

On the Complexity of Task Allocation

A. SCHONEVELD, J. F. DE RONDE, AND P. M. A. SLOOT

*Faculty of Mathematics, Computer Science, Physics & Astronomy, University of Amsterdam
Kruislaan 403, 1098 SJ Amsterdam
the Netherlands*

Received December 9, 1996; accepted August 21, 1997

A detailed study is presented on the combinatorial optimization problem of allocating parallel tasks to a parallel computer. Depending on two application/machine-specific parameters, both a sequential and a parallel optimal allocation phase are shown to exist. A sudden "phase" transition is observed if one of these parameters is varied. Simulated annealing is used to find the optimal allocations, which is justified by the self-similar structure of the task allocation energy landscape. It is shown that the difficulty of finding optimal allocations behaves anomalously near the transition, analogous to critical slowing down of simulated equilibration at second-order phase transitions. © 1997 John Wiley & Sons, Inc.

Key Words: critical behavior, combinatorial optimization, parallel task allocation, landscapes, auto-correlation functions

INTRODUCTION

An essential problem in the field of parallel computing is the so-called task allocation problem (TAP). Given a set of parallel communicating tasks (a parallel application) and a parallel distributed memory machine, find the optimal allocation of tasks onto the parallel system. The quality of an allocation is measured by the turnaround time of the application, which depends on communication and calculation components.

These two components cannot be regarded independently, but rather are strongly related. Equal distribution of the set of parallel tasks over the available parallel processors, without taking into account the inter-task communication leads to optimal work load balancing. On the other hand, if all tasks are placed on a single processor, the amount of communication is optimal.

We use the term "frustration" for the fact that optimization of one term conflicts with optimization of the other, in analogy to physical systems that exhibit frustration (e.g., spin glasses). Increasing dominance of either term reduces the amount of frustration in the system.

Many fundamental problems from natural sciences can be formulated in terms complex systems. A complex system can be described as a population of unique elements with well-defined attributes and interactions. In most cases, such systems are characterized by *quenched* disorder and *frustrated*, nonlinear interactions, between the set of elements constituting the system[1]. It is well known that these system ingredients result in unpredictable emergent behavior [2]. In general, the bulk properties of these systems are analytically intractable. Examples of such properties are asymptotic behavior and the exact location and value of the energetically

optimal states. The latter characteristic often causes the corresponding optimization problems to be NP-hard [3]. A recent example where parallelism is discussed in a complex systems context is the work by Macready et al. [4].

In this article, we define model representations for both the parallel application and the parallel computer. Additionally, an energy function, which quantifies the cost of a task allocation, is constructed. We explore the characteristics of the TAP in terms of phase space and optima structure. Consider a parallel computer consisting of identical processors and a finite speed communication network. Increasing the CPU performance continuously from 0 flop/s to ∞ flop/s induces a transition from optimal parallel to sequential allocation. It is shown that a paramount characteristic of the TAP is the presence of a *sudden* transition from sequential to parallel optimal allocation, for specific model constraints. In analogy with other combinatorial optimization problems that exhibit frustration and phase transitions, we expect that a phenomenon, known as “critical slowing down,” can be observed in the transition region of the TAP, that is, the difficulty of finding optimal solutions peaks near the transition region (see, e.g., [5]). The specific correlation structure of the corresponding energy landscape is used as a justification to select simulated annealing as the optimization heuristic [6].

MODELS

In order to facilitate a study on abstract parallel applications, a random graph representation as a model of n communicating parallel tasks is introduced. Each task (vertex) is assigned a workload and every pair of tasks in the task graph is connected with a probability γ . A message size is assigned to each link (edge) between two communicating tasks. The size of the message only contributes to the communication time if the two connected tasks are allocated to different processors. Workloads and message sizes are kept constant. The target parallel computer is assumed to be fully connected and homogeneous; that is, all P processors have identical constant performance. Moreover, the $P(P-1)/2$ communication channels are bi-directional and have equal bandwidths. An example corresponding to these models is a parallel molecular dynamics simulation with Coulomb interactions (long-range interactions, leading to global communication patterns) on the IBM SP2 (fully connected and homogeneous topology) [7]. The time evolution of the particles is always preceded by a data exchange phase. The total execution time is determined by the time spent in the communication phase and the calculation phase.

