



A 1,968-node coupled ring oscillator circuit for combinatorial optimization problem solving

William Moy , Ibrahim Ahmed, Po-wei Chiu, John Moy, Sachin S. Sapatnekar and Chris H. Kim  

Computational architectures that are optimized to solve non-deterministic polynomial-time hard or complete problems are of use in the development of machine learning, logistical planning and pathfinding. A range of quantum-, optical- and spintronic-based approaches have been explored for solving such combinatorial optimization problems, but they remain complicated to build and to scale. Here we report a scalable ring-oscillator-based integrated circuit for optimization problem solving. Our 1,968-node King's graph ring oscillator array has five levels of coupling strengths and can achieve up to 95% accuracy for randomly generated combinatorial optimization problems. The measured average power consumption of the Ising chip is 0.042 W and it takes less than 50 oscillation cycles to resolve to the ground state. Our device is resilient to environmental and variation effects. By using a multi-phase phase measurement circuit, we also capture the true phase behaviour within a coupled-oscillator integrated circuit.

Combinatorial optimization problems (COPs) occur in a variety of fields including logistics, resource allocation, communication network design, finance, drug discovery and transportation systems^{1,2}. Digital computers based on a von Neumann architecture often struggle to solve COPs in a timely and efficient manner due to the need for exponential computational power^{3–5}. Several COPs can be solved using a formulation known as the quadratic unconstrained binary optimization (QUBO) problem. The QUBO formulation covers problems ranging from max-cut to graph colouring and has parallels in machine learning and neuro-morphic computing^{6,7}. One way software algorithms solve a QUBO problem is by modelling it as the Ising spin glass model and updating the spin states in an iterative and probabilistic manner until a satisfactory ground state is found. The Ising model describes the behavior of ferromagnets and works on minimizing the energy level of coupled spins through a Hamiltonian cost function^{8,9}. Within the Ising model, s_i represents the spin value of node i with $s_i \in \{+1, -1\}$ and J_{ij} represents the coupling between spins s_i and s_j . Visualization of the Ising model is frequently represented by a mathematical graph structure with spins represented by nodes and coupling represented by edges (Fig. 1a)^{10–13}. The solution of a QUBO problem is the combination of spins that minimizes the Ising Hamiltonian cost function:

$$H(s) = \sum_{\langle ij \rangle} J_{ij} s_i s_j.$$

Although directly solving an Ising model can be inefficient with a digital computer, a physical representation of the Ising model can overcome these inefficiencies by simply letting a network of coupled oscillator circuits resolve to a low Hamiltonian state (Fig. 1a–e). The optimization problem, which can be defined by vertices (that is, spins) and edge weights (that is, coupling) of a graph, is converted to a locally connected architecture that is compatible with the Ising hardware. Once the weights of the Ising hardware are programmed, the oscillator network naturally resolves to the ground state with the oscillator phases representing the spin states. Simple error-correction and post-processing algorithms may be used to convert the phase information to the final solution of the original COP.

Alternative approaches that aim to solve COPs more efficiently include application-specific integrated circuit (ASIC) chips that directly implement the software algorithm in digital complementary metal–oxide–semiconductor (CMOS) hardware, as well as more unconventional devices that facilitate the coupling operation^{14,15}. However, ASICs suffer from the delay and power limitations of digital CMOS when solving exponentially large problems. Other more unconventional approaches, including quantum, spintronic and photonic methods, allow for the direct representation of the Ising model but depend on unproven technologies that are not mass manufacturable today. Quantum computers require near-absolute-zero temperature conditions requiring kilowatts of power per chip¹⁶. Spintronic devices have only been demonstrated at a smaller scale (less than ten devices) on a breadboard¹⁷. Optical Ising machines require kilometre-long optical fibre and multiple high-performance field-programmable gate array (FPGA) chips, limiting density and power efficiency¹⁸.

In this Article, we report an Ising chip based on a graph ring oscillator (ROSC) array with transmission-gate couplings. The approach allows conventional CMOS technology to be used and overcomes the typical von Neumann issues associated with digital ASICs. Our design shows resilience to process, voltage and temperature variations, and it can achieve an accuracy of up to 95% for random COPs. We also report realistic effects occurring in an integrated system that have previously not been captured within conventional simulation models.

ROSC Ising machine design

Figure 2a shows the circuit implementation of two ROSCs coupled via programmable transmission gates, or parallelly connected n-type and p-type transistors. Each transmission gate can be individually turned on or off to enable five coupling configurations. Local static random access memory (SRAM) cells store the coupling-weight information. The gate width and length of the coupling transistors are determined to provide adequate resistive values. Eight transmission gates of the same size are used to create programmable weights from -2 to $+2$. Coupling the ROSCs on the same nodes allows for a positive coupling. Taking advantage of the 180° phase difference

