

Vector Annealing for Improved Investment Models

Summary: Simulated Annealing is a classical heuristic search algorithm prevalently used for benchmarking and solving combinatorial optimization problems. Built on the principles of the Metropolis-Hastings algorithm, SA incorporates a monotonically decreasing temperature schedule which helps avoid local minima. For such heuristic methods, the time to find a candidate solution and the solution's quality quantify the performance of the discrete optimization approach. Throughout this text, we refer to an open-source implementation of the simulated annealing algorithm as SA¹.

Vector Annealing is NEC's optimized implementation of the simulated annealing algorithm using special purpose hardware. Vector Annealing, or VA, simulates quantum behavior using conventional computing technology without utilizing the actual quantum properties, hence is labeled as *quantum-inspired* computing. NEC Vector Annealing greatly reduces the computational complexity associated with traditional Simulated Annealers and accelerates the narrowing down of the candidate solutions by a factor of upto 300 times at problem sizes beyond the capabilities of conventional methods. This support for very large number of variables allows NEC VA to compute combinatorial optimization from a huge combination of variables having complex, real-world constraints.

Here we present a quantitative comparison of NEC's Vector Annealing (VA) solution against the simulated annealing algorithm on financial portfolio optimization. Markowitz's Modern Portfolio Theory was used to formulate portfolio management as a fully-connected discrete optimization problem and create a model that invests based on the solution². The number of stocks considered and our level of control over each stock determine the number of linear variables (i.e. the problem size). We found that Vector Annealing generally outperformed the Simulated Annealing algorithm in terms of solution quality and that its advantage over SA scales with problem size. As Figure 1 demonstrates, tuning the parametrization of portfolio management as a discrete optimization problem greatly affects the financial returns. We note that this is a simplified tuning example for pedagogical purposes, during a downtrend.




Figure 1: Both SA and VA perform better after tuning, VA's performance improvement is greater. This plot corresponds to a 5,243 linear variable problem.

At Icosa, we employ more sophisticated parametrization of portfolio management and approach tuning with a machine learning toolkit in real market conditions.

Methodology: Our formulation and testing of annealing approaches to improve portfolio optimization involves four stages. First, stock market data is obtained either from IEX's platform or, for some international stock prices, Yahoo Finance's platform. Second, using a tunable financial model, we deconstruct and reformulate the original problem into one of discrete optimization. Third, both SA and VA are used to find a candidate solution to this problem. Finally, the candidate with the lowest energy is considered the best solution and is formulated into portfolio decisions using the same model. To compare these discrete optimization approaches, we chose a historical period with known stock market information and test both annealing approaches. For our measure of performance, we subtracted the energy of the best VA solution from the best SA solution. We removed some of the highest energy differences - the outliers resulting mainly from SA's inability to produce any solution energy close to the averages produced by VA. Tests of SA were ran on x86 CPU hardware whereas VA tests were ran on NEC's Vector Engine. The temperature related parameters were selected using an internal function from SA software package and was used for both SA and VA¹. This selection of temperature range was best suited for SA.

¹D-Wave Systems, *dwave-neal* 0.6.0, Nov 2022

²Harry Markowitz. Portfolio selection. *The journal of finance*, 7(1):77–91, 1952.




Figure 2: Workflow for an ideal trading day. Financial data (historical stock prices) is collected and used to generate a Markowitz Model [panel (a)]. The hyper-parameters governing the Markowitz model are tuned using a machine learning model that runs continuously on a distributed computing platform [panel (b)]. The model is mapped to a quadratic unconstrained binary optimization (QUBO) problem, and solved using Quantum-Inspired (QI) hardware or algorithm [panel (c)]. The solution found by the solver represents the optimal investing strategy as defined by the constraints of the Markowitz Model [panel (d)]. In this work, we used stock price data from the previous 10 trading days to generate a daily trading strategy.

