Lossy Data Compression with Random Gates

Stefano Ciliberti,[1] Marc Mézard,[1] and Riccardo Zecchina[2]

[1]*Laboratoire de Physique Théorique et Modèles Statistiques, Université de Paris-Sud, Bâtiment 100, 91405, Orsay Cedex, France*
[2]*ICTP, Strada Costiera 11, I-34100 Trieste, Italy*

We introduce a new protocol for a lossy data compression algorithm which is based on constraint satisfaction gates. We show that the theoretical capacity of algorithms built from standard parity-check gates converges exponentially fast to the Shannon's bound when the number of variables seen by each gate increases. We then generalize this approach by introducing random gates. They have theoretical performances nearly as good as parity checks, but they offer the great advantage that the encoding can be done in linear time using the survey inspired decimation algorithm, a powerful algorithm for constraint satisfaction problems derived from statistical physics.

Introduction.—Constraint satisfaction problems (CSPs) are at the heart of an emerging field of research which is of interest to statistical physics, combinatorial optimization, statistical inference, and information theory [1]. Broadly speaking, these are problems involving a large number of variables, taking values in a finite set (hereafter we shall keep to binary variables). Each constraint involves K variables, and imposes a probability law on the 2^K possible assignments of the variables in this subset. Hard constraints just forbid some of the configurations. The spin glass problem [2], the satisfiability problem which lies at the heart of the theory of computational complexity in computer science [3], or the parity-check problems used in error correcting codes [4] all belong to this category.

A lot of progress has been made in recent years in the study of random constraint satisfaction problems where each constraint involves randomly chosen (with uniform distribution) variables [5]. This is the natural setting for spin glasses; it offers the possibility to study issues in typical case complexity in satisfiability, and it provides some of the most efficient codes for error correction. In several cases, it has been found that when the density of constraints increases the system enters first a "clustered" phase before it reaches the threshold of unsatisfiability where it cannot meet all the constraints. Above this threshold the configurations which violate the smaller number of constraints are also clustered. Clustering means that the configurations which satisfy all the constraints are grouped into many disconnected clusters which are distant from each other. Statistical physics methods originating from spin glass theory, like the replica and cavity methods, turn out to be very efficient to study these phenomena [6,7], and some of the results have been confirmed rigorously recently [8–11]. They have also led to a powerful algorithm (survey propagation) which is able to solve very large problems in the clustered region [7].

We show how one can take advantage of these clustered phases to address a classic problem in coding theory, lossy data compression. While a large amount of work has been done in this field [12], a number of challenging problems are open, among them the realization of a practical compression protocol for correlated sources or the exponential increasing time in the encoding-decoding step of typical algorithms. As for lossy compression schemes, it is worthwhile to mention, in particular, the good performance of algorithms based on the codes [13] developed in the context of channel coding. Here we propose an alternative strategy, and as a starting point we focus on the case of uncorrelated sources.

The problem of lossy data compression can be summarized as follows. We have a source alphabet $\mathcal{A}$, a source distribution $p(x_a)$, and a distortion measure $d(\cdot, \cdot)$ which takes values in $[0, 1]$. We start from an original message which is a sequence $\{x_a\}$ of M values independently drawn from the source distribution. The purpose of data compression is to map this message to a string of N bits, with $N < M$, in such a way that they can be, for example, easily transmitted or stored. Then, one wants to decode this N-bits string in order to reconstruct a sequence as close as possible to the original message. We call the decoded message $\{x_a^*\}$ and we want to minimize the expected value of the distortion $D \equiv \mathbb{E}\sum_{a=1}^{M} d(x_a, x_a^*)/M$. How small a distortion one can achieve depends on the rate $R = N/M$ at which the original message has been compressed.