We use the following Hamiltonian to quantify the quality or cost of a task allocation, which is inspired by a similar expression introduced by Fox et al. [8]:

$$H = (1 - \beta) \sum_{i>k}^n J_{ik} (1 - \delta(s_i, s_k)) + \beta \sum_l^P W_l^2 \quad (1)$$

$$\delta(s_i, s_j) = \begin{cases} 1 & \text{if } s_i = s_j \\ 0 & \text{otherwise} \end{cases}$$

The processor to which task i is allocated is denoted by $s_i \in \{1, \dots, P\}$ and P is the number of processors. J_{ik} is a contribution to the communication between the host processors of tasks i and k , resulting from the connection between these tasks. W_l the total calculation weight on processor l , following from the individual workloads of all allocated tasks. An optimization process that is steered by Eq. (1) implicitly minimizes both the variance in the workload distribution and the total communication “surface”. In nature, we can observe analogous processes, for example, the minimization of the surface/volume ratio in a droplet of water due to surface tension. In our case, the communication term can be compared to the surface of the droplet, while the workload variance is similar to its volume.

The β parameter can be varied in the range $[0,1]$, in order to tune the competition between the calculation and the communication terms. Variations of $\frac{\beta}{1-\beta}$ can be interpreted either as alterations in an application’s calculation-communication ratio or a computer’s processor speed-bandwidth ratio [9]. The connection probability γ in a random graph can be considered as a dual parameter for β . Also, γ can be increased in the range $[0,1]$, which is equivalent to augmenting the average communication load.

Additionally, Eq. (1) has the *locality property*, which means that local changes in a task allocation can be propagated into the Hamiltonian without recalculation. This is specifically useful if an optimization algorithm is applied that is based on incremental changes (e.g., simulated annealing [10]), and as such can exploit the direct consequence of these increments, reducing the computational cost associated with the optimization process.

TAP STRUCTURE

A random walk through some landscape can be used to characterize its structure [11]. For landscapes that are self-similar, it is known that the corresponding random walk auto-correlation function is a decaying exponential, with correlation length λ . Such landscapes are classified as AR(1) landscapes and have been identified in various fields, for instance, (bio)physics [11] and combinatorial optimization [12,13].

First, we will shortly discuss the relaxation function of random walks through the task allocation configuration space. This function indicates at what rate a random walk through the space deviates from the starting point, analogous to, e.g., relaxation of diffusion processes in physical systems. The previous function can be related to the auto-correlation function, which quantifies the ruggedness [11] of the TAP energy landscape. Using these functions, it can be shown that the landscape is AR(1) with a correlation length that is linearly proportional to the number of tasks n .

Configuration Space

The configuration space C of the TAP consists of all possible task allocations of the n tasks to the P processor machine. A

configuration can be encoded as a sequence of length n , which is composed of letters taken from the alphabet $\{1, 2, \dots, P\}$. The index of a sequence letter corresponds to a task identifier. The *Hamming* distance $d_H(A, B)$ (number of differing letter positions) between two sequences A and B is used as the distance metric on C . The corresponding *Hamming graph* Γ can be constructed by connecting every sequence pair (A, B) with $d_H(A, B) = 1$.

The relaxation functions $q_k(s)$ ($k=1, 2$) of a random walk through an arbitrary configuration space are given by (see [12]):

$$q_k(s) = 1 - \frac{\Delta_k(s)}{\Delta_k(\infty)} \quad (2)$$

with $\Delta_1(s)$ the average distance and $\Delta_2(s)$ the average squared distance of a random walk of length s .

In previous work, we have derived the following relaxation functions for Γ [3]:

$$q_1(s) = \left(1 - \frac{1}{n}\right)^s = e^{-s/\tau_1} \quad (3)$$

$$q_2(s) = \frac{(n-1)(1-P)}{1-n+nP} e^{-s/\tau_2} - \left(\frac{(n-1)(1-P)}{1-n+nP} - 1\right) e^{-s/\tau_1} \quad (4)$$

with $\tau_1 \approx n$ and $\tau_2 \approx \frac{n}{2}$.

Energy Landscape

The following auto-correlation function associated with a random walk through an energy landscape $H: \Gamma \rightarrow IR$ can be used to characterize its structure [11]:

$$r(d) = \frac{\langle (H(x) - \langle H \rangle) (H(y) - \langle H \rangle) \rangle_{d(x,y)=d}}{\sigma^2} \quad (5)$$

where d is the number of random walk steps between configurations x and y , and σ^2 is the variance of H over all possible allocations.

