Encoding arbitrary Ising Hamiltonians on Spatial Photonic Ising Machines

Jason Sakellariou,[1,2,*] Alexis Askitopoulos,[1,2,†] Georgios Pastras,[1] and Symeon I. Tsintzos[1,2,‡]

[1] *QUBITECH, Thessalias 8, GR 15231 Chalandri, Athens, Greece*
[2] *UBITECH Ltd., 95B Archiepiskopou Makariou, CY 3020 Limassol Cyprus*

(Dated: March 5, 2025)

Photonic Ising Machines constitute an emergent new paradigm of computation geared towards tackling combinatorial optimization problems that can be reduced to the problem of finding the ground state of an Ising model. Spatial Photonic Ising Machines have proven advantageous for simulating fully connected large-scale spin systems. However, fine control of a general interaction matrix J has only been accomplished so far through eigenvalue decomposition methods that either limit the scalability or increase the execution time of the optimization process. We introduce and experimentally validate a SPIM instance that enables direct control over the full interaction matrix, allowing the encoding of Ising Hamiltonians with arbitrary couplings and connectivity. We demonstrate the conformity of the experimentally measured Ising energy with the theoretically expected values and then proceed to solve both the unweighted and weighted graph partitioning problems, showcasing a systematic convergence to an optimal solution via simulated annealing. Our approach significantly expands the applicability of SPIMs for real-world applications without sacrificing any of the inherent advantages of the system. It paves the way to encoding the full range of NP problems that are known to be equivalent to Ising models on SPIM devices.

In recent years, a growing effort has been devoted toward implementing special-purpose physical machines that simulate the Ising Hamiltonian. These machines have garnered considerable attention for their potential to efficiently tackle optimization problems across diverse domains, as many non-deterministic polynomial-time hard (NP-hard) problems can be mapped to the Ising Hamiltonian [1–3]. Several physical implementations of Ising machines, employing either quantum or classical schemes, have been developed. Quantum Ising annealers have been realized by trapped atoms [4] and ions [5], single photons [6] and superconducting circuits [7]. In addition, their classical counterparts have been realized in various physical platforms including polariton condensates [8–11], stochastic magnetic junctions [12], memristors [13], coupled electrical oscillators [14], complementary metal oxide semiconductor technologies (CMOS) [15], networks of coupled optical pulses in a ring fiber [16, 17] and lately spatial photonic Ising machines (SPIMs) [18].

Photonic Ising machines are of particular interest as they exploit the advantages of mature optoelectronic technologies and have been shown to address large-scale combinatorial optimization problems [18, 19]. SPIMs, a newcomer to the field, utilize holographic optical phase modulation, leveraging 1) enhanced scalability, 2) all-to-all connectivity, 3) room temperature operation, 4) inherent parallelism, 5) low cost, and 6) low power consumption. Compared to other photonic Ising machines, SPIMs utilize a straightforward optical configuration to encode the physical parameters of the Ising Hamiltonian as distinct holographic phases on a discretized optical wavefront with the use of spatial light modulators (SLM) [20]. Since their introduction in 2019 [18], SPIMs have showcased their applicability to adiabatic evolution methods [21], while the inherent system noise was shown to be a valuable system resource guiding the convergence towards low-energy states [22]. Additionally, more advanced schemes utilizing optical non-linearities have also been implemented to include four-body interaction terms [23, 24]. Furthermore, they have been used to study the thermodynamics of Ising systems [25, 26], and for tackling various NP-hard problems such as the Number Partitioning Problem [27, 28], Max-Cut [29] and Knapsack problems [30]. While several computationally interesting problems can already be handled by SPIMs [31], the standard SPIM configuration is restricted to Mattis-type interactions [32], a fact that poses a significant limitation. Various approaches, such as vector-matrix multiplication [33], time division [34] and wavelength multiplexing schemes [35], have been suggested to overcome this constraint, by decomposing the interaction matrix J in a series of Mattis models that can then be independently treated with a SPIM, limiting however, the scalability of the system or significantly increasing the execution time of the optimization process.

In this letter, we formulate an alternative spin-interaction encoding for SPIMs that allows for manipulating arbitrary coupling matrices J. Moreover, this method offers scaling advantages in the case of sparse models. Each term $J_{ij}\sigma_i\sigma_j$ of the Ising Hamiltonian is directly encoded as an element of the phase matrix imprinted on the SLM. We validate our method by directly comparing the theoretically expected to the experimentally measured energy levels of a random spin glass Hamiltonian. Then, we apply this scheme to the Graph Partitioning Problem (GPP), showcasing persistently high-quality solutions for problems of arbitrary sparsity. Finally, we expand our approach to the weighted version of the problem, where the coupling matrix J_{ij} elements

take random positive values.

SPIN-PRODUCT-ENCODING OF ARBITRARY ISING HAMILTONIANS

The existing approach to treating Ising models by SPIMs [18] is based on the correspondence of each SLM pixel, a square with width L, to a spin of the Ising model. Let ζ_i be the amplitude of the incoming electric field to the i-th pixel. The SPIM adds a phase φ_i to the electric field. In the simplest approach, this phase is equal to the value of the corresponding spin, but in principle, it may also contain a constant angle θ_i different for each pixel, i.e., $\varphi_i = \sigma_i e^{i\theta_i}$. For the moment, let us assume that we do not use this extra freedom, i.e., $\theta_i = 0$.

For a given spin configuration $\sigma = \{\sigma_1, \ldots, \sigma_N\}$, the corresponding SLM pixels are set to the appropriate phases. The electric field is then Fourier-transformed by a Fourier lens and projected onto the camera. The captured image is the intensity

$$\widetilde{I}\left(\vec{k}\right) = \mathrm{sinc}^2 \frac{k_x L}{2} \mathrm{sinc}^2 \frac{k_y L}{2} \sum_{i,j} \zeta_i \zeta_j^* \varphi_i \varphi_j^* e^{i\vec{k}\cdot(\vec{r}_i - \vec{r}_j)}. \quad (1)$$

The Ising energy is computed as $H(\sigma) = -\widetilde{I}(\vec{0})$, which yields

$$H(\sigma) = -\operatorname{Re} \sum_{i,j} \zeta_i \zeta_j^* \varphi_i \varphi_j^*. \quad (2)$$

The relation between the phases φ_i and the values of the spins implies that the above Hamiltonian is simply the Ising model

$$H(\sigma) = -\sum_{i,j} J_{ij} \sigma_i \sigma_j, \quad (3)$$

with couplings

$$J_{ij} = \operatorname{Re} \zeta_i \zeta_j^*. \quad (4)$$

Further assuming that the phase of the incoming electric field is uniform, i.e., the amplitudes ζ_i are real, this corresponds to the Mattis model $H \propto -\sum_{i,j} \zeta_i \zeta_j \sigma_i \sigma_j$. This approach results in a restricted class of models since the number of controllable parameters (the amplitudes ζ_i) is linear to the number of spins N. In contrast, the general Ising model contains independent couplings whose number scales as $O(N^2)$.

The main idea of this work is the following: Being able to encode individual binary spins using an SLM implies that one can also encode products of spins, as they also take values ± 1. It follows that SLM pixels can be assigned to products of spins instead of single spins. In what follows, this method is called *Spin-Product-Encoding* (SPE). This idea can be beneficial for Ising models defined on sparse graphs because only non-zero couplings are allocated to SLM pixels.

Let the Ising model be defined on an interaction graph $G = (V, E)$, with $N = |V|$ spins identified by the index i, and non-vanishing couplings only between the pairs of spins $(i, j) \in E$. The Ising Hamiltonian reads

$$H(\sigma) = -\sum_{(i,j)\in E} J_{ij} \sigma_i \sigma_j, \quad (5)$$

We match each pixel to a pair of spins $(i, j) \in E$. These pixels add a phase delay to the electric field equal to $\varphi_{ij} = \sigma_i \sigma_j$. Furthermore, we employ an ancillary spin σ_0, assigned to one SLM pixel. Assuming real incoming electric field amplitudes ζ_{ij} and ζ for the set of pixels assigned to spin pairs and the ancillary spin, respectively, the Hamiltonian in Eq. (2) reads

$$\widetilde{H}(\sigma_0, \sigma) = -\zeta\sigma_0 \sum_{(i,j)\in E} \zeta_{ij} \sigma_i \sigma_j - \left(\sum_{(i,j)\in E} \zeta_{ij} \sigma_i \sigma_j \right)^2 - \zeta^2. \quad (6)$$

The first term of Eq. (6) is proportional to the Hamiltonian of a generic Ising model with couplings

$$J_{ij} = \zeta \zeta_{ij}. \quad (7)$$

The last term is just a non-dynamical constant. The second term, however, involves four-spin couplings. To cancel the four-spin contributions, we obtain two measurements of $\widetilde{H}$ for the two values of σ_0 while keeping the rest of the spins fixed. Then we have

$$H(\sigma) = \frac{\widetilde{H}(+1, \sigma) - \widetilde{H}(-1, \sigma)}{2} = -\zeta \sum_{(i,j)} \zeta_{ij} \sigma_i \sigma_j, \quad (8)$$

i.e., the generic Ising model Hamiltonian with couplings given by Eq. (7).

The treatment of an arbitrary Ising model with this method requires the modulation of the incoming laser beam since the couplings J_{ij} are determined by the amplitudes ζ_{ij}. We can trade amplitude modulation for phase modulation, which can be achieved by the SLM pixels, as shown in the supplementary material [36].

Furthermore, this encoding scheme allows for flexible partitioning of the Ising couplings, enabling a trade-off between space and time complexity, as discussed in the supplementary material [36].

EXPERIMENTAL CONFIGURATION

The optical configuration of our SPIM is depicted in Fig. 1(a). Light from a stabilised continuous wave He-Ne

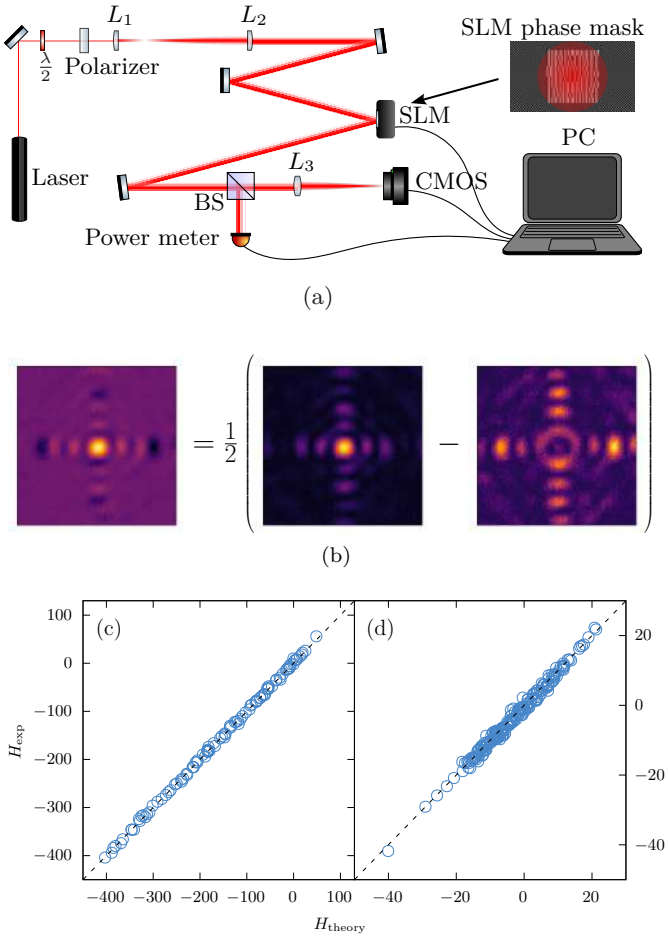

Figure 1. (a) Schematic of the experimental setup. (b) Images captured by the camera corresponding to the energy computation according to Eq. (8) for a given spin configuration. (c,d) The experimental energies versus the theoretical ones for a sparse ferromagnetic system (c), and a sparse spin glass system (d).

