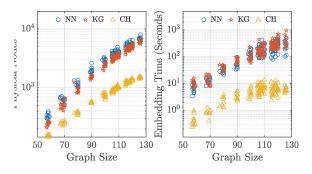


Fig. 9. Minor-embedding [40] mapping for diverse graphs. Y-axis is log scale. **Left:** Total physical nodes required for diverse logical node sizes. **Right:** Range of embedding time. *NN* is 20,000-spin Nearest-Neighbor [29], *KG* is 30,927-spin Kings-Graph [30], *CH* is 2048-spin Chimera.

2) In-depth comparison: Next, we look at a single graph problem (G22) to study ASA, OIM, and BRIM more indepth. In Fig. 10, we show the resulting cut value as different machines are given different annealing time. We add simulated annealing (SA) running on a workstation to this mix. Each dot's vertical coordinate represents the final cut value (higher is better) while its horizontal coordinate shows the annealing time. The range of annealing times are chosen such that the quality of the solution fits into roughly the same band. 10

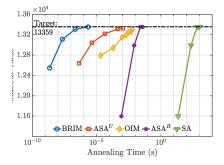


Fig. 10. Solution quality measured as max-cut value (higher is better) against total annealing time.  ${\rm ASA}^U$  means upper-bound, while  ${\rm ASA}^R$  means realistic.

An important point needs to be made about ASA. The 30,000 nominal spins can *not* accommodate G22's 2000 logical nodes. We thus create two variants instead, both better than the actual design. First, we create an idealized model where it supports all-to-all connection and can update all nodes sequentially, but all done in one annealing step. Note that this is intended as a very loose upper-bound of performance for the same *kind* of machine, and labeled as  $ASA^U$  in Fig. 10. Second, we take the reported improvement factor  $(2.6 \times 10^4 \text{ in reference [30]})$  over simulated annealing and draw a dotted line  $(ASA^R)$  to indicate what a more realistic performance roughly is. Even this is some idealization. The reported performance

improvement is over a much weaker graph that can be mapped onto the actual ASA. For G22, a much bigger system is needed and it will most definitely converge more slowly.

We can see the general trend that a longer annealing time provides a better answer for all machines. It is also clear that BRIM is orders of magnitude faster than other systems. However, rather than thinking about these curves as fixed and precise, it is helpful to think about them as a general shape that can move horizontally. Their position only represents the effect of current design parameters. And it is instructive to understand what it takes to move them to the left by, say, one order of magnitude. For ASA and SA, this is equivalent to making the machine 10x faster. We can imagine the challenge of this given the difficulty of scaling computation speed, not to mention memory access time. For the analog implementations, the question is more subtle. By reducing capacitance and/or inductance, the curve can shift to the left. The real question becomes the impact of noise and parasitics both to speed and solution quality. Such investigation is part of our future work.

TABLE II SOLUTION QUALITY OF THREE CONFIGURATIONS AND THEIR ANNEALING TIMES: BRIM  $(2.2\,\mu s)$ ; ASA  $(1\,ms)$ ; and SA  $(35\,min)$ . All solutions are best of 50 different runs shown as distance from best reported answers:  $\S$  [11],  $\ddag$  [38],  $\dag$  [39].

	Best	Ref.	BRIM	ASA	SA		Best	Ref.	BRIM	ASA	SA
G01	11624	§‡†	0	42	0	G16	3052	§‡†	6	27	0
G02	11620	§	0	38	0	G17	3047	‡†	6	24	1
G03	11622	§	0	1	0	G18	992	‡†	4	12	1
G04	11646	§‡†	0	23	0	G19	906	§	3	7	0
G05	11631	§‡†	0	27	0	G20	941	§‡†	0	0	0
G06	2178	§	0	2	0	G21	931	§	3	8	0
G07	2006	§	0	21	0	G43	6660	§	0	26	0
G08	2005	§	0	9	0	G44	6650	#	0	11	0
G09	2054	§	1	1	0	G45	6654	‡	1	5	0
G10	2000	§	1	1	1	G46	6649	§‡	2	16	0
G11	564	§‡†	6	8	0	G47	6657	‡	1	18	0
G12	556	§‡†	4	4	0	G51	3847	‡	3	21	0
G13	582	§‡†	4	6	0	G52	3851	‡	9	28	1
G14	3063	§	3	16	1	G53	3850	‡	10	20	1
G15	3050	‡†	4	7	0	G54	3851	‡	8	24	2

