

# UC San Diego

## UC San Diego Previously Published Works

### Title

Stress-Testing Memcomputing on Hard Combinatorial Optimization Problems

### Permalink

<https://escholarship.org/uc/item/2p1448pg>

### Journal

IEEE Transactions on Neural Networks and Learning Systems, 31(6)

### ISSN

2162-237X

### Authors

Sheldon, Forrest  
Cicotti, Pietro  
Traversa, Fabio L  
[et al.](#)

### Publication Date

2020-06-01

### DOI

10.1109/tnnls.2019.2927480

Peer reviewed

# Stress-testing memcomputing on hard combinatorial optimization problems

Forrest Sheldon, Pietro Cicotti, Fabio L. Traversa, Massimiliano Di Ventra

**Abstract**—Memcomputing is a novel paradigm of computation that utilizes dynamical elements with memory to both store and process information on the same physical location. Its building blocks can be fabricated in hardware with standard electronic circuits, thus offering a path to its practical realization. In addition, since memcomputing is based on non-quantum elements, the equations of motion describing these machines can be simulated efficiently on standard computers. In fact, it was recently realized that memcomputing, and in particular its *digital* (hence scalable) version, when simulated on a classical machine provides a significant speed-up over state-of-the-art algorithms on a variety of non-convex problems. Here, we stress-test the capabilities of this approach on finding approximate solutions to hard combinatorial optimization problems. These fall into a class which is known to require exponentially growing resources in the worst cases, even to generate approximations. We recently showed that in a region where state of the art algorithms demonstrate this exponential growth, simulations of digital memcomputing machines performed using the Falcon<sup>®</sup> simulator of MemComputing, Inc. only require time and memory resources that scale *linearly*. These results are extended in a stress-test up to  $64 \times 10^6$  variables (corresponding to about 1 billion literals), namely the largest case that we could fit on a single node with 128 GB of DRAM. Since memcomputing can be applied to a wide variety of optimization problems, this stress test shows the considerable advantage of non-combinatorial, physics-inspired approaches over standard combinatorial ones.

**Index Terms**—Memcomputing, Dynamical Systems, Optimization Problems

## I. INTRODUCTION

The increasing demand for computational power and efficiency is driving the scientific community and industry to explore new and unconventional ways to compute. In this respect, new ideas and radically different paradigms may be the key to solve or mitigate the computational bottlenecks that affect present computing technology.

A new computing paradigm has been recently proposed called *memcomputing* (which stands for computing *in* and *with* memory) [1], [2], with the potential to increase the computational efficiency in the solution of hard combinatorial/optimization problems. The formal description of memcomputing rests on the concept of Universal Memcomputing Machines (UMMs) [2]. UMMs are a collection of interconnected memprocessors (processors with memory) able to process and store information on the same physical location. UMMs are a class of non-Turing machines with specific properties as described in [2], [3]. In particular, they support

intrinsic parallelism, i.e., interconnected (mem-)processors act collectively on data using their collective state to perform computation [2], [4], [5]. Moreover, memprocessors are able to exploit the information available through the *topology* of their connections. Indeed, specific network topologies may be designed to embed problems in a one-to-one mapping. This last property has been named *information overhead* [2].

We have recently shown that the *digital* (hence scalable) subclass of UMMs we call *digital memcomputing machines* (DMMs) can deliver substantial benefits in the solution of hard combinatorial/optimization problems compared to traditional algorithmic approaches. For instance, in Ref. [6] we have shown that the simulation of DMMs on a classical computer solves the search version of the subset-sum problem in polynomial time for the worst cases. In Ref. [7] we have shown how to accelerate the pre-training of restricted Boltzmann machines using simulations of DMMs, and demonstrated their advantage over quantum-based machines (implemented in hardware), such as D-wave machines [8], as well as state-of-the-art supervised learning [9], [10]. Finally, in Ref. [11], we have employed simulations of DMMs and shown substantial advantages over traditional algorithms for a wide variety of optimization problems such as the Random 2 Max-SAT, the Max-Cut, the Forced Random Binary problem, and the Max-Clique [12]. In some cases, the memcomputing approach obtains the solution to problems on which the winners of the 2016 Max-SAT competition failed [11].