laser is expanded 10 times and impinges on a reflective spatial light modulator (SLM), Holoeye PLUTO-2.1-NIR. In addition to the dynamic spin phase encoding, a static holographic grating is applied to separate the 1st diffracted order from the unmodulated reflected light. The beam is then focused on a high QE Peltier-cooled CMOS camera, Atik Camera ACICS 7.1. The region of interest of the camera is selected around the 1st order of the holographic grating. We record the modulated laser light after impinging the SLM by placing a power meter (PM) to compensate for camera intensity variations due to laser fluctuations and SLM flickering. Then, the recorded image is normalized with the corresponding power value captured by the PM.

Since the SPIM Ising energy is directly related to the recorded light intensity, a calibration process takes place at the beginning of each experimental run by comparing the experimental energy values to the corresponding theoretical ones. First, we sample random spin configurations uniformly distributed throughout the configuration space. Then, we perform a linear fit between the experimental and theoretical energies and obtain the normalization factor and offset for the experimental ones.

Figure 1(c,d) presents the results of the above calibration process for a sparse ferromagnetic and a sparse spin glass system, showing a perfect match between the corresponding experimental and theoretical energies. The distribution of J_{ij} for the case of the sparse ferromagnet is $P(J) = p\delta(J-1) + (1-p)\delta(J)$, while for the spin glass is $P(J) = p\mathbb{1}_{[-1,1]}(J) + (1-p)\delta(J)$, where $\mathbb{1}_{[-1,1]}$ is the uniform distribution in the interval $[-1, 1]$ and p is the *edge probability* of the graph. We used $p = 0.05$ for both cases in Fig. 1(c,d).

GRAPH PARTITIONING

We apply our approach to the GPP, a member of the NP-hard complexity class. This problem is suitable for demonstrating our method's ability to tackle problems with arbitrary sparsity, overcoming the limitations of existing solutions.

The GPP considers an undirected graph $G = (V, E)$ with an even number $N = |V|$ of vertices. The problem asks for a partition of the set V into two subsets $\mathcal{A}$ and $\mathcal{B} = V \setminus \mathcal{A}$ of equal size, such that the number of edges connecting the two subsets is minimized. The *cut-set* is defined as $\mathcal{C} = \{(i, j) \in E \mid i \in \mathcal{A}, j \in \mathcal{B}\}$, and the *cost* for the GPP as the size of the cut-set $C = |\mathcal{C}|$.

Following [3], the GPP can be mapped to the Ising model using the following Hamiltonian

$$H = aH_a + bH_b, \text{ with} \tag{9a}$$

$$H_a = \left(\sum_{i=1}^{N} \sigma_i\right)^2, \tag{9b}$$

$$H_b = \sum_{(i,j)\in E} \frac{1 - \sigma_i\sigma_j}{2}. \tag{9c}$$

The term H_a penalises partitions into subsets of unequal size, while the term H_b corresponds precisely to the cost C. In this work, we use $a = b = 1$. The mapping between the GPP and the above Hamiltonian works as follows. Once the ground state of the Hamiltonian is found, the graph vertices are assigned to the two subsets according to the sign of the corresponding spins, i.e., $\mathcal{A} = \{i \mid \sigma_i = 1\}$ and $\mathcal{B} = \{i \mid \sigma_i = -1\}$.

To solve the GPP, we implement the Hamiltonian of Eq. (9) on the SPIM using a hybrid encoding scheme. We implement H_a using the existing approach found in [18] and H_b using the new encoding scheme proposed in this work. The term H_a has homogeneous couplings and is

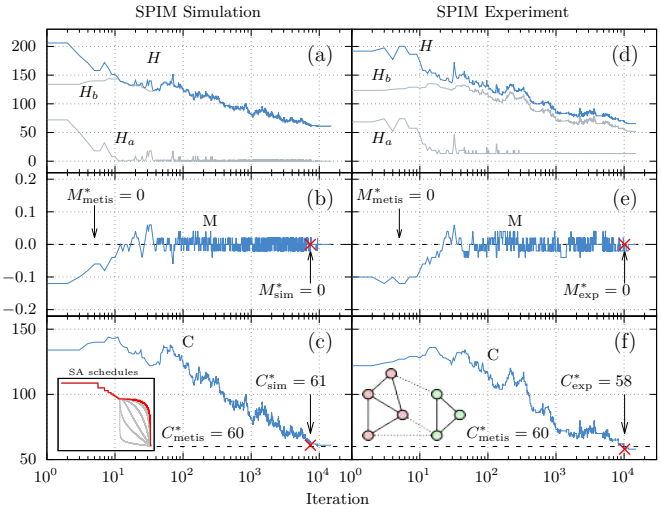

Figure 2. Simulated (a,b,c) and experimental (d,e,f) results for a graph partitioning Hamiltonian, showing the total (H) and individual (H_a and H_b) energies (a,d), the magnetization (b,e), and the cost (c,f) as a function of the optimization iterations. The inset in (c) depicts the different annealing schedules that were tried (grey) and the one that was used (red). Inset in (f) shows a graphical representation of a toy example of the graph partitioning problem. The dashed lines in (b,c,e,f) represent the solution obtained by the METIS package.