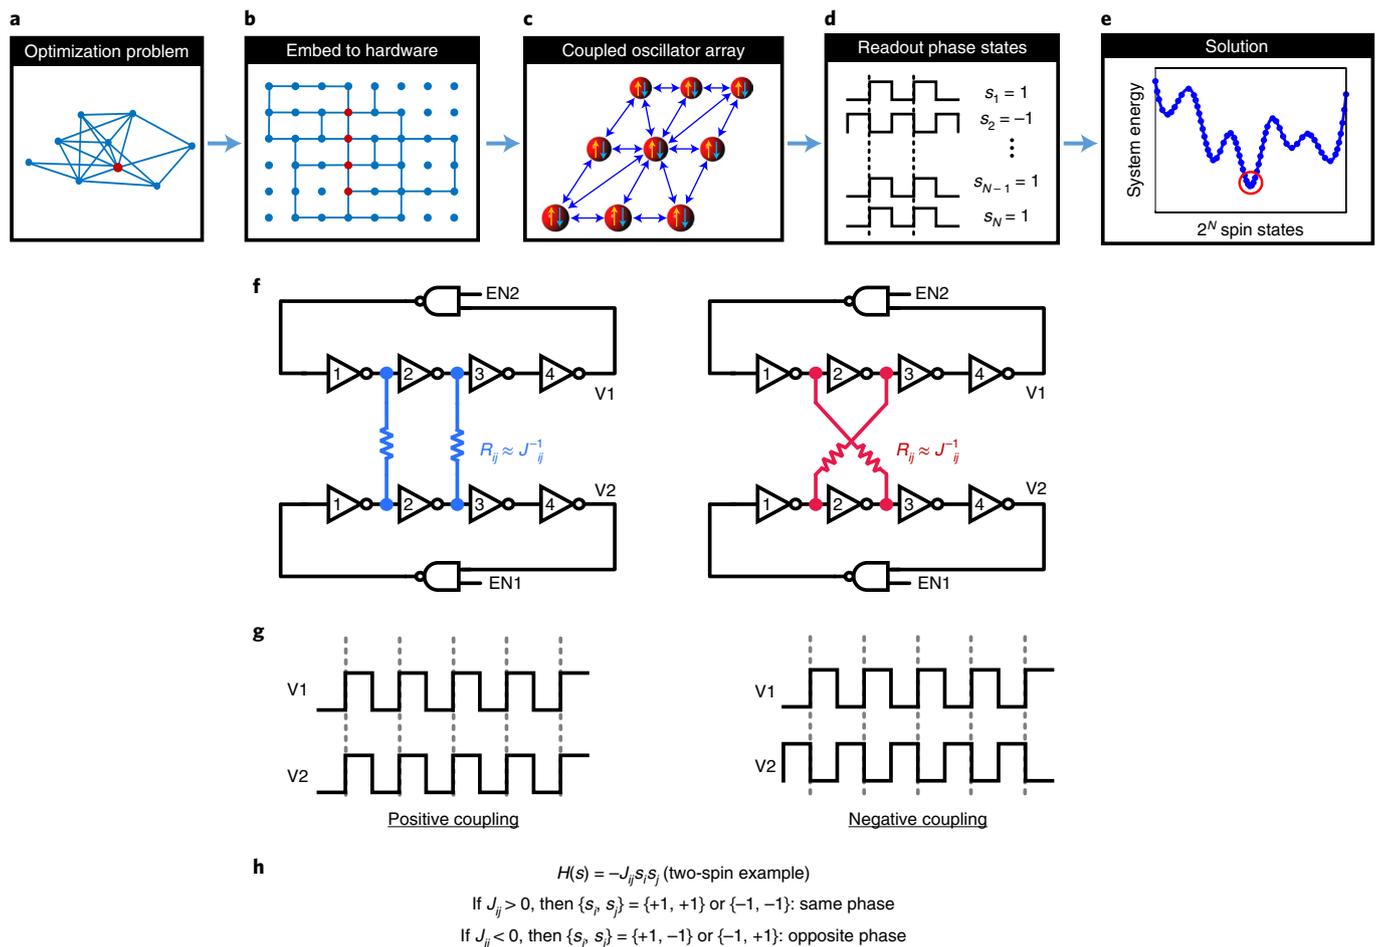


Fig. 1 | Computation steps and operating principle of the coupled-oscillator-based Ising machine. **a**, Optimization problem represented by a graph. **b**, Mapping the graph to a locally connected hardware architecture. **c**, Coupled oscillator hardware. **d**, Individual phase readout. **e**, Location of the coupled-oscillator-based Ising model solution within the Hamiltonian solution space. **f**, Two-node positive and negative coupled oscillator schemes. Same-node resistive connection provides a positive coupling. Cross-coupled resistive connections allow for a negative coupling. **g**, In-phase and out-of-phase ROSC spins due to the coupling type. Positive coupling aligns the phases together. Negative coupling creates out-of-phase spins. **h**, Mathematical representation of a two-node coupled ROSC system. The Hamiltonian is minimized when opposite phases are negatively coupled and same phases are positively coupled.

between the adjacent nodes, a cross-coupled configuration can be used to generate a negative coupling weight.

A 41×48 network of ROSCs was formed into a single Ising model utilizing a King's graph architecture where each oscillator is coupled to two vertical, two horizontal and four diagonal neighbours (Fig. 1c), recreating the movements of the king chess piece. The full-chip layout of the 2.1 mm^2 chip is shown in Fig. 2b. The chip was fabricated in a standard 65 nm low-power CMOS technology. To facilitate the array design, a unit tile was created and replicated with each tile forming a single unit of the 41×48 array. The tile consisted of a single ROSC, four coupling blocks (eight transmission gates per block), 20 SRAM cells, a phase sampler and storage circuit and a sub-harmonic injection locking (SHIL) circuit (Fig. 2b, right). When placed as a rectangular mosaic, the adjacent unit tiles create the eight couplings needed for the King's graph architecture. The majority of the unit-tile area is consumed by the coupling transmission gates and SRAM cells, with the ROSC circuit occupying an area less than 5%.

The ROSC block of the unit tile consists of a nine-stage ROSC (Fig. 3a) with a nominal frequency of 1.0 GHz (measured) at 1.2 V and room temperature. The global enable signal can simultaneously turn on or off all the ROSCs. The local enable signal allows for

individual ROSCs to be turned on and off. Four sets of transmission gates make up the coupling block (Fig. 3b). The coupling block in each unit tile connects the ROSC with the four nearest-neighbour ROSCs. The same set of ROSC nodes (namely, $A\langle 0 \rangle$ and $A\langle 1 \rangle$) are used for all the couplings within the coupling block, regardless of the directionality of the couplings. This is done to simplify the design and minimize the size of the ROSC. Transmission gates are controlled as pairs to provide a symmetric coupling. The local enable signal of each pair of transmission gates controls the coupling weight. A global enable signal is available to turn on or off all the gates at once. The coupling circuit contains an analogue voltage knob (omitted in the figure for simplicity) to manipulate the coupling strength relative to the core circuit voltage, which can be used to implement different annealing schedules. A SHIL block is employed to force the analogue phases of the coupled ROSC array into binary phases. The SHIL block works by perturbing the ROSCs at twice the base ROSC frequency¹⁹. A voltage-controlled oscillator is used to create the SHIL clock. The frequency of the voltage-controlled oscillator is controlled using an off-chip voltage supply. The ROSC is perturbed utilizing a set of pull-down n-type metal-oxide-semiconductor transistors. The duty cycle of the SHIL clock and the pull-down strength can be independently controlled