We have also performed scalability tests on finding approximate solutions to hard constructed instances of the Max-SAT problem [11]. Max-SAT possesses an *inapproximability gap* [13], meaning that even calculating an approximate solution beyond a certain fraction of the optimum will require resources which grow exponentially with the input size in the worst case. We observed this exponential growth on the tested instances with some of the best solvers from the 2016 Max-SAT competition. Instead, our simulations done for these hard cases, and up to a certain problem size succeed in overcoming that gap in linear time employing memory (of the single processor used for the simulations) that also scales linearly with input size [11].

Our goal for the present paper is to understand how far we can push these simulations with the available resources at our disposal. We focus again on hard combinatorial optimization problems since these arise in nearly every industrial and scientific discipline. These problems involve the minimization or maximization of a function of many independent variables, often called the cost function, whose value represents the quality of a given solution [14], [15]. In industrial settings this may be the wiring cost of a computer chip layout, or

F. Sheldon and M. Di Ventra are with the Department of Physics, University of California-San Diego, 9500 Gilman Drive, La Jolla, California 92093-0319, USA. P. Cicotti is with the San Diego Supercomputer Center, La Jolla, 92093 CA. F.L. Traversa is with MemComputing Inc., San Diego, CA, 92130 CA. e-mail: fsheldon@ucsd.edu, ftraversa@memcpu.com, pcicotti@sdsc.edu, diventra@physics.ucsd.edu

the length of a delivery route [16]. In scientific applications it may be searching for the ground state of a spin system or proteins [17].

We will show here that the simulations of DMMs on hard optimization problems can be pushed to tens of millions of variables, corresponding to about one billion literals, using a commercial and sequential MATLAB code (the Falcon<sup>®</sup> simulator provided by MemComputing, Inc.) running on a single thread of an Intel Xeon E5-2680 v3 with 128 Gb DRAM shared on 24 threads. These results show once more the power of physics-inspired approaches to computation over traditional algorithmic methods.

## II. THE PROBLEM

In this work, we focus on cost functions over a set of Boolean variables  $x_1, \dots, x_N$ ,  $x_i \in \{0, 1\}$ , into which a large number of problems may be cast.

A common formulation of these problems is given by a set of Boolean constraints where the cost function counts the number of constraints satisfied by an assignment. This is often expressed in conjunctive normal form (CNF) where each constraint (also called a *clause*) is expressed as the disjunction (logical OR denoted by  $\vee$ ) of a set of *literals* (a variable  $x_i$  or its negation  $\bar{x}_i$ ) which are then conjoined (logical AND denoted by  $\wedge$ ) together, e.g., in expressions of the type:

$$(x_1 \vee x_2) \wedge (x_1 \vee \bar{x}_2 \vee x_3) \wedge (x_1 \vee \bar{x}_2 \vee \bar{x}_3) \wedge (\bar{x}_1 \vee \bar{x}_2 \vee x_3) \\ \wedge (\bar{x}_1 \vee \bar{x}_3) \wedge (\bar{x}_2 \vee x_3) \wedge (\bar{x}_1 \vee x_2 \vee x_3).$$

The CNF representation may be considered general in the sense that any Boolean formula may be expressed in this form [12].

The problem of determining the assignment satisfying the maximum number of clauses is known as Max-SAT and is NP-hard, i.e., any problem in the class non-deterministic polynomial (NP) may be reduced to it in polynomial time [12]. As a result, algorithms for obtaining the solution will, in the worst cases, require a time that scales exponentially with the size of the instance, thus creating severe bottlenecks in the solution of complex industrial and scientific problems. The ubiquity and importance of this problem is exemplified by the yearly MAX-SAT competitions, where state-of-the-art solvers are tested on a variety of benchmarks to stimulate research and innovation [18].

In applications where the optimal solution is required *exact* algorithms must be used [19]. These *complete* solvers typically proceed by first obtaining bounds on the quality of the solution and then using these bounds to prune the search tree in a backtracking search. This systematic approach guarantees that the resulting solution will be the optimum, but typically scales poorly and is impractical for large instances.

In these cases, *heuristic* or *incomplete* algorithms must be used [19]–[21]. Rather than systematically searching the solution space, these solvers generate an initial assignment and then iteratively improve upon it, using a variety of strategies to boost the efficiency of their local search. After a specified number of steps, the algorithm returns its best assignment. As randomness is often used to drive the search, this procedure is referred to as *stochastic local search*. While

they can no longer guarantee optimality, these solvers have proven very effective at approximating, and sometimes solving difficult Max-SAT instances. For instance, in the 2016 Max-SAT competition [18], incomplete solvers performed 2 orders of magnitude faster than complete solvers on random and crafted benchmarks.

We might hope that if we seek an *approximation*, rather than a solution of a Max-SAT instance, we could avoid the exponential scaling of the run-time. Unfortunately, it turns out that even approximating the solution of many difficult optimization problems is NP-hard. More precisely, for a maximization problem with optimum  $O$  defined as the sum of the weights of all satisfied clauses, obtaining an assignment better than  $fO$  for the fraction  $f$  greater than some critical fraction,  $f_c$ , is an NP-hard problem [13], [22]. For example, obtaining an assignment for Max-E3SAT (a version of Max-SAT in which every clause has 3 literals) better than  $f_c = 7/8$  of the optimum is NP-hard, meaning that we cannot expect a polynomial algorithm to obtain the approximation for any instance of Max-E3SAT unless NP=P. Any known algorithm thus must show exponential scaling for a threshold past  $f_c$  in the worst case. In [11] we showed a region in which some of the best algorithms based on stochastic local search show an exponential growth with input size, but where the memcomputing approach based on deterministic dynamical systems requires only linearly growing time to achieve the same threshold.

## III. DIGITAL MEMCOMPUTING MACHINES

As mentioned in the Introduction, we present here a radically novel approach to these problems based on the simulation of DMMs on a standard classical computer [2], [4]. DMMs can be implemented in practice as specially-designed dynamical systems whose equilibrium (fixed) points represent the approximations to the computational problem at hand, and which can be realized with standard electrical components and those with memory.

The mathematical definition of DMMs is the eight-tuple [4]

$$DMM = (\mathbb{Z}_2, \Delta, \mathcal{P}, S, \Sigma, p_0, s_0, F), \quad (1)$$

where (although not strictly necessary) we consider the range of  $\mathbb{Z}_2 = \{0, 1\}$ . Generalization to any finite number of states is trivial.  $\Delta$  is a set of functions

$$\delta_\alpha : \mathbb{Z}_2^{m_\alpha} \setminus F \times \mathcal{P} \rightarrow \mathbb{Z}_2^{m'_\alpha} \times \mathcal{P}^2 \times S, \quad (2)$$

where  $m_\alpha < \infty$  is the number of memprocessors used as input of (read by) the function  $\delta_\alpha$ , and  $m'_\alpha < \infty$  is the number of memprocessors used as output (written by) the function  $\delta_\alpha$ ;  $\mathcal{P}$  is the set of the arrays of pointers  $p_\alpha$  that select the memprocessors called by  $\delta_\alpha$  and  $S$  is the set of indexes  $\alpha$ ;  $\Sigma$  is the set of the initial states written by the input device on the computational memory;  $p_0 \in \mathcal{P}$  is the initial array of pointers;  $s_0$  is the initial index  $\alpha$ , and  $F \subseteq \mathbb{Z}_2^{m_f}$  for some  $m_f \in \mathbb{N}$  is the set of final states.

The DMM approach to a specific Boolean problem may be summarized as follows [4], [6]:

- 1) The Boolean circuit representing the problem is constructed.




Fig. 1. Simulation time comparison between the incomplete solver DeciLS against the Falcon<sup>®</sup> solver of MemComputing, Inc. for the balanced and constrained delta-Max-E3SAT. A threshold of 1.5% of unsatisfiable clauses has been set. We have then tested how long DeciLS and our solver Falcon<sup>®</sup> take to overcome this limit with increasing number of variables. All calculations have been performed on a single thread of an Intel Xeon E5-2680 v3 with 128 Gb DRAM shared on 24 threads. The local solver requires an exponentially increasing time to reach that limit. Our memcomputing approach instead scales linearly even up to  $64 \times 10^6$  variables (corresponding to about 1 billion literals), which required a little over than  $10^5$  seconds on a single thread. We could not go beyond this limit because of memory resources (see text and Fig. 2).

- 2) Traditional logic gates are replaced with *self-organizing logic gates* (SOLGs) [4] which are circuit elements designed to interact and collectively organize into a logically-consistent state.
- 3) The resulting *self-organizing logic circuit* (SOLC) is fed voltages along its boundary, and allowed to evolve until it reaches equilibrium where the results of a computation may be read out.

A detailed account of SOLGs may be found in [4] (see also Ref. [6]). In essence, they may be understood as dynamical components whose equilibrium points encode the truth table of a logic gate, and can self-organize into a consistent logical relation of such truth table *irrespective* of the terminal to which a given truth value (a literal) is assigned. Each terminal is equipped with a *dynamic error correcting module* which reads the states of its neighbors, and sets the local current and voltage to enforce logical constraints. These elements are specially designed so that the resulting dynamical system is point dissipative [23] (in the context of functional analysis), avoids periodic orbits [24] and the chaotic behavior [25] that is typical of other non-linear dynamical systems, and utilizes components with memory to allow gates to efficiently correlate and *collectively* transition between states [26] (see also below).

The dynamics of the system are *deterministic* and, since they allow for collective transitions of large numbers of variables, they are also *non-local* in the same sense that complete combinatorial solvers are. Therefore, our approach contrasts sharply with those based on a stochastic local search such as simulated annealing [16], and many incomplete solvers [19], [20]. We

will comment further on the reasons for these differences and the role of long-range order later.

The procedure above may be followed to construct a hardware solver for a wide variety of combinatorial and optimization problems. However, since memcomputing machines are made of non-quantum elements, the resulting behavior is also captured by a set of nonlinear ordinary differential equations describing the circuit (see Ref. [4] for an example of DMMs equations of motion). These equations can be efficiently simulated, constituting an *algorithm* for the same problem. While this approach may seem indirect, as already noted, surprisingly the simulation of these circuits using standard numerical packages [27] is sufficient to outperform the state-of-the-art combinatorial approaches on many benchmarks [6], [7], [11].

In subsequent sections we demonstrate this and stress test these simulations by showing the results of a direct comparison on the same hardware between the Falcon<sup>®</sup> solver of MemComputing, Inc. and the solver DeciLS, an improved version of one of the best solvers of the 2016 Max-SAT competition [28]. In all cases our simulations considerably outperform the other method tested, by orders of magnitude. Our solver indeed scales *linearly* in time and memory compared to the expected exponential scaling of the other solver.

#### IV. MAX-SAT

As outlined above, we formulate our instances as Max-E3SATs, in which each clause contains exactly 3 literals and which has an inapproximability gap. A particular instance of




Fig. 2. Memory requirements of the Falcon<sup>®</sup> solver as a function of variables for the delta-Max-E3SAT. We provide both the input size memory (open circles) and the RAM used during computation.

Max-E3SAT may be characterized by the number of variables,  $N$ , and the number of clauses,  $M$ , or alternatively the clause density,  $\rho = M/N$ . As  $\rho$  increases, the relationships between the variables become increasingly constrained, and it becomes less likely that an assignment satisfying all clauses will exist, i.e., the problem is more likely UNSAT.

In the case of Random-3SAT, in which clauses are uniformly drawn from the set of possible clauses amongst  $N$  variables, instances undergo a SAT/UNSAT transition at  $\rho \approx 4.3$ , below which an instance is almost certainly satisfiable, and above which it is almost certainly unsatisfiable (the transition being sharply defined as  $N \rightarrow \infty$ ). Different methods of generating instances however, will lead to different transitions: Random-3XORSAT, in which clauses of 3 literals are formed using the exclusive OR,  $\oplus$ , undergoes a SAT/UNSAT transition at  $\rho \approx 0.918$  [29].

To stress test the capabilities of simulating DMMs we utilized a set of Max-E3SAT instances based on a variant of Random-3XORSAT in which each variable was constrained to occur the same number of times (or as nearly as possible while satisfying the specified  $N$  and  $\rho$ ) [11]. This particular set of instances were chosen for their low inter-instance variability in difficulty due to the constraint placed on variable occurrences, and the relevance of MAX-XORSAT to important problems in decoding [29]–[31]. The balanced structure of the resulting SAT makes them difficult for local solvers since, whenever a variable assignment is changed, it flips the state of all XORSAT clauses in which it was included. Thus, in order to find transitions which satisfy a larger number of clauses, the system is forced to flip increasingly large numbers of literals concurrently.

To generate these hard instances, we first generated a

random 3XORSAT instance with  $\rho = 1.25$  in which each variable was allowed to occur 3 or 4 times. These were then converted to a SAT instance by replacing each XORSAT clause with 4 CNF clauses which reproduce the truth table of the XORSAT clause. The resulting Max-E3SAT instances have a clause density of  $\rho = 5$ . We call this problem delta-Max-E3SAT [11].

To clearly show the superior performance of our approach compared to standard algorithms, we have set a threshold of 1.5% of unsatisfiable clauses and tested how long DeciLS and Falcon<sup>®</sup> take to overcome this limit with increasing number of variables. DeciLS is a compiled code that combines a unit propagation based decimation (Deci) and local search (LS) with restarts [28]. Past useful assignments are used to break conflicts in the unit propagation and may be considered “messages” from the past.

The solver Falcon<sup>®</sup> is a sequential MATLAB code that integrates forward the DMMs equations of motion using an explicit Euler method. All calculations have been performed on a single core of an Intel Xeon E5-2680 v3 with 128 GB DRAM. As expected the local solver requires an exponentially increasing time to reach that limit already at a few thousand variables (see squares symbols in Fig. 1). Instead, the simulations of DMMs scale linearly in time up to  $64 \times 10^6$  variables, which at a density of 5, corresponds to  $320 \times 10^6$  clauses, or about 1 billion literals. The largest case required a little over than  $10^5$  seconds to complete. In fact the time was not a limitation, rather the memory of the processor.

In Fig. 2 we present the memory used in the computation by Falcon<sup>®</sup> as a function of variables. It is clear that also the memory scales linearly up to  $64 \times 10^6$  variables. For comparison, we also display the memory size of the input

file, showing that the MATLAB implementation of DMMs equations of motion has about an order of magnitude overhead in memory. Since the memory of the Intel Xeon used for these simulations was only 128 GB DRAM, that limit has set a hard stop to the actual size we could fit on that processor. Of course, different implementations (e.g., using a compiled language rather than an interpreted one), different hardware, etc. may permit execution of even larger instances.

## V. COLLECTIVE DYNAMICS OF DMMs

The simulations surveyed in this paper and those already performed in relation to other problems [6] indicate that *collective* dynamics play an important role in the ability of DMMs to find good approximations to the optimum in non-convex landscapes. This contrasts sharply with the dynamics of stochastic-local-search solvers, and an analogy may be drawn to the co-tunneling events observed in quantum-annealing devices such as those manufactured by D-Wave [32], although DMMs are non-quantum machines.

Of particular interest is the fact that the transient dynamics of a DMM proceeds via an *instantonic phase* where the machine rids itself of logical defects [26], [33]. Instantons are families of trajectories of the non-linear equations of motion of DMMs, and connect different critical points in the phase space that have different indexes, namely different number of unstable directions.

In mathematical terms, if

$$\dot{x}(t) = F(x(t)) \quad (3)$$

is the equation of motion describing a DMM, with  $x$  the set of elements (e.g., voltages, currents and internal state variables), and  $F$  the flow vector field, then instantons are deterministic trajectories

$$\dot{x}_{cl}(t, \sigma) = F(x_{cl}(t, \sigma)); \quad x_{cl}(\pm\infty, \sigma) = x_{a,b}, \quad (4)$$

that connect two arbitrary critical points of  $F$ , say  $x_b$  and  $x_a$ . The parameters  $\sigma$  are the so-called moduli of instantons, and encode their non-locality.

