

Supplementary Material for Discretized Spin Interactions Enhance Higher-Order Analog Ising Machines

Robbe De Prins^{*1,2}, Guy Van der Sande², Peter Bienstman¹, and Thomas Van Vaerenbergh³

¹Photonics Research Group, Ghent University – imec, Technologiepark-Zwijnaarde 126, 9052 Gent, Belgium

²Applied Physics Research Group, Vrije Universiteit Brussel, Pleinlaan 2, 1050 Brussels, Belgium

³Large-Scale Integrated Photonics Lab, HPE Labs, HPE Belgium, Diegem, Belgium

Supplementary Note 1: Embedding SAT problems in a higher-order Ising model

The goal of a SAT problem is to find an assignment to a set of N Boolean variables such that a given logical formula evaluates to **True**. Every SAT problem can be expressed in conjunctive normal form (CNF) as a conjunction (logical AND, $\wedge$) of clauses, each of which is a disjunction (logical OR, $\vee$) of exactly L literals. A literal 1 is either a variable x_i or its negation $\neg x_i$. In the case of 3-SAT, we have $L = 3$. The derivation below trivially extends to the case where $L > 3$.

A 3-SAT problem with C clauses can be expressed as:

$$\bigwedge_{i=1}^C (l_{i1} \vee l_{i2} \vee l_{i3}), \quad (\text{S.1})$$

where l_{ij} denotes the j -th literal in the i -th clause.

To embed a 3-SAT problem into the higher-order Ising Hamiltonian defined in Eq. (2) from the main text, we first convert the logical formula into a real-valued Polynomial Unconstrained Binary Optimization (PUBO) energy function of binary variables $x_k \in \{0, 1\}$:

$$H(\mathbf{x}) = \sum_{i=1}^C (1 - l_{i1})(1 - l_{i2})(1 - l_{i3}), \quad (\text{S.2})$$

^{*}Robbe.DePrins@UGent.be

where

$$l_{ij} = \begin{cases} x_k & \text{if the literal } \mathbf{l}_{ij} \text{ is } \mathbf{x}_k, \\ 1 - x_k & \text{if the literal } \mathbf{l}_{ij} \text{ is } \neg \mathbf{x}_k, \end{cases} \quad (\text{S.3})$$

and $k \in \{1, \dots, N\}$. Note that every unsatisfied clause adds an energy penalty of +1 to Eq. (S.2).

We then express the energy function in terms of Ising spins $\sigma_i = \pm 1$ using the following transformation:

$$x_k = \frac{\sigma_k + 1}{2}. \quad (\text{S.4})$$

Supplementary Note 2: Simulating the analog Ising machine

The temporal evolution of the Ising machine is modeled using the Euler–Maruyama integration method:

$$\mathbf{s}_{t+1} = \mathbf{s}_t + \Delta t (-\mathbf{s}_t + \tanh(\alpha \mathbf{s}_t + \beta_t \mathbf{I}_t)) + \gamma \boldsymbol{\xi}_t, \quad (\text{S.5})$$

where $\mathbf{s}_t$ denotes the vector of spin amplitudes at time t , α is the linear gain, and $\Delta t = 0.01$ is the (dimensionless) time step (which is validated as detailed in the Supplementary Material). The local fields acting on the spins at time t are encoded in the vector $\mathbf{I}_t$, according to one of the local field models from the Methods section. γ denotes the noise strength, while each noise component $\xi_{t,i}$ is sampled independently from a normal distribution with zero mean and standard deviation $\sqrt{\Delta t}$. The interaction strength β_t is gradually increased over time using:

$$\beta_{t+1} = \beta_t + v_\beta \Delta t, \quad (\text{S.6})$$

starting from $\beta_0 = 0$, with v_β denoting the annealing rate.

The initial values of the spin amplitudes are sampled from the same Gaussian distribution as the noise process, i.e. $\gamma \boldsymbol{\xi}_0$. The iterative updates of Eqs. (S.5) and (S.6) continue until the IM either achieves the target energy or completes $10^4/\Delta t$ steps.

The TTS metric is defined as follows:

$$\text{TTS} = \begin{cases} T, & \text{if } P > 0.99, \\ T \frac{\log(0.01)}{\log(1-P)}, & \text{if } 0 < P \leq 0.99, \\ \infty, & \text{if } P = 0, \end{cases} \quad (\text{S.7})$$

where P is the probability of reaching the target state, and T is the (dimensionless) time window over which Eq. (4) is integrated. The reported values correspond to $\min_{0 < T < t_{\max}} \text{TTS}(T, P(T))$, which represents the total time needed to reach the optimal solution of a given problem with 99% probability.

For each problem instance, we perform a grid search over the hyperparameters, as detailed below. Each of the reported TTS values (e.g. in Figs. 1 and 2) corresponds to the optimal hyperparameter configuration identified through this process. For each hyperparameter setting, the IM evolution is repeated 100 times to estimate the TTS.

- Linear gain: $\alpha \in \{-10, -9, \dots, 1\}$
- Annealing speed: $v_\beta \in \{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1\}$
- Noise strength: $\gamma \in \{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\}$

Supplementary Note 3: Extending the local field models with quartic spin interactions

In the Methods section, we introduced multiple methods to incorporate higher-order spin interactions and provided the corresponding local field expressions for the third-order Ising model of Eq. (2), which we use for 3-SAT problems. Here, for completeness, we list the explicit extensions to fourth-order Ising Hamiltonians used for 4-SAT problems. The local field I_i for each method is summarized in Supplementary Table 1.

Supplementary Table 1: Local field expressions I_i for fourth-order Ising models using the different methods introduced in the Methods section. Here, $\langle |s_m| \rangle$ denotes the mean absolute value of the continuously-valued spin amplitudes.


Method	Local field I_i
Baseline	$J_i^{(1)} + \sum_j J_{ij}^{(2)} s_j + \sum_{j<k} J_{ijk}^{(3)} s_j s_k + \sum_{j<k<l} J_{ijkl}^{(4)} s_j s_k s_l$
Rescaling 1	$J_i^{(1)} + \sum_j J_{ij}^{(2)} \frac{s_j}{\langle s_m \rangle} + \sum_{j<k} J_{ijk}^{(3)} \frac{s_j s_k}{\langle s_m \rangle^2} + \sum_{j<k<l} J_{ijkl}^{(4)} \frac{s_j s_k s_l}{\langle s_m \rangle^3}$
Rescaling 2	$J_i^{(1)} \langle s_m \rangle + \sum_j J_{ij}^{(2)} s_j + \sum_{j<k} J_{ijk}^{(3)} \frac{s_j s_k}{\langle s_m \rangle} + \sum_{j<k<l} J_{ijkl}^{(4)} \frac{s_j s_k s_l}{\langle s_m \rangle^2}$
Rescaling 3	$J_i^{(1)} \langle s_m \rangle^2 + \sum_j J_{ij}^{(2)} s_j \langle s_m \rangle + \sum_{j<k} J_{ijk}^{(3)} s_j s_k + \sum_{j<k<l} J_{ijkl}^{(4)} \frac{s_j s_k s_l}{\langle s_m \rangle}$
Rescaling 4	$J_i^{(1)} \langle s_m \rangle^3 + \sum_j J_{ij}^{(2)} s_j \langle s_m \rangle^2 + \sum_{j<k} J_{ijk}^{(3)} s_j s_k \langle s_m \rangle + \sum_{j<k<l} J_{ijkl}^{(4)} s_j s_k s_l$
Spin sign	$J_i^{(1)} + \sum_j J_{ij}^{(2)} \text{sgn}(s_j) + \sum_{j<k} J_{ijk}^{(3)} \text{sgn}(s_j) \text{sgn}(s_k) + \sum_{j<k<l} J_{ijkl}^{(4)} \text{sgn}(s_j) \text{sgn}(s_k) \text{sgn}(s_l)$

Supplementary Note 4: Success rate and time-to-solution for individual problems


Supplementary Fig. 1 shows the maximal success rate achieved on each individual uniform random 3-SAT instance, for each of the local field models from the Methods section. This breakdown highlights the variability across instances that is not visible in the aggregated boxplots of Fig. 3 from the main text. Similarly, Supplementary Fig. 2 shows the maximal success rates for the individual uniform random 4-SAT instances that underly Fig. 4 from the main text. Further, Supplementary Figs. 3 and 4 show the TTS values for the individual 3-SAT and 4-SAT problems and local field methods, underlying Figs. 1 and 2, respectively.

