

Confirmation of Exponential Speed-up
Using MemComputing to Solve
Hard Optimization Problems



MemComputing, Inc.

Table of contents

1	Executive Summary
2	Introduction
3	MemComputing Advantage <ul style="list-style-type: none">- Self-Organizing Logic Gates- Nonlocal Collective State Computation- Scalability in Linear Time
5	Demonstrated Performance <ul style="list-style-type: none">- The Challenge- The Procedure- The Result
7	Conclusion
7	Reference

Executive Summary

Optimization is a core component of every data-based scientific discipline and industry. However, the explosive growth of Big Data combined with classical computing architecture can cause significant issues in solving optimization problems. The difficulty of computation is especially prominent in a set of problems classified as NP-hard. Currently, no algorithm exists for NP-hard problems that is able to find a solution in polynomial time.

Furthermore, even reaching approximations to an optimal solution is challenging to discover for the most difficult optimization problems. This is because as the number of inputs and constraints increase linearly, the computational time needed to find an approximation to the optimal solution can grow exponentially. It is common in industry for a single problem to take days, even weeks, to run. This is extreme, but obviously highlights the value of the information provided. There are also plenty of problems that could take months, years, decades or more. However, in most cases, it is not viable to wait this long or pay for that much computation time.

This document discusses how MemComputing exploits the physical properties of classical computer architectures to bypass these exceedingly long computations and provide results orders of magnitude faster than today's best algorithms. We also highlight results from a benchmark test where MemComputing was able to resolve a problem that would take today's best solvers many times the age of the universe. This demonstrates that MemComputing enables industry to approach computations that were previously impossible to effectively approximate.

Introduction

For industrial applications involving autonomous vehicles, delivery routing, traffic control, circuit design, drug discovery, and targeted clinical therapy (to name a few), improving the accuracy of approximations to recurring optimization problems at a fraction of the current computation time would result in significant improvements in operational efficiencies and related costs. It is not unreasonable to expect millions, if not hundreds of millions of dollars in cost savings or improved revenues annually as well as improved operational performance.

Currently, strategies for improving accuracy and overcoming this computational challenge can be separated into four distinct categories. Each, however, presents shortcomings in its approach:

- **Heuristic algorithm** research and development have historically been responsible for reducing the time required to solve complex optimization problems. While the implementation of heuristic algorithms is the most widely accepted solution across academia and commercial industries, the challenges facing heuristics (namely exponentially increasing computation time and difficulty refining candidate solutions) can limit their performance.
- **High Performance Computing (HPC)** can be applied to heuristics that can be broken down into a set of smaller problems that can be solved simultaneously and then merged into a composite solution. Exponential growth of run times is not eliminated but the speed at which the approximations can be calculated can be improved significantly. A downside is that the accuracy of the approximation can be reduced since the problem isn't solved as a whole.

- **Cloud computing** allows users to tap into shared system resources provided via the internet. This allows organizations to outsource the burden of maintaining a computer infrastructure and potentially spread the computational burden across many processors. Even though cloud computing can supplement an organization's processing power, it does not resolve the reliance on heuristic algorithms nor overcome certain exponential requirements to solve problems.
- **Quantum computing** in theory presents the ability to store and process huge amounts of information through quantum bits that can exist in a superimposition of 0s and 1s. However, there are significant physics, engineering, and scaling challenges hindering the realization of fully functional quantum computers. In fact, using quantum computers to solve optimization problems with more than a few dozen variables has shown to be very difficult. While it isn't possible to predict exactly when and if these breakthroughs will occur, many estimates assume availability is at least 5-10 years away.

MemComputing is not a heuristic approach. It introduces a new paradigm where these problems are converted from a combinatorial problem into a physics problem. Solving the physical problem avoids the exponential growth associated with related heuristic approaches. This approach allows MemComputing to provide performance enhancements that are dramatically faster than HPC, Cloud Computing and Quantum Computing.

MemComputing Advantage

The unprecedented performance of MemComputing technology is based on the pioneering theoretical concept of **universal memcomputing machines** (UMMs). UMMs are a class of scalable memcomputing machines built with interconnected memory units (memprocessors) capable of performing computation.