While Fig. 10 shows the comparison of one benchmark, the trend is very similar across all benchmarks. In Table II, we provide the detail of solution quality for benchmarks of Gset with no more than 1000 nodes. We compare BRIM, an ASA with unlimited nodes to map the problem, and SA. For BRIM, the distances range from 0 to 10, with a mean of 2.6 and a median of 1.5. In contrast, the distance for ASA (1ms) ranges from 0 to 42 with a mean of 15.1, and a median of 14. Additionally, the distance for SA (35 Minutes) ranges from 0 to 2 with a mean of 0.3, and a median of 0. SA obtains slightly better results, but note that it takes 9 orders of magnitudes longer than BRIM. Increasing annealing time for BRIM may improve quality but becomes extremely expensive to simulate.

3) Effects of annealing control: As already mentioned earlier, this paper highlights a different approach to Ising machine. The specific example used so far is but one design point in the space and much of the space remains to be explored in more detail. For example, annealing control plays a crucial role in obtaining good solutions. Fig. 11 compares

<sup>&</sup>lt;sup>10</sup>Note that each additional dot to the right takes 10x longer to obtain.

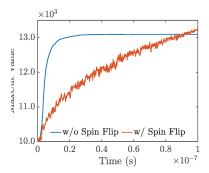


Fig. 11. Max-cut values as annealing evolves. Both trajectories start at the same initial condition.

the system with and without spin flips. It clearly shows the effect and advantage of exploring more than a single basin of attraction.

## E. Discrete-component prototype

We now showcase a limited prototype of BRIM using discrete electronics with operational amplifiers and passive components such as capacitors as well as resistors. The prototype is meant to demonstrate the key operation principle and therefore does not have some features such as programming interface and advanced annealing control to escape local minima.

The six nodes are implemented each with a capacitor and a ZIV diode made out of an op-amp. These are seen in the bottom right of Fig. 12. The coupling can be seen above the nodes in the figure. In this prototype, the nodes are directly connected with resistors without any buffer. Thus the coupling is bidirectional.

There is no state initialization in this prototype and each node starts at a voltage of 0. Noise intrinsic to the setup helps each node break out of the saddle point  $(v_i=0)$  and bifurcate to one of the stable states of  $\pm 1.15V$  determined by the specific set up of the ZIV diodes. Each run from power-up thus gives a different waveform on the scope and a potentially different solution. A waveform from simulation as well as the waveform on the scope can be see in Fig. 12. In this particular setup, the energy landscape is relatively simple and the prototype always reaches the ground state.

### V. Conclusions

Ising machines can be programmed to map an abstract problem and let physics naturally guide the dynamical system towards some kind of optimal state. This state can then be translated back to be a heuristic solution to a combinatorial optimization problem. The use of nature suggests the possibility of significant speed and/or efficiency. Consequently, exploration of Ising machines has been gaining attentions. Quantum annealers and optical coherent Ising machines are prominent examples of such machines which drew particular interests from the physics community. However, these Ising machines are generally bulky and energy intensive. Continued exploration is certainly a worthwhile endeavor for scientific

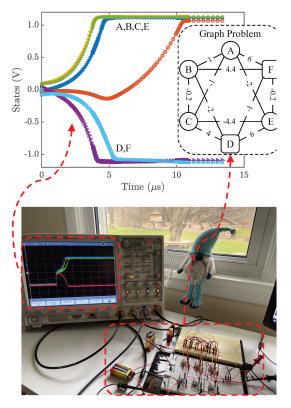


Fig. 12. Photo of the discrete-component prototype (with bi-directional coupling) connected to a 4-channel scope capturing the bifurcation of 6 nodes (A to F) after powering up (bottom). The graph representation of the problem being solved and a Spice-simulated waveform are also shown (top). In this particular simulation, nodes A, B, C, E converge to +1.15V while the other two nodes converge to -1.15V representing an opposite spin.

