Feedback and constraints in physical optimizers

Niharika Gunturu[a], Hideo Mabuchi[a], Edwin Ng[b], Daniel Wennberg[a], Ryotatsu Yanagimoto[b]
[a]Edward L. Ginzton Laboratory, Stanford University; [b]NTT Research, Inc. (United States)

ABSTRACT

Extremizing a quadratic form can be computationally straightforward or difficult depending on the feasible domain over which variables are optimized. For example, maximizing $E = x^T V x$ for a real-symmetric matrix V with x constrained to a unit ball in R^N can be performed simply by finding the maximum (principal) eigenvector of V, but can become computationally intractable if the domain of x is limited to corners of the ± 1 hypercube in R^N (*i.e.*, x is constrained to be a binary vector). Many gain-loss physical systems, such as coherently coupled arrays of lasers or optical parametric oscillators, naturally solve minimum/maximum eigenvector problems (of a matrix of coupling coefficients) in their equilibration dynamics. In this paper we discuss recent case studies on the use of added nonlinear dynamics and real-time feedback to enforce constraints in such systems, making them potentially useful for solving difficult optimization problems. We consider examples in both classical and quantum regimes of operation.

Keywords: optimization, nonlinear dynamics, feedback, constraints, nonlinear optics, quantum optics

1. INTRODUCTION

Recent benchmarking studies[1] of the empirical performance of Coherent Ising Machine (CIM)-type architectures and simulation algorithms[2] has sparked interest in understanding the underlying principles of these approaches to quadratic unconstrained binary optimization (QUBO), in which we seek an assignment of the binary variables σ_i that maximizes the cost function,

$$E = \sum_{i,j} J_{ij} \sigma_i \sigma_j,$$

or, equivalently, minimizes the equivalent quadratic form with $J_{ij} \to -J_{ij}$. The canonical CIM algorithm is understood[2] to operate by first finding a linearized solution in the form of the (real) eigenvector $v_{\max}$ of J_{ij} with largest eigenvalue $\varepsilon_{\max}$, and then gradually imposing soft constraints (through the activation of an additional quadratic energy potential) that deform $v_{\max}$ into a quasi-binary vector whose entries cluster tightly around real values $\pm\alpha$. Finally, the proposed QUBO solution is taken as the corresponding sign vector with entries ± 1. Key aspects of the performance of canonical CIM as a QUBO method can be understood in the setting of gaussian random J_{ij} in high dimensions[3] and modifications of the CIM with improved performance have been formulated. The most notable CIM variant, known as CIM-CAC (Coherent Ising Machine with chaotic amplitude control), utilizes dynamic feedback[4] in place of, or in addition to, the nonlinear potential to induce binarization of the solution vector.

While quadratic optimization over real variables (in the form of principal eigenvector-finding) is generally considered to be tractable even for problems with billions of variables[5], the imposition of nonconvex constraints on the domain of the optimization creates difficulty for practical applications[6] even for variable counts $N \sim 10^3$. Typical constraints arising in applications include restriction of optimization variables to binary or integer values, as well as inequality constraints on variables (e.g., box constraints) or functions of variables (e.g., knapsack problems).

For non-convex energy landscapes (cost functions), the broad class of real optimization methods based on modified gradient descent suffers from obstruction by local minima and low-index saddle points. CIM-type algorithms appear to avoid high-lying traps[3] by gradually interpolating from the unconstrained (principal eigenvector) potential to a potential that includes binarizing nonlinear components. It is interesting to note that these algorithms thus seem to keep a set of optimization parameters (the J_{ij} matrix) fixed while interpolating between two distinct domains for the optimization variables – starting with real values allowed but ending with a restriction to binary values. Recently, more general domain interpolations[7,8] have been investigated for CIM variants.