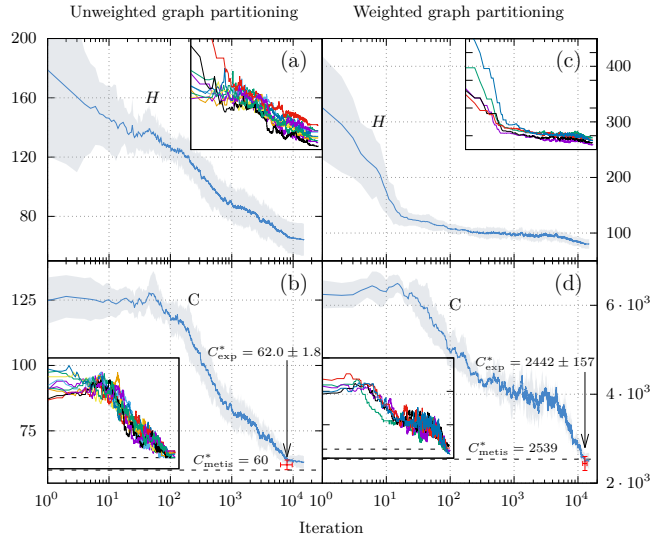

Figure 3. Multiple runs obtained experimentally for the Unweighted Graph Partitioning Problem (GPP) (a,b) and the Weighted Graph Partitioning Problem (WGPP) (c,d). Both instances have $N = 100$ and $p = 0.05$. Each column shows the total energy (a,c) and the cost (b,d) for 10 (a,b) and 5 (c,d) different initial spin configurations, showing that the SPIM systematically finds solutions with comparable cost to that obtained by METIS. The insets show all individual runs.

fully connected, which is the ideal case for the existing method, as the number of pixels required scales linearly with the number of spins. On the other hand, the term H_b has sparse couplings, which is the ideal case for the new encoding scheme, as the number of pixels required scales linearly with the number of edges of the interaction graph $|E|$. The two terms are computed sequentially in two steps, and the final energy is obtained by summing the two contributions.

Figure 2 illustrates the simulation and experimental results by applying the proposed optical encoding scheme to the GPP. In this case, we considered a graph with $N = 100$ vertices and an edge density of $p = 0.05$. In addition, we used the software package METIS [37] that implements a state-of-the-art algorithm for GPP [38] to compare with our results.

Both the simulation and experimental results are obtained by running a Simulated Annealing (SA) optimization algorithm to minimize the Hamiltonian in Eq. (9). Various annealing schedules were implemented and tested, as shown in the inset of Fig. 2(c). Among them, the *linear additive* (red line) was selected because we empirically found that it provides the best results.

The simulation results were obtained by developing a model based on Fresnel diffraction theory. The model incorporates the optical encoding of the spins and interactions among them as additional phase masks, mimicking the role of the SLM.

The top panel of Fig. 2(a,c) presents the evolution of the total energy, H, and the two individual energy terms, H_a and H_b, as defined in Eq. (9), as a function of the iteration number of the SA algorithm. After several iterations, the SA algorithm converges to a minimum for the total energy. The cost C, obtained by the simulated and experimental solution of the GPP, is shown in Fig. 2(c,f). The red cross on the graphs denotes the *optimal cost* C^* obtained from the minimization procedure, while the dashed line shows the solution of the METIS package. In particular, the simulation and experimental values are $C^*_{\text{sim}} = 61$ and $C^*_{\text{exp}} = 58$, respectively, while $C^*_{\text{metis}} = 60$.

The simulation above and experimental values were obtained by selecting the spin configuration that minimizes the cost while having vanishing magnetization, see Fig. 2(b,e), implying that we are only interested in solutions that split the graph into two equal subsets, as required by the GPP. Both values are comparable to the METIS solution. The experimental value slightly outperforms both simulation and METIS, highlighting the versatility of our method to address arbitrary graphs using the SPIM optical architecture.

For the graph appearing in Fig. 2, we performed 10 distinct experimental trials starting from different initial spin configurations to showcase the systematic convergence to a low-energy state. The average energy and cost obtained from this process are depicted in Fig. 3(a,b). The average optimal cost is $C^*_{\text{exp}} = 62 \pm 1.8$ and is depicted by the red cross. For reference, the METIS solution is $C^*_{\text{metis}} = 60$.

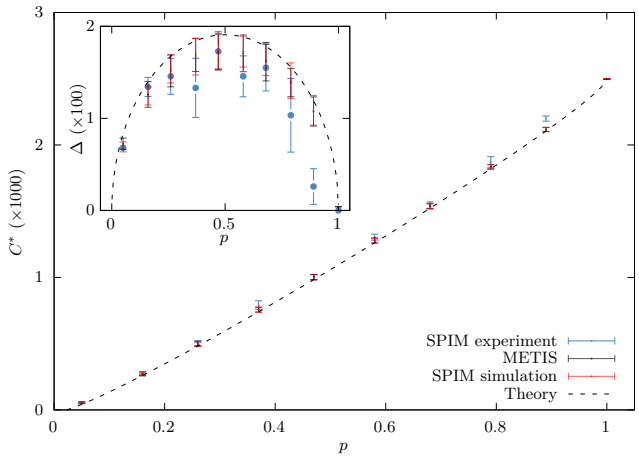

Figure 4. Optimal cost C^* of the GPP for varying edge probability p from the SPIM experiment, METIS, SPIM simulation, and theory. The SPIM experiment points have been averaged over 20 random problem instances. The METIS and SPIM simulation points have been averaged over 50 problem instances. Inset: the improvement due to optimization $\Delta(p)$.