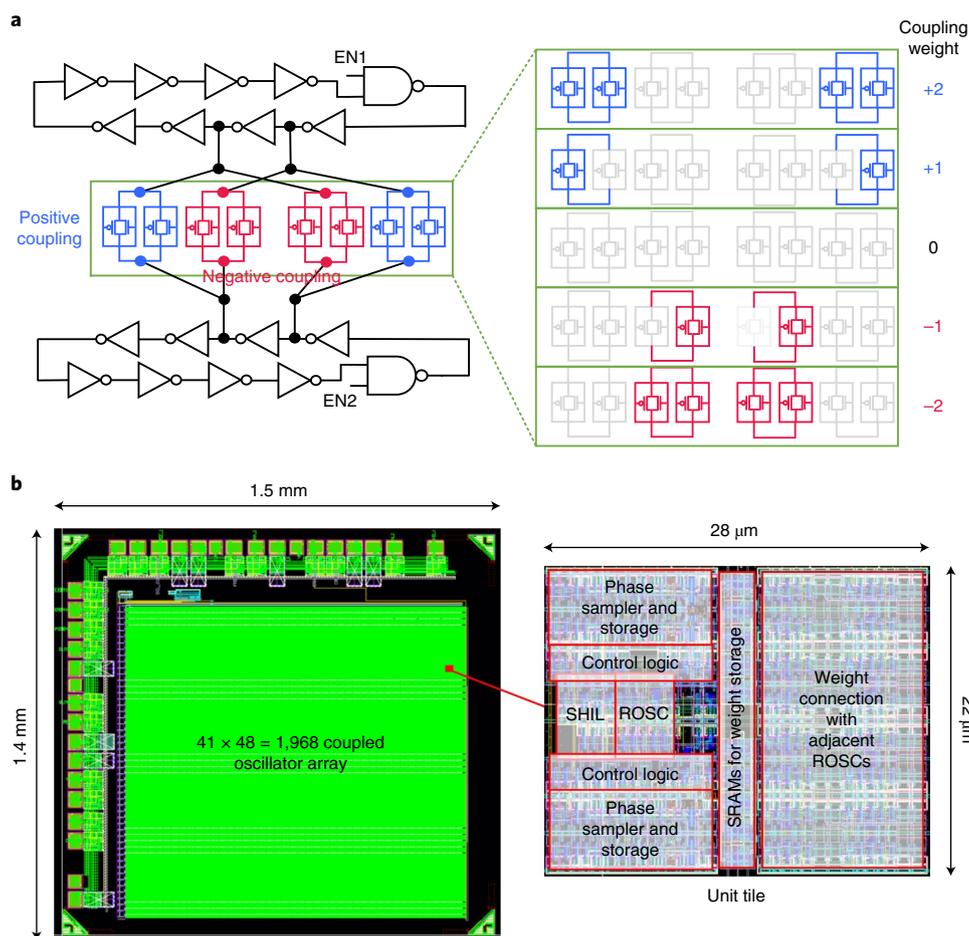


Fig. 2 | Schematic of a coupled oscillator and 1,968-node Ising chip layout. **a**, Five-level transmission-gate coupling between two ROSC nodes. A positive coupling is achieved through like-node connections. A negative coupling is achieved through phase-shifted connections. A bank of programmable switches enables five distinct coupling strengths. **b**, Full-chip and unit-tile layouts of the Ising chip. Due to the supporting circuits, the ROSC occupies less than 5% of the total chip area.

to vary the SHIL perturbation strength. The SRAM was used in lieu of flip-flops to reduce the area of the design. Out of the 20 SRAM cells, 16 are used to control the coupling weights, three are used to control the SHIL signal and one is used to enable the ROSC.

A four-stage delayed sampling circuit is used to capture the spin information of each ROSC. The master and slave stages of the sampling flip-flops are clocked using separate signal clocks, namely, CLKA and CLKB (Fig. 3e). CLKA controls the master stage and samples the delayed ROSC signal. CLKB controls the slave stage and stores the sampled data. Pseudo-analogue phase measurements of the ROSC are taken by sampling the ROSC output delayed by 0-, 2-, 4- and 6-inverter-stage delays. The flip-flops act as both sampling and serial readout circuits. The scan chain is connected to adjacent unit tiles for serial readout. All the 1,968 ROSC signals are simultaneously sampled by the same sampling clock edge.

Phase measurements and post-processing methodology

A portion of the ROSC period is captured by the phase-sampling circuit and the analogue phase information is decoded (Fig. 4a). The raw phase maps (Fig. 4b) reveal the non-idealities (namely, noise, process variation and nonlinearity) of a large-scale coupled oscillator array implemented in an integrated circuit chip. The three measured phase maps correspond to uniform, chequerboard and random problems, respectively. The uniform problem sets all the available couplings to a uniform positive weight (+2), creating a uniform final spin pattern. The chequerboard problem utilizes only

horizontal and vertical negative couplings (−2) to create a chequerboard spin pattern. The random problem was generated using all the available couplings weights (−2, −1, 0, +1 and +2). An equal probability (that is, 20%) was given to each of the five coupling weights during the random weight selection.

The phase drift is the most noticeable in the uniform and chequerboard problems. The uniform graph slowly transitions from a positive to negative phase on the horizontal axis of the phase map. Although the absolute phase between the nodes differs between the right and left sides of the array, the nearest-neighbour couplings maintain the same or almost the same phase. The chequerboard problem has a spiralling pattern in which the absolute phase of the design slowly drifts, whereas the nearest-neighbour phases maintain the expected 180° phase differences fairly consistently throughout the map. The random problem experiences the same phase drift, but simply lacks a perceptible pattern.

As the relative phase difference between the adjacent nodes is nearly always correct, a lightweight post-processing algorithm (Fig. 4c) is used to eliminate the global phase drift and extract the correct phase information. The post-processing algorithm works by condensing multiple samples into a single solution as follows. First, multiple chip samples are taken and recorded for a given problem. The average phase difference for each coupling is then found and ranked based on a confidence interval. A coupling is ranked with more confidence based on how close the average phase difference is to either 0° or 180°. By ranking the couplings in this way, spin