UMMs are practically realized in the form of **digital memcomputing machines** (DMMs). DMMs harness the power of **self-organizing logic circuits** (SOLCs) which are differentiated from traditional circuits through the unique properties of the **self-organizing logic gates** (SOLGs) used in their construction. These SOLGs, in turn, can be realized in hardware with available technology (non-quantum) or simulated efficiently in software.

The proprietary components of the MemComputing system enable a non-combinatorial approach to solve combinatorial optimization problems. Thus, it removes the widespread dependency on heuristic algorithms and stochastic local searches. The features of SOLGs, the foundational components of MemComputing capabilities, provide the ultimate computational technology that is cost-effective, reduces computation time by orders of magnitude, and increases the quality of the solution to complex optimization problems.

Self-Organizing Logic Gates

Logic gates are the building blocks for digital circuits. Traditional logic gates are essentially miniature electric circuits that receive incoming electrical currents (inputs) and send an outgoing electrical current (output) based on what came in. The role of logic gates is to perform a logical

operation (e.g., AND, OR, NOT) on its one or more binary inputs to produce its singular binary output.

Unlike conventional logic gates, SOLGs can receive signals from *both* the traditional input and output terminals. In other words, SOLGs are “terminal-agnostic” and multi-directional. They accomplish this by *self-organizing* into their own logical proposition as well as in logical relations to another gate. For example, if a SOLG represents the logical operation OR it has an internal mechanism that propels its terminal states to fulfill the relationship defined by

$x_0 = x_1 \vee x_2$ (x_0 being the state of the traditional output terminal and x_1, x_2 being the states of the traditional input terminals). Then, setting x_0 will initiate the logic gate’s self-organization to produce x_1 and x_2 states *consistent* with x_0 and the truth table of an OR gate. The ability to set the state of the output terminal is not available in traditional logic gates.

As previously mentioned, SOLGs can be realized through standard, non-quantum electronic components. This capability is due to its dynamic correction module which is designed to correct inconsistent logic gate configurations. The error correction module reads voltages at the logic gate’s terminals and injects current when the gate is in an inconsistent configuration.

When SOLGs replace traditional logic gates, a self-organizing logic circuit (SOLC) is created. The dynamic, *collective* self-organization of all the SOLGs in the circuit allows SOLCs to solve complex optimization problems from any state selected at random by evolving to converge into equilibrium points. The equilibria represent approximations that come closer to the global optimum than current best solutions.

Nonlocal Collective State Computation

The most significant feature of SOLGs are their manifested *long-range order*. Long-range order describes physical systems which demonstrate correlated behavior across remote particles. In other words, systems with long-range order contain components that correspond to the states of other components regardless of distance. This simultaneous collective responsiveness of individual parts describes the temporal and spatial non-locality of the system.

The capability of SOLGs to realize long-range order is due to the existence of instantons. Instantons connect topologically inequivalent critical points in the phase space. They are the classical analogue of quantum tunneling. Instantons create non-locality in the system which generates the collective, dynamic behavior of SOLGs to correlate at an arbitrary distance. In effect, this collective behavior allows SOLGs to efficiently adapt their truth value to satisfy the logical proposition of another gate without violating their own internal logical proposition. The nonlocality of SOLGs allows for simultaneous variable flips which is a necessary task that combinatorial approaches cannot accomplish once they reach a certain number of satisfied constraints.

It is precisely the long-range order of SOLCs that produces computation acceleration by orders of magnitude. As discussed in greater detail in the next section on the demonstrated performance of Falcon[®], the system converges quickly to the equilibrium points which represent current closest approximations to the global optimum for complex optimization problems.

Scalability in Linear Time

Another key feature of the MemComputing solution is the ability of DMMs to scale polynomially and particularly linearly as demonstrated for the problem sets that map efficiently to a Max SAT problem space. It is important to note that the linear scaling is independent of the input size because the number of logic gates grows linearly with each step requiring only a linear number of floating-point operations and linearly growing memory. In other words, the number of variables can be increased without an exponential growth in computation time which resolves the primary issue with conventional computation solutions.

The configuration of DMMs outlined here is able to support infinite-range correlations in the infinite size limit. This support allows for an ideal scale-free behavior of the SOLC in which the correlations do not decay. This was derived analytically using topological field theory.²

Demonstrated Performance

In order to demonstrate the extraordinary performance of MemComputing technology, we compared Falcon[®] with the best solvers of the 2016 MaxSAT competition. Max-SAT problems are notoriously harder than traditional satisfiability (SAT) problems. Max-SAT problems cannot be satisfied. Instead, a solution must be found that identifies the maximum/largest number of clauses/constraints that can be satisfied.