To demonstrate the ability of the method to handle general Ising models as well, we proceed with the *weighted* version of the GPP (WGPP). In WGPP, the graph is defined as the triplet $G = (V, E, w)$, where $w_{ij} \in \mathbb{N}^*$ is the weight of the edge (i, j). The objective in this case is to partition the vertices into two subsets, minimizing the cost $C = \sum_{(i,j)\in\mathcal{C}} w_{ij}$. For the WGPP the H_b term of Eq. (9) takes the form

$$H_b^{(\text{weighted})} = \sum_{(i,j)\in E} w_{ij} \frac{1 - \sigma_i\sigma_j}{2}. \tag{10}$$

Following the theoretical description provided in the supplementary material [36], we encode the w_{ij} using two adjacent pixels with $\theta_{ij} = \arccos(w_{ij}/w_{\max})$. Figure 3(c,d) depicts the results obtained from a graph with $|V| = 100$, $p = 0.05$. The w_{ij} are uniformly random integers in the range $[1, w_{\max}]$ with $w_{\max} = 100$.

The results are consistent with the unweighted case. Our method successfully finds solutions comparable to those obtained by the METIS package.

Apart from the method's ability to address arbitrary Ising models, it also scales well for sparse graphs: The number of free parameters of the Ising model scales as $O(N^2)$. However, for sparse graphs, many of these parameters are zero; hence, they do not contribute to the system's energy. By encoding only the non-zero couplings, our method requires an SLM with a number of pixels that scales as $O(|E|)$.

To quantify the method's ability to address sparse graphs, we performed experiments to solve the GPP for various values of the edge probability p. The results of those experiments are presented in Fig. 4 where the optimal cost is plotted as a function of p. In addition to the experimental realizations, Fig. 4 includes the results obtained from the simulation model, the METIS package, and the theoretical prediction for the optimal cost obtained by the replica method [1, 39], namely

$$C^*(p) = \frac{N^2}{4}p - 0.38N^{3/2}\sqrt{p(1-p)}. \tag{11}$$

Since the first term in Eq. (11) does not depend on the optimization [39], we present in the inset of Fig. 4 the improvement due to optimization $\Delta = pN^2/4 - C^*$. Notably, the agreement between the simulation and METIS is perfect and almost matches the theoretical optimum. The agreement of the experimental results with the rest is satisfactory, especially for sparse graphs, i.e., small values of p.

CONCLUSION

We have introduced and experimentally validated a novel encoding scheme for SPIMs, effectively augmenting their functionality and applicability without sacrificing scalability or speed. This scheme directly encodes any Ising Hamiltonian's full interaction matrix J through a two-step iteration process. The method requires $O(|E|)$ pixels; hence, it is particularly advantageous in the case of sparse graphs in comparison to existing implementations [36]. In the worst case (fully connected graphs), $|E| \sim N^2$, the method becomes equivalent to the worst case (full-rank graphs) of existing implementations. Nevertheless, its ability to trade space complexity for time complexity by partitioning the Ising couplings allows the method to be implemented on any SLM, regardless of its size, unlike existing approaches.

We applied this method to our SPIM instance to solve the unweighted and weighted graph partitioning problems, demonstrating comparable quality solutions with a GPP-specific algorithm for a varying degree of sparsity. Notably, our method can be easily expanded to Hamiltonians having p-spin interactions [36], which play an equally important role in combinatorial optimization [40]. As the SLM technology enters the MHz range [41], SPIMs are poised to become a powerful and versatile technology for real-world applications.

We acknowledge useful discussions with C. Conti and N. Berloff. J.S., A.A., G.P., and S.T. acknowledge support from HORIZON EIC-2022-PATHFINDERCHALLENGES-01 HEISINGBERG Project 101114978. J.S., A.A., and S.T. acknowledge support from HoloCIM (CODEVELOP-ICTHEALTH/0322/0047).