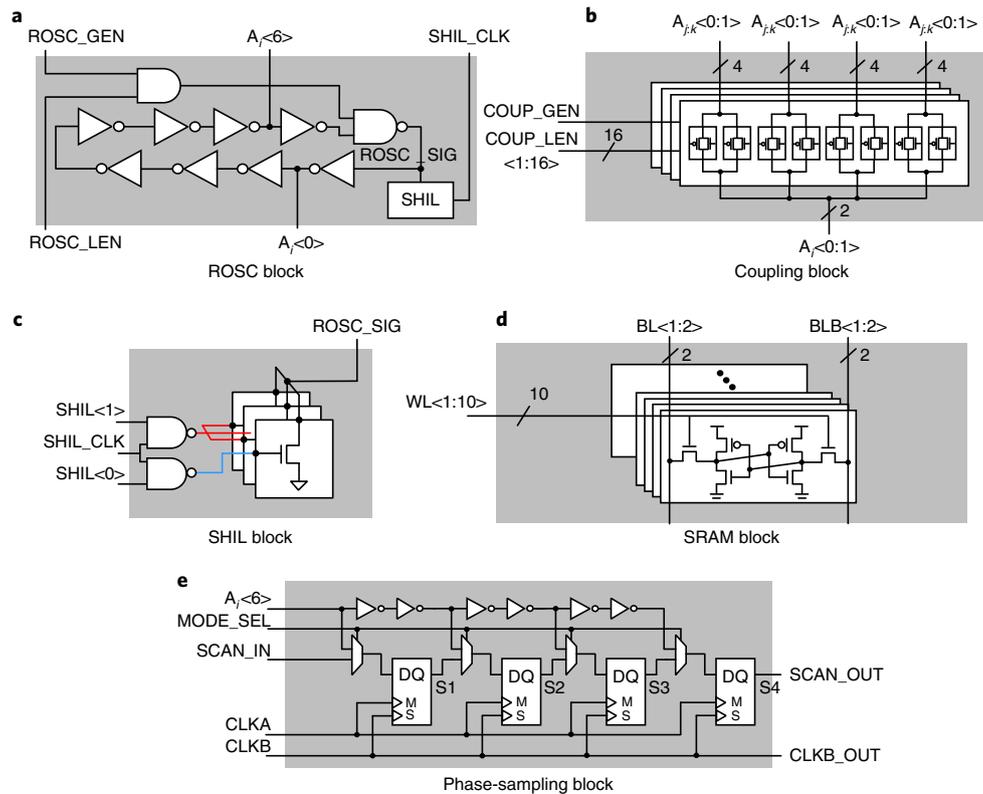


Fig. 3 | Unit-tile oscillator and support circuits. **a**, ROSC circuit with local and global ROSC-enable signals. **b**, Transmission-gate coupling block circuit with local and global coupling-enable signals. **c**, SHIL circuit for binarizing the ROSC phases. **d**, SRAM memory that is used to store the weight and configuration information for the unit tile. **e**, Phase-sampling circuit that measures and stores the ROSC phase information.

couplings that contain noise and other imperfections are ignored, leading to a more accurate final solution. Once all the spins have been determined, the algorithm ends, potentially ignoring any couplings that could have been detrimental to the overall solution. The lightweight post-processing solution creates highly accurate final solutions (100% for the uniform and checkerboard graphs) based on only the most confident couplings (Extended Data Fig. 1).

Statistical data from Ising chip

The accuracy of the fabricated chip was compared with a Tabu search software solution^{7,20,21}. The Tabu solution is treated as the ground truth with regard to normalized accuracy results. The reported design was programmed with numerous random problems that utilized all the available couplings and coupling weights. All the measured results reported were generated using the maximally sized graph of 41×48 (that is, 1,968) nodes. By randomizing the problems, contentious problems of constant movement are generated, challenging the reported Ising machine during solving. A simple example of a contentious problem would be an all-to-all negatively coupled graph. The reported design was able to achieve 90–95% accuracy compared with the ground-truth solution (Fig. 5a, red distribution). To illustrate the high accuracy of solutions from our chip, we also plotted the Hamiltonian values based on 100,000 random guesses. That is, we randomly assigned +1 or -1 to the 1,968 spins for a given graph, computed each Hamiltonian value and plotted the accuracy distribution in blue. The best random guesses rarely achieved greater than 10% accuracy. A low accuracy is expected from the random-guess approach due to the solution space size of $2^{1,968} \approx 10^{592}$ of the 1,968-node graph.

The time to solution is defined as the number of ROSC cycles needed for the phases to resolve to the good ground state. The global coupling control and the sampling clock edge are external and

independent signals that are used to measure the time to solution (Fig. 5b). The time between the external sampling clock edge and the global coupling-control clock edge was varied until a steady-state accuracy was found. Before the global coupling control is enabled, the reported design is not solving the programmed problem. Additionally, although the couplings are off, the ROSCs are still on, allowing for the spins to randomly drift apart. Once the couplings are enabled, the chip begins to solve the programmed problem. The sampling edge was then activated at a consistent time interval after the couplings were enabled to generate the time-to-solution graph. Multiple samples were needed at each time interval to perform the post-processing described previously. Checkerboard, uniform and random problems were used for the time-to-solution testing. These problems were chosen to illustrate how the problem difficulty affects the time-to-solution results. Although the time to solution can slightly vary between problems, the Ising chip was able to find the ground state within 50 ROSC cycles. Due to the solution-space landscape, the random problem is far harder to resolve to the ground state during testing than the uniform or checkerboard problems as seen by the 100% normalized accuracy of the latter two problems compared with the approximately 90% normalized accuracy of the random problem. The shape of the time-to-solution curve varies with the problem type. In Fig. 5b, it is evident that unlike the uniform problem, the checkerboard and random problems show hesitation before rapidly reaching a steady-state solution.

Ising problems with random weights were tested on multiple chips to study if chip-to-chip variation caused by manufacturing affects the overall accuracy of the reported design (Fig. 5c). The random weights applied to the chips were generated in the same way as the time-to-solution testing. The accuracy range between the chips varies very little (~1% difference in the maximum accuracy), indicating that the reported design is highly resistant to

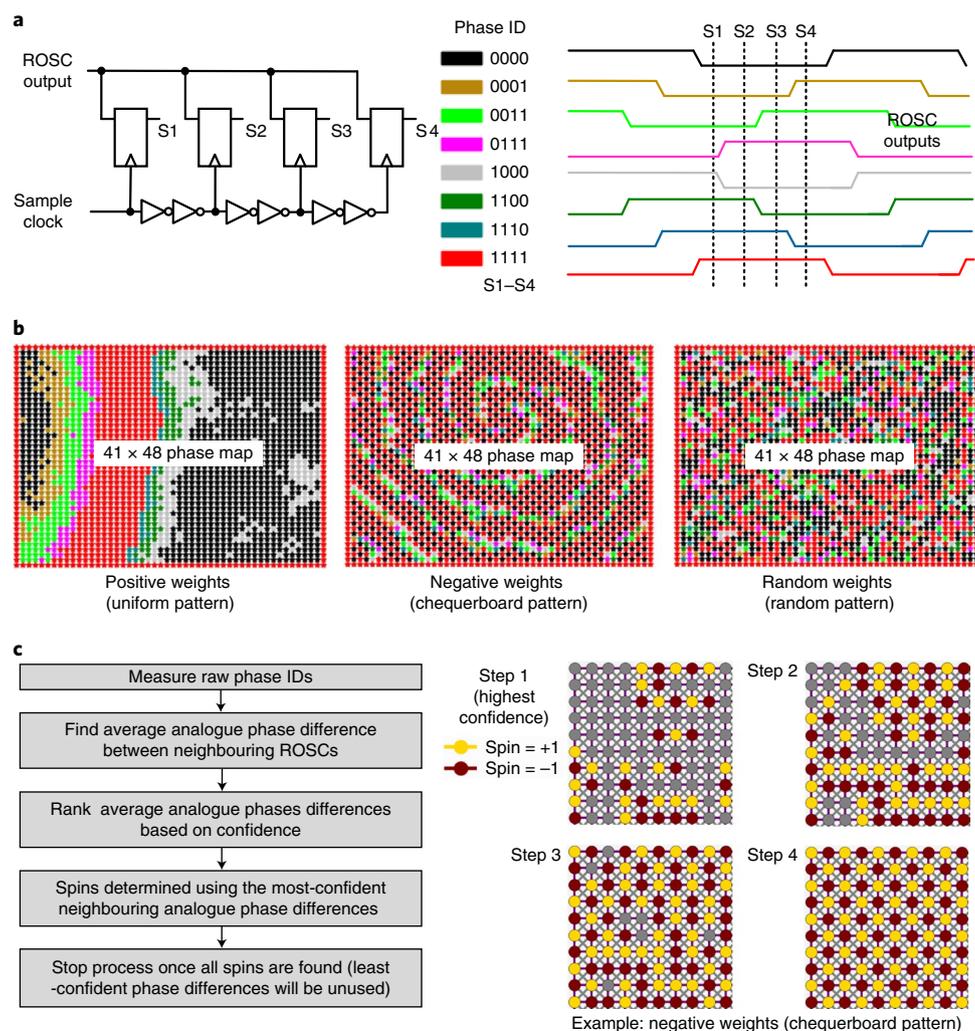


Fig. 4 | Phase-sampling circuit, measured raw phase maps and post-processing steps. **a**, Sampling circuit with phase IDs. The sampling circuit captures a portion of the ROSC phase and produces a four-bit phase ID for each ROSC. **b**, Measured phase maps from 1,968 coupled oscillators for three distinct graphs: positive weight, negative weight and random weight. Non-idealities of a real integrated circuit chip are revealed. **c**, Post-processing algorithm condenses multiple chip samples into a final spin solution using a lightweight nearest-neighbour phase-averaging scheme. Steps 1–4 show the post-processing algorithm in action on a 10×10 portion of a negative-weight chequerboard pattern.

chip-to-chip variation. Other Ising machines based on exotic devices (spintronics and optics) may not have the same degree of manufacturing repeatability.

CMOS transistor characteristics are, by nature, nonlinear with respect to the supply voltage. Due to this, the interval between the coupling levels can vary with the voltage of the core circuitry. Additionally, varying the core voltage of the chip allows for the noise level of the reported design to be manipulated. Although the interval between the coupling levels and noise level within the chip cannot be directly measured, by sweeping the supply voltage of the Ising chip, the optimal voltage level that balances the above two design aspects can be found (Fig. 5d). A core voltage of 1.8 V gives marginally (<2%) better average accuracy results than the other voltages. The overall accuracy difference, however, is minimal between the voltages, showing resilience to supply-voltage variation.

We repeated the accuracy testing at three different chip temperatures. The phase interaction dynamics were expected to change as both device noise and transistor parameters (for example, threshold voltage and mobility) are impacted by temperature. A previous study²² reported a lower accuracy at higher temperatures possibly due to the increased thermal noise that may have prevented the

phases from stabilizing. In this work, however, we observed a weak dependence on temperature, with the overall accuracy of the design practically staying the same (<1%) at -40 , 25 and 90°C (Fig. 5e). Several factors make the temperature dependence different in this work. First, we sample the phases of all the oscillators using one clock edge, whereas the previous design could only measure the phase difference between two nearest-neighbour oscillators at a time. The improved phase readout scheme prevents measurement error due to phase drifts, which could make the solution quality less sensitive to temperature changes. Second, the transmission-gate-based coupling circuits are weaker in the reported design to support five coupling levels. The lower transistor threshold voltage at higher temperatures could have increased the coupling strength, compensating for the increased thermal noise.

Our chip includes two global analogue biases for separately tuning the p-type and n-type transistors in the coupling block. These external control voltages determine the gate-voltage levels of the transmission-gate-based coupling transistors (Fig. 3b). Using the two global biases, ‘external annealing’ was applied, wherein the coupling strengths of all the weights in the chip were gradually increased in the hopes of finding a better ground state. Contrary

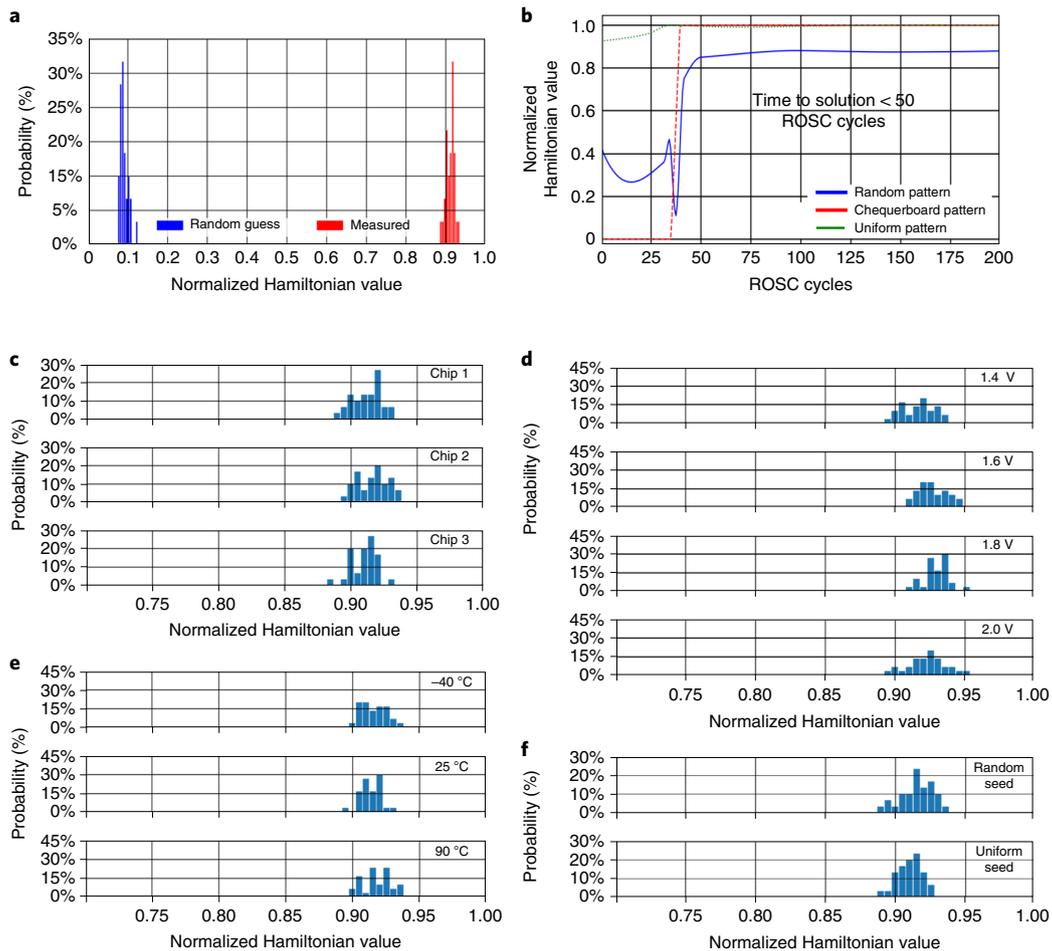


Fig. 5 | Measured results from the 1,968 oscillator chips. a, Measured Ising Hamiltonian results normalized to the best software solution and compared with 100,000 random guesses. **b**, Time to solution for different Ising problems, showing a relaxation time of less than 50 ROSC cycles. **c–f**, Accuracy fluctuation across different chips, supply voltages, chip temperatures and initial spin states. Each distribution consists of 30 problems with five solutions per problem (namely, 150 samples). The uniform seed initialized all the spins to the same phase before the problem is allowed to be solved by the reported chip. The random initial seed is uncontrolled and varies each time the chip solves the Ising problem.

to conventional wisdom, improvement in the average and best (5 problems \times 40 iterations = 80 solutions) Hamiltonian values were only 1.00% and 0.04%, respectively, which suggests that external annealing may not provide the accuracy improvements predicted by previous simulation-based studies and that the inherent annealing already occurring inside the chip may be sufficient for finding a competitive solution. Annealing has been shown to be useful for a diverse set of architectures and Ising implementations. However, with respect to the King's graph architecture implemented with the reported ROSC Ising machine, the additional effort of annealing was not warranted. Further research into the effectiveness of annealing as it applies to non-deterministic Ising machines with fixed architectures should be done.

We also report how the solution accuracy is affected by the initial phases of the oscillators. Our hardware allows us to initialize the phases to either a random seed or a uniform seed. A random seed was created by allowing the ROSCs to drift apart as the couplings are disabled, before the global coupling-enable signal is fired. A uniform seed was created by simultaneously disabling and enabling the global ROSC control, causing all the spins to start in phase. By comparing the two distributions in Fig. 5f, the random initial seed marginally improves the overall accuracy of the coupled oscillator array. This can be explained by the fact that a random initial seed may allow more of the solution space to be explored, leading to

better overall results. However, the actual accuracy improvement is relatively small (1–2%) on average. Hence, even without a randomized initial seed, the reported device appears to explore the solution space sufficiently and find a good solution.

Comparison with earlier work

In Table 1, we compared our ROSC integrated-circuit Ising design to previous designs that use emerging technologies, digital logic, superconductors and oscillators^{18,19,22–27}. The earlier work can be split into categories of deterministic and non-deterministic Ising machines. The digital and software-based approaches of refs. ^{23,24} and ref. ¹⁹ are deterministic and seek to simulate the Ising model instead of directly implementing the Ising model through spins and couplings. This work and refs. ^{22,25–27} directly implement the Ising model and can be considered non-deterministic.

The most relevant comparisons given the size, power and performance metrics of the reported device would be with refs. ^{23,25,22}. Although ref. ¹⁹, ref. ²⁶ and ref. ²⁷ are valuable contributors to the body of work involving Ising machines, they lack spin size, hardware measurement and scalability, respectively, to be considered mature Ising machines. In comparison to ref. ²³, the reported device in this paper has a major disadvantage when directly comparing the spins as ref. ²³ has 60,000 spins. However, when the overall die size and technology node is taken into account, the spin advantage

Table 1 | Comparison with earlier Ising works

	This work	Digital ²³	Digital ²⁴	Qubit ²⁵	CMOS oscillator ²²	Phase ²⁶	OIM ¹⁹	OIM ²⁷	CIM ¹⁸
Architecture	King's graph	King's graph	King's graph	Chimera	Hexagonal	All to all	All to all	Chimera	All to all
Technology	CMOS 65 nm	CMOS 40 nm	CMOS 40 nm	Superconductor	CMOS 65 nm	Phase transition	Simulation	PCB/discrete components	Optical
Coupling	Oscillator interaction	Digital logic	Digital logic	Qubit interaction	Oscillator interaction	Resistive or capacitive	Equation	Oscillator interaction	Optical + digital (FPGA)
Bias term	No	Yes	Yes	Yes	No	No	Yes	No	No
Operating temperature	Room temperature	Room temperature	Room temperature	−273.14 °C	Room temperature	Room temperature	N/A	Room temperature	Room temperature
Weight resolution	5 levels	3 bits	5 bits	Not reported	1 bit	N/A	N/A	3 levels	−1 or +1
Time to solution	50 ns (50 cycles)	22 μs	6.5 ms	Not reported	200 ns	1 ms	N/A	1 ms	593 μs
Average power	42 mW	Not reported	Not reported	25 kW	23 mW	Not reported	N/A	5 W	Not reported
Measured accuracy	89–100%	98.8%	Not reported	Not reported	82–100%	97.6%	99.6%	Not reported	96.3%
Spins	1,968	60,000	147,000	2,048	560	4	N/A	240	100,000
Die size	2.1 mm ²	47.3 mm ²	216 mm ²	Not reported	1.44 mm ²	N/A	N/A	N/A	N/A

Note: die size for refs. ^{23,24} is the total die size of all the chips used in their respective systems.

is put into more reasonable terms. Ref. ²³ is over 22 times the size of this work and is at a more advanced technology node (40 nm); additionally, the time to solution in that work is 2,200 clock cycles or 22 μs. This is much more than the sub-50 cycles reported in this work and could be evidence supporting non-deterministic Ising machines. Power figures are not provided for ref. ²³; therefore, no accurate comparison can be made. As accuracy is only reported for non-contentious problems in ref. ²³, accuracy comparisons provide little value. It is important to note the resolution is three bits compared with this work's five levels. The follow-up paper to refs. ^{23,27} expands the body of work again to 147,000 spins in a 40 nm process. However, the area per spin increases between ref. ²³ and ref. ²⁷, with the caveat that the weight resolution increases to five bits. It is also seen that the time to solution increases to 6.5 ms. Accuracy numbers have not been reported at this time for ref. ²⁷. It is difficult to compare the proposed work here with ref. ²³ and ref. ²⁷ due to the immaturity of the CMOS Ising field compared with those of ref. ²³ and ref. ²⁷, but preliminary results show that ROSC-based Ising machines could prove to be more area efficient and hence more power efficient based on the rough trend of increasing power with increasing die size. Without transparency in the problem sets used, accuracy and time-to-solution conclusions are preliminary.

Another study²⁵ provides a non-CMOS comparison to the device reported in this paper. That study consists of a non-deterministic Ising machine that uses superconductor loops to represent spins. Although the spin counts are the same, the power consumption and operating temperature shown in ref. ¹⁹ present distinct disadvantages due to the magnitudes of difference in power consumption and operating temperature between the two devices. Accuracy numbers are unavailable for a direct comparison.

In addition to ref. ¹⁸, ref. ²⁵ is another non-CMOS, non-deterministic Ising machine with 100,000 spins. However, ref. ¹⁸ has limited weight resolution despite being a relatively mature Ising machine technology. Additionally, the device needs 56 peripheral FPGAs, which would be a limitation to scalability¹⁸. All-to-all coupling enables problem mapping without an embedding step which is required for the King's graph architecture used in this work. Ref. ¹⁸ needs a 5 km loop of optical cable and an extensive

temperature-controlled test setup—something that an integrated circuit approach avoids. The portability and space requirements of ref. ¹⁸ are reasons for concern with the design.

Comparing this work to a previous integrated oscillator-based Ising machine (OIM)¹⁹ shows the advancements made in this paper. As can be seen, spin count and accuracy have been improved, whereas the power per spin is decreased and spin density is increased. Additionally, the time-to-solution results and weight resolution for this work are better than ref. ¹⁹. Comparison of this work to ref. ²⁷ is similar to the comparison with ref. ¹⁹ with respect to power, area, spin count, weight resolution and time-to-solution results.

Conclusions

We have reported a 1,968-node coupled-ROSC-based Ising machine with five-level coupling for solving optimization problems. Our design has the density and energy efficiency of a CMOS integrated circuit, as well as using the natural spin-updating behaviour of a physics-based computer. Using a pseudo-analogue phase measurement circuit, we captured the true phase behaviour within a coupled-oscillator integrated circuit. A lightweight post-processing methodology was developed to address the realistic phase interactions observed in the measured data. The design achieves an accuracy of up to 95% for random COPs, and can relax to the ground state in less than 50 ROSC cycles and consuming an average power of only 0.042 W. The solution quality was also consistent across a wide range of processes, voltages and temperatures.

Methods

Circuit simulation. Synopsys HSPICE was used for the simulation of the multibit coupled oscillator array. A Taiwan Semiconductor Manufacturing Company (TSMC) 65 nm low-power CMOS compact model was used as the basis for the simulations. Synopsys Cosmoscope was used to view the simulated waveforms.

Integrated circuit design. Cadence Virtuoso was used to design the schematic and layout of the coupled oscillator design in TSMC 65 nm low-power technology. The layout was design-rule-checking and layout-versus-schematic verified using Mentor Graphics' Calibre software. One hundred dies were manufactured by TSMC and bonded in a dual in-line package with 40 pins by Advotech.

Test platform. Printed circuit board (PCB) as well as breadboard were used in conjunction with Agilent benchtop power supplies. A Raspberry Pi 3 Model B was used to interface with the fabricated integrated circuit chip. The RPi.GPIO library was used to control the Raspberry Pi's general-purpose input/output pins. The NumPy Python library was used to generate the random graphs and bit streams. For temperature testing, we placed the test boards in a TestEquity Model 107 temperature chamber.

Post-processing software. The post-processing software was developed using Python and the NumPy Python package.

Software COP solver. The Python package called qbsolv, which was used to generate the software Hamiltonian solutions, was obtained from D-Wave Systems. The solver was not time restricted or resource restricted when solving COPs.

Data availability

The problem set and test data presented in this paper are available at https://figshare.com/projects/A_1_968_node_coupled_ring_oscillator_circuit_for_combinatorial_optimization_problem_solving/134594.

Code availability

The code used to generate the optimal Hamiltonians is available at https://figshare.com/projects/A_1_968_node_coupled_ring_oscillator_circuit_for_combinatorial_optimization_problem_solving/134594.

Received: 15 August 2021; Accepted: 18 March 2022;

Published online: 02 May 2022

References

- Yang, G. *Industrial Applications of Combinatorial Optimization* (Springer, 1998).
- Vangelis, T. P. *Applications of Combinatorial Optimization* 2nd edn (Wiley, 2014).
- Markov, I. L. Limits on fundamental limits to computation. *Nature* **512**, 147–154 (2014).
- Markov, I. L. Know your limits. *IEEE Des. Test* **30**, 78–83 (2013).
- Hoover, H. J. Greenlaw, R. & Ruzzo, W. L. *Limits to Parallel Computation: P-Completeness Theory* (Oxford Univ. Press, 1995).
- Date, P., Arthur, D. & Pusey-Nazzaro, L. QUBO formulations for training machine learning models. *Sci. Rep.* **11**, 10029 (2021).
- Glover, F. Kochenberger, G. & Du, Y. A tutorial on formulating and using QUBO models. Preprint at <https://arxiv.org/abs/1811.11538> (2019).
- Lucas, A. Ising formulations of many NP problems. *Front. Phys.* **2**, 5 (2014).
- Ising, E. Beitrag zur theorie des ferromagnetismus. *Z. Phys.* **31**, 253–258 (1925).
- Neogy, T. & Roychowdhury, J. Analysis and design of sub-harmonically injection locked oscillators. In *2012 Design, Automation & Test in Europe Conference & Exhibition (DATE)* 1209–1214 (IEEE, 2012).
- Böhm, F., Verschaffel, G. & Van der Sande, G. A poor man's coherent Ising machine based on opto-electronic feedback systems for solving optimization problems. *Nat. Commun.* **10**, 3538 (2019).
- Vadlamani, S. K., Xiao, T. P. & Yablonoitch, E. Physics successfully implements Lagrange multiplier optimization. *Proc. Natl Acad. Sci. USA* **117**, 26639–26650 (2020).
- Mallick, A. et al. Using synchronized oscillators to compute the maximum independent set. *Nat. Commun.* **11**, 4689 (2020).
- Yamaoka, M. et al. A 20k-spin Ising chip to solve combinatorial optimization problems with CMOS annealing. *IEEE J. Solid-State Circuits* **51**, 303–309 (2016).
- Hsu, J. How much power will quantum computing need? *IEEE Spectrum* **5** (2015).
- Yamamoto, Y. et al. Coherent Ising machines—optical neural networks operating at the quantum limit. *npj Quantum Inf.* **3**, 49 (2017).
- Borders, W. et al. Integer factorization using stochastic magnetic tunnel junctions. *Nature* **573**, 390–393 (2019).
- Honjo, T. et al. 100,000-spin coherent Ising machine. *Sci. Adv.* **7**, eabh0952 (2021).
- Wang, T. & Roychowdhury, J. Oscillator-based Ising machine. Preprint at <https://arxiv.org/abs/1709.08102> (2017).
- Glover, F. Tabu search—part I. *ORSA J. Comput.* **1**, 190–206 (1989).
- Glover, F. Tabu search—part II. *ORSA J. Comput.* **2**, 4–32 (1990).
- Ahmed, I., Chiu, P.-W., Moy, W. & Kim, C. H. A probabilistic compute fabric based on coupled ring oscillators for solving combinatorial optimization problems. *IEEE J. Solid-State Circuits* **56**, 2870–2880 (2021).
- Takemoto, T. et al. A 2 × 30k-spin multi-chip scalable CMOS annealing processor based on a processing in-memory approach for solving large-scale combinatorial optimization problems. *IEEE J. Solid-State Circuits* **55**, 145–156 (2020).
- Takemoto, T. et al. 4.6 A 144Kb annealing system composed of 9 × 16Kb annealing processor chips with scalable chip-to-chip connections for large-scale combinatorial optimization problems. In *2021 IEEE International Solid-State Circuits Conference (ISSCC)* 64–66 (IEEE, 2021).
- Bian, Z. et al. The Ising model: teaching an old problem new tricks. *D-Wave Systems* **2**, 1–32 (2010).
- Dutta, S. et al. Experimental demonstration of phase transition nano-oscillator based Ising machine. In *2019 IEEE International Electron Devices Meeting (IEDM)* 37.8.1–37.8.4 (IEEE, 2019).
- Wang, T., Wu, L. & Roychowdhury, J. New computational results and hardware prototypes for oscillator-based Ising machines. In *Proc. 56th Annual Design Automation Conference 2019* **56**, 1–2 (IEEE, 2019).