AI and Optical Data Sciences V, edited by Ken-ichi Kitayama, Volker J. Sorger, Proc. of SPIE
Vol. 12903, 129030K · · 0277-786X · doi: 10.1117/12.3005007

In this paper we discuss two ongoing studies that further extend this approach. In the first we analyze how interpolating between binary and phase-like variable domains impacts the performance of CIM-type algorithms. For the smallest QUBO instances ($N = 4$) that are non-trivial for canonical CIM, we find that allowing the optimization variables to explore a phase-like domain before being strictly binarized (e.g., allowing $\sigma_i \sim \alpha e^{i\varphi}$ with gradual restriction of φ to $\pm\pi$) modifies critical bifurcation structures that set boundaries between easy and hard instances. In the second study we explore the use of non-demolition measurements with real-time feedback to enforce constraints in a network of coherently-coupled optical parametric amplifiers (OPAs) operating in a quantum regime, in which the (signal) photon number in each OPA represents an integer variable in a linear programming problem. Our preliminary findings suggest that sequential quantum nondemolition measurements with feedback may provide a way to simultaneously enforce multiple equality constraints on linear combinations of integer variables represented by the photon numbers in coherently-coupled OPAs (or optical parametric oscillators).

2. INTERPOLATION BETWEEN PHASE-LIKE AND BINARY DOMAINS

In the canonical CIM based on a model of coupled OPOs, the binary optimization variables are represented by optical modes whose amplitudes are intrinsically complex numbers, but whose imaginary parts are strongly driven to zero by the phase sensitivity of the gain. However, one can also inject phase-insensitive gain that only imposes a soft constraint on the mode amplitudes and does not affect the phases. By combining the two types of gain, the interpolation between the principal eigenvector and the QUBO solution can pass through a regime where the amplitude constraint takes precedence over the phase constraint, lowering the barrier to flipping the sign of a principal eigenvector component in response to the amplitude constraint before the spins are fully binarized by the phase constraint.

As a testbed for this strategy, we consider a one-parameter family of N = 4 QUBO instances where the ground state and all but one coupling are ferromagnetic, and the antiferromagnetic coupling is the tunable parameter. For a range of values, one component of the principal eigenvector has the wrong sign, but the canonical CIM successfully recovers from this sign error in the initial estimate for a subset of this range. However, by adding phase-insensitive gain, we can significantly extend the range of couplings for which the machine succeeds. This is mediated by a bifurcation that destabilizes the all-real equilibrium of the canonical CIM and creates a bridge through the complex plane to the branch with the lower Ising energy. An example of this is shown in Figure 1.

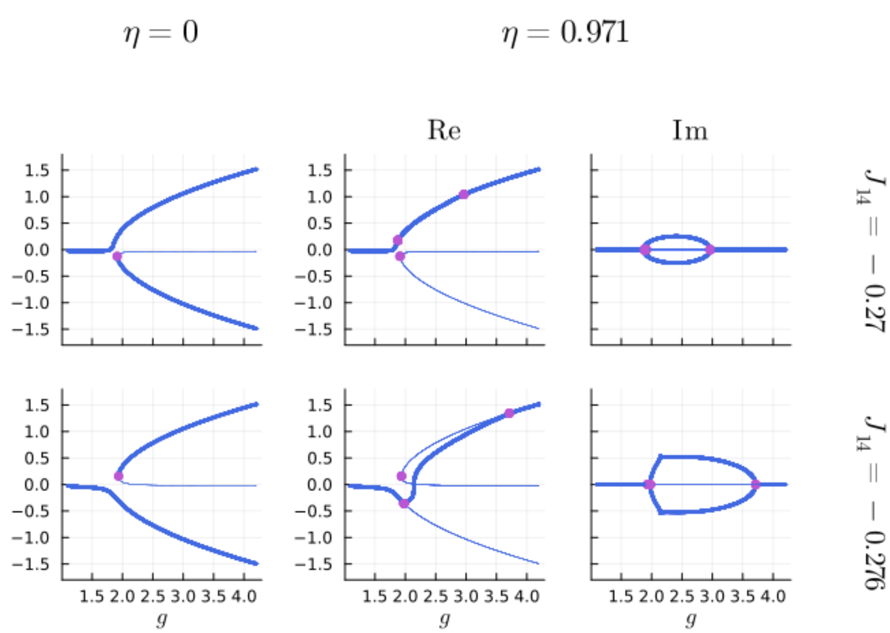

Figure 1. Bifurcation diagrams for the CIM with and without phase insensitive gain. The diagrams show equilibrium values of the least committed spin as a function of overall gain g in a frustrated $N = 4$ instance where the ground state and all but

