www.sciencemag.org/cgi/content/full/science.aah4243/DC1

Supplementary Materials for

A coherent Ising machine for 2000-node optimization problems

Takahiro Inagaki,* Yoshitaka Haribara, Koji Igarashi, Tomohiro Sonobe, Shuhei Tamate, Toshimori Honjo, Alireza Marandi, Peter L. McMahon, Takeshi Umeki, Koji Enbutsu, Osamu Tadanaga, Hirokazu Takenouchi, Kazuyuki Aihara, Ken-ichi Kawarabayashi, Kyo Inoue, Shoko Utsunomiya, Hiroki Takesue*

*Corresponding author. E-mail: inagaki.takahiro@lab.ntt.co.jp (T.I.); takesue.hiroki@lab.ntt.co.jp (H.T.)

Published 20 October 2016 on *Science* First Release
DOI: 10.1126/science.aah4243

This PDF file includes:

Materials and method

Stochastic differential equations

The cooperative oscillation process of networked degenerate optical parametric oscillators (DOPO) with a measurement and feedback (MFB) scheme is described by the following c-number stochastic differential equations (*9*, *23*):

$$dc_i = \left[(-1 + p - c_i^2 - s_i^2)c_i + \sum_j \xi_{ij} \tilde{c}_j \right] dt + \frac{1}{A_s} \sqrt{c_i^2 + s_i^2 + \frac{1}{2}} dW_1 , \tag{S1}$$

$$ds_i = (-1 - p - c_i^2 - s_i^2)s_i dt + \frac{1}{A_s} \sqrt{c_i^2 + s_i^2 + \frac{1}{2}} dW_2 , \tag{S2}$$

where c_i and s_i are the in-phase and quadrature phase amplitudes of the ith DOPO, respectively. $\tilde{c}_i$ is the measured in-phase amplitude of the ith DOPO, which can differ from the original state due to the noise induced in the measurement process. Only the in-phase component c_i is amplified by the phase sensitive amplification of a DOPO, and the sign of c_i corresponds to the binary oscillation mode. $\xi_{ij} = -rJ_{ij}$ is a coupling coefficient from the jth to ith DOPO, where r is a constant that determines the coupling strength. The pump rate p is normalized by the pump threshold of a single isolated DOPO, and A_S is defined as the saturation amplitude. dW_1 and dW_2 are independent Gaussian noise processes that originate from vacuum fluctuations from the open port and the pump field fluctuation for in-phase and quadrature phase components, respectively. The 2nd term in the square brackets on the right hand side of Eq. (S1) gives the spin-spin coupling term, which is obtained by a calculation undertaken in the FPGA module.

Maximum cut (MAX-CUT) problem

MAX-CUT is the problem of dividing the nodes of a given graph into two subsets $\{S_1, S_2\}$ to maximize the sum of the weights w_{ij} ($i \in S_1, j \in S_2$) of the edges between the divided subsets. The cut value $C(\sigma)$ can be counted as follows:

$$C(\sigma) = \sum_{i \in S_1, j \in S_2} w_{ij} = \sum_{i<j} w_{ij} \frac{(1 - \sigma_i \sigma_j)}{2} = \frac{1}{2} \left(\sum_{i<j} w_{ij} - \sum_{i<j} H(\sigma) \right) , \tag{S3}$$

where we assign binary variables $\sigma \in \{-1, +1\}$ to represent the divided subsets $\{S_1, S_2\}$. Here, H is the Ising Hamiltonian defined in Eq. (1) by $J_{ij} = -w_{ij}$, and the cut value is maximized when the Ising energy is minimized. This means that a MAX-CUT problem is intrinsically equivalent to the Ising problem, and we can solve a MAX-CUT problem with the coherent Ising machine (CIM) by finding the lowest-energy spin configuration of the corresponding Ising model. In the experiment, the quantities we observed with the BHD were the in-phase amplitudes of the signal DOPOs, $\{\tilde{c}_i\}$, which are analog values. Therefore, when calculating cut values or Ising energies, we converted $\{\tilde{c}_i\}$ into $\{\sigma_i\}$ as $\sigma_i = \mathrm{sgn}\, \tilde{c}_i$.

