Copositive optimization via Ising solvers

Robin Brown*†‡, David Bernal†, Ankit Sahasrabudhe*,
Aaron Lott†‡, Davide Venturelli†‡, and Marco Pavone*

*Stanford University,
†NASA Quantum Artificial Intelligence Laboratory (QuAIL)
‡USRA Research Institute for Advanced Computer Science (RIACS)

Recent years have seen significant advances in physics-inspired technologies capable of approximately searching for the ground state of Ising Hamiltonians, in other words, solving optimization problems of the form: $\min_{z\in\{-1,1\}^n} \sum_{i,j} J_{i,j} z_i z_j + \sum_i h_i z_i$, where $J_{i,j}$, h_i are real coefficients and $z_i \in \{-1,1\}$ are discrete variables to be optimized over [1, 2]. The promise of leveraging quantum technology to accelerate the solution of difficult optimization problems has spurred an increased interest in developing mappings from such optimization problems to Ising Hamiltonians. A standard approach for constructing these mappings involves discretizing continuous variables and passing constraints into the objective through a penalty function. However, due to the heuristic nature of many state-of-the-art Ising solvers, such an approach is insufficient in scenarios that require guarantees of global optimality, such as neural network verification. Moreover, these Ising solvers are often treated as a black box, eluding analysis on the optimality gap of the returned solution.

In this work, we propose to use Ising solvers to solve copositive programs—these are linear optimization problems over the cone of copositive matrices. Despite being convex, copositive programs are generally intractable because checking copositivity of a matrix is co-NP-complete. However, overlooking the challenges of checking copositivity, convexity implies a problem structure that is amenable to efficient optimization algorithms. Specifically, we propose to use cutting-plane/localization algorithms, a broad class of convex optimization algorithms that alternate between checking feasibility of a test point, updating an outer approximation of the feasible region, and judiciously selecting the next test point. Checking feasibility (i.e., copositivity) is most naturally posed as a quadratic minimization problem; critically, this is precisely the class of problems that Ising solvers excel at. We propose to check feasibility using an Ising solver and to carry out the remainder of the algorithm with a classical computer.

There are many instantiations of cutting plane algorithms, differing in how the feasible set is approximated and how test points are selected. However, they all require only a polynomial number of outer iterations in the dimension of the problem, and the selection of each test point is also polynomial in the problem dimension [3]. This means that the proposed algorithm necessarily shifts the complexity of the copositive program onto the copositivity checks, taking full advantage of any speedups from a (potentially quantum) Ising solver.

Moreover, one can use any certificate of non-copositivity to refine the outer approximation of the feasible region, even if it is not a global minimizer of the copositivity check. Thanks to convexity, upper and lower bounds can be readily computed from these outer approximations. Accordingly, the proposed approach offers resilience to the heuristic nature of Ising solvers and transforms black box solutions into ones with rigorous optimality guarantees.

Finally, the choice to focus on copositive optimization problems is non-restrictive. In a landmark result, Burer demonstrated an exact mapping from mixed-binary quadratic optimization problems to linear optimization problems over the cone of completely positive matrices—dualizing these problems yields a linear optimization problem with a single copositivity constraint [4]. Combining Burer's mapping with our approach yields a generic framework for solving mixed-binary quadratic optimization problems that encompasses many problems tackled with Ising solvers thus far.

We conducted a preliminary investigation of the proposed method on the maximum clique problem, which finds the largest complete subgraph of a graph. Given a graph, the maximum clique problem can be formulated as the copositive program $\min_{\lambda \in \mathbb{R}}\{\lambda \mid (\lambda(I + \overline{A}) - \mathbb{1}\mathbb{1}^\top \in \mathcal{C}\}$, where $\overline{A}$ is the adjacency matrix of the graph's complement, I the identity, and $\mathcal{C}$ the cone of copositive matrices [5]. To study the scaling of the proposed approach, we considered random max-clique problems with $10, 20, \ldots, 100$ vertices. For each graph size, we generated 25 random Erdős-Renyi instances with edge density 0.5 and solved to global optimality using the proposed copositive cutting plane algorithm. The copositivity checks were conducted by solving Anstreicher's mixed-integer linear programming characterization of copositivity, [6], using Gurobi version 9.0.3 [7]. Figure 1 plots the time the copositive cutting plane algorithm spent on the copositivity checks versus other operations (updating the outer approximation and computing test points). The time spent on the copositivity checks scales exponentially with the number of vertices in the graph, while the time spent on other operations stays constant. This confirms that the proposed approach shifts the complexity of the copositive program onto the copositivity checks.

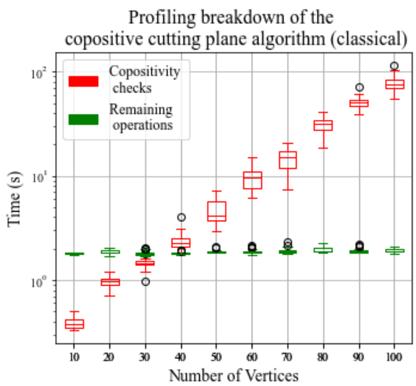

Fig. 1. Time spent on the copositivity checks versus all other operations in the proposed method. The copositivity checks grow exponentially with the number of vertices, while the other operations time remains constant.