Supplementary Note 5: Distribution of time-to-solution across problem sizes


In Figs. 1 and 2 of the main text, we observed that data points increasingly diverge from the diagonal as problem size grows (indicated by darker dots). To make this trend more explicit, we show the TTS distribution for all local field models from the Methods section, grouped by problem size. As seen in Supplementary Fig. 5 for 3-SAT problems and in Supplementary Fig. 6 for 4-SAT problems, the spin sign




Supplementary Figure 1: Instance-specific success rates for the uniform random 3-SAT instances, underlying the boxplots in Fig. 3. We compare the various methods to incorporate higher-order interactions, as described in the Methods section. Problem names follow the format $uf\{N\}-\{i\}$, where N is the number of spin variables (indicated in the subplot title) and i is the instance identifier (shown along the x-axis).




Supplementary Figure 2: Instance-specific success rates for the uniform random 4-SAT instances, underlying the boxplots in Fig. 4. We compare the various methods to incorporate higher-order interactions, as described in the Methods section (cf. Supplementary Table 1 for explicit local field expressions). Problem names follow the format $uf\{N\}-\{i\}$, where N is the number of spin variables (indicated in the subplot title) and i is the instance identifier (shown along the x-axis).



Supplementary Figure 3: Instance-specific time-to-solution values for the uniform random 3-SAT instances, underlying Fig. 1. We compare the various methods to incorporate higher-order interactions, as described in the Methods section. The dashed line denotes the resolution limit of the TTS metric. All bars surpassing this limit denote $TTS = \infty$, meaning they were not solved by the given method for any of the tested hyperparameter values. Problem names follow the format $uf\{N\}-\{i\}$, where N is the number of spin variables (indicated in the subplot title) and i is the instance identifier (shown along the x-axis).



Supplementary Figure 4: Instance-specific time-to-solution values for the uniform random 4-SAT instances, underlying Fig. 2. We compare the various methods to incorporate higher-order interactions, as described in the Methods section (cf. Supplementary Table 1 for explicit local field expressions). The dashed line denotes the resolution limit of the TTS metric. All bars surpassing this limit denote $TTS = \infty$, meaning they were not solved by the given method for any of the tested hyperparameter values. Problem names follow the format $uf\{N\}-\{i\}$, where N is the number of spin variables (indicated in the subplot title) and i is the instance identifier (shown along the x-axis).




Supplementary Figure 5: Time-to-solution distributions across problem sizes for all local field models from the Methods section. Each boxplot represents TTS values over ten 3-SAT instances of the same size. Outliers are shown as dots. The advantage of the spin sign method over the alternative methods grows with increasing problem size.


method maintains lower TTS values than the alternative methods, with the performance gap widening as the number of spins increases.

Supplementary Note 6: Validation of time step size

As detailed under Eq. (S.5), all results in the main text were obtained using a time step of $dt = 0.01$. To ensure this choice is sufficiently small for accurately integrating the governing equations of the IM (Eq. (4)), we repeat the TTS analysis for 3-SAT problems with a smaller time step of $dt = 0.001$. The results, shown in Supplementary Fig. 7, closely match those in Fig. 1, confirming that $dt = 0.01$ yields reliable integration.



Supplementary Figure 6: Time-to-solution distributions across problem sizes for all local field models from the Methods section. Cf. Supplementary Table 1 for explicit local field expressions. Each boxplot represents TTS values over ten 4-SAT instances of the same size. Outliers are shown as dots. The advantage of the spin sign method over the alternative methods grows with increasing problem size.



