The TSP Phase Transition

Ian P. Gent
Department of Computer Science
University of Strathclyde
Glasgow G1 1XH
United Kingdom
ipg@cs.strath.ac.uk

Toby Walsh
Mechanized Reasoning Group
IRST, Trento &
DIST, University of Genoa,
Italy
toby@itc.it

Keywords: NP-complete problems; complexity; traveling salesman problem; search phase transitions; finite size scaling; easy and hard instances

Abstract

The traveling salesman problem is one of the most famous combinatorial problems. We identify a natural parameter for the two-dimensional Euclidean traveling salesman problem. We show that for random problems there is a rapid transition between soluble and insoluble instances of the decision problem at a critical value of this parameter. Hard instances of the traveling salesman problem are associated with this transition. Similar results are seen both with randomly generated problems and benchmark problems using geographical data. Surprisingly, finite-size scaling methods developed in statistical mechanics describe the behaviour around the critical value in random problems. Such phase transition phenomena appear to be ubiquitous. Indeed, we have yet to find a NP-complete problem which lacks a similar phase transition.

1 Introduction

There are many useful connections between statistical mechanics and a wide variety of combinatorial problems. For example, an analogy between annealing in solids and combinatorial optimization suggested simulated annealing, a very general and powerful optimization procedure [14]. A phase transition occurs in the quality of the solution returned by this procedure as the "temperature" control parameter is lowered. This phase transition in the algorithm behaviour has been described with the aid of ideas from statistical mechanics [3, 23]. More recently, statistical mechanics has been used not to model algorithm performance as in [3, 23] but the properties of ensembles of problems (for example, the probability that a solution exists). In particular, phase transition phenomena in Boolean satisfiability problems have been described by finite-size scaling methods [15]. We show here that phase transitions and finite-size scaling occur more widely than predicted in [15].

Phase transition phenomena occur in the solubility of many combinatorial problems. If such problems are very loosely constrained, there are likely to be many solutions. Almost all problems will therefore be soluble. If problems are very tightly constrained, almost all problems will be insoluble. With randomly generated problems, there is often a sharp transition between the two regions as a "control parameter" is varied. Cheeseman *et al.* [4] showed that computational difficulty is often associated with this phase transition. In the loosely constrained region, the very large number of solutions means that it is typically easy to find one. In the tightly constrained region, it is usually comparatively easy to show that no solution exists. Problems from the phase transition in between are typically hard since they cannot easily be proved soluble or insoluble. The correlation between a phase transition in solubility and a peak in search cost has been observed in Boolean satisfiability [15, 4, 18], graph colouring [4], constraint satisfaction [19, 22], independent set [9], and Hamiltonian circuit problems [4]. Many algorithms are now benchmarked using problems from such phase transitions.

2 Traveling salesman problem

The traveling salesman problem is one of the simplest and most famous combinatorial problems in computer science. Given a list of n cities, and a matrix giving the distance between each pair of cities, the traveling salesman problem is to determine if you can visit all n cities and return to the starting city in a distance l or less. We focus on the decision version of the problem (does a tour of length l or less exist?) and not the closely related optimization problem (what is the minimal length tour?). The decision problem can be used to solve the optimization problem, and vice versa. However, as we argue later, the decision problem can be easier than the optimization problem. Indeed, many optimization procedures solve a series of increasingly constrained decision problems. The result of the decision problem may also be of primary importance. For example, if the post office discovers there is no tour of length less than or equal to the daily travel allowance, two people will be needed to deliver the mail.

As we focus on the decision problem, as opposed to the optimization problem, the tour length l is a parameter we can fix and is not the result returned. A soluble phase will thus correspond to a region where the decision question is always answered "Yes" (a tour of this length or less exists). And an insoluble phase will correspond to a region where the decision question is always answered "No" (no tour of this length or less exists). In this paper, we focus on the symmetric, 2-dimensional Euclidean traveling salesman problem. However, we have seen similar phase transition behaviour in the asymmetric traveling salesman problem [8].