Apart from totally uncorrelated landscapes, $r(d) = \delta(d, 0)$, the simplest class consists of the nearly fractal or self-similar landscapes. For such landscapes, it is known that the corresponding random walk auto-correlation ($r(d)$) function is a decaying exponential, with correlation length λ :

$$r(d) = r(1)^d = e^{-d/\lambda}, d \ll n \quad (6)$$

Such landscapes are classified as AR(1) (or elementary) landscapes and have been identified in various fields, for example, in (bio)physics and combinatorial optimization [11-13].

According to Eq. (6) the auto-correlation function of an AR(1) landscape can be determined from the 1-step auto-correlation. Let t and t' be two configurations with $d(t, t') = 1$, with the Hamiltonian having values H and H' :

$$r(1) = 1 - \frac{\langle (H - H')^2 \rangle}{2\sigma^2} := 1 - \xi \quad (7)$$

If ξ is sufficiently small:

$$\lambda = -\frac{1}{\ln(r(1))} = -\frac{1}{\ln(1-\xi)} \approx \frac{1}{\xi} \quad (8)$$

or equivalently,

$$\lambda = \frac{2\sigma^2}{\langle (H - H')^2 \rangle} \quad (9)$$

Three important parameters that can be analytically determined are the average value of the Hamiltonian $\langle H \rangle$, the variance σ^2 , and $\langle (H - H')^2 \rangle$, for the target application and machine model presented earlier. For technical details, refer to [3]. We obtain:

$$\langle H \rangle = \beta n \left(\frac{n}{P} + 1 - \frac{1}{P} \right) + (1 - \beta) \gamma \frac{(P-1)n(n-1)}{P} \quad (10)$$

$$\sigma^2 = \frac{2(n-1)n(P-1)(\beta^2 + \gamma - 4\beta\gamma + 3\beta^2\gamma)}{P^2} \quad (11)$$

$$\langle (H - H')^2 \rangle = \frac{8(n-1)(P-1)(\beta^2 + \gamma - 4\beta\gamma + 3\beta^2\gamma)}{P^2} \quad (12)$$

Using Eqs (9), (11), and (12), it follows that

$$\lambda = \frac{n}{2} \quad (13)$$

λ is linearly proportional to the number of tasks n , which corresponds to the relaxation time τ_2 . Summarizing, it has been established that the TAP energy landscape is AR(1), with correlation length $n/2$. Since the TAP landscape has a self-similar structure, we use simulated annealing to find (sub) optima, which is known to be an efficient search method for such landscapes [6]. The SA process effectively increases the resolution at which the space is searched by lowering the temperature. At the same time, the size of the region of space that is explored decreases. Due to the fractal structure of the TAP search space, SA is able to effectively use its ability to “zoom” into regions of increasingly deeper local minima (see [6]).

TAP PHASE TRANSITION

TAP Extremes

Although the task allocation problem is NP-hard [14], the two extremes, $\beta = 0$ and $\beta = 1$, are easy to solve. For $\beta = 0$ (infinitely fast CPUs), the only relevant term in the Hamiltonian is an attracting communication term, which will cause all *connected* tasks to be allocated to one processor. For this extreme (with a corresponding lowest energy state of value zero), the number of optima is equal to P . In the case of $\beta = 0$, the P optima are at maximum distance in terms of the defined distance metric. The P -ary inversion operation (analogous to spin-flipping in spin glass models) and arbitrary permutations, applied to a given optimal configuration, leave the value of the Hamiltonian invariant. Note that, in this case, the

TAP landscape is highly symmetrical. The entire landscape consists of P identical sub-landscapes. Each sub-landscape has only one optimum, which is automatically the global optimum.

For $\beta = 1$ (infinitely fast network), only a repulsive workload term is present, which will force the variance in the workload distribution to be minimized. This results in an equal partitioning of the total workload over all available processors. It can easily be shown that the total number of optima in this case equals:

$$\prod_{k=1}^P \binom{n}{k(n/P)} = \frac{n!}{(n/P)!^P} \quad (14)$$

where it has been assumed that n/P is integer. The corresponding optimal value of the Hamiltonian is equal to n^2/P . In case of $\beta = 1$, the optima are relatively close to one another. Again, two types of operations can be distinguished that leave the value of the Hamiltonian invariant, which are rotation of the sequence and permutation of two arbitrary tasks.

Locating the Transition

A transition from sequential to parallel allocation can be observed when β is increased from 0 to 1 (or equivalently, if γ is decreased from 1 to 0). In order to quantify this (phase) transition, we define an *order parameter*, which expresses the *degree of parallelism* present in an optimal allocation. Since all tasks and connection weights are unity, the order parameter P , quantifying the *parallelism* in a given optimal allocation, can be defined as follows:

$$P = 1 - \frac{(\langle W^2 \rangle - \langle W \rangle^2) P^2}{n^2 (P-1)} \quad (15)$$

where W is the time spent in calculation and $n^2(P-1)/P^2$ is the maximal possible variance in W . Eq. (15) takes the value 1 in the case of optimal parallelism ($\beta = 1$ or $\gamma = 0$) and the value 0 ($\beta = 0$ or $\gamma = 1$) in the case of a sequential allocation.

Using Eq. (10), we can calculate whether the average value Eq. (1) either increases or decreases by using more processors. Using a mean field argument, the transition from sequential to parallel allocation will approximately occur for those values of β and γ for which Eq. (1) will change from a monotonically decreasing function to a monotonically increasing function of P . In other words, setting

$$\frac{\partial \langle H \rangle}{\partial P} = 0 \quad (16)$$

with the additional constraint that either γ or β is fixed, we obtain the following transition values:

$$\beta_c = \frac{\gamma}{1+\gamma} \quad (17)$$

$$\gamma_c = \frac{\beta}{1-\beta} \quad (18)$$

The values of β_c and γ_c are interpreted as the critical values of β and γ in analogy with the critical temperature T_c in thermal phase transitions or the percolation threshold p_c .

Critical Slowing Down

Many search methods show anomalous behavior for certain critical parameters of combinatorial search problems [5, 15-17]. For example, in the case of graph coloring, it has been observed that the “difficulty” of determining if a graph can be colored increases abruptly when the average connectivity in the graph is gradually increased to some critical value [5]. In Ising model simulations, the difficulty of equilibrating increases when the critical temperature is approached (critical slowing down).

In analogy, we expect that in the TAP comparable phenomena can be found in a critical region of the β and γ domain. For both β extremes, the optima are known in advance. The difficulty to find these optima is therefore trivial. If the calculation and the communication term in the Hamiltonian (Eq. (1)) are of comparable magnitude, the system is said to be in a critical area. Moving away from this critical region, one term becomes small noise for the other.

We will use the following empirical method to estimate the computational cost of finding optima. The number of local optima are measured, in which independent steepest descent (SD) runs get stuck. A specific search space is considered to be “simple” if it contains a relatively small number of local optima; otherwise it is classified as “difficult.” The distinction between local optima is based on the value of the Hamiltonian of the corresponding task allocations. That is, two local optima i and j are called distinct if:

$$H(i) \neq H(j) \quad (19)$$


Due to the fact that the TAP energy landscape is AR(1), we do not expect large plateaus in which SD can get stuck in its search for a true local minimum (a minimum with the lowest value of the Hamiltonian in its local neighborhood). Therefore, it is to be expected that plateau states will not have a major contribution to the cost of the search.

Of course, one may wonder how this cost heuristic relates to a standard search measure, such as the number of convergence steps taken in SA. In section 5, we will present empirical evidence that both measures are strongly correlated.

Finite Size Scaling


In many physical systems, the sharpness and the location of transition points depend on the system size. This dependence can be analyzed by a method from statistical physics called “finite size scaling.” The existence of scaling parameters for transitions at different system sizes is a direct evidence for critical behavior at the transition. The system is

FIGURE 1



Analytical (dashed lines) and experimental values for the auto-correlation function $r(s)$ with $n = 100$, $P = 8$ and $\gamma = 0.0$ (diamonds) and $n = 64$, $P = 4$ and $\gamma = 0.5$ (pluses). The experimental auto-correlation functions are generated from a random walk of 640000 steps.

FIGURE 2



Two-phase transitions with fixed γ and increasing β , with $\gamma = 0.2$, $n = 256$ and $P = 32$, and with $\gamma = 0.5$, $n = 128$ and $P = 8$ respectively. The vertical lines indicate the location of the transition as predicted by Eq. (17). The β domain is scanned with steps of $\Delta\beta = 0.01$. For each value of β , P is estimated by averaging over 25 simulated annealing runs.

indistinguishable at all sizes, except for a change of scale. The scale change can be found by analyzing the shift of the TAP transition (with fixed γ):

$$\beta_c(n) - \beta_c(\infty) \propto n^{-1/\nu} \quad (20)$$

where $\beta_c(n)$ is the location of the transition for a TAP instance of size n and ν is called a “critical exponent.” The parameter ν can be used to rescale β as follows:

$$\beta^* = n^{1/\nu}(\beta - \beta_c(\infty)) / \beta_c(\infty) \quad (21)$$

Critical behavior has also been shown to occur at transitions of other combinatorial optimization problems[4,18].