Supplementary Figure 7: Comparison of time-to-solution across methods for incorporating higher-order interactions on uniform random 3-SAT instances. The spin sign method of Eq. (9) is compared to (a) the baseline method of Eq. (5) and (b-d) rescaling methods 1-3 of Eqs. (6) to (8). This figure mirrors Fig. 1 from the main text but uses a smaller timestep of $dt = 0.001$ instead of $dt = 0.01$. The similarity with Fig. 1 indicates that $dt = 0.01$ is sufficient for accurately integrating the IM's governing equations (Eq. (4)).

Supplementary Note 7: Aligning interaction strengths with the binary energy landscape

The baseline method, defined in Eq. (5) and Supplementary Table 1, is susceptible to magnitude imbalances between spin interactions of different orders. For instance, if $|s_i| \ll 1$, the linear terms dominate the quadratic terms, which in turn dominate the cubic terms. To address this, the rescaling methods apply continuously varying correction factors to the interaction terms. While this reduces the imbalance, it does not fully eliminate it. In contrast, the spin sign method preserves the exact magnitudes of all interactions as specified in the binary Hamiltonian of Eq. (2).


To formalize this comparison, we introduce the error metrics $\epsilon_i^{(2)}$ and $\epsilon_i^{(3)}$. The metric $\epsilon_i^{(2)}$ quantifies the deviation of the effective ratio between spin couplings and external fields, as experienced by spin i , from the target ratio defined by the higher-order Ising Hamiltonian. Similarly, $\epsilon_i^{(3)}$ measures the deviation of the effective ratio between third-order spin interactions and external fields from the corresponding target ratio. More precisely, we define:

$$\epsilon_i^{(2)} = \frac{\left| \frac{\Gamma^{(2)} \sum_j J_{ij}^{(2)} s_j}{\Gamma^{(1)} J_i^{(1)}} \right|}{\left| \frac{\sum_j J_{ij}^{(2)} \text{sgn}(s_j)}{J_i^{(1)}} \right|} = \left| \frac{\Gamma^{(2)} \sum_j J_{ij}^{(2)} s_j}{\Gamma^{(1)} \sum_j J_{ij}^{(2)} \text{sgn}(s_j)} \right|, \quad (\text{S.8})$$

$$\epsilon_i^{(3)} = \frac{\left| \frac{\Gamma^{(3)} \sum_{j < k} J_{ijk}^{(3)} s_j s_k}{\Gamma^{(1)} J_i^{(1)}} \right|}{\left| \frac{\sum_{j < k} J_{ijk}^{(3)} \text{sgn}(s_j) \text{sgn}(s_k)}{J_i^{(1)}} \right|} = \left| \frac{\Gamma^{(3)} \sum_{j < k} J_{ijk}^{(3)} s_j s_k}{\Gamma^{(1)} \sum_{j < k} J_{ijk}^{(3)} \text{sgn}(s_j) \text{sgn}(s_k)} \right|, \quad (\text{S.9})$$

where $\Gamma^{(k)} = \langle |s_m| \rangle^{(p-k)}$ for rescaling method p , and $\Gamma^{(k)} = 1$ ($\forall k$) for the baseline method. By construction, $\epsilon_i^{(2)} = \epsilon_i^{(3)} = 1$ under the spin sign method.

Supplementary Fig. 8 illustrates the temporal evolution of the uniform random 3-SAT instance `uf20-4` from the SATLIB benchmark set [1] across all methods. Similar behavior was observed for other instances, confirming the generality of the observations. Panels (a–e) show the spin amplitudes, which typically converge to ± 1 , although some spins remain frustrated. Panels (f–j) show the number of violated clauses in the SAT problem. All methods find a solution that satisfies all clauses, but they do so at different rates. Frustrated spins with vanishing amplitudes can persist even when the solution is reached, indicating that the solution is invariant under their spin-sign reversal. Panels (k–o) and (p–t) present $\epsilon_i^{(2)}$ and $\epsilon_i^{(3)}$, respectively. For the baseline method, both metrics start near zero, indicating that external fields $J_i^{(1)}$ dominate the quadratic and cubic interactions. As the IM evolves, both metrics fluctuate irregularly and deviate from unity, reflecting distortions in the effective energy landscape. When the spins eventually converge to ± 1 , the metrics approach unity—consistent with Eqs. (S.8) and (S.9)—but only after the system has effectively frozen. The rescaling methods (panels l–n and q–s) avoid a field-dominated start but still exhibit erratic deviations from unity. This shows that, despite adjusting interaction magnitudes dynamically, they cannot enforce the exact intended ratios for each spin. The spin-sign method, by contrast, preserves perfect consistency between all interaction orders and external fields, providing a more faithful realization of the binary Hamiltonian throughout the