Implementation of simulated annealing

Here we describe the implementation of simulated annealing used in our benchmark study, which was initially introduced by Kirkpatrick et al. (*2*). Since the benchmark instance of interest was a complete graph with edge weights $w_{ij} \in \{-1, +1\}$, graph data could be represented by a binary array of size $N(N-1)/2$. Spin values for $N = |V|$ vertices were also stored in a binary array of size N, which is known as multi-spin coding (*30*). In the case of N=2,000, the size of graph data was ~244KB, which fitted in the cache of modern CPU so that we could efficiently access to the graph data. For the efficient calculation, we used a single instruction multiple data (SIMD) operation (bitwise XOR and pop count) to boost an energy difference calculation. The scheduling was based on a logarithmic function: the inverse temperature is set as

$$\beta = \beta_0 \log(1 + t/T), \tag{S4}$$

where β_0 and T are the temperature and time scaling factor, respectively. We implemented single-threaded C++ program whose pseudo-code is shown in Algorithm 1. The code is available in at (https://github.com/haribara/SA-complete-graph/releases/tag/WK2000)

Algorithm 1 Simulated Annealing

Read graph adjacency matrix $\boldsymbol{J}$ — Binary when solving ±1-weighted complete graph
Initialize spin variable σ — Binary array
$E \leftarrow H(\sigma)$ — Calculate initial Ising energy using SIMD bit-count
loop
  $\beta \leftarrow \beta_0 \log(1 + t/T)$ — Inverse temperature with normalized time $t \in [0, 1]$
  Randomly choose vertex $v \in V$
  $\sigma' \leftarrow \mathrm{flip}(\sigma, v)$ — State whose vertex v is flipped
  $\Delta E \leftarrow H(\sigma') - E$ — Calculate energy difference (SIMD)
  if $\Delta E < 0$ or $\exp(-\beta \Delta E) \geq \mathrm{Uniform}(0, 1)$ **then**
    $\sigma \leftarrow \sigma'$ — Update spin
    $E \leftarrow E + \Delta E$
  end if
end loop — When the value reaches the target or time runs out
return E, σ

The graph we used in the benchmarking was generated by rudy.c provided by Helmberg which implements a machine independent random number generator (*27*). For example, the following command

./ rudy −clique 2000 −random 0 1 20001 −times 2 −plus −1

generates a ±1-weighted complete graph of the order $N = 2000$.

The results in Fig. 3 and Fig. 4 were obtained after the parameter optimization. Figure S1A shows the parameter (β_0, T) dependence of the ±1-weighted K_{2000} in a fixed time of 50 ms. Figure S1B shows the parameter dependence on the same graph when the target accuracy was changed. The benchmarking was performed on an Ubuntu 14.04.4 Linux server with two 6 core 2.67 GHz Intel Xeon X5650 processors (L1 data and L2 cache size were 32KB and 256KB respectively) and a 96 GB memory. The codes were

compiled with a GNU C++ compiler (version4.9.3). An option to reduce the computation time further is to parallelize on a multi-core CPU. However, in this N=2,000 case, the data size is small enough to fit in a cache of a single core, and each iteration is very fast. Therefore, because of the synchronization and communication overheads between threads, we benefit little from paralleling SA on a multi-core CPU.

Details of experimental setup

The experimental setup is shown in Fig. 1. A 1536-nm continuous wave from a laser output is split into two paths, one of which is used as a local oscillator for the balanced homodyne detector (BHD), and the other is launched into an intensity modulator (IM1), which produces a pulse train with a 40-ps width and a 1-GHz repetition frequency. The pulse train is amplified by an erbium-doped fiber amplifier (EDFA1), and the resulting pulse train is further split into two parts, one of which is launched into a push-pull modulator to prepare a feedback optical signal and the other is launched into a second intensity modulator (IM2), which controls the amplitude of the pump pulses. The pulses are then amplified by EDFA2, filtered to suppress amplified spontaneous emission noise, and then launched into a periodically poled lithium niobate (PPLN) waveguide module (PPLN1) to obtain 768-nm pump pulses via the second harmonic generation (SHG) process. The 768-nm pump pulses are then input into another PPLN waveguide module (PPLN2), which is placed in a 1-km fiber cavity as a phase sensitive amplifier (PSA).

To realize the pump schedule shown in Fig. S2A, the transmittance of IM2 was controlled in a 5-ms period.

Periodically poled lithium niobate waveguides

PPLN1 and PPLN2 are fiber-coupled modules that contain 5-cm-long ZnO-doped $LiNbO_3$ ridge waveguides fabricated using the direct bonding method (*31*). The doping of the LiNbO3 core layer with ZnO made the waveguides highly resistant to photorefractive damage. By using the dry etching technique, we obtained a ridge waveguide with fine uniformity. As shown in Fig. 1, the PPLN waveguide modules are equipped with dichroic mirrors that reflect the 0.8-μm light while transmitting the 1.5-μm light in front of the input and output facets of the waveguides. The temperatures of the PPLN modules are tuned and stabilized so that the SHG power and the DOPO powers are maximized.