Acknowledgements

I.A. and C.H.K. acknowledges the initial support of the project from the National Science Foundation under ECCS 1739635 and the Semiconductor Research Corporation (SRC) under 2759.007. For the chip testing and data analysis works, W.M. and C.H.K. have been supported in part by the SRC under 3024.001. We would like to thank SRC's industry liaisons for their technical feedback.

Author contributions

W.M., I.A., P.-W.C., S.S.S. and C.H.K. participated in the circuit and architecture design of the integrated circuit. W.M., I.A. and P.-W.C. created the layout for the fabrication of the integrated circuit. W.M. and C.H.K. performed the testing and measurement of the chip. J.M. created the cloud-testing interface and testing automation setup for the chip.

Competing interests

The authors declare the following competing interest: US patent application 17/213,396 (Probabilistic compute engine using coupled ROSCs).

Additional information

Extended data is available for this paper at <https://doi.org/10.1038/s41928-022-00749-3>.

Correspondence and requests for materials should be addressed to Chris H. Kim.

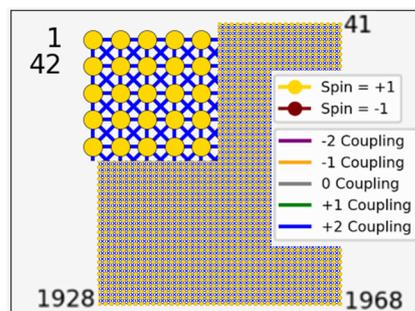
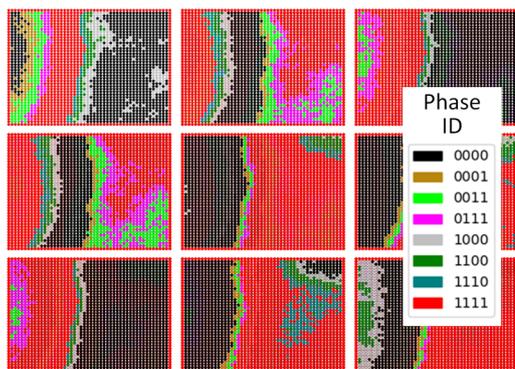
Peer review information *Nature Electronics* thanks Michael Huang and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

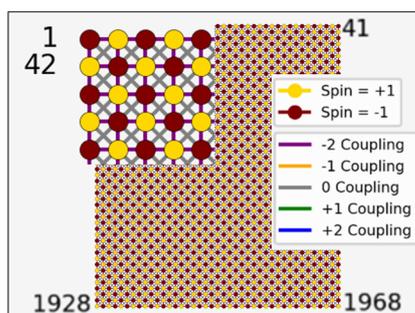
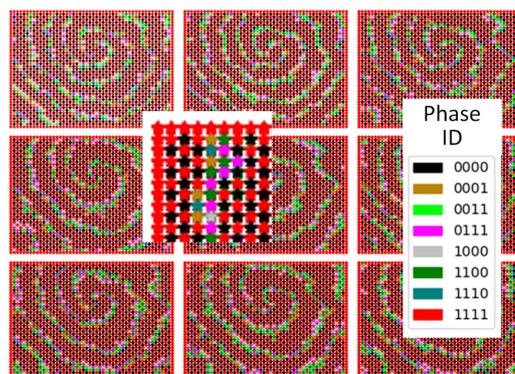
Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature Limited 2022

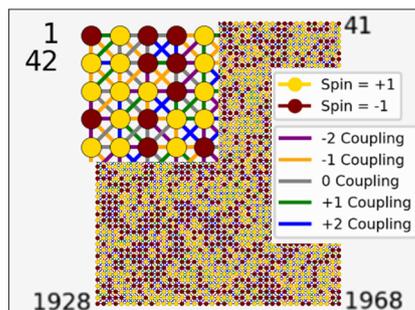
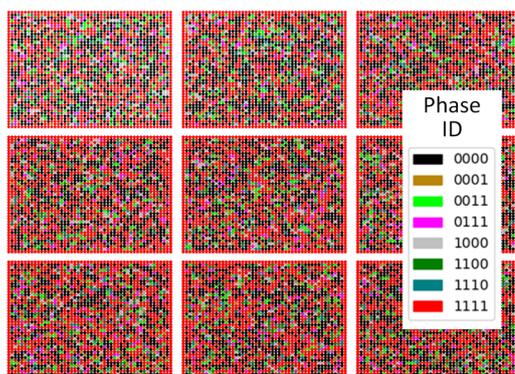
Positive weights (all weights = +2)



Negative weights (all weights except for diagonals = -2)



Random weights (weights = -2, -1, 0 +1, +2)



Raw phase maps (9 measurements)

After light-weight post-processing

Extended Data Fig. 1 | Raw phase map data with post-processed results. Raw phase maps are shown for uniform, chequerboard, and random pattern graphs. A lightweight post-processing method that emphasizes the relative phase difference between adjacent nodes can eliminate non-ideal effects and produce the expected spin value maps on the right. The weights and spin values are colour coded and overlaid for better visualization. ROSCs are indexed from 1 (upper left corner) to 1,968 (lower right corner), with ROSC 1 being directly above ROSC 42.