Results: Our results are summarized in Figure 3:

Linear Variables	Energy Gap	Trading Days	Market	Stocks	Granularity	VA Mean Energy	SA Mean Energy
2,430	0.0283	301/341	SP500	486	5	-0.1200	-0.0917
5,243	0.3522	212/230	U.S.	5,243	1	-8.2677	-7.9155
10,486	2.0646	31/31	U.S.	5,243	2	0.0549	2.1195
17,833	4.2298	23/23	International	17,833	1	-0.5460	3.6838
25,034	4.5050	3/3	International	25,034	1	-1.6194	2.8856
5,243 (Tuned)	0.4390	230/230	U.S.	5,243	1	-0.2198	0.2192

Figure 3: VA and SA performance for different problems. Trading days column shows number of days VA generated a better solution than SA out of total number of days

The problem sizes were dictated by granularity chosen and the equity market that we selected to evaluate. In the scope of this text, granularity corresponds to the number of linearly independent variables associated with one stock. Higher granularity results in a higher number of independent variables. The first test was conducted using price data of equities in the S&P 500 between 3/12/2018 and 8/1/2019. Some of the equities had missing stock price data, so we limited our scope to 486 equities that had sufficient data. We applied a granularity of 5, which increased the problem size to 2,430 independent variables over 342 trading days.

We then considered all U.S. equities who had price data available on IEX, yielding 5,243 stocks. We ran a test on this data between 3/18/2022 and 3/2/2023 with a granularity of 1, resulting in 5,243 independent variables over 231 trading days. In addition, we ran a second test on the same data but with a granularity of 2 and a date range of 3/18/2022 to 5/16/2022, resulting in 10,486 independent over 32 trading days.

To create problems with large numbers of variables, we made 2 different tests which included many international equities. The first test included 17,833 equities traded in France, Germany, U.K., and the U.S. between 3/18/2022 and 5/4/2022, yielding 17,833 independent variables with a granularity of 1 over 24 trading days. The second test included 25,034 equities from Canada, France, Germany, Japan, Turkey, U.K., and the U.S. between 3/17/2022 and 4/1/2022, yielding 25,034 equities with a granularity of 1 across 3 trading days.

VA consistently produced better quality solutions than SA in our tests, this is despite the temperature parameters being chosen to optimize SA. We also found that the the magnitude of the advantage of VA over SA — the average energy gap between VA and SA — increased as the number of variables grew for untuned models (as seen in Figure 4), demonstrating that VA has a nontrivial scaling advantage.




Figure 4: Growth of Vector Annealing's advantage with problem size. The errors plotted reflect the outlier excluded 95% confidence interval.

Market Performance: We tested the performance of VA and SA in US and international markets by converting their solutions to investment decisions based on a financial model. In general, this model can have many tunable parameters. Throughout our analysis, we focused on the solution quality difference between VA and SA and disregarded actual financial returns as the model’s parameters are not tuned. However, both the discrete optimization energies and the financial returns are quite sensitive to these parameters.

Figure 1 represents an instance with 5,243 tradeable stocks in the US equity market. Here, tuned the return-to-risk ratio coefficient and both the energy difference and financial market performance between SA and VA changed significantly, where VA’s superior performance grew in both cases. Accurate fine-tuning for real world trading is beyond what’s demonstrated in this paper. At Icosa, we have built a proprietary machine-learning toolkit to fine-tune problem parameters to generate profit in financial markets.

Future Work: Our comparison of NEC’s improved implementation of simulated annealing (VA) with an open-source non hardware-accelerated simulated annealing showed that VA has a clear advantage that scales with the number of independent variables in the prob-

lem. To the extent of our knowledge, this is the first test where such a large number of stocks are considered as part of a portfolio optimization problem and NEC’s VA solution provides a competitive advantage over the open-source method. We recognize that these large-scale tests are not indicative of how this solver would perform in real-world scenarios, as it is infeasible for most institutions to regularly trade 25,034 individual stocks, and in addition, we used an untuned MPT model, which is itself a too simple for real-world trading.

Instead, a more suitable application for financial institutions is to focus on highly tuned trading algorithms with high granularity focused on mid-sized markets, ones involving 1000-5000 stocks. For example, we demonstrated a significant improvement in the market performance of the trading algorithm in the US equity market with 5,243 tradable stocks with just a minor tuning. At Icosa, we use more nuanced trading strategies and tune them with machine-learning to achieve performance improvements. VA’s ability to handle large problems makes it practical to increase the granularity of these mid-sized market models. Combining this with Icosa’s proprietary model tuning, future work will include testing fine tuned, hardware accelerated models with real trades to achieve financial advantage with quantum-inspired methods.