It is worth noting that while Falcon[®] was tested as a single threaded interpreted MatLab application it was compared to the best solvers from the MaxSAT competition, CCLS and DeciLS. These solvers were designed using high level languages that are compiled software solutions designed for the greatest efficiency. Interpreted MatLab is notoriously slow when compared to compiled software. However, despite the anticipated performance gaps between the two, MemComputing offered better approximations in the least amount of time regardless of the increasing number of variables. Further, Falcon[®] was shown to scale linearly vs. exponentially, which only becomes more evident as the number of inputs increases.

The Challenge

Three types of problems tested (in order of increasing difficulty) were the random-Max-E3SAT, hyper-Max-E3SAT, and delta-Max-E3SAT. Two competing solvers were used as benchmarks for Falcon's[®] performance: CCLS and DeciLS. These were the best solvers in the 2016 Max-SAT competition.


Since these problems are considered NP-hard, Max-SAT algorithms can experience an inapproximability gap which says that no approximation can surpass a fraction of the global

optimal solution without incurring exponential overhead. Particularly for Max-E3SAT, the best approximation, under the conditions associated with classical computer algorithms where it is assumed that $NP \neq P$, is that a maximum of $\frac{7}{8}$ of the optimal number of constraints can be satisfied.³

The Procedure

Step 1: Construct a Boolean circuit that represents the optimization problem

- All optimization problems can be written in Boolean format and any standard Boolean formula can be written in conjunctive normal form (CNF)



Step 2: Replace traditional uni-directional Boolean gates with SOLGs

Step 3: Feed the appropriate terminals with the required output in CNF form (for instance, the logical 1 if we are interested in checking its satisfiability)

Step 4: The circuit is represented by non-linear ordinary differential equations which can be solved numerically to find the equilibrium points. Given the collective behavior of the circuit, the equilibria come extremely close to the global optimum

The Result

With a threshold of 1.5% of unsatisfiable clauses, all three solvers were tested for the length of time it would take to surpass the limit with an increasing number of clauses for the hardest cases, namely those of the delta-Max-E3SAT.

The exponential blowup of both CCLS and DeciLS is immediately apparent in Fig. 1. In comparison,

Falcon's non-combinatorial approach scaled linearly for up to 2×10^6 variables which necessitated 10^4 seconds (a little more than a couple of hours) to reach the unsatisfiable clause limit at 1.5%. Had CCLS and DeciLS been required to compute 2×10^6 variables, the computation time would, at best, run for 10^{2500} seconds which is 10^{2480} times the estimated age of the universe.




FIG. 1. Time comparison between CCLS, DeciLS, and Falcon for the delta-Max-E3SAT problem with a set threshold of unsatisfiable clauses at 1.5%. The projected times for CCLS and DeciLS were estimated using a linear regression of the $\text{Log}_{10}(\text{time})$ versus the number of variables.

Conclusion

In conclusion, we have shown empirical evidence that a non-combinatorial approach utilizing DMMs to solve hard combinatorial optimization problems exponentially outperforms heuristics designed to solve such problems.

Specifically, compared to the best solvers in the 2016 Max-SAT competition, DMMs:

- Found closer approximations to the global optimum
- Accelerated computation to a few hours (vs. the age of the universe) on a single thread of an Intel Xeon E5-2680
- Scaled linearly in time and memory

It is important to note that the linear scalability of DMMs does not prove the availability of polynomial solutions to NP-hard problems. Nevertheless, the outstanding performance of a physics-based approach to solving complex optimization problems demonstrates a promising trajectory for advancements in optimization computation.

Reference

[1] Traversa, Fabio L., Pietro Cicotti, Forrest Sheldon, and Massimiliano Di Ventra. Evidence of an exponential speed-up in the solution of hard optimization problems. *arXiv preprint arXiv:1710.09278* (2017).

[2] Di Ventra, M., Traversa, F. L. & Ovchinnikov, I. V. *Topological field theory and computing with instantons*. Ann. Phys. (Berlin) 1700123 (2017).

[3] Hastad, J. Some optimal inapproximability results. *Journal of the ACM* 48, 798-859 (2001).

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT.