

Phase Transitions of the Asymmetric Traveling Salesman*

Weixiong Zhang

Department of Computer Science and Engineering
Washington University in St. Louis
St. Louis, MO 63130, USA
Email: zhang@cse.wustl.edu

Abstract

We empirically study phase transitions of the asymmetric Traveling Salesman. Using random instances of up to 1,500 cities, we show that many properties of the problem, including the backbone and optimal tour cost, experience sharp transitions as the precision of intercity distances increases across a critical value. We also show that the average computational cost of the well-known branch-and-bound subtour elimination algorithm for the problem also exhibits a threshold behavior, transitioning from easy to difficult as the distance precision increases. These results provide strong positive evidences to a decade-long open question regarding the existence of phase transitions of the Traveling Salesman.

1 Introduction

Phase transition refers to such a phenomenon of a system in which some global properties change rapidly and dramatically when a control parameter crosses a critical value. A simple example of phase transition is water changing from the liquid phase to the solid phase when the temperature drops below the freezing point. Phase transitions of combinatorial problems and threshold behavior of combinatorial algorithms have drawn much attention in recent years [12]. It has been shown that many combinatorial decision problems have phase transitions, such as Boolean satisfiability [7; 21; 22], graph coloring [7] and number partitioning [4].

Another useful concept for characterizing combinatorial problems is that of backbones [17; 22]. A backbone variable refers to such a variable that has a fixed value in all solutions of a problem. All such backbone variables are collectively referred to as the backbone of the problem. The fraction of backbone, the percentage of variables in the backbone, reflects the constrainedness of the problem and directly affects an algorithm searching for a solution. The larger the backbone, the more tightly constrained the problem becomes.

*This research was supported in part by NSF grants IIS-0196057 and ITR/EIA-0113618, and in part by DARPA Cooperative Agreements F30602-00-2-0531 and F33615-01-C-1897. Thanks to Sharlee Climer for the joint work on an algorithm for finding backbones and Alan Frieze for a question on a bin-ball problem. Thanks also to Sharlee Climer, David Johnson, Scott Kirkpatrick and the anonymous reviewers for very helpful comments and suggestions to early versions of the paper.

In contrast to numerous phase-transition studies on decision problems, the research on the phase transitions and backbones of *optimization problems* is limited. An early work on the symmetric Traveling Salesman Problem, which is an optimization problem, introduced the concept of backbones and left an open question of whether there exists a phase transition of the TSP [17]. However, this question has not been addressed since 1985. One of the most rigorous phase-transition results was obtained on number partitioning [4], an optimization problem. However, the phase transition analyzed was the existence of a perfect partition of a set of integers, which is in essence a decision problem. In our early work, we studied the relationship between the phase transitions of satisfiability, a decision problem, and maximum satisfiability, an optimization problem [24]. In addition, the relationship between backbones and average-case algorithmic complexity has also been considered [23].

In this paper, we study the phase transitions of the asymmetric Traveling Salesman Problem (TSP), which has many real-world applications such as scheduling and routing. The TSP [13; 18] is an archetypical combinatorial optimization problem, and also very often a touchstone for combinatorial algorithms. In this paper, we consider the asymmetric TSP (ATSP), where the distance from one city to another may not be necessarily the same as the distance in the reverse direction. The ATSP is more general and difficult than the symmetric TSP (STSP). Our results provide strong positive evidences to the long-standing open question of [17] using the more general form of the problem. Another reason that we choose the ATSP rather than the symmetric TSP is that many properties of the assignment problem [19], a problem closely related to the ATSP, are known (Section 3.2) and useful to our analysis.

Specifically, using random problem instances of up to 1,500 cities, we empirically reveal that the average optimal tour length, the accuracy of the best cost function (the assignment problem), and the backbone of the ATSP undergo sharp phase transitions. The control parameter is the precision of intercity distances which is typically represented by the number of digits for the distances. Note that these results are algorithm independent and properties of the problem. Furthermore, we show that the average computational cost of the well-known branch-and-bound subtour elimination algorithm [1; 3] for the ATSP exhibits a threshold behavior, in which the computational cost grows abruptly and dramatically as the distance precision increases.