one coupling are ferromagnetic. Heavy and thin curves represent stable and unstable fixed points, respectively, and the markers show the locations of bifurcations. The parameter η determines the fraction of gain that is phase insensitive; thus, the left column, where $\eta = 0$, represents the canonical CIM. The two right columns, where $\eta = 0.971$, represent a case where almost all the gain is phase insensitive, and we must plot both the real and imaginary parts of the amplitudes, shown in the middle and right columns, respectively. Each row shows an instance with a particular value of the antiferromagnetic coupling J_{14} as shown on the right.

For both instances, the principal eigenvector has large positive values for the first three components and a small negative value for the fourth component, such that the first three spins commit early and take large positive values on all the plotted branches (not shown). The remaining spin, however, has a small and negative value until it is driven to commit by a displaced pitchfork around $g = 1.8$, and its eventual sign depends on the value of the antiferromagnetic coupling. The top row shows an instance where the canonical CIM succeeds in recovering the ferromagnetic ground state, and the CIM with phase-insensitive gain performs similarly albeit with an inconsequential excursion into the complex plane. (Though note that it also destabilizes the incorrect branch, thus increasing robustness against noise and nonequilibrium dynamics.) The bottom row shows an instance with a stronger antiferromagnetic coupling, where the canonical CIM fails but the CIM with phase-insensitive gain succeeds through a bifurcation that connects back to the ground state branch via the complex plane.

3. LINEAR INTEGER CONSTRAINTS VIA QND MEASUREMENT AND FEEDBACK

It can be shown that a coherent pump field driving a frequency-detuned optical parametric amplifier (OPA) experiences displacements conditioned on the number of signal Bogoliubov excitations[9]:

$$\hat{H} \approx -\tilde{g}\left(\hat{N}_A + \frac{1}{2}\right) + \Delta\hat{N}_A + const,$$

where $\tilde{g}$ is the effective nonlinear coupling strength. The form of this Hamiltonian thus allows us to perform QND measurements on the signal photon number. This information can be accessed through homodyne measurement of the p-quadrature of the pump mode.

We can extend this PNR-QND system to include two OPAs driven by a single pump, as shown in Figure 2. In such a case, the pump encodes information about the photon number *sum* of the two signal modes. We can similarly compose a photon number *difference* measurement scheme by using the same setup but inducing a π phase shift of the pump between the two OPAs.

In our ongoing study, we investigate strategies for computing the solution to a *system* of constraints by feeding forward the solution of each constraint to the next. In a preliminary scheme we feedforward post-selected states that are conditioned upon the sum or difference being the desired value. In continuing work, we are studying the use of real-time feedback control to enforce the constraints autonomously[10] (without post-selection).

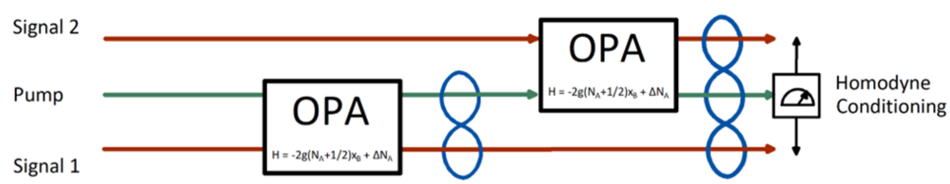

Figure 2. We compose a system wherein two OPAs share the same pump mode. Each OPA encodes in the pump mode information about the number of Bogoliubov excitations in the respective signal mode, thus entangling the two modes. After both OPA interactions, the system exists in a three-mode entangled state, where the pump has encoded the sum (or difference) of the Bogoliubov excitations in the two signal modes. For each OPA we set the system parameters as follows: We use the system parameters of $\Delta/\tilde{g} = 150$ and the total interaction time of $\tilde{g}t = 1$.

