Comparison of annealing methods for photonics Ising machines

Toon Sevenants,[1] Guy Verschaffelt,[1] and Guy Van der Sande[1]

[1] Vrije Universiteit Brussel, Applied Physics research group, Pleinlaan 2, 1050 Brussels, Belgium

Ising machines (IMs) are a novel natural computing system that could solve complex optimization problems more efficiently than traditional digital computers. A problem of interest is encoded into a network of non-linear elements, where the optimal solution is found by minimizing the cost function through an iterative feedback process. In our setup, the non-linearity is created through an opto-electronic loop, including an MZM, photodiode and FPGA. However, IMs do not always reach this optimal solution due to trapping in sub-optima. To address this challenge, annealing schemes are being investigated to navigate around these local minima and increase the overall success rate of finding the optimal solution. In this study, we compare (based on numerical simulations) the success rate and time-to-solution of two promising schemes: linear annealing (LA) and linear adiabatic annealing (LAA). In the LA scheme, we slowly change the feedback strength of the iterative update process over time. In the LAA scheme, the system starts in the known optimum solution to an easy problem. Subsequently, the network connection strengths are gradually changed towards that of the more complex target problem.

Introduction

Combinatorial optimization problems are very common in today's society. They are for example related to airline scheduling, Very Large Scale Integration or even gene-sequencing [1]. However, traditional digital computers are unable to efficiently compute the solution of these types of optimization problems, because the number of possible solutions scales very strongly with the number of variables in the problem. To solve these types of optimization problems more efficiently, a renewed interest in analog computing platforms has emerged, particularly in natural computing. This field of research tries to harness the computing power that is inherently present in nature to solve real-life problems. We focus here on one of these novel promising computing platforms, namely the Ising machine (IM), with the goal being to solve combinatorial optimization problems more efficiently than traditional digital computers.

IMs exploit the fundamental physical principle that a dissipative system evolves to the configuration that minimizes its energy, i.e., its ground state. More specifically, an IM uses a network of analog artificial spins to emulate the Ising model of which the energy is given by $H_{Ising} = -\frac{1}{2}\sum_{ij} J_{ij}\sigma(x_i)\sigma(x_j)$, where $\sigma(x_i)$ represents the sign of the i-the spin amplitude x_i and J_{ij} is the matrix that describes how the spins are interconnected. An optimization problem can be encoded through this interconnection matrix and driven by the spins' coupling, the spin network will tend to evolve towards its ground state energy. The network configuration in this ground state can then be translated to the optimal solution of the encoded problem [2]. In our experimental IM-setup, artificial spins are created through a time-multiplexed opto-electronic loop, including an Mach-Zehnder modulator, photodiode, and FPGA [3]. The time evolution of a spin's amplitude x_i in our setup is described by the following transfer equation:

$$\dot{x}_i = \cos^2(ax_i + b\sum_j J_{ij}\, x_j + \gamma - \pi/4) - 1/2$$

where $\pi/4$ and 1/2 are two bias terms, a and b represent the self-feedback strength and mutual-feedback strength respectively, J_{ij} represents again the interconnection matrix and γ the gaussian white noise.

IMs do not always reach the optimal solution due to trapping in local energy levels, meaning that solving a problem on an IM is inherently probabilistic. Therefore, it is possible to associate a certain success rate of finding the optimal solution with each problem. Another consequence is that numerous runs are required to ensure that at least one instance reaches the ground state energy of the problem being considered. This unfortunately strongly decreases the efficiency of the IM. To address this challenge, annealing schemes are being investigated to navigate around these local minima and increase the overall success rate of finding the optimal solution [1, 4]. In this contribution, we use numerical simulations to study and compare the performance two promising schemes: linear annealing (LA) and linear adiabatic annealing (LAA). Fundamentally, both schemes gradually alter the energy-landscape of the IM over time to avoid trapping in local energy levels.