* isakellariou@q.ubitech.eu

† aaskitopoulos@q.ubitech.eu
‡ stsintzos@q.ubitech.eu

[1] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin glass theory and beyond: An Introduction to the Replica Method and Its Applications*, Vol. 9 (World Scientific Publishing Company, 1987).
[2] M. Mézard and A. Montanari, *Information, physics, and computation* (Oxford University Press, 2009).
[3] A. Lucas, Frontiers in Physics **2**, 10.3389/fphy.2014.00005 (2014), publisher: Frontiers.
[4] P. Scholl, M. Schuler, H. J. Williams, A. A. Eberharter, D. Barredo, K.-N. Schymik, V. Lienhard, L.-P. Henry, T. C. Lang, T. Lahaye, A. M. Läuchli, and A. Browaeys, Nature **595**, 233 (2021), publisher: Nature Publishing Group.
[5] K. Kim, M.-S. Chang, S. Korenblit, R. Islam, E. E. Edwards, J. K. Freericks, G.-D. Lin, L.-M. Duan, and C. Monroe, Nature **465**, 590 (2010), publisher: Nature Publishing Group.
[6] X.-s. Ma, B. Dakic, W. Naylor, A. Zeilinger, and P. Walther, Nature Physics **7**, 399 (2011), publisher: Nature Publishing Group.
[7] M. W. Johnson, M. H. S. 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. S. Truncik, S. Uchaikin, J. Wang, B. Wilson, and G. Rose, Nature **473**, 194 (2011), publisher: Nature Publishing Group.
[8] N. G. Berloff, M. Silva, K. Kalinin, A. Askitopoulos, J. D. Töpfer, P. Cilibrizzi, W. Langbein, and P. G. Lagoudakis, Nature Materials **16**, 1120 (2017), publisher: Nature Publishing Group.
[9] H. Ohadi, A. Ramsay, H. Sigurdsson, Y. del Valle-Inclan Redondo, S. Tsintzos, Z. Hatzopoulos, T. Liew, I. Shelykh, Y. Rubo, P. Savvidis, and J. Baumberg, Physical Review Letters **119**, 067401 (2017), publisher: American Physical Society.
[10] K. P. Kalinin, A. Amo, J. Bloch, and N. G. Berloff, Nanophotonics **9**, 4127 (2020), publisher: De Gruyter.
[11] S. Alyatkin, J. Töpfer, A. Askitopoulos, H. Sigurdsson, and P. Lagoudakis, Physical Review Letters **124**, 207402 (2020), publisher: American Physical Society.
[12] W. A. Borders, A. Z. Pervaiz, S. Fukami, K. Y. Camsari, H. Ohno, and S. Datta, Nature **573**, 390 (2019), publisher: Nature Publishing Group.
[13] F. Cai, S. Kumar, T. Van Vaerenbergh, X. Sheng, R. Liu, C. Li, Z. Liu, M. Foltin, S. Yu, Q. Xia, J. J. Yang, R. Beausoleil, W. D. Lu, and J. P. Strachan, Nature Electronics **3**, 409 (2020), publisher: Nature Publishing Group.
[14] N. Shukla, A. Parihar, E. Freeman, H. Paik, G. Stone, V. Narayanan, H. Wen, Z. Cai, V. Gopalan, R. Engel-Herbert, D. G. Schlom, A. Raychowdhury, and S. Datta, Scientific Reports **4**, 4964 (2014), publisher: Nature Publishing Group.
[15] P. A. Merolla, J. V. Arthur, R. Alvarez-Icaza, A. S. Cassidy, J. Sawada, F. Akopyan, B. L. Jackson, N. Imam, C. Guo, Y. Nakamura, B. Brezzo, I. Vo, S. K. Esser, R. Appuswamy, B. Taba, A. Amir, M. D. Flickner, W. P. Risk, R. Manohar, and D. S. Modha, Science **345**, 668 (2014), publisher: American Association for the Advancement of Science.
[16] 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, Science **354**, 614 (2016), publisher: American Association for the Advancement of Science.
[17] 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, Science **354**, 603 (2016), publisher: American Association for the Advancement of Science.
[18] D. Pierangeli, G. Marcucci, and C. Conti, Physical Review Letters **122**, 213902 (2019), publisher: American Physical Society.
[19] T. Honjo, T. Sonobe, K. Inaba, T. Inagaki, T. Ikuta, Y. Yamada, T. Kazama, K. Enbutsu, T. Umeki, R. Kasahara, K.-i. Kawarabayashi, and H. Takesue, Science Advances **7**, eabh0952 (2021), publisher: American Association for the Advancement of Science.
[20] W. Sun, W. Zhang, Y. Liu, Q. Liu, and Z. He, Optics Letters **47**, 1498 (2022), publisher: Optica Publishing Group.
[21] D. Pierangeli, G. Marcucci, and C. Conti, Optica **7**, 1535 (2020), publisher: Optica Publishing Group.
[22] D. Pierangeli, G. Marcucci, D. Brunner, and C. Conti, Nanophotonics **9**, 4109 (2020), publisher: De Gruyter.
[23] S. Kumar, H. Zhang, and Y.-P. Huang, Communications Physics **3**, 1 (2020), publisher: Nature Publishing Group.
[24] S. Kumar, Z. Li, T. Bu, C. Qu, and Y. Huang, Communications Physics **6**, 1 (2023), publisher: Nature Publishing Group.
[25] Y. Fang, J. Huang, and Z. Ruan, Physical Review Letters **127**, 043902 (2021), publisher: American Physical Society.
[26] M. Leonetti, E. Hörmann, L. Leuzzi, G. Parisi, and G. Ruocco, Proceedings of the National Academy of Sciences **118**, e2015207118 (2021), publisher: Proceedings of the National Academy of Sciences.
[27] J. Huang, Y. Fang, and Z. Ruan, Communications Physics **4**, 1 (2021), publisher: Nature Publishing Group.
[28] A. Prabhakar, P. Shah, U. Gautham, V. Natarajan, V. Ramesh, N. Chandrachoodan, and S. Tayur, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences **381**, 20210409 (2022), publisher: Royal Society.
[29] X. Ye, W. Zhang, S. Wang, X. Yang, J. Du, and Z. He, in *CLEO 2023 (2023), paper JTh2A.32* (Optica Publishing Group, 2023) p. JTh2A.32.
[30] H. Yamashita, K.-i. Okubo, S. Shimomura, Y. Ogura, J. Tanida, and H. Suzuki, Physical Review Letters **131**, 063801 (2023), publisher: American Physical Society.
[31] R. Z. Wang, J. S. Cummins, M. Syed, N. Stroev, G. Pastras, J. Sakellariou, S. Tsintzos, A. Askitopoulos, D. Veraldi, M. C. Strinati, S. Gentilini, D. Pierangeli, C. Conti, and N. G. Berloff, Efficient Computation Using Spatial-Photonic Ising Machines: Utilizing Low-Rank and Circulant Matrix Constraints (2024), arXiv:2406.01400 [cond-mat, physics:physics].
[32] D. C. Mattis, Physics Letters A **56**, 421 (1976).
[33] J. Ouyang, Y. Liao, Z. Ma, D. Kong, X. Feng, X. Zhang, X. Dong, K. Cui, F. Liu, W. Zhang, and Y. Huang, Communications Physics **7**, 1 (2024), publisher: Nature Publishing Group.