Two related results are worth mentioning. The research in [26] revealed that the average complexity of the subtour elim-

ination algorithm for the ATSP is controlled by the number of distinct intercity distances. We will further extend this result in Section 5. However, these results are algorithm specific, which may not necessarily reflect intrinsic features of the underlying problem. The research in [11] studied the decision version of the symmetric TSP. Specifically, it analyzed the probability that a tour of length less than a specific value exists for a random symmetric euclidean TSP, showing that the probability has a one-to-zero phase transition as the length of the targeting tour increases. Note that the phase-transition result in [11] does not address the open question of [17] which is on the optimization version of the problem.

2 The Problem and Algorithms

Given n cities and a matrix $D = (d_{i,j})$ that defines the distances or costs of pairs of cities, the *Traveling Salesman Problem* (TSP) is to find a minimum-cost complete tour visiting each city once. When the cost matrix is asymmetric, i.e., $d_{i,j}$ is not necessarily equal to $d_{j,i}$, the problem is the asymmetric TSP (ATSP), which is more difficult than the STSP, for both optimization and approximation [14].

The branch-and-bound (BnB) subtour elimination algorithm [1; 3] solves the ATSP using the assignment problem (AP) as a lower-bound cost function. The AP, which can be solved in $O(n^3)$, is to assign to each city i another city j with cost equal to distance $d_{i,j}$, so that the total cost of all such assignments is minimized [19]. If the AP solution happens to be a complete tour, it is also an ATSP solution.

The BnB search takes the ATSP as the root of the state space and solves the AP to the root node. If the AP solution is not a complete tour, decompose it into subproblems by selecting a subtour from the AP solution and generate subproblems by excluding some edges in the subtour to eliminate the subtour. We used the Carpaneto-Toth subtour-elimination heuristic [6] in our implementation, which generates no duplicate subproblem, so that the search space is a tree. Next, select as the current problem a new subproblem that has been generated but not yet expanded. This process continues until there is no unexpanded problem, or all unexpanded problems have costs greater than or equal to the cost of the best complete tour found so far.

Our algorithm is, in principle, the same as that of [5], which is probably the best known complete algorithm for the ATSP. The main difference between the two is that we use depth-first branch-and-bound (DFBnB), due to its low space requirement, while [5] used best-first search. We further extended our DFBnB algorithm to finding all optimal solutions and backbones.

3 Previous, Related Results

Two previous lines of work influenced this research and helped reveal the properties of the ATSP.

3.1 Phase transitions in tree search

To capture the tree search of the DFBnB algorithm for the ATSP, we first introduce an abstract tree model. An *incremental tree* [15; 20; 25], $T(m, d)$, is a tree with depth d , and independent and identically distributed (i.i.d.) random branching factors with mean m . Edges are assigned costs that are finite and nonnegative i.i.d. variables. The cost of a node is the sum of the edge costs along the path from the root to that node. An edge cost is nothing but the difference between a node cost and

the cost of its parent. An *optimal goal* is a node of minimum cost at depth d . The overall goal is to find an optimal goal node.

Let p_0 be the probability of a node having the same cost as its parent. The expected number of child nodes of a node that have the same cost as their parent is mp_0 . The following results have been proven [15; 20; 25]. Let C^* be the expected optimal goal cost of $T(m, d)$ with $m > 1$. Let N_{BFS} and N_{DFS} be the expected numbers of nodes expanded by best-first search and DFBnB, respectively. As $d \rightarrow \infty$, (1) $C^*/d \rightarrow \alpha_1$ and $N_{BFS} = N_{DFS} = \theta(\alpha_2^d)$ almost surely when $mp_0 < 1$, where α_1 and α_2 are constants; (2) $C^*/(\log \log d) \rightarrow 1$ and $N_{BFS} = \theta(d^2)$ and $N_{DFS} = O(d^3)$ almost surely when $mp_0 = 1$; (3) C^* remains bounded by a constant and $N_{BFS} = \theta(d)$ and $N_{DFS} = O(d^2)$ almost surely when $mp_0 > 1$.

The above results mean that the cost of an optimal goal node almost surely undergoes a phase transition from a linear function of depth d to a constant when mp_0 increases beyond one. Meanwhile, the expected complexity of best-first search and DFBnB changes dramatically from exponential to polynomial in d as mp_0 is reduced below one.