The traveling salesman problem is of considerable theoretical importance as it is representative of the large class of NP-complete (non-deterministic polynomial time complete) problems. No polynomial time algorithm is known to exist for any NP-complete problem, so this class can be considered to mark the start of computational intractability. All the combinatorial problems mentioned earlier are examples of NP-complete problems. Although any NP-complete problem can be transformed into any other NP-complete problem [6], this mapping usually does not map the problem space uniformly. The existence of a phase transition in, say, Boolean satisfiability does not therefore guarantee the existence of a similar phase transition in the traveling salesman problem.

Random instances of the traveling salesman problem can be easily generated by placing n cities on a square of area A uniformly at random. For such problems, the expected optimal tour length l_{opt} approaches the limit,

$$\lim_{n \rightsquigarrow \infty} \frac{l_{opt}}{\sqrt{n}} = k.\sqrt{A} \qquad (1)$$

where k is a constant [2]. The best current estimate for k is 0.7124 ± 0.0002 [13]. This asymptotic result suggests that a natural control parameter is the dimensionless ratio $l/\sqrt{n.A}$. At large values of this parameter, random problems are likely to be soluble as the tour length required is large compared to the number of cities to visit and the average inter-city distance. At small values of this parameter, random problems are likely to be insoluble as the tour length required is small compared to the number of cities to visit and the average inter-city distance. At intermediate values around 0.7, random problems can be either soluble or insoluble.

To solve the traveling salesman problem, we use a branch and bound algorithm with the well known Hungarian heuristic for branching [17]. Branch and bound is a backtracking algorithm which searches through the space of partial tours. Search is halted on a given branch when the length of the current partial tour plus a lower bound on the length of the rest of the tour exceeds the shortest tour found to that point. Branch and bound is one of the best complete algorithms for the traveling salesman problem. In Table 1, we give the mean and standard deviation in the optimal tour length and the parameter $l_{opt}/\sqrt{n.A}$ for random traveling salesman problems with up to 30 cities. As the number of cities increases, the mean and standard deviation in the parameter $l_{opt}/\sqrt{n.A}$ decrease. As is typical in other problems, the optimal tour length is significantly displaced from the asymptotic value at small n. Equation (1) therefore provides a poor approximation for the mean optimal tour length for small n. In addition, it provides no indication of how tour lengths are distributed around the mean.

n	mean of l_{opt}	s.d. of l_{opt}	mean of $\frac{l_{opt}}{\sqrt{n.A}}$	s.d. in $\frac{l_{opt}}{\sqrt{n.A}}$
6	2302	390	0.940	0.159
12	3102	331	0.895	0.095
18	3678	307	0.867	0.072
24	4158	293	0.849	0.060
30	4590	283	0.838	0.052

Table 1. Mean and standard deviation in the optimal tour length and the parameter $l_{opt}/\sqrt{n.A}$ for 1000 instances of n-city traveling salesman problems. Cities are placed at random on a square of side 1000.

3 Phase transition

To explore finite problem sizes in more detail, we plot in Figure 1A the probability of a tour of length l existing for different n. As n increases, the transition between soluble and insoluble problems occurs at larger l. In Figure 1B, we plot the same data but with the parameter, $l/\sqrt{n.A}$ along the x-axis. As expected from (1), there is a rapid transition in solubility around a fixed value of $l/\sqrt{n.A}$. As n increases, the transition occurs over a smaller range of the control parameter. We next applied finite-size scaling methods [1] to determine more precisely how the distribution of tour lengths scales with problem size. Around some critical point, averages for sets of problems of different sizes ought to be indistinguishable except for a change of scale. This suggests that the probability of a tour existing is a simple function of the single parameter,

$$(\frac{l}{\sqrt{n.A}} - \alpha) \,.\, n^{1/\nu}$$

where α is the critical point, and $n^{1/\nu}$ provides the change of scale. The size dependency around the critical point is perhaps a little surprising. In finite-size scaling, the size dependency is explained by means of the correlation length, the distance over which behaviour is correlated. This diverges at the critical point in an infinite system. If two points are separated by more than the correlation length, they behave independently. Despite the fact that lengths appear in the traveling salesman problem, the size dependency of the crossover obeys a simple power law not with a length but with n, the number of cities. Similar power law behaviour was seen with finite-size scaling in Boolean satisfiability [15] where there are no lengths and the number of variables replaces the usual scaling with respect to a length.

If finite-size scaling holds, then there will be a critical value, α, of the control parameter, where the probability of a tour existing will be invariant as n changes. To find this critical value, we measured the spread of observed probabilities. That is, for given $l/\sqrt{n.A}$ we measured the probability of a tour existing for each n, interpolating between observed values of l where necessary. We then noted the spread between the largest and smallest such probabilities over all values of n. This spread was minimised at $l/\sqrt{n.A} = 0.78$, giving $\alpha = 0.78 \pm 0.01$, where the errors indicate the range over which the probabilities differ by less than 5% . To compute ν, we then set $\alpha = 0.78$ and compared median tour lengths. If l_1 and l_2 are the median tour lengths for n_1 and n_2 cities respectively ($n_1 \neq n_2$) then,

$$(\frac{l_1}{\sqrt{n_1.A}} - \alpha)\,.\,n_1^{1/\nu} \quad = \quad (\frac{l_2}{\sqrt{n_2.A}} - \alpha)\,.\,n_2^{1/\nu}$$

Rearranging this equation gives,

$$\nu \quad = \quad \frac{\log(n_1/n_2)}{\log((\frac{l_2}{\sqrt{n_2.A}} - \alpha)/(\frac{l_1}{\sqrt{n_1.A}} - \alpha))}$$

Every distinct pair of values n_1 and n_2 thus provide an estimate for ν. We computed all possible estimates from our data and took the median. This gives $\nu = 1.5 \pm 0.2$ where errors represent the upper and lower quartiles of the estimates made from pairs n_1, n_2.

We therefore define a rescaled parameter $\gamma = (l/\sqrt{n.A} - 0.78)\,.\,n^{1/1.5}$. In Figure 1C, we plot the probability of a tour existing against γ. Viewing our data for the distributions of optimal tour lengths l_{opt} in terms of the rescaled parameter γ, we observe a mean of $\gamma = 0.58$, median of 0.60 and a standard deviation of 0.50. We tested the hypothesis that the distribution of γ is normally distributed with median 0.6 and standard deviation 0.5, using a goodness of fit test with 50 intervals. Using data for n from 8 to 30 in steps of 1, we could not reject this hypothesis at the 5% significance level, but it could be rejected for $n = 6$ and 7. We have shown that using finite-size scaling we can model the entire distribution of tour lengths using a single normal distribution, from $n = 8$ to 30. While it does not follow that tour lengths are in fact normally distributed, providing a significantly better model will require much more experimentation with larger sample sizes, number of cities, or both.

Whilst finite-size scaling fits our data extremely well, caution is required if we wish to extrapolate to larger n. The asymptotic value of the control parameter, at the phase transition reported in [13] is 0.7124 ± 0.0002. This is significantly smaller than the value of α used here. If α is set to 0.7124, a critical value like, for example, the mean or median tour length can be modelled using finite-size scaling. However, this rescaling fails to describe accurately the complete tour length distribution. To obtain a good fit for the whole distribution, we must scale around not a fixed value but a critical and varying value (for example, the value of the control parameter at the median optimal tour length). A similar issue appears to occur in random Boolean satisfiability problems. In [15] the critical value used for finite-size scaling is 4.17. This gives a crossover point where 50% of problems with N variables and L clauses are satisfiable at $L \approx 4.17N + 3.1N^{1/3}$. By contrast, the crossover point is empirically determined to be $L = 4.258N + 58.68N^{-2/3}$ in [5]. The differences between these two models remain to be resolved.

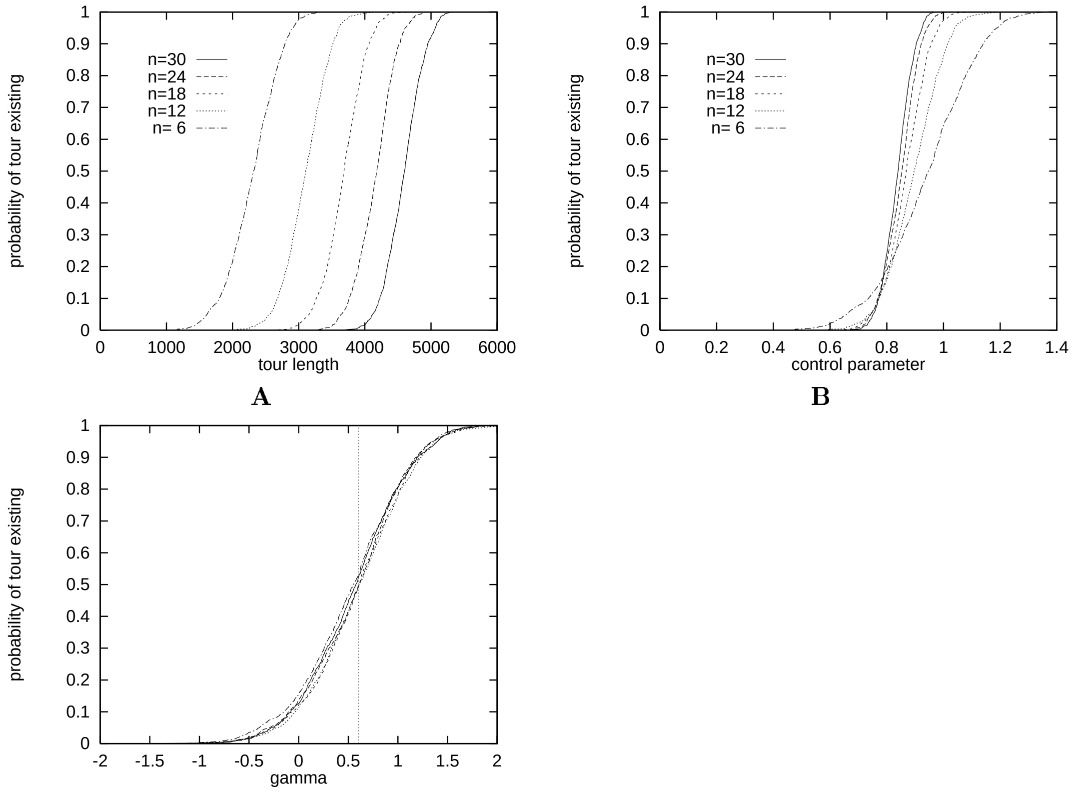

Fig. 1. **(A)** The probability of a tour existing plotted against tour length, l for 6, 12, 18, 24, and 30 city random problems. 1000 instances were used at each problem size giving roughly 3% accuracy. **(B)** The same data plotted against the control parameter $l/\sqrt{n.A}$. The curves become sharper as n increases. **(C)** The same data plotted against the rescaled parameter γ. Only the curve for $n = 6$ can be distinguished statistically from a normal distribution with parameters independent of n. The vertical line at $\gamma = 0.6$ indicates where 50% of instances are soluble.

4 Search cost

As in other combinatorial problems, a peak in search cost is associated with the phase transition in solubility. With a very large bound on the required tour length, many tours satisfy the bound and it is typically easy to find one. With a very small bound, almost all tours are too long and many are quickly ruled out, so again problems are typically easy. Problems at the phase boundary are often hard, as it is difficult to determine if a tour of the required length exists without exhaustive search. Surprisingly, search cost appears to be directly correlated with the rescaled parameter γ. In Figure 2, we plot the median number of nodes searched by the branch and bound algorithm as we vary γ. For 6 and 12 city problems, median search cost is never large, but for 18 and more cities there is a well defined peak. This is closely associated with the phase transition in probability of a tour existing. Median search cost for $n = 18$, 24, and 30 peaks at $\gamma = 0.6 \pm 0.1$, exactly the point where we observe that 50% of problems have no tour of the required length. The connection between search cost and the rescaled parameter used in finite-size scaling has also been noted in satisfiability and constraint satisfaction problems [21, 7]. Whilst the results here are restricted to a branch and bound algorithm, experience with other NP-complete problems suggest that search cost is likely to peak in the phase transition for a wider variety of procedures.

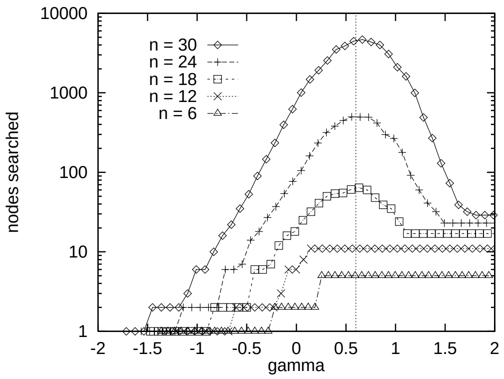

Fig. 2. Median number of nodes searched by the branch and bound algorithm in solving the decision problem, plotted against γ for 6, 12, 18, 24, and 30 city random problems. 1000 random instances were used at each point. The vertical line indicates where 50% of instances are soluble. Other probabilities can be interpolated from Fig. 1C.

It is important to study measures other than median search cost. In Figure 3A we plot further percentiles of search cost for 24 city problems against the control parameter, including the best and worst cases seen in each of 1000 problems. Search cost in the worst case can be orders of magnitude more than the median. We see similar behaviour if we fix the tour length and vary the number of cities visited [11]. Occasional underconstrained problems can be very hard for the branch and bound algorithm even where median behaviour is almost trivial. A similar phenomenon has been observed for graph-colouring and Boolean satisfiability problems [10, 12]. It appears to be the result of commitment to poor branching decisions early in search.

Random problems may be different to problems derived from geographical data. We therefore compared our results with some standard benchmark data from TSPLIB[20]. We use the capitals of the 48 contiguous states of the U.S.A. In Figure 3B, we fix the number of capitals visited at 24 and vary the tour length required. An ensemble of problems was generated by picking 1000 random subsets of 24 out of the 48 capitals. Behaviour is similar to that seen with random problems although the phase transition occurs at a smaller value of the control parameter, and search cost is typically greater. We have also observed similar behaviour using other procedures like simulated annealing [14] with both random and geographical data. Since simulated annealing cannot prove problems insoluble, we can only look at soluble traveling salesman problems. Nevertheless, search cost grows rapidly as we approach the phase boundary from the soluble side [11].

5 Conclusions

Our results correct the claim of Kirkpatrick and Selman in [15] that "there are other NP-complete problems (for example, the traveling salesman problem or max-clique) that lack a clear phase boundary at which 'hard problems' cluster". We have shown here that there is a sharp phase boundary for the traveling salesman decision problem at which search cost peaks. The max-clique problem is an optimization problem rather than a decision problem. The related NP-complete decision problem is deciding if a graph contains a k-clique. That is, if there exists k nodes of the

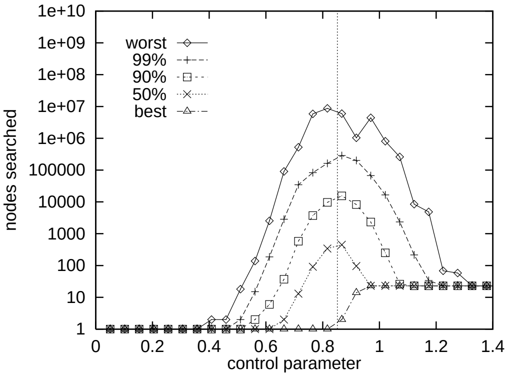

A

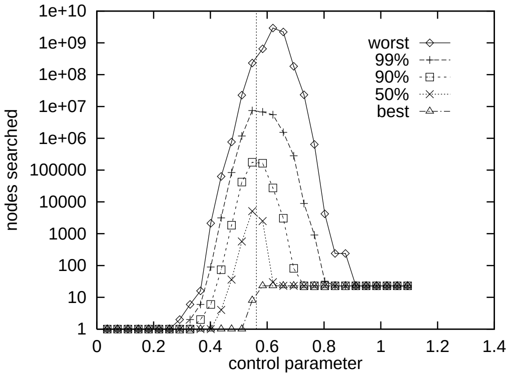

B

Fig. 3. Percentiles in the number of nodes searched by the branch and bound algorithm in solving 1000 decision problems at each point, plotted against the control parameter $l/\sqrt{n.A}$. The vertical line indicates where 50% of instances are soluble. **(A)** Random problems with 24 cities. **(B)** Random sets of 24 capitals of the contiguous states in USA. A is taken to be the surface area of the 48 states.

graph which are all connected to each other. Again, there is a sharp boundary between soluble and insoluble random k-clique problems, and hard problems cluster around this point. We showed this by investigating the (isomorphic) independent set problem [9].

Our results help to explain the hardness of optimization problems. To solve an optimization problem you *effectively* solve the related decision problem at the optimal tour length. That is, by solving the optimization problem, you have also shown that there exists a tour of length l_{opt} but no tour of length $l_{opt} - 1$, thus solving the decision problem immediately either side of the optimal tour length. Especially at large n, such decision problems are likely to be close to the phase boundary. They are thus usually very difficult. Indeed, the search process in an algorithm such as branch and bound represents precisely the solution of increasingly more constrained decision problems until finally the decision problems either side of the optimal tour length are solved. These results may therefore have implications in understanding the behaviour of other optimization procedures. In particular, they justify our earlier claim that traveling salesman optimization problems are typically more difficult than traveling salesman decision problems. However, traveling salesman optimization problems do also display transitions in behaviour. Using branch and bound, Zhang and Korf [24] have observed a transition in the cost of finding an optimal tour as inter-city distances are changed. They associated this transition with the probability that inter-city distances are the same length. Cheeseman *et al.* [4] have also observed that the difficulty of the optimization problem depends on the standard deviation in inter-city distances.

The phase transition phenomena described here appear to be widespread. It may seem surprising that methods like finite-size scaling developed in statistical mechanics to describe systems like spin glasses should be so widely applicable to combinatorial problems. There are, however, several important analogies to be made. For instance, interactions between atoms in a spin glass are both ferromagnetic (promoting alignment of spins) and anti-ferromagnetic (promoting opposite spins). As a result, a spin glass is a "frustrated" system with a large number of near optimal equilibrium configurations which cannot be locally improved by flipping a spin. It is difficult therefore to get the spin glass into a state of least energy. A traveling salesman problem is also usually a frustrated system, having a large number of near optimal tours which cannot be locally improved by making small changes. Frustration arises from the tension between visiting all cities (*i.e.* making a tour) and minimizing the distance travelled. Because of this frustration, it is very difficult to find the optimal tour from amongst the many near optimal tours. Similar remarks apply to many other NP-complete problems [16]. It remains an open question therefore if phase transition behaviour will be seen in all NP-complete problems.

Acknowledgments

We thank Alan Bundy and the Mathematical Reasoning Group at Edinburgh University (supported by SERC grant GR/H/23610), the MRG group at IRST (Trento) and the Department of Computer Science at the University of Strathclyde for donating a large number of CPU cycles to these experiments. The second author is supported by a HCM fellowship from the European Commission. We especially thank Robert Craig for his implementation of the TSP branch and bound algorithm, and referees of this journal for helpful comments.

References

[1] Michael N. Barber. Finite-size scaling. In *Phase Transitions and Critical Phenomena*, **8**, 145–266. Academic Press, 1983.

[2] J. Beardwood, J.H. Halton and J.M. Hammersley. The shortest path through many points. In *Proceedings of the Cambridge Philosophical Society*, **55**, 299–327, 1959.

[3] E. Bonomi and J-L. Lutton. The n-city travelling salesman problem: statistical mechanics and the Metropolis algorithm. *SIAM Review*, **26**, 551–568, 1984.

[4] P. Cheeseman, B. Kanefsky, and W.M. Taylor. Where the really hard problems are. In *Proceedings of the 12th IJCAI*, 331–337. International Joint Conference on Artificial Intelligence, 1991.

[5] J.M. Crawford and L.D. Auton Experimental Results on the Crossover Point in Random 3SAT. *Artificial Intelligence*, **81**, 31–57.

[6] M. R. Garey and D. S. Johnson. *Computers and intractability : a guide to the theory of NP-completeness.* W H Freeman, 1979.

[7] I.P. Gent, E. MacIntyre, P. Prosser, and T. Walsh. Scaling effects in the CSP phase transition. In U. Montanari and F. Rossi, editors, *Principles and Practice of Constraint Programming - CP95*, pages 70-87, Springer, 1995.

[8] I.P. Gent, E. MacIntyre, P. Prosser, and T. Walsh. The Constrainedness of Search. Technical Report, IRST, Trento, Italy, 1996. To appear in *Proceedings AAAI-96.*

[9] I.P. Gent and T. Walsh. The hardest random SAT problems. In B. Nebel and L. Dreschler-Fischer, editors, *KI-94: Advances in Artificial Intelligence. 18th German Annual Conference on Artificial Intelligence*, pages 355–366. Springer-Verlag, 1994.

[10] I.P. Gent and T. Walsh. Easy problems are sometimes hard. *Artificial Intelligence*, **70**, 335–345, 1994.

[11] I.P. Gent and T. Walsh. The TSP phase transition. Research report/95/178, Department of Computer Science, University of Strathclyde, 1995.

[12] T. Hogg and C. Williams. The hardest constraint problems: A double phase transition. *Artificial Intelligence*, **69**, 359–377, 1994.

[13] D.S. Johnson, L.A. McGeoch and E.E Rothberg. Asymptotic Experimental Analysis for the Held-Karp Traveling Salesman Bound. *Proc. of the 7th Annual ACM-SIAM Symposium on Discrete Algorithms*, 1996.

[14] S. Kirkpatrick, C.D. Gelatt, and M.P. Vecchi. Optimization by simulated annealing. *Science*, **220**, 671–680, 1983.

[15] S. Kirkpatrick and B. Selman. Critical behavior in the satisfiability of random boolean expressions. *Science*, **264**, 1297–1301, 1994.

[16] S. Kirkpatrick and R. Swendsen. Statistical Mechanics and Disordered Systems. *Communications of the ACM*, **28:4**, 363–373, 1985.

[17] E.L. Lawler, J.K. Lenstra, A.H.G. Rinnooy Kan, and D.B. Shmoys, editors. *The Traveling Salesman Problem.* John Wiley & Sons, Chichester, 1985.

[18] D. Mitchell, B. Selman, and H. Levesque. Hard and easy distributions of SAT problems. In *Proceedings, 10th National Conference on Artificial Intelligence*, 459–465. AAAI Press/The MIT Press, July 12-16 1992.

[19] P. Prosser. Binary constraint satisfaction problems: Some are harder than others. In A.G. Cohn, editor, *Proceedings of ECAI-94*, 95–99. John Wiley, Amsterdam, 1994.

[20] G. Reinelt. TSPLIB - a traveling salesman problem library. *ORSA Journal on Computing*, **3**, 376–384, 1991.

[21] B. Selman and S. Kirkpatrick. Critical behavior in the computational cost of satisfiability testing. *Artificial Intelligence*, **81**, 273–295.

[22] B. M. Smith. Phase transition and the mushy region in constraint satisfaction problems. In A.G. Cohn, editor, *Proceedings of ECAI-94*, 100–104. John Wiley, Amsterdam, 1994.

[23] J. Vannimenus and M. Mézard. On the statistical mechanics of optimization problems of the travelling salesman type. *J. Physique Lett.*, **45**, L1145–L1152, 1984.

[24] W. Zhang and R.E. Korf. An Average-Case Analysis of Branch-and-Bound with Applications: Summary of Results. In *Proceedings, 10th National Conference on Artificial Intelligence*, 545–550. AAAI Press/The MIT Press, July 12-16 1992.