[34] S. Wang, W. Zhang, X. Ye, and Z. He, Applied Optics **63**, 2973 (2024), publisher: Optica Publishing Group.
[35] L. Luo, Z. Mi, J. Huang, and Z. Ruan, Science Advances **9**, eadg6238 (2023), publisher: American Association for the Advancement of Science.
[36] See Supplemental Material at [URL will be inserted by publisher].
[37] KarypisLab/METIS (2024), original-date: 2020-03-02T03:37:17Z.
[38] D. LaSalle and G. Karypis, in *2016 45th International Conference on Parallel Processing (ICPP)* (2016) pp. 236–241, iSSN: 2332-5690.
[39] Y. Fu and P. W. Anderson, Journal of Physics A: Mathematical and General **19**, 1605 (1986).
[40] M. K. Bashar and N. Shukla, Scientific Reports **13**, 9558 (2023), publisher: Nature Publishing Group.
[41] J. Park, B. G. Jeong, S. I. Kim, D. Lee, J. Kim, C. Shin, C. B. Lee, T. Otsuka, J. Kyoung, S. Kim, K.-Y. Yang, Y.-Y. Park, J. Lee, I. Hwang, J. Jang, S. H. Song, M. L. Brongersma, K. Ha, S.-W. Hwang, H. Choo, and B. L. Choi, Nature Nanotechnology **16**, 69 (2021), publisher: Nature Publishing Group.

Supplemental material: Encoding arbitrary Ising Hamiltonians on Spatial Photonic Ising Machines

Jason Sakellariou,[1,2] Alexis Askitopoulos,[1,2] Georgios Pastras,[1] and Symeon I. Tsintzos[1,2]

[1] *QUBITECH, Thessalias 8, GR 15231 Chalandri, Athens, Greece*
[2] *UBITECH Ltd., 95B Archiepiskopou Makariou, CY 3020 Limassol Cyprus*

TRADING AMPLITUDE MODULATION FOR PHASE MODULATION

The treatment of an arbitrary Ising model requires the modulation of the incoming laser beam, since the couplings J_{ij} are determined by the amplitudes ζ_{ij}. Here we show how we can trade amplitude modulation for phase modulation. This is convenient since the phase modulation can be achieved by a single SLM. This can be done by doubling the set of pixels corresponding to the spin couplings. We distinguish the pixels corresponding to the same spin coupling using the upper index $(\pm)$. It is essential to take advantage of the extra spin-independent phase delay that can be introduced by the SLM pixels. Namely, we introduce them so that obey $\theta_{ij}^{+} = -\theta_{ij}^{-} \equiv \theta_{ij}$. In other words, $\varphi_{ij}^{\pm} = \sigma_i \sigma_j e^{\pm i\theta_{ij}}$. The last part of the SLM is again composed by one pixel that causes no phase delay to the electric field. No modulation of the incoming electric field is applied, i.e. $\zeta_{ij}^{+} = \zeta_{ij}^{-} = \zeta$. It is a matter of simple algebra to show that the Hamiltonian in Eq. (7) in the main text reads

$$\widetilde{H}(\sigma_0, \sigma) = -\zeta^2 \sigma_0 \sum_{(i,j)\in E} \cos\theta_{ij} \sigma_i \sigma_j - 4\zeta^2 \left(\sum_{(i,j)\in E} \cos\theta_{ij} \sigma_i \sigma_j \right)^2 - \zeta^2. \tag{S.1}$$

Repeating the same procedure as described in the main text, we obtain two measurements of $\widetilde{H}$, then using Eq. (8) in the main text we find

$$H(\sigma) = -\zeta^2 \sum_{(i,j)\in E} \cos\theta_{ij} \sigma_i \sigma_j, \tag{S.2}$$

which is the generic Ising model with couplings

$$J_{ij} = \zeta^2 \cos\theta_{ij}, \tag{S.3}$$

i.e. we managed to tune all couplings J_{ij} without modulating the laser beam, but solely with the use of the SLM, via the choice of appropriate values for the angles θ_{ij}.

EXTENDING THE ENCODING SCHEME TO THE p-SPIN MODEL

Let us consider a model with p-spin interactions, i.e.

$$H(\sigma) = -\sum_{1\le i_1<\cdots<i_p\le N} J_{i_1\ldots i_p} \sigma_{i_1} \ldots \sigma_{i_p}. \tag{S.4}$$

This Hamiltonian can be encoded in the same way as the Hamiltonian of Eq. (3) in the main text that contains only 2-spin interactions. Each pixel is associated with a unique product of p spins, $\sigma_1\sigma_2\ldots\sigma_p$, that has a non-vanishing coefficient in the Hamiltonian (S.4). Let the incoming field amplitude to this pixel be $\zeta_{i_1,i_2,\ldots,i_p}$. We also introduce an ancillary spin σ_0 that is associated with a single pixel of the SLM, where the incoming electric field amplitude is ζ. Finally, we configure the SLM pixels to add a phase delay to the electric field equal to $\sigma_1\sigma_2\ldots\sigma_p$. It directly follows that the cost functional, i.e., the intensity measured at the center of the Fourier space, reads

$$\widetilde{H}(\sigma_0, \sigma) = -\zeta\sigma_0 \sum_{1\le i_1<\cdots<i_p\le N} \zeta_{i_1\ldots i_p} \sigma_{i_1} \ldots \sigma_{i_p} - \left(\sum_{1\le i_1<\cdots<i_p\le N} \zeta_{i_1\ldots i_p} \sigma_{i_1} \ldots \sigma_{i_p} \right)^2 - \zeta^2, \tag{S.5}$$

Fixing the amplitudes so that

$$J_{i_1,i_2,\ldots,i_p} = \zeta \zeta_{i_1,i_2,\ldots,i_p} \tag{S.6}$$

allows the encoding of the Hamiltonian (S.4) via two iterations, using Eq. (8) in the main text.