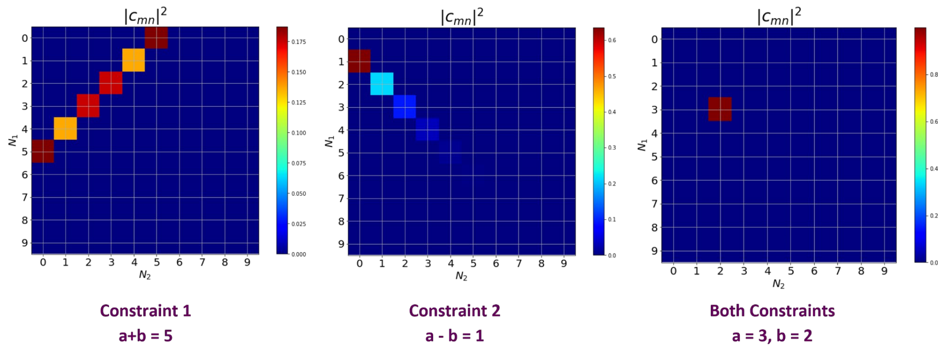

Figure 3. Examining the signal state $|\psi\rangle = \sum c_{mn}|m\rangle|n\rangle$ post homodyne conditioning. We observe the signal state exists in a superposition of valid solutions to the constraints (i) $a + b = 5$, (ii) $a - b = 1$. In (iii), we feedforward the solution of the adder system to the subtractor system and observe that the signal state post homodyne conditioning is the solution to the system, *i.e*, $a = 3, b = 2$.

REFERENCES

[1] Don, S. H. G. M., Inui, Y., Kako, S., Yamamoto, Y. and Aonishi, T., "Mean-field coherent Ising machines with artificial Zeeman terms," Journal of Applied Physics 134, 234901 (2023).

[2] Yamamoto, Y., Leleu, T., Ganguli, S. and Mabuchi, H., "Coherent Ising Machines—Quantum optics and neural network Perspectives," Applied Physics Letters 117, 160501 (2020).

[3] Yamamura, A., Mabuchi, H. and Ganguli, S., "Geometric landscape annealing as an optimization principle underlying the coherent Ising machine," arXiv:2309.08119 [cond-mat.dis-nn] (2023).

[4] Kako, S., Leleu, T., Inui, Y., Khoyratee, F., Reifenstein, S. and Yamamoto, Y., "Coherent Ising Machines with Error Correction Feedback," Advanced Quantum Technologies 3, 2000045 (2020).

[5] Zbikowski, R. M. and Johnson, C. W., "Boostrapped block Lanczos for large-dimension eigenvalue problems," Computer Physics Communications 291, 108835 (2023).

[6] Singh, A. K., Venturelli, D. and Jamieson, K., "Perturbation-based Formulation of Maximum Likelihood MIMO Detection for Coherent Ising Machines," *GLOBECOM 2022 – 2022 IEEE Global Communications Conference, Rio de Janiero, Brazil, 2022*, 2523 (2022).

[7] Strinati, M. C. and Conti, C., "Multidimensional hyperspin machine," Nature Communications 13, 7248 (2022).

[8] Strinati, M. C. and Conti, C., "Hyperscaling in the Coherent Hyperspin Machine," Physical Review Letters 132, 017301 (2024).

[9] Yanagimoto, Y., Nehra, R., Hamerly, R., Ng, E., Marandi, A. and Mabuchi, H., "Quantum Nondemolition Measurements with Optical Parametric Amplifiers for Ultrafast Universal Quantum Information Processing," PRX Quantum 4, 010333 (2023).

[10] Stockton, J. K., Van Handel, R. and Mabuchi, H., "Deterministic Dicke-state preparation with continuous measurement and control," Physical Review A 70, 022106 (2004).