Indeed, the presence of instantons creates long-range order in the system, both in time and in space [26]. Spatial long-range order means that logic gates can correlate at arbitrary distances from each other. Temporal long-range order implies that the system can follow different paths to obtain the same solution.

It is the presence of these instantons that renders these machines efficient in the solution search of complex problems. In fact, by transforming a combinatorial or optimization problem (a Boolean problem) into a physical (dynamical) one as described above, its simulation is also efficient. The reason is that the original Boolean problem is no longer solved combinatorially, as it is usually done. A combinatorial approach requires a search in a space that grows exponentially for a complex problem. Rather, the solution search is accomplished, via instantons, by a physical, dynamical system and the latter can be simulated efficiently by solving its corresponding differential equations. In addition, since instantons are topological objects, the solution search is robust against noise and structural disorder, a fact that was also shown explicitly in Ref. [33].

## VI. CONCLUSIONS

The performance of digital memcomputing machines on the benchmarks presented in this paper demonstrates the substantial advantages of our approach, based on the *simulation* of non-linear dynamical systems, compared to traditional combinatorial ones. While we have focused on the maximum-satisfiability problem, the methods we have illustrated readily generalize to a wide variety of combinatorial optimization problems. It would then seem that physics-based approaches offer a lot to the world of computing, and we believe these ideas may form the basis for the next generation of computational devices.

## ACKNOWLEDGMENTS

We sincerely thank Dr. Shaowei Cai for providing us with the binary compiled code DeciLS. We also thank Haik Manukian and Robert Sinkovits for helpful discussions. M.D. and F.L.T. acknowledge partial support from the Center for Memory Recording Research at UCSD. M.D. and F.S. acknowledge partial support from MemComputing, Inc. All calculations reported here have been performed by one of us (P.C.) on a single processor of the Comet cluster of the San Diego Supercomputer Center, which is an NSF resource funded under award #1341698. The authors would be delighted to provide, upon request, all instances of the constrained delta-Max-E3SAT used to generate Figs. 1 and 2.

## REFERENCES

- [1] M. Di Ventra and Y. V. Pershin, "The parallel approach," *Nature Physics*, vol. 9, pp. 200–202, 2013.
- [2] F. L. Traversa and M. Di Ventra, "Universal memcomputing machines," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 26, no. 11, p. 2702, 2015.
- [3] Y. Pei, F. L. Traversa, and M. Di Ventra, "On the universality of memcomputing machines," *arXiv:1712.08702*, 2017.
- [4] F. L. Traversa and M. Di Ventra, "Polynomial-time solution of prime factorization and np-complete problems with digital memcomputing machines," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 27, p. 023107, 2017.
- [5] F. L. Traversa, C. Ramella, F. Bonani, and M. Di Ventra, "Memcomputing NP-complete problems in polynomial time using polynomial resources and collective states," *Science Advances*, vol. 1, no. 6, p. e1500031, 2015.
- [6] M. D. Ventra and F. L. Traversa, "Memcomputing: Leveraging memory and physics to compute efficiently," *J. Appl. Phys.*, vol. 123, p. 180901, 2018.
- [7] H. Manukian, F. L. Traversa, and M. Di Ventra, "Accelerating deep learning with memcomputing," *arXiv:1801.00512*, 2018.
- [8] S. H. Adachi and M. P. Henderson, "Application of quantum annealing to training of deep neural networks."
- [9] X. Glorot, A. Bordes, and Y. Bengio, "Deep sparse rectifier neural networks," in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, 2011, pp. 315–323.
- [10] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proceedings of the 32nd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, F. Bach and D. Blei, Eds., vol. 37. Lille, France: PMLR, 07–09 Jul 2015, pp. 448–456.
- [11] F. L. Traversa, P. Cicotti, F. Sheldon, and M. Di Ventra, "Evidence of exponential speed-up in the solution of hard optimization problems," *Complexity (in press)*; *arXiv:1710.09278*.
- [12] M. R. Garey and D. S. Johnson, *Computers and Intractability; A Guide to the Theory of NP-Completeness*. New York, NY, USA: W. H. Freeman & Co., 1990.
- [13] J. Håstad, "Some optimal inapproximability results," *Journal of the ACM*, vol. 48, no. 4, pp. 798–859, jul 2001.