Phase stabilization of optical system

The output DOPO pulses extracted from 9:1 coupler 1 were divided into two using another coupler (not shown in Fig. 1). One part was launched into the BHD and the other was used for monitoring and cavity locking. The DOPO pulse train was further divided into two paths; one was used to monitor optical power and the other was launched into a photodetector after being amplified with an EDFA. The signal from the photodetector was used to obtain the error signal for cavity locking. We employed a dither-and-lock scheme (*20*), and we controlled the cavity length by feedback-controlling piezo-based fiber stretcher (FS1) so that the photocurrent was maximized. As a result, we observed that the phase pattern of the dummy DOPO pulses, which is inherently random but that should repeat that same pattern for every round trip, was kept the same on the

oscilloscope for an average of several tens of seconds. This confirms that the cavity phase was successfully locked to the pump phase.

In addition to the phase locking of the 1-km fiber cavity, the phases of the local oscillator for the BHD and the coupling pulses were also stabilized by employing the dither-and-locking scheme using FS2 and FS3 shown in Fig. 1. To phase-lock the local oscillator, we obtained the error signal by dividing the detection signal from the BHD. The error signal for locking the coupling pulse phase was obtained from a photodetector placed at the output port of 9:1 coupler 2 (see Fig. 1).

FPGA modules for measurement and feedback scheme

We used two cascaded Xilinx Virtex-7 (XC7VX690T) FPGAs for MFB. The FPGA contains 693,120 logic cells and 3,600 digital signal processor (DSP) slices, and has a 52,920-Kb memory. The 2,048 signal DOPO amplitudes measured by the BHD were digitized by a 12-bit analog-to-digital converter (ADC). The FPGA could handle input data at up to 1 GS/s from the ADC through 8-channel parallel processing at an operation frequency of 125 MHz, and identify the number of signal DOPOs thanks to an external header signal. Since the number of logic cells in the FPGA was not large enough to allow us calculate the feedback signals for the 2,048×2,048 matrix of spin-spin interactions J_{ij}, we divided the calculation tasks between two FPGAs and connected them in tandem with four 12.5 Gbps transceivers. To reduce the number of logic cells required for the feedback signal calculations, we used the logic cells as the input data selectors instead of multipliers to implement the interaction weights of three values $J_{ij} \in \{-1, 0, +1\}$. In addition, the resolution of the input data was reduced from 12 bits (at the ADC output) to 5 bits when calculating the feedback signals, so that we could further save the logic cell resources. To ensure the accuracy of the feedback signal calculations, the amplitude range for the 5 bits was changed during the temporal evolution of the DOPO amplitudes so that we could obtain adequate resolutions in the DOPO amplitude measurements. The 2,048 calculation results were converted to feedback signals by a 14-bit DAC at a sampling frequency of 2 GHz, which were then imposed on the coupling pulses with a push-pull optical modulator. The coupling pulses that conveyed the feedback signals were injected into the corresponding signal DOPOs. The coupling pulses and the signal DOPOs were synchronized by adjusting the temporal delay of the FPGA output and the length of the optical path for the coupling pulses using a variable optical delay line placed in front of the push-pull modulator (not shown in Fig. 1). We could accomplish this MFB sequence within 3 μs including the measurement time of 2,048 DOPO amplitudes and the 2,048×2,048 matrix calculation time in FPGA modules. The FPGA could store up to 512 sets of 2,048 DOPO amplitudes. This means that we could store the final spin states of 512 sequential measurements, or temporal evolutions of the 2,048 DOPO signals for up to 512 circulations. Therefore, when we observed a temporal evolution as shown in Figs. 2 and 4, we recorded the signal DOPO amplitudes every 2 circulations since a computation trial took 1,000 circulations of DOPOs in the cavity.

In the current FPGA module, the host computer sets J_{ij} and retrieves c_i one after the other via a slow serial communication interface (RS-232C), which makes the data transfer between the host computer and the FPGA module very slow (~60s). By implementing the data transfer between the FPGA module and the host computer with a high speed interface such as PCI-Express, we can significantly improve the transfer time.

Under the experimental condition described in this paper, the host computer sends ~244KB of J_{ij} data, and retrieves ~4KB of c_i data. If we use PCI-Express Gen3 (x8), the J_{ij} and c_i transfer times are estimated to be ~37 and 0.6 μsec, respectively, which are much shorter than the typical operation time of the CIM.