EXPERIMENTAL RESULTS

In this section experimental results regarding correlation length, phase transition, and search cost for the TAP are presented.

In Figure 1, measured and predicted correlation functions are displayed, with parameters $n = 100$, $P = 8$, and $\gamma = 0$. In the second experiment, a TAP instance with a nonzero connection probability ($\gamma = 0.5$), $n = 64$, and $P = 4$ is used. The theoretical correlation functions with correlation lengths 50 and 32, respectively ($n/2$), are plotted as dashed lines. Clearly, the predicted correlation functions match the experimental data.

Next, several experiments are conducted to demonstrate the existence of a phase transition. Furthermore, the location of the transition, as predicted by Eqs. (17) and (18), is checked.


In Figure 2, β is varied in the range $[0,1]$ and γ is fixed at two different values (0.2 and 0.5). In Figure 3, the dual experiment is performed, where γ is varied in the range $[0,1]$ and β is fixed at the value 0.25. The results presented are comparable with those found for arbitrary parameter values (data not

shown). The mean field transition points (Eqs. (17) and (18)) are plotted as vertical lines. As shown in Figures 2 and 3, the approximate location of the phase transition that is induced by variation of β or γ can be predicted by a mean field argument (Eq. (16)).

Figure 4 displays the search cost and the order parameter P , for task graphs with $n = 32$, $P = 4$, γ fixed to 0.5, and β varied in the range $[0,1]$. In Figure 5, the same experiment is carried out, now with γ varied in the range $[0,1]$, $n = 64$, $P = 8$, and β fixed to 0.2. The divergence of the search cost near the transition point can be observed in both plots. The method described in the previous section is used to quantify the search cost. In Figure 4, the standard deviation over the order parameter (i.e., the order parameter susceptibility) is depicted. It can be observed that the order parameter susceptibility increases at the phase transition, which can be compared with increasing magnetic susceptibility of Ising spin systems near the critical temperature.


To illustrate the correlation between the cost heuristic, introduced in the previous section, and the average number of SA convergence steps, an experiment has been carried out with the following parameter setting: $n = 64$, $P = 4$, and $\gamma = 0.5$, while β is varied between 0 and 1. We applied the following convergence criterion for SA: If the optimal allocation does not improve with more than 5 percent during 50 consecutive temperature lowerings, SA has reached convergence. Figure 6 shows that both measures peak near the predicted transition point β_c . The results indicate that it takes more temperature steps to find an optimal allocation in the critical β region. In other words, the “freezing temperature” of SA appears to behave anomalously near the TAP phase transition.

FIGURE 3




A phase transition with $\beta = 0.25$, $n = 64$, $P = 8$. The vertical solid line indicates the location of the transition as predicted by Eq. (18). The γ domain is scanned with steps of $\Delta\gamma = 0.025$. For each value of γ , P is estimated by averaging over 10 simulated annealing runs.

FIGURE 4




A phase transition and the search cost (with standard deviations) with $\gamma = 0.5$, $n = 32$, and $P = 4$ with β varied with steps of $\Delta\beta = 0.025$ (The cost is determined for $\beta \in \{0.1, \dots, 0.5\}$.) The vertical line indicates the location of the transition as predicted by Eq. (17). The values for P are estimated by averaging over 10 simulated annealing runs. Each point in the search cost is estimated over 10 random graph instances, where for each instance 10n steepest descent runs are conducted. The cost value is scaled to fit in the range $[0,1]$.

FIGURE 5



A phase transition (with standard deviations) and the search cost with $\beta = 0.2$, $n = 64$, and $P = 8$ with γ varied with steps of $\Delta\gamma = 0.025$ (The cost is determined cost for $\gamma \in \{0.1, \dots, 0.5\}$.) The vertical line indicates the location of the transition as predicted by Eq. (18). The values for \mathcal{P} are estimated by averaging over 10 simulated annealing runs. Each point in the search cost is estimated over 10 random graph instances, where for each instance $10n$ steepest descent runs are conducted. The cost value is scaled to fit in the range $[0, 1]$.

