

Equilibration through local information exchange in networks

K. Y. Michael Wong¹ and David Saad²

¹*Department of Physics, The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong, China*

²*The Neural Computing Research Group, Aston University, Birmingham B4 7ET, United Kingdom*

(Received 19 December 2005; published 17 July 2006)

We study the equilibrium states of energy functions involving a large set of real variables, defined on the links of sparsely connected networks, and interacting at the network nodes, using the cavity and replica methods. When applied to the representative problem of network resource allocation, an efficient distributed algorithm is devised, with simulations showing full agreement with theory. Scaling properties with the network connectivity and the resource availability are found.

DOI: [10.1103/PhysRevE.74.010104](https://doi.org/10.1103/PhysRevE.74.010104)

PACS number(s): 02.50.-r, 02.70.-c, 89.20.-a

Many theoretically challenging and practically important problems involve interacting variables connected by network structures [1]. Statistical mechanics of disordered systems makes contributions towards the understanding of such systems at two levels. Macroscopically, it describes the typical behavior of the networks, using techniques such as the replica method. Microscopically, it analyzes the relation between variables, using techniques such as the cavity method, that give rise to efficient computational algorithms. In computer science, probabilistic inference based on graphical structures has been developed and applied [2,3], mainly for providing approximate solutions to specific instances of limited size. Examples of recent success included the belief propagation algorithm for error-correcting codes [4] and the survey propagation algorithm for the satisfiability problem [5].

Most analyses so far have focused on networks of discrete variables. However, many typical problems, such as network resource allocation, involve continuous variables. Compared with discrete variables, analyses for continuous variables were much less explored. The main obstacle comes from the need to pass among the nodes entire free energy functions as messages. This is much more complex than cases of discrete values, where the messages are countable sets of conditional probability estimates of *discrete values*. Previous work in the computer science literature focused on modeling these functions for getting good approximations in feasible time scales [6]. There have been attempts to simplify the messages for continuous variables, for example, to parametrize them using eigenfunction decomposition for special cases, but the general feasibility remains an open question [7].

In this Rapid Communication we study a system of real variables that can be mapped onto a sparse graph. Based on the analysis, we demonstrate the close relationship between belief propagation algorithms and the Bethe approximation in statistical physics [8], and propose a message-passing approximation method, generally applicable to problems of continuous variables. The method is efficient since the messages consist of only the first and second derivatives of the vertex free energies derived from our analysis. The key to the successful simplification, not needed for the simpler case of discrete variables, is that the messages passed to a target node are accompanied by information-provision messages from the target node, to first determine the state at which the derivatives should be calculated.

We first formulate the problem at a general temperature, and then focus on a prototype for optimization. The traditional approach for optimization on networks is to adopt computationally demanding global optimization techniques, such as linear or quadratic programming [9]. In contrast, message-passing approaches have the potential to solve global optimization problems via local updates, thereby reducing the computational complexity. An even more important advantage, relevant to practical implementation, is its distributive nature. Since it does not require a global optimizer, it is particularly suitable for distributive control in large or evolving networks.

We consider a sparse network with N nodes, labeled $i=1, \dots, N$. Each node i is randomly connected to c other nodes. The connectivity matrix is given by $\mathcal{A}_{ij}=1, 0$ for connected and unconnected node pairs respectively. A link variable y_{ij} is defined on each connected link from j to i . We consider an energy function (cost) $E = \sum_{(ij)} \mathcal{A}_{ij} \phi(y_{ij}) + \sum_i \psi(\lambda_i, \{y_{ij} | \mathcal{A}_{ij}=1\})$, where λ_i is a quenched variable defined on node i . In the context of probabilistic inference, y_{ij} may represent the coupling between observables in nodes j and i , $\phi(y_{ij})$ may correspond to the logarithm of the prior distribution of y_{ij} , and $\psi(\lambda_i, \{y_{ij} | \mathcal{A}_{ij}=1\})$ the logarithm of the likelihood of the observables λ_i . In the context of resource allocation, $y_{ij} \equiv -y_{ji}$ may represent the current from node j to i , $\phi(y_{ij})$ the transportation cost, and $\psi(\lambda_i, \{y_{ij} | \mathcal{A}_{ij}=1\})$ the performance cost of the allocation task on node i , dependent on the node capacity λ_i .

We are interested in the case of intensive connectivity $c \sim O(1) \ll N$. Since the probability of finding a loop of finite length on the network is low, the cavity method well describes the local environment of a node. A node is connected to c branches in a tree structure, and the correlations among the branches of the tree are neglected. In each branch, nodes are arranged in generations. A node is connected to an ancestor node of the previous generation, and other $c-1$ descendent nodes of the next generation. Considering node i as the ancestor of node j , the descendents of node j form a tree structure \mathbf{T} with $c-1$ branches, labeled by $k \neq i$ for $\mathcal{A}_{jk}=1$. At a temperature $T \equiv \beta^{-1}$, the free energy $F(y_{ij} | \mathbf{T})$ can be expressed in terms of the free energies $F(y_{jk} | \mathbf{T}_k)$ of its descendents. The free energy can be considered as the sum of two parts, $F(y | \mathbf{T}) = N_{\mathbf{T}} F_{\text{av}} + F_{\text{v}}(y | \mathbf{T})$, where $N_{\mathbf{T}}$ is the number of nodes in the tree \mathbf{T} , F_{av} is the average free energy per

node, and $F_V(y|\mathbf{T})$ is referred to as the *vertex free energy*. This leads to the recursion relation

$$F_V(y_{ij}|\mathbf{T}) = -T \ln \left\{ \prod_{k \neq i} \left(\int dy_{jk} \right) \exp \left[-\beta \psi(\lambda_j, \{y_{jk}\}) - \beta \sum_{k \neq i} [F_V(y_{jk}|\mathbf{T}_k) + \phi(y_{jk})] \right] \right\}_{A_{jk}=1} - F_{av}, \quad (1)$$

$$F_{av} = -T \left\langle \ln \left\{ \prod_k \left(\int dy_{jk} \right) \exp \left[-\beta \psi(\lambda_j, \{y_{jk}\}) - \beta \sum_k [F_V(y_{jk}|\mathbf{T}_k) + \phi(y_{jk})] \right] \right\}_{A_{jk}=1} \right\rangle_{\lambda}, \quad (2)$$

where \mathbf{T}_k is the tree terminated at node k , and $\langle \cdots \rangle_{\lambda}$ represents the average over the distribution of λ . Interestingly, the recursive relation of Eq. (1) can be directly linked to probabilistic message passing (belief propagation), where the logarithms of messages passed between nodes are proportional to the vertex free energies.

For more concrete discussions, we focus on a prototype for optimization, termed resource allocation and well known in the areas of computer science and operations management [10,11]. The analysis of the problem is applicable to typical situations where a large number of nodes are required to balance loads and/or resources, such as reducing internet traffic congestion and streamlining network flows of commodities [12]. In computer science, many practical solutions are usually heuristic and focus on practical aspects (e.g., communication protocols). Here we study a more generic version of the problem represented by nodes of some computational power that should carry out tasks. Both computational powers and tasks will be chosen at random from some arbitrary distribution. The nodes are located on a randomly chosen sparse network of some connectivity. The goal is to allocate tasks on the network such that demands will be satisfied while the migration of (sub-)tasks is minimized.

We focus here on the satisfiable case where the total computing power is greater than the demand, and where the number of nodes involved is very large; the unsatisfiable case can be investigated using a similar approach [13]. Each node on the network has a capacity (computational capability minus allocated tasks) λ_i randomly drawn from a distribution $\rho(\lambda_i)$. With the aim to satisfy the capacity constraints, we have $\psi(\lambda_i, \{y_{ij} | A_{ij}=1\}) = \ln[\Theta(-\sum_j A_{ij} y_{ij} - \lambda_i) + \epsilon]$, where $\epsilon \rightarrow 0$. The problem reduces to the load balancing task of minimizing the energy function (cost) $E = \sum_{(ij)} A_{ij} \phi(y_{ij})$, subject to the capacity constraints $\sum_j A_{ij} y_{ij} + \lambda_i \geq 0$.

When $\phi(y)$ is a general even function of the current y , we may also derive Eq. (1) using the replica method. We first introduce the *chemical potentials* μ_i of nodes i , and approximate the current y_{ij} as driven by the potential differences between nodes $y_{ij} = \mu_j - \mu_i$. Since sparse networks are locally treelike, the probability of finding short loops is vanishing in large networks, and the approximation works well.

Considering the optimization problem in the space of chemical potentials, we calculate the replicated partition function $\langle Z^n \rangle_{A,\lambda}$ averaged over the connectivity matrix and capacity distribution, and take the limit $n \rightarrow 0$. Assuming replica symmetry, the saddle point equations yield a recursion relation for a two-component function R dependent on the tree structure \mathbf{T} , given by

$$R(z, \mu|\mathbf{T}) = \frac{1}{\mathcal{D}} \prod_{k=1}^{c-1} \left(\int d\mu_k R(\mu, \mu_k|\mathbf{T}_k) \right) \times \Theta \left(\sum_{k=1}^{c-1} \mu_k - c\mu + z + \lambda_{V(\mathbf{T})} \right) \times \exp \left(-\frac{\beta\epsilon}{2} \mu^2 - \beta \sum_{k=1}^{c-1} \phi(\mu - \mu_k) \right), \quad (3)$$

where \mathcal{D} is a constant, \mathbf{T}_k represents the tree terminated at the k th descendent, and $\lambda_{V(\mathbf{T})}$ the capacity of the vertex of the tree \mathbf{T} . The term $\beta\epsilon\mu^2/2$, with $\epsilon \rightarrow 0$, is introduced to break the translational symmetry of the chemical potentials, since the energy function is invariant under the addition of a constant to all chemical potentials.

Equation (3) expresses $R(z, \mu|\mathbf{T})$ in terms of $c-1$ functions $R(\mu, \mu_k|\mathbf{T}_k)$ ($k=1, \dots, c-1$), a characteristic of the tree structure. Furthermore, except for the factor $\exp(-\beta\epsilon\mu^2/2)$, R is a function of $y \equiv \mu - z$, which is interpreted as the current drawn from a node with chemical potential μ by its ancestor with chemical potential z . One can then express the function R as the product of a *vertex partition function* Z_V and a normalization factor W , that is, $R(z, \mu|\mathbf{T}) = W(z)Z_V(y|\mathbf{T})$. In the limit $\epsilon \rightarrow 0$, the dependence on μ and y decouples, enabling one to derive a recursion relation for the *vertex free energy* $F_V(y|\mathbf{T}) \equiv -T \ln Z_V(y|\mathbf{T})$ and arrive at Eq. (3).

The current distribution and the average free energy per link can be derived by integrating the current y' in a link from one vertex to another, fed by the trees \mathbf{T}_1 and \mathbf{T}_2 , respectively; the obtained expressions are $P(y) = \langle \delta(y - y') \rangle_{\star}$ and $\langle E \rangle = \langle \phi(y') \rangle_{\star}$ where $\langle \cdot \rangle_{\star} = \langle \int dy' \exp[-\beta E(y')] (\cdot) / \int dy' \exp[-\beta E(y')] \rangle_{\lambda}$, and $E(y') = F_V(y'|\mathbf{T}_1) + F_V(-y'|\mathbf{T}_2) + \phi(y')$.

Figure 1(a) shows results of the iteration of Eq. (1), in the case of optimization ($T=0$) based on discretizing $F_V(y|\mathbf{T})$ as a vector. The capacity distribution $\rho(\lambda)$ is Gaussian of variance 1 and average $\langle \lambda \rangle$. Each iteration corresponds to adding one extra generation (1000 new nodes in our simulations) to the tree structure, such that the iterative process corresponds to approximating the network by an increasingly extensive tree. We observe that after an initial rise with iteration steps, the average energies converge to steady-state values, at a rate which increases with the average capacity.

To study the convergence rate of the iterations, we fit the average energy at iteration step t using $\langle E(t) - E(\infty) \rangle \sim \exp(-\gamma t)$ in the asymptotic regime. As shown in the inset of Fig. 1(a), the relaxation rate γ increases with the average capacity. A cusp appears at the average capacity of about 0.45, below which convergence is slow due to a plateau that

develops in the average energy curve before the final stage. The slowdown is probably due to the appearance of increasingly large clusters of nodes with negative resources, which draw currents from increasingly extensive regions of nodes with excess resources to satisfy the demand.

The local nature of the recursion relation (1) points to the possibility that the network optimization can be solved by message-passing approaches. Instead of passing the functions $F_V(y|\mathbf{T})$ of the current y as messages, we simplify each message to two parameters, namely, the first and second derivatives of the vertex free energies. Let $(A_{ij}, B_{ij}) \equiv [\partial F_V(y_{ij}|\mathbf{T}_j)/\partial y_{ij}, \partial^2 F_V(y_{ij}|\mathbf{T}_j)/\partial y_{ij}^2]$ be the message passed from node j to i . Using Eq. (1), the recursion relations lead to the message (A_{ij}, B_{ij})

$$A_{ij} \leftarrow -\mu_{ij}, \quad B_{ij} \leftarrow \frac{\Theta(-\mu_{ij} + \epsilon)}{\sum_{k \neq i} \mathcal{A}_{jk} (\phi''_{jk} + B_{jk})^{-1}}, \quad (4)$$

$$\mu_{ij} = \min \left\{ \left[\sum_{k \neq i} \mathcal{A}_{jk} [y_{jk} - (\phi'_{jk} + A_{jk})(\phi''_{jk} + B_{jk})^{-1}] + \lambda_j - y_{ij} \right] \times \left[\sum_{k \neq i} \mathcal{A}_{jk} (\phi''_{jk} + B_{jk})^{-1} \right]^{-1}, 0 \right\}, \quad (5)$$

with ϕ'_{jk} and ϕ''_{jk} representing the first and second derivatives of $\phi(y)$ at $y=y_{jk}$, respectively.

The algorithm is complete with the determination of the drawn current y_{ij} at which the derivatives comprising the messages should be computed. Two methods are proposed. In the first, when messages are sent from node j to the ancestor node i , backward messages y_{jk} computed from the same optimization step are sent from node j to the descendent nodes k , informing them of the particular arguments to be used for calculating subsequent messages. In the second, node j first receives the messages (A_{ji}, B_{ji}) and current y_{ji} from the ancestor node i , and update the current y_{ij} by minimizing the total cost. Both methods work well for the quadratic cost functions.

For comparison, an independent exact optimization is available at zero temperature. The chemical potentials turn out to be the Lagrange multipliers of the capacity constraints, and the relation between the currents and the chemical potentials turns out to be exact. The Kuhn-Tucker conditions for the optimal solution yield

$$\mu_i = \min \left[\frac{1}{c} \left(\sum_j \mathcal{A}_{ij} \mu_j + \lambda_i \right), 0 \right]. \quad (6)$$

Like in the message-passing algorithm, this condition also provides a local iterative solution to the optimization problem. Simulations show that it yields excellent agreement with Eqs. (1), (4), and (5).

To study the dependence on the connectivity, we first consider the limit of large $K \equiv c-1$, in which Eq. (4) converges to the steady-state results of $A_{ij} = \max[K^{-1}(\sum_{k \neq i} \mathcal{A}_{jk} A_{jk} - \lambda_j), 0]$ and $B_{ij} \sim K^{-1}$. Then $\sum_{k \neq i} \mathcal{A}_{jk} A_{jk}$ becomes self-averaging and equal to $K m_A$, where





FIG. 1. Results for $N=1000$ and $\phi(y)=y^2/2$. (a) $\langle \phi \rangle$ obtained by iterating Eq. (1) as a function of t for $\langle \lambda \rangle = 0.1, 0.2, 0.4, 0.6, 0.8$ (top to bottom) and $c=3$. Dashed line: The asymptotic $\langle \phi \rangle$ for $\langle \lambda \rangle = 0.1$. Inset: γ as a function of $\langle \lambda \rangle$. (b) $K^2 \langle \phi \rangle$ as a function of $\langle \lambda \rangle$ for $c=3$ (\circ), 4 (\square), 5 (\diamond), large K (line). Inset: $K^2 \langle \phi \rangle$ as a function of time for random sequential update of Eqs. (4) and (5). Symbols: same as (a).

$m_A \sim K^{-1}$ is the mean of the messages A_{ij} . Thus, $y_{ij} \sim \mu_i \sim K^{-1}$. The physical picture of this scaling behavior is that the current drawn by a node is shared among the K descendent nodes. After rescaling, quantities such as $K^2 \langle \phi \rangle$, $P(Ky)/K$, and $P(K\mu)/K$ become purely dependent on the average capacity $\langle \lambda \rangle$.

For increasing finite values of K , Fig. 1(b) shows the common trend of $K^2 \langle \phi \rangle$ decreasing with $\langle \lambda \rangle$ exponentially, and gradually approaching the large K limit. The scaling property extends to the optimization dynamics [Fig. 1(b) inset]. As shown in Fig. 2(a), the current distribution $P(Ky)/K$ consists of a δ function component at $y=0$ and a continuous component, whose breadth decreases with $\langle \lambda \rangle$. Remarkably, the distributions for different connectivities collapse almost perfectly after the currents are rescaled by K^{-1} , with a very mild dependence on K and gradually approaching the large K limit. As shown in the inset of Fig. 2(a), the fraction of idle links increases with $\langle \lambda \rangle$. The fraction has a weak dependence on the connectivity, confirming the almost universal distributions rescaled for different K .

Since the current on a link scales as K^{-1} , the allocated




FIG. 2. Results for $N=1000$ and $\phi(y)=y^2/2$. (a) The current distribution $P(Ky)/K$ for $\langle\lambda\rangle=0.02, 0.5, 1$, and $c=3$ (solid lines), 4 (dotted), 5 (dot-dashed), large K (long dashed). Inset: $P(y=0)$ as a function of $\langle\lambda\rangle$ for $c=3$ (\circ), 4 (\square), 5 (\diamond), large K (line). (b) The resource distribution $P(r)$ for $\langle\lambda\rangle=0.02, 0.1, 0.5$, large K . Symbols: as in (a). Inset: $P(r>0)$ as a function of $\langle\lambda\rangle$. Symbols: as in the inset of (a).

resource of a node should have a weak dependence on the connectivity. Defining the resource at node i by $r_i \equiv \lambda_i + \sum_j A_{ij} y_{ij}$, the resource distribution $P(r)$ shown in Fig. 2(b) confirms this behavior even at low connectivities. The

fraction of nodes with unsaturated capacity constraints increases with the average capacity, and is weakly dependent on the connectivity [Fig. 2(b) inset]. Hence the saturated nodes form a percolating cluster at a low average capacity, and breaks into isolated clusters at a high average capacity. It is interesting to note that at the average capacity of 0.45, below which a plateau starts to develop in the relaxation rate of the recursion relation, Eq. (1), the fraction of saturated nodes is about 0.47, close to the percolation threshold of 0.5 for $c=3$.

In summary, using the example of the resource allocation problem on sparsely connected networks, we studied the use of message-passing methods for equilibration using both replica and cavity based analyses. A local algorithm was devised and successfully applied to this task. The study also reveals the scaling properties of this model, showing that the resource distribution on the nodes depends principally on the networkwide availability of resources, and depends only weakly on the connectivity. Links share the task of resource provision, leading to current distributions that are almost universally dependent on the resource availability after rescaling.

While the analysis focused on fixed connectivity and zero temperature, it can accommodate any connectivity profile and temperature parameter and may be used for analyzing a range of inference problems. For instance, we have considered the effects of adding anharmonic and frictional terms to the quadratic cost function. The message-passing function can be adapted to these variations, and the results will be presented elsewhere [13]. Both analysis and algorithm extend the use of current message-passing techniques to inference in problems with continuous variables, opening up a rich area for further investigations with many potential applications.

We thank David Sherrington for useful comments. This work is partially supported by the Research Grant Council of Hong Kong (Grants No. HKUST6062/02P, No. DAG04/05.SC25, and No. DAG05/06.SC36) and EU Grants EVERGROW, IP 1935 in FP6 and STIPCO in FP5.

-
- [1] H. Nishimori, *Statistical Physics of Spin Glasses and Information Processing* (Oxford University Press, Oxford, UK, 2001).
 [2] J. Pearl, *Probabilistic Reasoning in Intelligent Systems* (Morgan Kaufmann, San Mateo, CA, 1988).
 [3] D. J. C. Mackay, *Information Theory, Inference and Learning Algorithms* (Cambridge University Press, Cambridge, UK, 2003).
 [4] M. Opper and D. Saad, *Advanced Mean Field Methods* (MIT Press, Cambridge, MA, 2001).
 [5] M. Mézard, e-print cond-mat/0401237.
 [6] S. L. Lauritzen, *Graphical Models* (Oxford University Press, New York, 1996).
 [7] N. Skantzos, I. P. Castillo, and J. P. L. Hatchett, *Phys. Rev. E* **72**, 066127 (2005).
 [8] J. S. Yedidia, W. T. Freeman, and Y. Weiss, *IEEE Trans. Inf. Theory* **51**, 2282 (2005).
 [9] D. Bertsekas, *Linear Network Optimization* (MIT Press, Cambridge, MA, 1991).
 [10] L. Peterson and B. S. Davie, *Computer Networks: A Systems Approach* (Academic Press, San Diego, CA, 2000).
 [11] Y. C. Ho, L. Servi, and R. Suri, *Large Scale Syst.* **1**, 51 (1980).
 [12] S. Shenker, D. Clark, D. Estrin, and S. Herzog, *Comput. Commun. Rev.* **26**, 19 (1996).
 [13] K. Y. M. Wong and D. Saad (unpublished).