Pump schedule of CIM in solving MAX-CUT

To obtain the results shown in Fig. 3, we employed the 1.5-μm pump schedule for SHG shown in Fig. S2A. Here, the pump power was normalized at the DOPO threshold without optical couplings. The 1.5-μm pump power, which is proportional to the amplitude of the 0.8-μm pump for the PSA, was gradually increased from the first to the 950th cavity circulation, and then abruptly increased at the 951st circulation to saturate the DOPOs so that their phases were discretized to either 0 or π. The pump and injection pulses were turned off after the 1000th circulation to initialize the DOPO states. On the other hand, in the time evolution measurement shown in Fig 4, the 1.5-μm pump power was abruptly increased at the first circulation to the maximum normalized pump power of ~1.2 so that the CIM could obtain the benchmark Ising energy faster than the schedule shown in Fig. S2A. The DOPO output power as a function of the normalized 1.5-μm pump power is shown in Fig. S2B. At the DOPO threshold, the average 1.5-μm pump power was ~270 mW.

Role of quantumness in CIM

The balanced homodyne detection in our system is operated at nearly the shot-noise limit. Therefore, the DOPO oscillation is initiated by the quantum noise originating from a vacuum, and it undergoes a squeezed vacuum state formed by the nonlinear process in the phase sensitive amplification (PSA), before reaching a coherent state as a result of the DOPO phase transition. According to a numerical study in which optically-coupled DOPOs were analyzed, the formation of a superposed state is observed during the DOPO phase transition (Fig. 7 in (*14*)). It is also shown numerically that the solution searching assisted by the formation of a superposed state is severely perturbed by injecting thermal noise into the DOPO cavity (Fig. 13 in (*14*)). From these numerical results, we consider that the CIM is being operated in a partially quantum regime, but we have yet to confirm the quantum behavior of our system experimentally. Moreover, it is also important to clarify experimentally whether the existence of a superposed state in the CIM contributes to better performance in terms of computation time and accuracy.

The lower-energy searching process of the CIM occurs during the DOPO phase transition, which is caused by the nonlinear interactions of photons in a PSA with an initial state originating from quantum noise. To the best of our knowledge, it is still an open question whether such a nonlinear phenomenon seeded by quantum noise can be efficiently emulated with a purely classical system, such as special purpose electronics without any quantum effects.

Supplementary Text

Visualization of 2000-node graph

Figure S3A shows the graph structure of G22, which is a sparsely connected random graph with 2,000 nodes and 19,900 edges. A solution to the MAX-CUT problem on G22 graph obtained with the CIM is shown in Fig. S3B. The red and blue nodes correspond to up and down spin states of Ising model, respectively. The green lines show the edges between the divided subsets. The cut value was 13,313.

Temporal evolutions of cut value and Ising energy

Figure S4 shows the temporal evolutions of the cut value and the Ising energy for the data shown in Fig. 2. In this measurement, we employed the same pump schedule as that shown in Fig. S2A. At the 1,000th circulation, the cut value reached 32,612 and the Ising energy decreased to −66,264.

Average temporal evolutions of Ising energy for solving complete graph K_{2000}

In addition to the best results obtained in 100 trials with the CIM and SA shown in Fig. 4, we evaluated the average performance of the CIM and SA. The temporal evolutions of the Ising energy averaged for 26 trials with the CIM and those averaged for 100 trials with SA are shown by red and black curves, respectively, in Fig. S5. In the current CIM system, the transfer of the data to a computer takes approximately 60 seconds, which is much longer than the average stabilized time of the optical setup in the CIM. Therefore, we performed temporal evolution measurements 100 times sequentially using the CIM, and then eliminated the results where DOPO instabilities were observed. Such instabilities include phase flips of the local oscillators for the BHD or the coupling pulses and abrupt phase changes possibly caused by mechanical vibrations of the optical setup. The times required to reach the benchmark Ising energy were 270 μs for the CIM and 5.5 ms for SA.

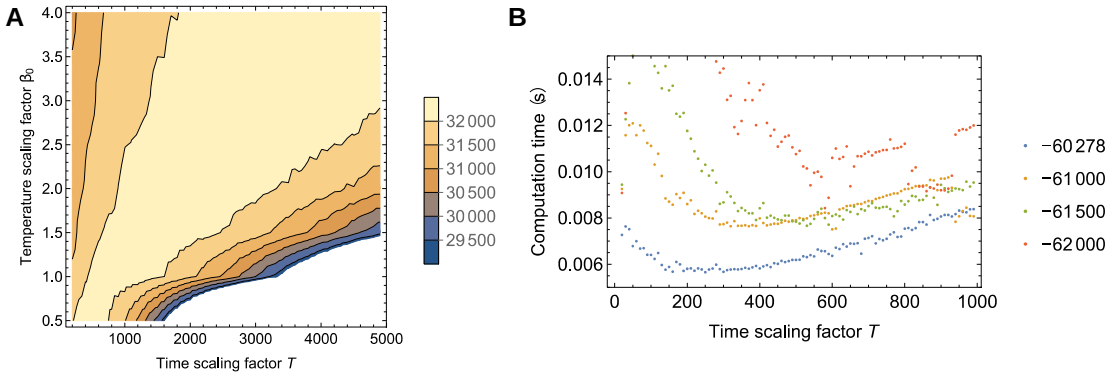

Fig. S1

Parameter optimization of SA in solving MAX-CUT for the complete graph K_{2000}. (**A**) Average cut value of 100 trials obtained in a fixed time of 50 ms with various temperature scaling factors β_0 and time scaling factors T. We used $\beta_0 = 4.0$ and $T = 4{,}200$ for obtaining the results shown in Fig. 3. (**B**) Computation time as a function of time scaling factor T for various target values. We used $T = 200$ to reach the target Ising energy of −60,278 as soon as possible for the time comparison in Fig. 4.

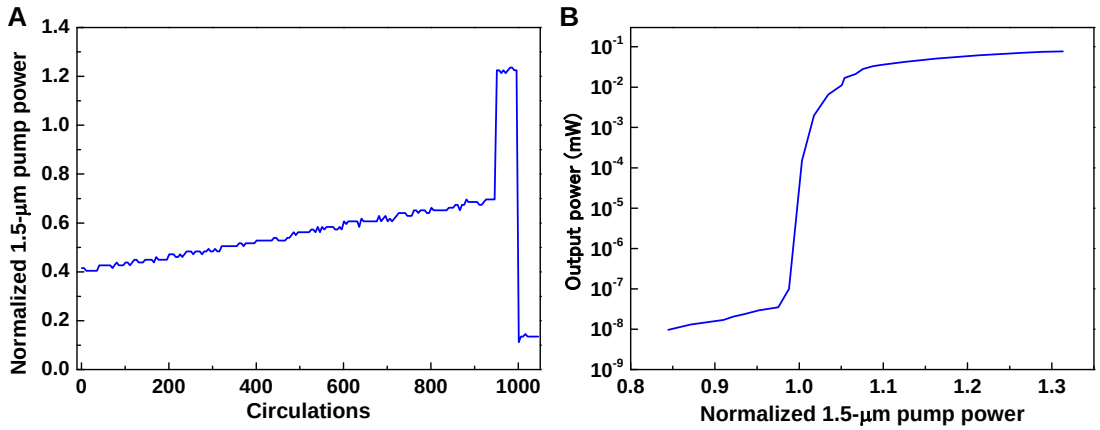

Fig. S2

(**A**) Pump schedule for CIM when solving MAX-CUT in Fig. 3. (**B**) DOPO output power as a function of a normalized pump rate without optical couplings.