purposes. But for now, integrated circuit designs allow more immediate applications. We propose one such design we call BRIM that uses bistable nodes resistively coupled with programmable and variable strengths. The design is fully CMOS compatible for chip-scale implementations. Through our experimental analysis, we show that the machine is significantly smaller, faster, and more energy-efficient relative to some other designs. Compared to the room-sized Ising machines, it is about 4 orders of magnitude faster and consuming 7-8 orders of magnitude less energy. Compared to a recent chip-scale simulated annealing accelerator, it is also orders of magnitude better in speed, area, and power. While there is some degree of uncertainty with these statistics, we also envision continued improvement and optimizations. On the other hand, there are also many challenges facing machines like BRIM. Scalable architecture, area-efficient connection networks, and circuit scaling in the presence of Process-Voltage-Temperature (PVT) variations are but a few examples of the challenges.

## VI. ACKNOWLEDGEMENT

This work has been supported by a Pump Primer award from the University of Rochester and also in part by DARPA under Agreement No. HR00112090012.

#### REFERENCES

- [1] C. Johnson, D. H. Allen, J. Brown, S. Vanderwiel, R. Hoover, H. Achilles, C. Cher, G. A. May, H. Franke, J. Xenedis, and C. Basso, "A wire-speed power<sup>TM</sup> processor: 2.3GHz 45nm SOI with 16 cores and 64 threads," in 2010 IEEE International Solid-State Circuits Conference - (ISSCC), 2010, pp. 104–105.
- [2] B. Erbagci, N. E. C. Akkaya, C. Teegarden, and K. Mai, "A 275 Gbps AES encryption accelerator using ROM-based S-boxes in 65nm," in 2015 IEEE Custom Integrated Circuits Conference (CICC), 2015, pp. 1–4.
- [3] Y. Chen, T. Luo, S. Liu, S. Zhang, L. He, J. Wang, L. Li, T. Chen, Z. Xu, N. Sun, and O. Temam, "DaDianNao: A Machine-Learning Supercomputer," in 2014 47th Annual IEEE/ACM International Symposium on Microarchitecture, Dec 2014, pp. 609–622.
- [4] N. P. Jouppi, C. Young, N. Patil, D. Patterson, G. Agrawal, R. Bajwa, S. Bates, S. Bhatia, N. Boden, A. Borchers, R. Boyle, P.-I. Cantin, C. Chao, C. Clark, J. Coriell, M. Daley, M. Dau, J. Dean, B. Gelb, T. V. Ghaemmaghami, R. Gottipati, W. Gulland, R. Hagmann, C. R. Ho, D. Hogberg, J. Hu, R. Hundt, D. Hurt, J. Ibarz, A. Jaffey, A. Jaworski, A. Kaplan, H. Khaitan, D. Killebrew, A. Koch, N. Kumar, S. Lacy, J. Laudon, J. Law, D. Le, C. Leary, Z. Liu, K. Lucke, A. Lundin, G. MacKean, A. Maggiore, M. Mahony, K. Miller, R. Nagarajan, R. Narayanaswami, R. Ni, K. Nix, T. Norrie, M. Omernick, N. Penukonda, A. Phelps, J. Ross, M. Ross, A. Salek, E. Samadiani, C. Severn, G. Sizikov, M. Snelham, J. Souter, D. Steinberg, A. Swing, M. Tan, G. Thorson, B. Tian, H. Toma, E. Tuttle, V. Vasudevan, R. Walter, W. Wang, E. Wilcox, and D. H. Yoon, "In-Datacenter Performance Analysis of a Tensor Processing Unit," SIGARCH Comput. Archit. News, vol. 45, no. 2, p. 1–12, Jun. 2017.