Supplementary Figure 8: System dynamics for the uniform random 3-SAT problem uf20-4 from the SATLIB benchmark set [1] across different methods to incorporate higher-order spin interactions. (a–e) Each color denotes a different spin amplitude. The majority of spins reach ± 1 , while some remain frustrated. (f–j) All approaches successfully solve the problem, with progressively longer runtimes from left to right. Both error metrics, $\epsilon_i^{(2)}$ (k–o) and $\epsilon_i^{(3)}$ (p–t), fluctuate for the baseline and rescaling methods but remain exactly 1 for the spin-sign method. Hence, only the spin-sign method maintains the intended balance between spin couplings and external fields ($\epsilon_i^{(2)} = 1$) and between third-order interactions and external fields ($\epsilon_i^{(3)} = 1$), faithfully representing the binary Hamiltonian.

entire evolution.


Supplementary Note 8: Flips-to-solution

In the main text, we used time-to-solution and success rate as performance metrics. Prior work has shown that the number of spin flips required to reach a solution provides an additional informative measure [2]. This metric reflects how efficiently an Ising machine explores the binary energy landscape of Eq. (2). In some Ising machine implementations, spin flips may also contribute significantly to the energy consumption, which further motivates its study.

In this section, we adopt the flips-to-solution (FTS) metric as defined in Ref. [2]:

$$\text{FTS} = \begin{cases} \tilde{n}_{\text{flips}}, & \text{if } P > 0.99, \\ \tilde{n}_{\text{flips}} \frac{\log(0.01)}{\log(1-P)}, & \text{if } 0 < P \leq 0.99, \\ \infty, & \text{if } P = 0, \end{cases} \quad (\text{S.10})$$

where P is the probability of reaching the target state within the (dimensionless) time window T , and $\tilde{n}_{\text{flips}}$ denotes the average number of flips within that same time window. We further report the values $\min_{0 < T < t_{\max}} \text{FTS}(P(T), \tilde{n}_{\text{flips}}(T))$, which corresponds to the average total number of flips required to reach the optimal solution of a given problem with 99% probability.



Supplementary Figure 9: Comparison of flips-to-solution between the methods to incorporate higher-order interactions for uniform random 3-SAT problems. The spin sign method of Eq. (9) is compared to (a) the baseline method of Eq. (5) and (b-d) rescaling methods 1-3 of Eqs. (6) to (8). Dots in the grey area on the right denote COPs that could be solved by the spin sign method within the allocated compute time of $t_{\max} = 10^4$, but not by the method on the x-axis ($\text{TTS} = \infty$, $\text{SR} = 0$).

Supplementary Fig. 9 provides pairwise comparisons between the local field models from the Methods section for 3-SAT problems. Panels (c) and (d) show that the spin sign method outperforms rescaling methods 2 and 3 in terms of FTS. Both for the baseline method (panel a) and rescaling method 1 (panel b), the spin-sign method yields lower FTS on roughly half of the 60 tested COPs (28 and 33 instances, respectively). However, we observe that it generally yields lower relative FTS as the problem size increases. Moreover, several COPs could only be solved by the spin-sign method. In these cases, the baseline method and rescaling method 1 yield $\text{FTS} = \infty$ (8 and 6 instances, respectively).

References

- [1] H. Hoos and T. Stutzle. Satlib: An online resource for research on sat. In H. van Maaren I. P. Gent and T. Walsh, editors, *SAT2000*, pages 283–292. IOS Press, 2000.
- [2] Anshujit Sharma, Matthew Burns, and Michael C Huang. Combining cubic dynamical solvers with make/break heuristics to solve sat. In *26th International Conference on Theory and Applications of Satisfiability Testing (SAT 2023)*, pages 25–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2023.