FIGURE 6



The average number of SA steps to reach convergence and the search cost, with $n = 64$, $P = 4$, $\gamma = 0.5$, and $\beta \in [0, 1]$. The number of SA steps is averaged over 25 runs. The vertical line indicates the location of the transition as predicted by Eq. (17). Each point in the search cost is estimated over 10 random graph instances, where for each instance $10n$ steepest descent runs are conducted. Both measures are scaled to fit in the range $[0, 1]$.

The results of a finite size scaling experiment are summarized in Figure 7 for task graphs with $P=8$ and $\gamma=0.2$. For increasing task graphs sizes ($n \in \{80, 160, 240, 320, 400, 480, 560, 640\}$), the order parameter \mathcal{P} has been estimated for ten different values of β near the predicted phase transition. Each point in the figure is averaged over 25 SA runs. It can be observed that all functions intersect in one point. Assuming that all task graph order parameter functions intersect at a common point, this must be the exact value of β_c . The critical value of β found experimentally, closely corresponds to the value as predicted by Eq. (17). For experiments with different parameters, also common intersection points, close to the predicted critical value, have been found (data not shown). In order to test the finite size scaling hypothesis, we have plotted the data from Figure 7 against the rescaled β parameter, β^* . We define $\beta_c(n)$ to be the value of β at which $\mathcal{P} = 0.5$ and $\beta_c(\infty)$ is given by Eq. (17). In Figure 8, we observe that all rescaled plots fall on a n -independent curve.


The parallel and sequential phases and the separation as predicted by Eqs. (17) and (18) are depicted in a phase diagram (see Figure 9).

CONCLUSIONS

The results presented in this article clearly show that the task allocation problem exhibits a variety of interesting properties. For specific parameter sets, the task allocation problem only exists in a small parameter range. Outside this range the problem is trivial. The problem becomes complex in the region where the calculation and communication terms are of comparable magnitude. The location of this complex region is marked by the presence of a transition from sequential to parallel allocation.


Different allocation regimes are summarized in Figure 10. The *sequential al-*

FIGURE 7



For task graphs sizes $n \in \{80, 160, 240, 320, 400, 480, 560, 640\}$, $P=8$, and $\gamma=0.2$, \mathcal{P} is calculated around the mean field prediction of β_c for β varied in the range $[0.1, 0.19]$ and a step size of $\Delta\beta = 0.01$. The vertical line indicates the location of the transition as predicted by Eq. (17).

FIGURE 8

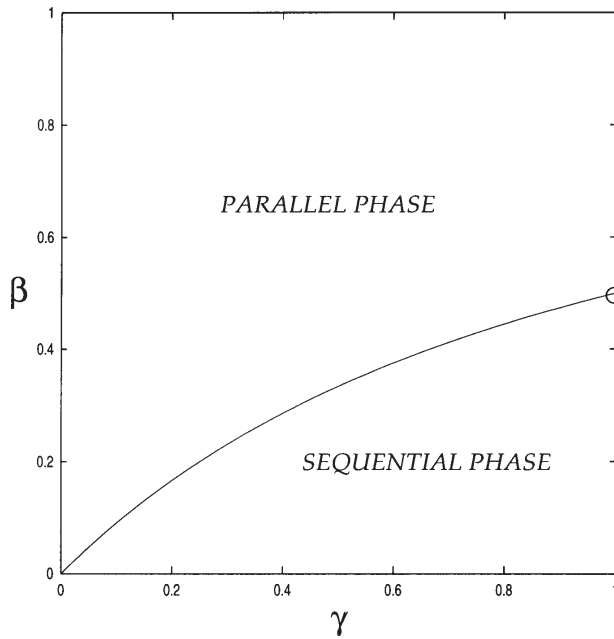


Finite size scaling. The curves of Fig. 7 are plotted against the rescaled $\beta^* = n^{1/\nu}(\beta - \beta_c(\infty)) / \beta_c(\infty)$, where $\nu \approx 1.83 \pm 0.04$.

location region only contains optima where all tasks are allocated to one processor. The *semi-parallel allocation* region corresponds to the situation, where not-all-available processors are necessarily used due to the high competition between the calculation and communication terms. Also, the locality in the task graph has its consequences for the allocation sequence. Tasks that are connected to one another “desire” to be grouped onto the same processor. The last region, *parallel allocation*, corresponds to the mode where the inter task connectivity has become insignificant. This may either be due to a high-speed communication network or a weakly connected task graph. For increasing task graph sizes, the transition region narrows. This implies the existence of exactly two regions of task allocation order in the limit $n \rightarrow \infty$. Hence, for large task graphs that display long-range interactions, the TAP is trivial for allocation on fully connected parallel computers.