The rate-distortion theorem proved by Shannon [14] provides a bound for the minimum rate at which a compression is possible once we fix the average distortion we tolerate. The analytic expression of this *rate-distortion function* $R(D)$ is not known explicitly in the most general case of correlated (memory) sources, and it is most often obtained by means of some numerical algorithm (see, e.g., [15]). On the other hand, for an uncorrelated unbiased binary source (i.e., with $p(x_a = 0) = 1 - p(x_a = 1) = 1/2$), the rate-distortion function in the large N limit has the simple expression

$$R(D) = 1 + D\log_2 D + (1 - D)\log_2(1 - D). \quad (1)$$

As a first requirement, a good lossy data compression

algorithm must be able to approach this theoretical limit. One should mention that in the lossless case, that is, the $D \to 0$ limit, practical algorithms that asymptotically saturate the Shannon's bound were discovered a long time ago [16]. The lossy case turns out to be more difficult from the algorithmic point of view. Recently, a perceptron with a nonlinear transfer function has been proposed [17] as a lossy compressor and it has been analytically shown to achieve theoretically optimal performance, but its practical use is strongly reduced by the fact that there is no known polynomial algorithm in this case, and the typical string lengths that can be compressed in reasonable time are thus rather short.

The general idea.—We are interested here in developing an approach to the problem which is based on CSPs, as suggested in [18]. Our CSP uses M constraints between N Boolean variables taking values in $\{0, 1\}$. Each constraint, say a, is actually a gate controlled by the value x_a of the ath bit of the original message (see Fig. 1). The gate a is connected to K_a randomly chosen variables. The 2^{K_a} possible configurations of these variables are partitioned into two equal size subsets, S_a and U_a. When the control bit is $x_a = 0$, the configurations in S_a satisfy the constraint; the configurations in U_a do not. When the control bit is $x_a = 1$, the configurations in U_a satisfy the constraint; the configurations in S_a do not. A simple example is provided by parity checks. S_a consists of the configurations with an even number of 0's. The gate then performs a linear operation: it checks whether the sum of all variables and b_a is equal to 0 modulo 2. In this case the CSP is nothing but the well-known XORSAT problem in computer science [19]; it can also been seen as a spin glass problem with three-spin interactions.

In our procedure, the initial word of M bits is used to build a CSP with M constraints for $N(<M)$ variables $\{y_1, \ldots y_N\}$. We then look for a configuration of variables $\mathbf{y}^*$ that minimizes the number of violated constraints, that is, the ground state configuration. This is the encoded (compressed) word. Of course, this step is nontrivial since one must be able to handle a CSP which is in its "unsat" phase. We note that the rate R of the process is simply related to the density of constraints $\alpha = M/N$ by $R =$
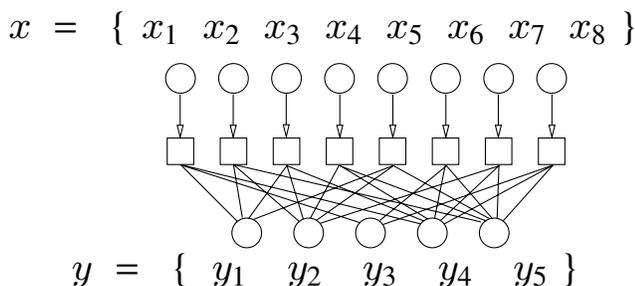

FIG. 1. A Tanner graph is a convenient representation for a CSP. We emphasize in this cartoon the topological support of our protocol.

$1/\alpha$. Once we have a ground state configuration, the decoding is easy: for each constraint a, one considers the configuration of the subset of K_a variables appearing in a. If it lies in S_a, the reconstructed bit is $x_a^* = 0$, otherwise it is $x_a^* = 1$. The number of bits of the original message which are wrongly reconstructed is nothing but the number of constraints violated in a ground state configuration.

We shall measure the distortion as $D = \sum_{a=1}^{M} |x_a - x_a^*|/M$. We define the total "energy" of a configuration $\mathbf{y}$ of the CSP as $E(\mathbf{y}) = \sum_{a=1}^{M} \varepsilon_a$, where $\varepsilon_a(\mathbf{y}) = 0$ if the constraint a is satisfied by the global configuration $\mathbf{y}$, and $\varepsilon_a(\mathbf{y}) = 1$ otherwise. The distortion is then related to the ground state energy E_0 of the CSP through

$$D = E/M. \qquad (2)$$

We are interested in the thermodynamic limit $N, M \to \infty$ at fixed density of constraints α. Shannon's theorem provides a lower bound $E_{\mathrm{Sh}}(\alpha)$ to the ground state energy, and a good compression algorithm should be based on a CSP with a very low ground state energy, as near as possible to Shannon's value.

The coder based on parity-check gates (XORSAT-CSP) is a good candidate. A general strategy for computing the ground state energy of this problem has been developed in [8]. When all checks involve K variables and K becomes large, a computation based on this strategy [20] shows that $E_0(\alpha) - E_{\mathrm{Sh}}(\alpha)$ decreases exponentially with K. So the theoretical capacity of these gates rapidly approach Shannon's limit when K increases. Unfortunately there is no known algorithm which matches this theoretical capacity. This is in contrast to the use of low density parity-check codes for channel coding, where message passing techniques are known to perform quite well. However, we shall see below that message passing does perform well on some other classes of gates.

Message passing.—Useful techniques for solving random CSPs are based on local iterative updates of some "messages" sent along the graph. For example, applying the "min-sum" algorithm [21] to CSPs, one obtains the warning propagation algorithm: each constraint a sends a warning message $u_{a \to i}$ to one neighbor variable i according to the values of the other variables attached to it: This message can be 0—meaning that the variable is free to assume any value [22]—or 1—meaning that, in order to satisfy that constraint, the variable should assume a certain value. This algorithm is very powerful and efficient in many CSPs where the underlying factor graph is locally treelike, when the density of constraints is small enough. However it is limited in two aspects: (1) it stops to converge when the density of constraint is such that the system is in a clustered phase, and, in particular, in the unsat regime which is of interest for our compression scheme. (2) It does not work for parity checks because of the basic symmetry of these gates.

The first limitation can be handled by going to a more sophisticated message passing algorithm, survey propagation (SP). The SP algorithm is the direct implementation of

the one-step replica symmetry breaking (1RSB) cavity
equations on a single sample. In this case, one works
with the probability distribution $Q_{a \to i}(u_{a \to i})$ that, if we
pick one cluster at random, the warning $u_{a \to i}$ is sent along
the link from constraint a to variable i. This is the general
object which is needed in order to cope with the appear-
ance of disconnected clusters of solutions. Once we know
the probabilities of all the warnings, we estimate the
probability of the total *bias* H_i on the variable i, defined
as $H_i = \Sigma_{a \in V(i)} u_{a \to i}$, where $V(i)$ is the set of constraints
attached to the variable i. The idea behind the survey
inspired decimation [7] is to take advantage of this infor-
mation to fix the most biased variable in the system. Once
this is done, one has a simplified CSP with $N - 1$ variables
and we can thus repeat this step until one is left with an
under-constrained problem solvable by some standard
local algorithm. This algorithm has been shown to be
very useful in CSP problems like the coloring or the K-
SAT, where one is able to find efficiently a SAT assignment
in the difficult region. It can also be used for problems with
higher constraint density in order to find configurations of
low energy (small number of violated constraints) [23],
which makes it a very useful tool for compression.

Nonlinear nodes.—While it performs well on many
CSPs, the SP algorithm is useless for the "parity source
coder." The problem here comes from the fact that the
distribution of the total bias is always symmetric, so that
the messages obtained after convergence give no hint of
how to decimate the problem. In order to solve this prob-
lem, we propose a compression scheme based on some
other gates, different from parity checks. These turn out to
have a theoretical capacity close to the parity checks, and a
generalized version of the SP algorithm [20] leads to a
convergent decimation scheme which is an efficient coding
algorithm. Among the several types of constraints we have
examined, the "random" nodes have been found to be the
more efficient. They are defined as follows: The subset S_a
is a randomly chosen subset of size 2^{K-1}. While parity
checks just implement linear constraints on Z_2, these ran-
dom nodes are nonlinear functions of their inputs. However
from the point of view of the cavity method (used to
compute their theoretical performance) and of the SP
message passing procedure (used as encoding algorithm),
they can be handled by relatively straightforward general-
izations of the methods used for parity checks [20]. In this
note we summarize the results.

We build up a list of r random checks, and each con-
straint a picks up one check randomly in this list [24]. This
allows memorization of the truth table of all nodes and thus
speeds up the algorithms. All the results quoted below are
for $r = 30$. The theoretical capacity of this system, which
is proportional to the ground state energy according to (2),
is illustrated in Fig. 2. As we clearly see, the ground state
quickly approaches the Shannon's bound of Eq. (1) as K
increases. Thus, this particular CSP is very promising from
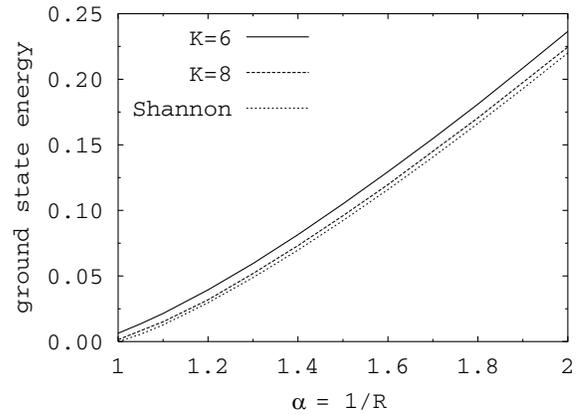the point of view of lossy data compression. In Fig. 3 we

FIG. 2.   The ground state energy of the CSP based on nonlinear
nodes versus the constraint density α. Shannon's bound is also
plotted for comparison.

show the phase diagram for $K = 6$. The *static* energy is the
ground state energy per variable also plotted in Fig. 2; from
the algorithmic point of view, this is the performance of the
best possible algorithm which minimizes the number of
violated constraints. The *dynamical* energy marks the ap-
pearance of a regime where solutions group in many differ-
ent well separated clusters (that is, in order to go from one
cluster to another one we should flip an extensive number
of variables); any local algorithm, as, for example, warning
propagation, will be trapped at this dynamical threshold.
The dynamical energy computed here in the 1RSB ap-
proximation is believed to be an upper bound to the exact
one. Finally, the stability curve [9] indicates the range of
validity of the 1RSB formalism used to determine this
phase diagram (in particular, the ground state energy which
we compute should be the exact one for $\alpha < \alpha_{1\text{RSB}}$).

So the theoretical properties of the random-node CSP
are quite similar to the parity-check CSP (the XORSAT
problem discussed in [8]). The good point with respect to
the parity-check CSP is that in this case the survey inspired
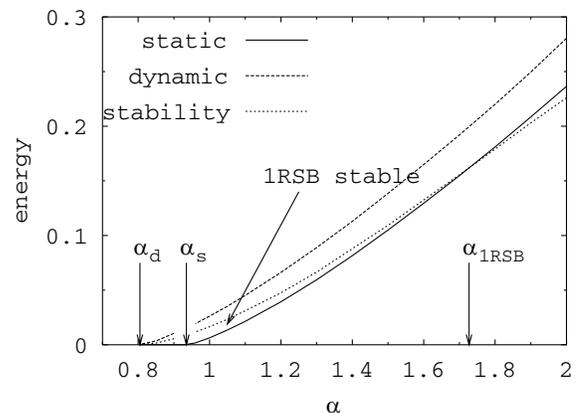decimation algorithm does converge in the unsat phase in a

FIG. 3.   The phase diagram of a system with 30 different
random nodes with $K = 6$. The values for the thresholds are:
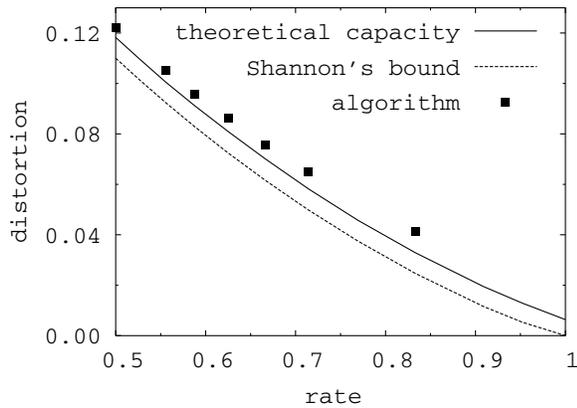$\alpha_d = 0.803$, $\alpha_s = 0.935$ and $\alpha_{1\text{RSB}} = 1.727$.