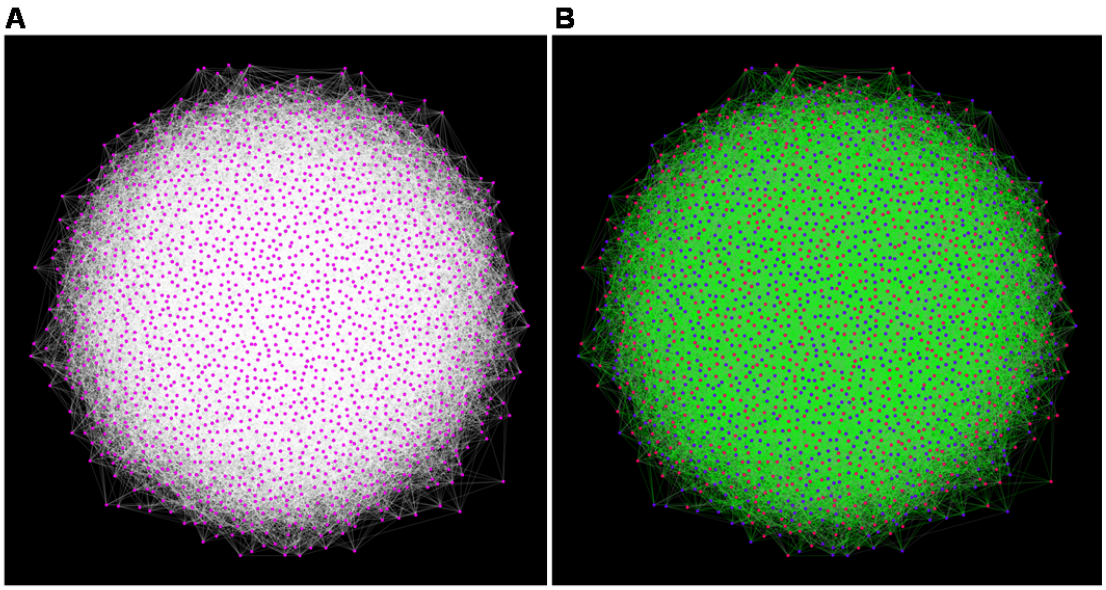

Fig. S3

Visualization of 2000-node graph. (**A**) G22: a sparsely connected random graph with 2,000 nodes and 19,900 edges. (**B**) A solution to MAX-CUT problem on G22 graph obtained by the CIM.

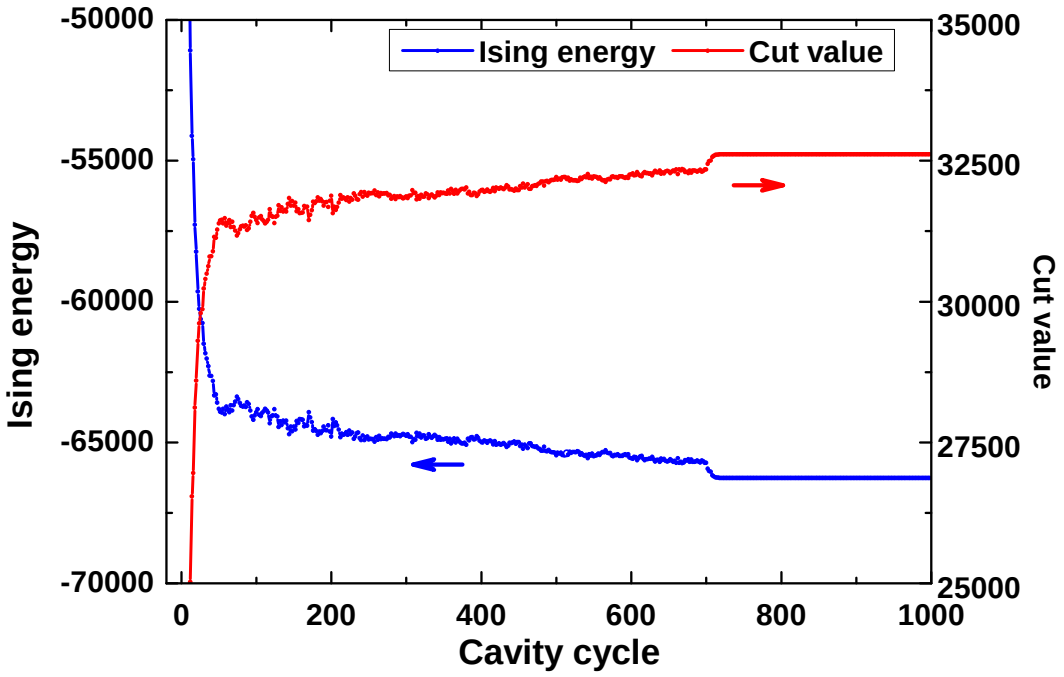

Fig. S4.

Time evolutions of the MAX-CUT value (red) and Ising energy (blue) when solving the complete graph K_{2000}.