Linear Annealing (LA)

The idea behind the LA scheme, which is schematically illustrated in Fig.1, comes from the fact that the ground state only exists above a specific self- and mutual-feedback value. We will refer to this critical point as the operation threshold of the IM. Although the ground state energy is the only energy minimum right at the operation threshold, it is also a very shallow energy minimum meaning the system can easily jump out of it due to noise that is inherently present in the system, e.g., thermal noise.

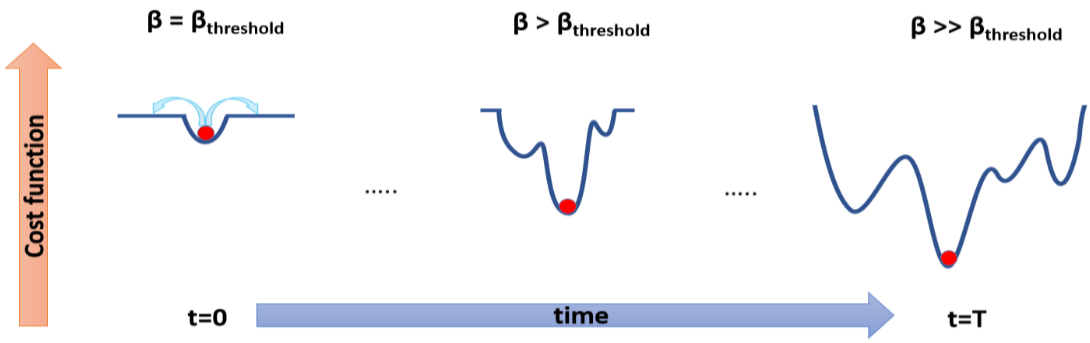

Figure 1. This figure illustrates the concept of LA. At the operation threshold, only the ground state solution exists but the system can easily escape it. At large coupling-strength values the ground state energy is better defined but there are now also local minima in which the system could get trapped. By slowly changing the coupling strength, and keeping the noise low, the system could remain in the ground state.

To better define the system's global minimum compared to its surrounding energy landscape, i.e. to make the ground state level deeper, we can further increase the mutual feedback. However, in doing this, many local minima will emerge in which the system can get trapped. In LA, we therefore ramp the mutual feedback over time from a start value below the operation threshold to a value above the operation threshold while fixing the self-feedback strength throughout this evolution. This would give the system a better chance of remaining in the ground state pit as it does not suddenly have to navigate around the many local minima as it would have if initiated far above the operation threshold.

Linear Adiabatic Annealing (LAA)

In the LAA scheme, we do not change the mutual-feedback strength but rather the Hamiltonian of the system itself. An example of an LAA evolution is schematically illustrated in Fig.2. At the start, the system is initialized in the known ground state configuration of an easy-to-solve problem, e.g., a network of spins only coupled anti-ferromagnetically to their nearest-neighbors. Subsequently, the network connection strengths J_{ij} are gradually changed towards that of the more complex target problem, effectively changing the Hamiltonian of the system. The time-dependent Hamiltonian can be written as follows:

$$H_{Ising}(t) = \left(1 - \frac{t}{T}\right) H_{start} + \frac{t}{T} H_{target}$$

where T is the runtime of the system. The idea behind the LAA technique is that if the system is initialized in the ground state of the easy Hamiltonian, the system will remain in the ground state of an interpolated Hamiltonian if the transition happens slow enough and in a low noise regime. Note that the principle of adiabatic evolution has already been proven in quantum mechanics but not in classical mechanics, which is the regime in which our IM operates.

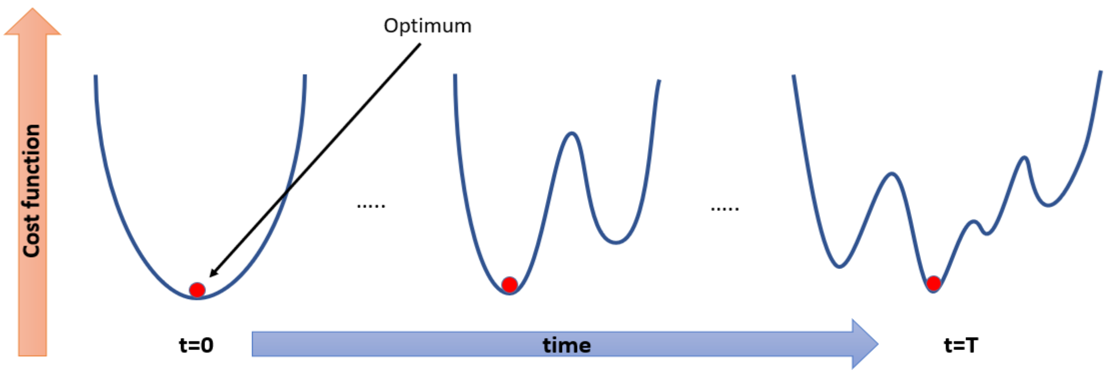

Figure 2. This figure illustrates the concept of LAA. The system starts at t = 0 from the known optimal solution of a simple problem (red dot). As time increases, we change the network connections towards that of the target problem, which has a more complex cost function, i.e., more local minima.

Results and summary

We use the 'g05' problem set of the biqmac benchmark library. These ten problems all have 100 spins and an edge probability of 50%. Table 1 shows the success rate for four of these problems. We used two different annealing schemes, namely LA and LAA. For the LAA scheme, the system is initiated in an anti-ferromagnetic configuration with nearest-neighbor coupling. For each problem, the table also contains the success rate of finding the ground state energy without annealing. These values are used as a reference to compare against the success rates obtained through annealing. The noise strength γ, the amount of runs and the number of iterations remained fixed for all the reference and annealing simulations to 0.001, 100 and 21.000 respectively.

Looking at Table 1, the first row contains the problems names. Underneath each problem, the success rate without annealing, the success rate with LA and the success rate with LAA is given. The green colored cells are used to indicate if an annealing scheme results in an increased success rate with respect to the reference value. For problem 100.0, both LA and LAA result in an significant increase in success rate. For problems 100.1 and 100.9, only one of the two annealing schemes seem to improve the success rate, which are LAA and LA respectively. For problem 100.5, neither LA or LAA seem to increase the success rate as both return a 0% success rate.

Overall, our data shows that is possible to increase the success rate of some benchmark problems using annealing schemes. However, there does not seem to be an annealing scheme that works for all the problems equally well. The LAA technique increases the success rate of problem 100.0 and 100.1 while LA seem to be beneficial for problems 100.0 and 100.9. Additionally, there also seems to be problems in the biqmac benchmark library that do not improve under the considered annealing schemes such as 100.5, which already has a significantly lower reference success rate.

Table 1. This table contains the success rates of four different biqmac problems under three different annealing schemes.

Problem name	**100.0**	**100.1**	**100.5**	**100.9**
Reference value	43%	39%	3%	55%
LA	89%	34%	0%	85%
LAA (anti-ferromagnetic)	100%	54%	0%	0%

In summary, the IM is a natural computing system that finds the solution of a problem by minimizing the energy of its spin network. To increase the probability of ending up in the lowest energy state, i.e., the success rate, our research uses linear annealing (LA) and linear adiabatic annealing (LAA) to slowly modify the IMs energy landscape. Our results suggest that different problems require different annealing schemes, although further research is needed to reveal what the exact underlying reasons are.

References

[1] S. Kumar, T. Van Vaerenbergh, and J.P. Strachan, "Classical Adiabatic Annealing in Memristor Hopfield Neural networks for Combinatorial Optimization," 2020 International Conference on Rebooting Computing (ICRC), pp. 76-79, 2020.

[2] N. Mohseni, P. L. McMahon, and T. Byrnes, "Ising machines as hardware solvers of combinatorial optimization problems," ArXiv, 2022.

[3] F. Böhm, G. Verschaffelt, and G. Van der Sande, "A poor man's coherent Ising machine based on opto-electronic feedback systems for solving optimization problems," Nat. Comm., 2019.

[4] D. Pierangeli, G. Marcucci, and C. Conti, "Adiabatic evolution on a spatial-photonic Ising machine," Optic., 2020