The incoming amplitude modulation can be avoided by doubling the size of the SLM via the same mechanism as the one described in the previous section. Let us note that a model with local magnetic fields of the form

$$H(\sigma) = \sum_{(i,j)\in E} J_{ij}\sigma_i\sigma_j - \sum_{i=1}^{N} h_i\sigma_i \tag{S.7}$$

can be encoded similarly by treating the local field terms as p-spin interactions with $p = 1$.

The generic model containing p-spin interactions has $\binom{N}{p} \sim N^p$ independent couplings. Thus, it is clear that its encoding in a SPIM via SPE requires an SLM with that many pixels. As for the $p = 2$ case, the technique will have better scaling properties in the case of models with sparse tensors J, as only the non-vanishing elements are allocated.

SPACE AND TIME COMPLEXITY

The two main resources that come into play in the context of SPIMs are the number of SLM pixels and the number of time steps required to determine the energy. Here, one SLM/CMOS cycle is considered as a single time step. In principle, loading an array of $O(n)$ elements onto the SLM requires at least $O(n)$ time steps due to the memory copy operation. However, in the context of SPIMs, the memory copy operation is orders of magnitude faster than the SLM rise/fall time, which is of the order of milliseconds. Therefore, the computation time is dominated by the SLM rise/fall time; hence, we consider the memory copy operation as instantaneous. We refer to these as the space and time complexity, respectively. This section discusses our method's space and time complexity and compares it to existing methods.

At the time of writing, the only available single-SLM methods for encoding arbitrary Ising models on SPIMs rely on *time-division multiplexing* (TDM) [1] and *wavelength-division multiplexing* (WDM) [2]. Both methods decompose the coupling matrix into a sum of rank-1 matrices and encode these simpler models separately—either at different time steps (TDM) or different wavelengths (WDM). Moreover, in both these approaches, the SLM pixels are mapped to individual spins.

In TDM, each of the simpler models is encoded separately in the SLM, and its contribution to the Ising Hamiltonian is measured independently. This kind of encoding allows the use of a smaller SLM, as it only needs to encode one of the simpler component models at a time. However, this advantage comes at the cost of multiple measurements—one for each component model—equal to the rank of the coupling matrix. WDM follows a similar decomposition approach but encodes all component models simultaneously in the SLM, each at a different wavelength. This kind of multiplexing enables the total Ising Hamiltonian contribution to be measured in a single step. In essence, TDM and WDM achieve the same goal, with WDM effectively trading time complexity for space complexity.

In contrast, our method, *Spin-product-encoding* (SPE), takes a different approach. Instead of decomposing the coupling matrix, it directly encodes only nonzero couplings, mapping SLM pixels to spin products. It always requires two time steps to determine the energy. For an Ising model with N spins and random couplings—such a model naturally is fully connected ($|E| = N(N-1)/2$) and full rank ($\mathrm{rank}(J) = N$)—SPE appears to have a space complexity of $N(N-1)/2$ and a time complexity of 2, similar to WDM. In fact, the two methods are identical up to a factor of two, trading time complexity for space complexity.

However, SPE offers a key flexibility advantage: Since coupling matrix elements (and thus the Ising Hamiltonian terms) are directly encoded in the SLM, SPE can naturally incorporate the core idea of TDM. One can partition the couplings into subsets, encode only one subsets at a time, and sum the measurements to obtain the total energy. In other words, SPE enables a trivial decomposition of the Ising Hamiltonian into any set of components containing a subset of the original couplings. Consequently, the naive assumption that SPE always requires $N(N-1)/2$ space complexity is incorrect—one can trade space complexity for time complexity, achieving a space complexity as low as 1.

The space and time complexity of the three methods are summarized in the table below. Notice that space complexity has to be doubled in all methods if one desires to trade amplitude modulation for phase modulation as described in Section

	Space complexity	**Time complexity**
TDM [1]	N	$\mathrm{rank}(J)$
WDM [2]	$N\,\mathrm{rank}(J)$	1
SPE (minimum time)	$\lvert E\rvert$	2
SPE (minimum space)	1	$2\lvert E\rvert$

The above arguments lead us to several conclusions:

1. Since SPE allows trading space complexity for time complexity, the product of time and space complexities is a useful measure of efficiency. This is supported by the fact that TDM can be viewed as WDM with time complexity exchanged for space complexity.
   (a) The product of time and space complexities for TDM and WDM is $N\,\mathrm{rank}(J)$
   (b) The product of time and space complexity for SPE is $|E|$
2. The above complexities suggest that SPE is advantageous for sparse graphs, whereas TDM or WDM is for low-rank graphs. This fact makes these approaches complementary since low-rank graphs tend to be fully connected, while sparse graphs are typically full-rank.
3. Both WDM and SPE (configured for minimal time complexity) have comparable time complexity, $O(1)$. However, in terms of space complexity:
   (a) SPE is preferable for sparse graphs whenever $|E| < N\,\mathrm{rank}(J)$.
   (b) SPE allows space complexity to be traded for time complexity. If $|E|$ exceeds the available SLM size, the couplings can be partitioned into subsets smaller than the SLM capacity and processed sequentially—an option unavailable in WDM.

[1] S. Wang, W. Zhang, X. Ye, and Z. He, Applied Optics **63**, 2973 (2024), publisher: Optica Publishing Group.
[2] L. Luo, Z. Mi, J. Huang, and Z. Ruan, Science Advances **9**, eadg6238 (2023), publisher: American Association for the Advancement of Science.