FIG. 4.   The performance of the algorithm is plotted versus the rate $R = 1/\alpha$ of the compression (here $K = 6$, $N = 1000$), together with the theoretical capacity and Shannon's bound.

time which scales as $2^K N \log N$, and gives very low energy states, i.e., nearly optimal global configurations. In Fig. 4 we show the performance of the compressor based on random gates, for $K = 6$. The distortion achieved in practice by the algorithm with $N = 1000$ is close to the theoretical capacity, which brings it a few percent above Shannon's bound. As shown in Fig. 2, we expect the performance to improve with increasing K (at the price of an increase in computer time).

Conclusions.—We have shown, by using techniques borrowed from the statistical physics of disordered systems, how one can use CSPs as a tool for compressing data. In particular, the algorithmic performance of the random gates—CSPs based on nonlinear nodes—is found to be nearly optimal, since the Shannon bound is reached at large K. The generalization of the present approach to compression of data from a larger alphabet (beyond binary input) looks like an interesting perspective.

[1] Special Issue, edited by O. Dubois, R. Monasson, B. Selman, and R. Zecchina [Theor. Comput. Sci. **265,** 1 (2001)].

[2] M. Mézard, G. Parisi, and M. A. Virasoro, (World Scientific, Singapore, 1987);*Spin Glasses and Random Fields*, edited by A. P. Young (World Scientific, Singapore, 1998).

[3] C. H. Papadimitriou, *Computational Complexity* (Addison-Wesley, Reading, MA, 1994).

[4] R. C. Gallager, *Information Theory and Reliable Communication* (Wiley, New York, 1968).

[5] M. Mézard, Science **301,** 1685 (2003).

[6] R. Monasson *et al.*, Nature (London) **400,** 133 (1999); G. Biroli, R. Monasson, and M. Weigt, Eur. Phys. J. B **14,** 551 (2000); M. Mézard, G. Parisi, and R. Zecchina, Science **297,** 812 (2002).

[7] M. Mézard and R. Zecchina, Phys. Rev. E **66,** 056126 (2002).

[8] M. Mézard, F. Ricci-Tersenghi, and R. Zecchina, J. Stat. Phys. **111,** 505 (2003).

[9] A. Montanari and F. Ricci-Tersenghi, Eur. Phys. J. B **33,** 339 (2003).

[10] S. Cocco, O. Dubois, J. Mandler, and R. Monasson, Phys. Rev. Lett. **90,** 047205 (2003).

[11] M. Mézard, T. Mora, and R. Zecchina, Phys. Rev. Lett. **94,** 197205 (2005).

[12] T. Berger and J. D. Gibson, IEEE Trans. Inf. Theory **44,** 2693 (1998).

[13] G. Caire, S. Shamai, and S. Verdù, in *Proceedings of IEEE Information Theory Workshop, Paris, 2003* (IEEE, New York, 2003), p. 291; Y. Matsunaga and H. Yamamoto, in *Proceedings of IEEE International Symposium on Information Theory, Lausanne, 2002* (IEEE, New York, 2002), Vol. 461.

[14] C. E. Shannon, IRE Natl. Conv. Record **4,** 142 (1959).

[15] R. E. Blahut, IEEE Trans. Inf. Theory **18,** 460 (1972).

[16] D. A. Huffman, Proc. of IRE **40,** 1098 (1952); J. Ziv, A. Lempel, IEEE Trans. Inf. Theory **23,** 337 (1977).

[17] T. Hosaka, Y. Kabashima, and H. Nishimori, Phys. Rev. E **66,** 066126 (2002).

[18] E. Martinian and J. S. Yedidia, MERL, Report No. TR2003-120 2003.

[19] T. J. Schaefer, in *Proceedings of the 10th Annual ACM Symposium on Theory of Computing, San Diego* (ACM, New York, 1978), pp 216–226.

[20] S. Ciliberti and M. Mézard (to be published).

[21] F. R. Kschischang, B. J. Frey, and H. A. Loeliger, IEEE Trans. Inf. Theory **47,** 498 (2001).

[22] A variable is free when the status of the constraint ($\varepsilon_a = 0$ or 1) is independent of the actual value of the variable.

[23] D. Battaglia, M. Kolàr, and R. Zecchina, Phys. Rev. E **70,** 036107 (2004).

[24] We forbid "fully canalizing" nodes that just depend on one variable.