We intend to investigate the effects of introducing short-range locality in both the task graph as well as the processor topology. We intend to study whether other system-specific properties, such as order parameter susceptibility, scale into universal curves. Furthermore, we will adapt our formalization of the TAP such that it can be applied to load balancing of dynamic and heterogeneous parallel applications on dynamic and heterogeneous parallel computers.

FIGURE 9



The phase diagram of the TAP. The circle corresponds to the situation $(\beta, \gamma) = (1/2, 1)$. In this case $\gamma = \frac{0}{0}$ (indeterminate). The separating line between the two phases corresponds to $\beta(\gamma) = \frac{\gamma}{1+\gamma}$.

FIGURE 10



The different allocation regimes in the task allocation problem for varying β .

REFERENCES

1. M. Mezard, G. Parisi, and M.A. Virasoro: Spin glass theory and beyond. World Scientific, Singapore, 1987.
2. D. Sherrington: Complexity due to disorder and frustration. In E. Jen (Ed.) 1989 Lectures in complex systems, volume II of Santa Fe Institute studies in the sciences of complexity, Addison-Wesley, Redwood City, 1990.
3. P. M. A. Sloot, A. Schoneveld, J. F. de Ronde, and J. A. Kaandorp: Large-scale simulations of complex systems, part I: conceptual framework. Working Paper 97-07-070, Santa Fe Institute, 1997.
4. W. G. Macready, A. G. Siapas, and S. A. Kauffman: Criticality and parallelism in combinatorial optimization. Science, 271: 56-59, 1996.
5. C. P. Williams and T. Hogg: Using deep structure to locate hard problems. In Proceedings of the 10th National Conference on Artificial Intelligence, pp. 472-477, San Jose, CA, 1992.
6. G. B. Sorkin. Simulated annealing on fractals: Theoretical analysis and relevance for combinatorial optimisation. In W.J. Dally (ed.), Advanced research in VLSI, Proceedings of the 6th MIT Conference, pp. 331-351, 1990.
7. T. Agerwala, J. L. Martin, J. H. Mirza, D. C. Sadler, D. M. Dias, and M. Snir: Sp2 system architecture. IBM Systems Journal, 32(2): 152-184, 1995.
8. G. C. Fox, R. D. Williams, and P. Messina. Parallel computing works! Morgan Kaufmann Publishers, San Francisco, 1994.
9. G. Fox, M. Johnson, G. Lyzenga, S. Otto, J. Salmon, and D. Walker: Solving problems on concurrent processors, vol. 1, Prentice-Hall, Englewood Cliffs, NJ, 1988.
10. S. Kirkpatrick, C. D. Gelatt Jr., and M. P. Vecchi: Optimization by simulated annealing. Science, 220 (4598): 671-680, May 1983.
11. E. D. Weinberger: Correlated and uncorrelated fitness landscapes and how to tell the difference. Biological Cybernetics, 63: 325, 1990.
12. P. F. Stadler and R. Happel: Correlation structure of the landscape of the graph-bipartitioning problem. Journal of Physics A, 25: 3103-3110, 1992.
13. P. F. Stadler and W. Schnabl. The landscape of the traveling salesman problem. Physics Letters A, 161:337-344, 1992.
14. D. Bernstein, M. Rodeh, and I. Gertner: On the complexity of scheduling problems for parallel/pipelined machines. IEEE Transaction on Computers, C-38(9): 1308-1313, 1989.
15. C.P. Williams and T. Hogg: Phase transitions and coarse-grained search. In Workshop on Physics and Computation, Dallas, Texas, 1994.
16. T. Hogg: Statistical mechanics of combinatorial search. In Workshop on Physics and Computation, Dallas, Texas, 1994.
17. P. Cheeseman, B. Kanefsky, and W. M. Taylor: Computational complexity and phase transitions. In Workshop on Physics and Computation. IEEE Computer Society, 1992.
18. S. Kirkpatrick and B. Selman: Critical behavior in the satisfiability of random boolean expressions. Science, 264: 1297-1301, 1994.