- [14] K. S. Christos H. Papadimitriou, *Combinatorial Optimization*. Dover Publications Inc., 1998.
- [15] S. H. Z. Edwin K. P. Chong, *An Introduction to Optimization*. JOHN WILEY & SONS INC, 2013.
- [16] S. Kirkpatrick, C. D. Gelatt, M. P. Vecchi *et al.*, “Optimization by simulated annealing,” *Science*, vol. 220, no. 4598, p. 671, 1983.
- [17] S. Kobe and J. Krawczyk, “Ground states, energy landscape, and low-temperature dynamics of  $\pm j$  spin glasses,” in *Computational complexity and statistical physics*, A. Percus, G. Istrate, and C. Moore, Eds. OUP USA, 2006.
- [18] [Online]. Available: <http://www.maxsat.udl.cat/16/index.html>
- [19] C. P. Gomes, H. Kautz, A. Sabharwal, and B. Selman, “Satisfiability solvers,” in *Handbook of knowledge representation*, F. Van Harmelen, V. Lifschitz, and B. Porter, Eds. Elsevier, 2008, vol. 1, ch. 2, pp. 89–134.
- [20] H. A. Kautz, A. Sabharwal, and B. Selman, “Incomplete algorithms,” in *Handbook of satisfiability*, A. Biere, M. Heule, and H. van Maaren, Eds. IOS press, 2009, vol. 185, pp. 185–203.
- [21] J. Hromkovic, *Algorithmics for Hard Problems: introduction to combinatorial optimization, randomization, approximation, and heuristics*. Springer, 2010.
- [22] U. Feige, “A threshold of  $\ln n$  for approximating set cover,” *Journal of the ACM*, vol. 45, no. 4, pp. 634–652, jul 1998.
- [23] J. Hale, *Asymptotic Behavior of Dissipative Systems*, 2nd ed., ser. Mathematical Surveys and Monographs. Providence, Rhode Island: American Mathematical Society, 2010, vol. 25.
- [24] M. Di Ventra and F. L. Traversa, “Absence of periodic orbits in digital memcomputing machines with solutions,” *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 27, p. 101101, 2017.
- [25] M. Di Ventra and F. L. Traversa, “Absence of chaos in digital memcomputing machines with solutions,” *Phys. Lett. A*, vol. 381, p. 3255, 2017.
- [26] M. Di Ventra, F. L. Traversa, and I. V. Ovchinnikov, “Topological field theory and computing with instantons,” *Ann. Phys. (Berlin)*, vol. 529, p. 1700123, 2017.
- [27] [Online]. Available: <https://www.mathworks.com/>
- [28] S. Cai, C. Luo, and H. Zhang, “From decimation to local search and back: A new approach to maxsat,” in *to appear in proceedings of International Joint Conference on Artificial Intelligence*, 2017.
- [29] F. Ricci-Tersenghi, M. Weigt, and R. Zecchina, “Simplest random k-satisfiability problem,” *Physical Review E*, vol. 63, no. 2, p. 026702, 2001.
- [30] H. Jia, C. Moore, and B. Selman, “From spin glasses to hard satisfiable formulas,” in *International Conference on Theory and Applications of Satisfiability Testing*. Springer, 2004, pp. 199–210.
- [31] S. Cocco, R. Monasson, A. Montanari, and G. Semerjian, “Analyzing search algorithms with physical methods,” 2006.
- [32] V. S. Denchev, S. Boixo, S. V. Isakov, N. Ding, R. Babbush, V. Smelyanskiy, J. Martinis, and H. Neven, “What is the computational value of finite-range tunneling?” *Phys. Rev. X*, vol. 6, p. 031015, Aug 2016. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevX.6.031015>
- [33] S. R. B. Bearden, H. Manukian, F. L. Traversa, and M. Di Ventra, “Instantons in self-organizing logic gates,” *Physical Review Applied*, vol. 9, p. 034029, 2018.