- [5] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, "Characterizing quantum supremacy in near-term devices," *Nature Physics*, vol. 14, no. 6, pp. 595–600, 2018.
- [6] E. Pednault, J. A. Gunnels, G. Nannicini, L. Horesh, and R. Wisnieff, "Leveraging Secondary Storage to Simulate Deep 54-qubit Sycamore Circuits," arXiv: Quantum Physics, 2019.
- [7] P. I. Bunyk, E. M. Hoskinson, M. W. Johnson, E. Tolkacheva, F. Altomare, A. J. Berkley, R. Harris, J. P. Hilton, T. Lanting, A. J. Przybysz, and J. Whittaker, "Architectural Considerations in the Design of a Superconducting Quantum Annealing Processor," *IEEE Transactions on Applied Superconductivity*, vol. 24, no. 4, pp. 1–10, 2014.
- [8] H. Yu, Y. Huang, and B. Wu, "Exact Equivalence between Quantum Adiabatic Algorithm and Quantum Circuit Algorithm," *Chinese Physics Letters*, vol. 35, no. 11, p. 110303, Oct 2018.
- [9] W. van Dam, M. Mosca, and U. V. Vazirani, "How Powerful is Adiabatic Quantum Computation?" *Proceedings 2001 IEEE International Conference on Cluster Computing*, pp. 279–287, 2001.
   [10] T. Inagaki, Y. Haribara, K. Igarashi, T. Sonobe, S. Tamate, T. Honjo,
- [10] T. Inagaki, Y. Haribara, K. Igarashi, T. Sonobe, S. Tamate, T. Honjo, A. Marandi, P. L. McMahon, T. Umeki, K. Enbutsu, O. Tadanaga, H. Takenouchi, K. Aihara, K.-i. Kawarabayashi, K. Inoue, S. Utsunomiya, and H. Takesue, "A coherent ising machine for 2000-node optimization problems," *Science*, vol. 354, no. 6312, pp. 603–606, 2016.
- [11] T. Wang and J. Roychowdhury, "OIM: Oscillator-Based Ising Machines for Solving Combinatorial Optimisation Problems," 2019.
- [12] C. Roques-Carmes, Y. Shen, C. Zanoci, M. Prabhu, F. Atieh, L. Jing, T. Dubček, C. Mao, M. R. Johnson, V. Čeperić, J. D. Joannopoulos, D. Englund, and M. Soljačić, "Heuristic recurrent algorithms for photonic Ising machines," *Nature Communications*, vol. 11, no. 1, p. 249, 2020
- [13] A. Messiah, Quantum Mechanics. North Holland, 1976, vol. 2.
- [14] R. M. Karp, *Reducibility among Combinatorial Problems*. Boston, MA: Springer US, 1972, pp. 85–103.
- [15] A. Lucas, "Ising formulations of many np problems," Frontiers in Physics, vol. 2, p. 5, February 2014.
- [16] J. Cai, W. G. Macready, and A. Roy, "A practical heuristic for finding graph minors," *ArXiv*, vol. abs/1406.2741, 2014.
  [17] "The D-Wave 2000Q Quantum Computer." [On-
- [17] "The D-Wave 2000Q Quantum Computer." [Online]. Available: https://www.dwavesys.com/sites/default/files/D-Wave%202000Q%20Tech%20Collateral\_0117F.pdf
- [18] Y. Yamamoto, K. Aihara, T. Leleu, K.-i. Kawarabayashi, S. Kako, M. Fejer, K. Inoue, and H. Takesue, "Coherent ising machines optical neural networks operating at the quantum limit," npj Quantum Information, vol. 3, no. 1, p. 49, 2017.