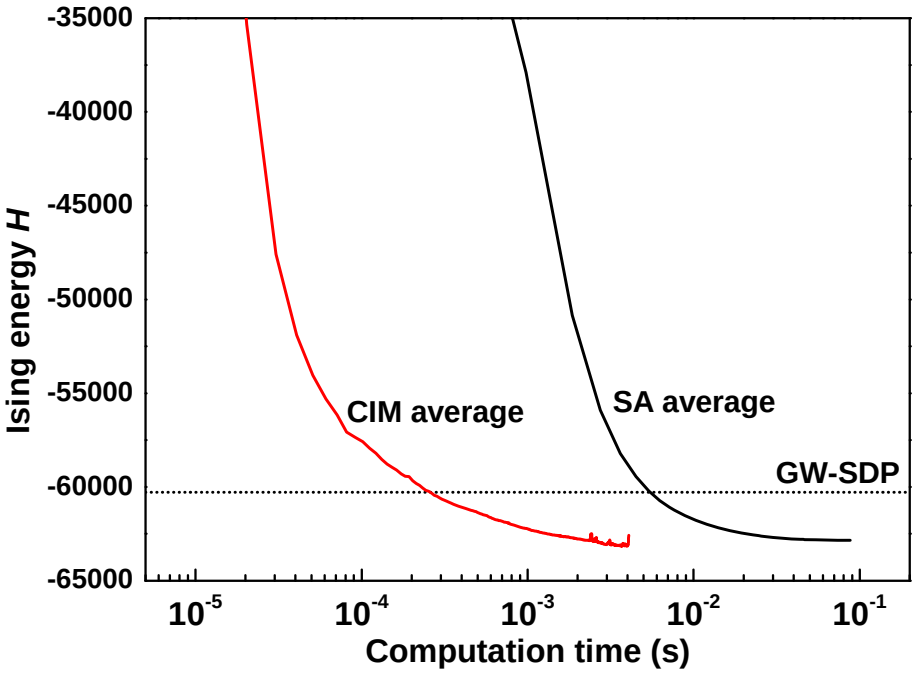

Fig. S5

Time evolutions of Ising energy obtained with CIM and SA when solving the complete graph K_{2000}. The dotted line corresponds to the benchmark energy H = –60,278 obtained with the GW-SDP algorithm. The temporal evolution of the Ising energy averaged for 26 trials with the CIM and that averaged for 100 trials with SA are shown by the red and black curves, respectively. The times required to reach the benchmark energy were 270 μs (CIM) and 5.5 ms (SA).

References

1. R. M. Karp, *Complexity of Computer Computations* (Plenum, 1972), pp. 85–103.

2. S. Kirkpatrick, C. D. Gelatt Jr., M. P. Vecchi, Optimization by simulated annealing. *Science* **220**, 671–680 (1983). Medline doi:10.1126/science.220.4598.671

3. F. Barahona, On the computational complexity of Ising spin glass models. *J. Phys. A* **15**, 3241–3253 (1982). doi:10.1088/0305-4470/15/10/028

4. M. W. Johnson, M. H. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk, E. M. Chapple, C. Enderud, J. P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, C. J. Truncik, S. Uchaikin, J. Wang, B. Wilson, G. Rose, Quantum annealing with manufactured spins. *Nature* **473**, 194–198 (2011). Medline doi:10.1038/nature10012

5. K. Kim, M. S. Chang, S. Korenblit, R. Islam, E. E. Edwards, J. K. Freericks, G. D. Lin, L. M. Duan, C. Monroe, Quantum simulation of frustrated Ising spins with trapped ions. *Nature* **465**, 590–593 (2010). Medline doi:10.1038/nature09071

6. I. Mahboob, H. Okamoto, H. Yamaguchi, An electromechanical Ising Hamiltonian. *Sci. Advances* **2**, e1600236 (2016). doi:10.1126/sciadv.1600236

7. M. Yamaoka, C. Yoshimura, M. Hayashi, T. Okuyama, H. Aoki, H. Mizuno, A 20k-spin Ising chip to solve combinatorial optimization problems with CMOS annealing. *IEEE J. Solid-State Circuits* **51**, 303–309 (2015).

8. S. Utsunomiya, K. Takata, Y. Yamamoto, Mapping of Ising models onto injection-locked laser systems. *Opt. Express* **19**, 18091–18108 (2011). Medline doi:10.1364/OE.19.018091

9. Z. Wang, A. Marandi, K. Wen, R. L. Byer, Y. Yamamoto, Coherent Ising machine based on degenerate optical parametric oscillators. *Phys. Rev. A* **88**, 063853 (2013). doi:10.1103/PhysRevA.88.063853

10. M. Wolinsky, H. J. Carmichael, Quantum noise in the parametric oscillator: From squeezed states to coherent-state superpositions. *Phys. Rev. Lett.* **60**, 1836–1839 (1988). Medline doi:10.1103/PhysRevLett.60.1836

11. L. Krippner, W. J. Munro, M. D. Reid, Transient macroscopic quantum superposition states in degenerate parametric oscillation: Calculations in the large-quantum-noise limit using the positive P representation. *Phys. Rev. A* **50**, 4330–4338 (1994). Medline doi:10.1103/PhysRevA.50.4330