To investigate potential speedups from using a stochastic Ising solver, we resolved each copositivity check that yielded a certificate of non-copositivity using Simulated Annealing (SA)[8]. For each copositivity check solved, we considered discretizations corresponding to $\min\{x^\top Qx \mid x \in \frac{1}{2^{i+1}}\{2^0, \ldots, 2^i\}^n\}$ for $i = 1, 2, 3, 4$. The performance of SA, which depends on the number of sweeps, was optimized using Hyperopt [9]. Figure 2 plots the best times to solution from SA against the time to solution from Gurobi. We see that for all discretization sizes, SA can consistently find

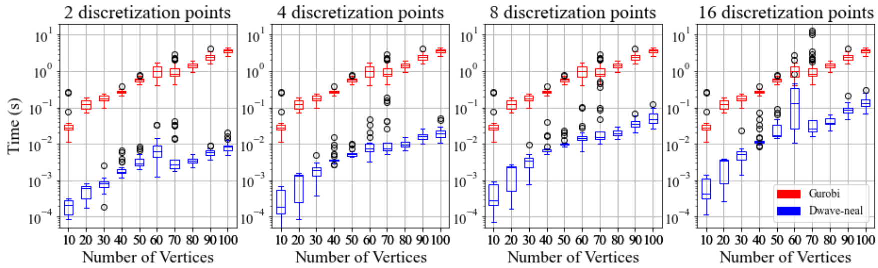

Fig. 2. Optimized time to solution of simulated annealing (Dwave-neal) against solution time of Gurobi for 2, 4, 8, and 16 discretization points in the copositivity checks. Both methods scale exponentially with the number of vertices in the graph; however, SA is several orders of magnitude faster than Gurobi.

certificates of non-copositivity in orders of magnitude less time than Gurobi. Notably, SA and Gurobi demonstrate similar scaling with respect to the number of vertices.

Conclusions: In this work, we propose to leverage Ising solvers to solve copositive programs to global optimality. To take advantage of potential speedups from quantum, or quantum-inspired Ising solvers, such as quantum annealers [10], coherent Ising machines [11], or quantum approximate optimization circuits [12], the proposed approach shifts the complexity of solving copositive programs onto the copositiviy checks—these are carried by the Ising solver. The convex structure of copositive programs allows for rigorous optimality guarantees even if the Ising solver used is a black box and/or heuristic in nature. In conjunction with the approach of [4], the proposed algorithm is a generic framework suitable for a broad class of mixed-binary quadratic programs.

Acknowledgements

This work was supported by NSF CCF (grant #1918549), NASA Academic Mission Services (contract NNA16BD14C – funded under SAA2-403506). R.B. acknowledges support from the Feynman Academy Internship program.

References

1. C. McGeoch and P. Farré, "The D-wave advantage system: An overview," *D-Wave Systems Inc., Burnaby, BC, Canada, Tech. Rep*, 2020.
2. Y. Yamamoto, K. Aihara, T. Leleu, K.-i. Kawarabayashi, S. Kako, M. Fejer, K. Inoue, and H. Takesue, "Coherent Ising machines—Optical neural networks operating at the quantum limit," *npj Quantum Information*, vol. 3, no. 1, pp. 1–15, 2017.
3. S. Boyd and L. Vandenberghe. Localization and cutting-plane methods. Stanford University. Available at https://see.stanford.edu/materials/lsocoee364b/05-localization_methods_notes.pdf.
4. S. Burer, "On the copositive representation of binary and continuous nonconvex quadratic programs," vol. 120, no. 2, p. 479–495, 2009.

5. E. De Klerk and D. V. Pasechnik, "Approximation of the stability number of a graph via copositive programming," *SIAM Journal on Optimization*, vol. 12, no. 4, pp. 875–892, 2002.
6. K. M. Anstreicher, "Testing copositivity via mixed–integer linear programming," *Linear Algebra and its Applications*, vol. 609, pp. 218–230, 2021.
7. *Gurobi Optimizer Reference Manual*, 2022, available at http://www.gurobi.com.
8. *dwave-neal Documentation*, D-Wave Systems Inc, 2021, available at https://docs.ocean.dwavesys.com/_/downloads/neal/en/latest/pdf/.
9. J. Bergstra, D. Yamins, and D. Cox, "Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures," in *International conference on machine learning.* PMLR, 2013, pp. 115–123.
10. M. W. Johnson, M. H. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk *et al.*, "Quantum annealing with manufactured spins," *Nature*, vol. 473, no. 7346, pp. 194–198, 2011.
11. T. Honjo, T. Sonobe, K. Inaba, T. Inagaki, T. Ikuta, Y. Yamada, T. Kazama, K. Enbutsu, T. Umeki, R. Kasahara *et al.*, "100,000-spin coherent ising machine," *Science advances*, vol. 7, no. 40, p. eabh0952, 2021.
12. E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," *arXiv preprint arXiv:1411.4028*, 2014.