- [19] P. L. McMahon, A. Marandi, Y. Haribara, R. Hamerly, C. Langrock, S. Tamate, T. Inagaki, H. Takesue, S. Utsunomiya, K. Aihara, R. L. Byer, M. M. Fejer, H. Mabuchi, and Y. Yamamoto, "A fully programmable 100-spin coherent ising machine with all-to-all connections," *Science*, vol. 354, no. 6312, pp. 614–617, 2016.
- [20] K. Takata, A. Marandi, R. Hamerly, Y. Haribara, D. Maruo, S. Tamate, H. Sakaguchi, S. Utsunomiya, and Y. Yamamoto, "A 16-bit coherent ising machine for one-dimensional ring and cubic graph problems," *Scientific Reports*, vol. 6, no. 1, p. 34089, 2016.
- [21] F. Böhm, G. Verschaffelt, and G. Van der Sande, "A poor man's coherent ising machine based on opto-electronic feedback systems for solving optimization problems," *Nature Communications*, vol. 10, no. 1, p. 3538, 2019. [Online]. Available: https://doi.org/10.1038/s41467-019-11484-3
- [22] R. Hamerly, T. Inagaki, P. L. McMahon, D. Venturelli, A. Marandi, T. Onodera, E. Ng, C. Langrock, K. Inaba, T. Honjo, K. Enbutsu, T. Umeki, R. Kasahara, S. Utsunomiya, S. Kako, K. ichi Kawarabayashi, R. L. Byer, M. M. Fejer, H. Mabuchi, D. Englund, E. G. Rieffel, H. Takesue, and Y. Yamamoto, "Experimental investigation of performance differences between coherent ising machines and a quantum annealer," in Science advances. 2019.
- [23] H. Takesue, T. Inagaki, K. Inaba, T. Ikuta, and T. Honjo, "Large-scale coherent ising machine," *Journal of the Physical Society of Japan*, vol. 88, no. 6, p. 061014, 2019.
- [24] M. G. Rosenblum, A. S. Pikovsky, and J. Kurths, "Phase synchronization of chaotic oscillators," *Physical review letters*, vol. 76, no. 11, p. 1804, March 1996.
- [25] M. Vidyasagar, Nonlinear System Analysis. SIAM, 1978.
- [26] T. Wang, L. Wu, and J. Roychowdhury, "New computational results and hardware prototypes for oscillator-based ising machines," in *Proceedings* of the 56th Annual Design Automation Conference 2019, ser. DAC '19. New York, NY, USA: Association for Computing Machinery, 2019.
- [27] J. Chou, S. Bramhavar, S. Ghosh, and W. Herzog, "Analog coupled oscillator based weighted ising machine," *Scientific Reports*, vol. 9, no. 1, p. 14786, 2019.
- [28] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," science, vol. 220, no. 4598, pp. 671–680, 1983.
- [29] M. Yamaoka, C. Yoshimura, M. Hayashi, T. Okuyama, H. Aoki, and H. Mizuno, "24.3 20k-spin ising chip for combinational optimization problem with cmos annealing," in 2015 IEEE International Solid-State Circuits Conference - (ISSCC) Digest of Technical Papers, Feb 2015, pp. 1–3.
- [30] T. Takemoto, M. Hayashi, C. Yoshimura, and M. Yamaoka, "2.6 a 2 by 30k-spin multichip scalable annealing processor based on a processingin-memory approach for solving large-scale combinatorial optimization problems," in *IEEE International Solid- State Circuits Conference*, February 2019.
- [31] Z. Varty, "Simulated annealing overview," 2017. [Online]. Available: https://www.lancaster.ac.uk/pg/varty/RTOne.pdf
- [32] R. Harris, M. W. Johnson, T. Lanting, A. J. Berkley, J. Johansson, P. Bunyk, E. Tolkacheva, E. Ladizinsky, N. Ladizinsky, T. Oh, F. Cioata, I. Perminov, P. Spear, C. Enderud, C. Rich, S. Uchaikin, M. C. Thom, E. M. Chapple, J. Wang, B. Wilson, M. H. S. Amin, N. Dickson, K. Karimi, B. Macready, C. J. S. Truncik, and G. Rose, "Experimental investigation of an eight-qubit unit cell in a superconducting optimization processor," *Phys. Rev. B*, vol. 82, p. 024511, Jul 2010.
- [33] "dwave-system." [Online]. Available: https://docs.ocean.dwavesys.com/en/stable
- [34] "Breakdown of qpu access time." [Online]. Available: https://docs.dwavesys.com/docs/latest/c\_timing\_2.html
- [35] T. G. J. Myklebust, "Solving maximum cut problems by simulated annealing," 2015.
- [36] Y. Ye. G-set benchmark. [Online]. Available: https://web.stanford.edu/~yyye/yyye/Gset
- [37] C. Helmberg and F. Rendl, "A spectral bundle method for semidefinite programming," SIAM J. on Optimization, vol. 10, no. 3, p. 673–696, Jul. 1999.
- [38] Q. Wu and J.-K. Hao, "Memetic search for the max-bisection problem," Computers & Operations Research, vol. 40, no. 1, pp. 166 – 179, 2013.
- [39] F. Ma, J.-K. Hao, and Y. Wang, "An effective iterated tabu search for the maximum bisection problem," *Computers & Operations Research*, vol. 81, pp. 78 – 89, 2017.
- [40] J. Cai, W. G. Macready, and A. Roy, "A practical heuristic for finding graph minors," 2014.