12. P. D. Drummond, K. Dechoum, S. Chaturvedi, Critical quantum fluctuations in the degenerate parametric oscillator. *Phys. Rev. A* **65**, 033806 (2002). doi:10.1103/PhysRevA.65.033806

13. K. Takata, A. Marandi, Y. Yamamoto, Quantum correlation in degenerate optical parametric oscillators with mutual injections. *Phys. Rev. A* **92**, 043821 (2015). doi:10.1103/PhysRevA.92.043821

14. D. Maruo, S. Utsunomiya, Y. Yamamoto, Truncated Wigner theory of coherent Ising machines based on degenerate optical parametric oscillator network. *Phys. Scr.* **91**, 083010 (2016). doi:10.1088/0031-8949/91/8/083010

15. A. Marandi, Z. Wang, K. Takata, R. L. Byer, Y. Yamamoto, Network of time-multiplexed optical parametric oscillators as a coherent Ising machine. *Nat. Photonics* **8**, 937–942 (2014). doi:10.1038/nphoton.2014.249

16. T. Inagaki, K. Inaba, R. Hamerly, K. Inoue, Y. Yamamoto, H. Takesue, Large-scale Ising spin network based on degenerate optical parametric oscillators. *Nat. Photonics* **10**, 415–419 (2016). doi:10.1038/nphoton.2016.68

17. J. King, S. Yarkoni, M. M. Nevisi, J. P. Hilton, C. C. McGeoch, http://arxiv.org/abs/1508.05087 (2015).

18. E. Ising, Beitrag zur Theorie des Ferromagnetismus. *Z. Phys. A* **31**, 253–258 (1925). doi:10.1007/BF02980577

19. C. D. Nabors, S. T. Yang, T. Day, R. L. Byer, Coherence properties of a doubly-resonant monolithic optical parametric oscillator. *J. Opt. Soc. Am. B* **7**, 815–820 (1990). doi:10.1364/JOSAB.7.000815

20. A. Marandi, N. C. Leindecker, V. Pervak, R. L. Byer, K. L. Vodopyanov, Coherence properties of a broadband femtosecond mid-IR optical parametric oscillator operating at degeneracy. *Opt. Express* **20**, 7255–7262 (2012). Medline doi:10.1364/OE.20.007255

21. D. K. Serkland, G. D. Bartolini, A. Agarwal, P. Kumar, W. L. Kath, Pulsed degenerate optical parametric oscillator based on a nonlinear-fiber Sagnac interferometer. *Opt. Lett.* **23**, 795–797 (1998). Medline doi:10.1364/OL.23.000795

22. Y. Okawachi, M. Yu, K. Luke, D. O. Carvalho, S. Ramelow, A. Farsi, M. Lipson, A. L. Gaeta, Dual-pumped degenerate Kerr oscillator in a silicon nitride microresonator. *Opt. Lett.* **40**, 5267–5270 (2015). Medline doi:10.1364/OL.40.005267

23. Y. Haribara, S. Utsunomiya, K. Kawarabayashi, Y. Yamamoto, https://arxiv.org/abs/1501.07030 (2015).

24. See supplementary materials on *Science* Online.

25. M. E. Marhic, C. H. Hsia, J. M. Jeong, Optical amplification in a nonlinear fiber interferometer. *Electron. Lett.* **27**, 210–211 (1991). doi:10.1049/el:19910136

26. R.-D. Li, P. Kumar, W. L. Kath, Dispersion compensation with phase-sensitive optical amplifiers. *J. Lightwave Technol.* **12**, 541–549 (1994). doi:10.1109/50.285338

27. C. Helmberg, F. Rendl, A spectral bundle method for semidefinite programming. *SIAM J. Optim.* **10**, 673–696 (2000). doi:10.1137/S1052623497328987

28. G. Rinaldy, rudy graph generator, www-user.tu-chemnitz.de/˜helmberg/rudy.tar.gz (1996).

29. M. X. Goemans, D. P. Williamson, Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM* **42**, 1115–1145 (1995). doi:10.1145/227683.227684

30. S. V. Isakov, I. N. Zintchenko, T. F. Rønnow, M. Troyer, Optimised simulated annealing for Ising spin glasses. *Comput. Phys. Commun.* **192**, 265–271 (2015). doi:10.1016/j.cpc.2015.02.015

31. T. Umeki, O. Tadanaga, A. Takada, M. Asobe, Phase sensitive degenerate parametric amplification using directly-bonded PPLN ridge waveguides. *Opt. Express* **19**, 6326–6332 (2011). Medline doi:10.1364/OE.19.006326