3.2 Relationship between AP and ATSP

The assignment problem (AP) cost function and its relationship with the ATSP have been a research interest for a long time [8; 9; 10; 16]. The relationship between the AP cost, $AP(D)$, and the ATSP cost, $ATSP(D)$, has remarkably different characteristics under different distance distributions of a random matrix D . The $AP(D)$ and $ATSP(D)$ can be the same with a high probability, or they can differ from each other, with a high probability, by a function of n . If the expected number of zeros in a row of D approaches infinity when $n \rightarrow \infty$, $AP(D) = ATSP(D)$ with a probability tending to one [9]. However, if the distances are uniform over $[0, 1, \dots, \lfloor c_n n \rfloor]$, $AP(D) = ATSP(D)$ with a probability going to zero, where c_n grows to infinity with n [9]. Indeed, when the entities of D are uniform over $[0, 1]$, $E(ATSP(D) - AP(D)) \geq c_0/n$, where c_0 is a positive constant [10].

4 Phase Transitions of the ATSP

The results in the previous section indicate that the quality of the AP function varies significantly, depending on the underlying distance distribution. Precisely, the difference between the AP cost, $AP(D)$, and the ATSP cost, $ATSP(D)$, has two phases, controlled by the number of zero distances in the distance matrix D . In one phase, the difference is zero with high probability, while in the other phase, the expectation of the difference is a function of the problem size n .

How does the difference between $ATSP(D)$ and $AP(D)$ change phases? Does it have a sharp phase transition, or does it follow a slow process? Do other properties of the ATSP, such as the backbone, also have phase transitions?

The two-phase result on the accuracy of the AP cost function discussed in Section 3.2 is in principle consistent with the phase-transition result of incremental random trees in Section 3.1. The root of the search tree has a cost equal to the AP cost $AP(D)$ to the problem and an optimal goal node has the ATSP tour cost $ATSP(D)$. If we subtract the AP cost to the root from every node in the ATSP search tree, the root node has cost zero and an optimal goal node has cost equal to $ATSP(D) - AP(D)$. When there are a large number of zero distances in D , the AP cost of a child node in a search tree is

more likely to be the same as the AP cost of its parent, since AP will tend to use the zero distances. Therefore, it is expected that more nodes in the search tree will have more than one child node having the same cost as their parents. In fact, the concept of zero distances can be extended to minimal-cost distances. If the cost of the minimal distance is known, we can simply subtract it from every distance without affecting the accuracy of the AP function.

4.1 The control parameter

Furthermore, the concept and role of zero and minimal distances can be extended to that of equal distances. The AP to a child node is obtained by excluding one arc, and possibly including some arcs, of a subtour of the AP solution to the parent node. With this new restriction, the child AP is then computed by constructing a new augmenting path from the starting point to the end point of the excluded arc [19]. The child and parent AP solutions usually differ by only a relatively small number of arcs. In other words, the child AP solution is derived by replacing some arcs in the parent AP solution. If the numbers of distances of equal values are large, it is more likely that the child AP cost is equal to the parent AP cost. Conversely, when the number of distinct distances is large, it is unlikely the costs of two APs will be the same. When the distances are uniform over $[0, 1, \dots, R - 1]$, the probability that the child AP cost is equal to the parent AP cost will depend on the range R . If R is small, relative to the problem size n , this probability will be high. In short, for a given problem size, the accuracy of the AP cost function is controlled by the number of distinct distances in matrix D . More precisely, the accuracy of the AP is determined by the fraction, denoted as $\rho(n)$, of distinct distances.

In practice, however, we do not directly control the number or the fraction of distinct distances. In addition to the actual structures of the “layouts” of the cities, the precision of the distances also affects the number of distinct distances. The precision of a number is usually represented by the maximal number of digits allowed for the number. As a result, the number of digits for distances is naturally a good choice for the control parameter.

The effect of a given number of digits on the fraction of distinct distances is relative to the problem size n . Consider a matrix D with distances uniformly over $\{0, 1, 2, \dots, R - 1\}$, where the range R is determined by the number of digits b . For a fixed b , the fraction of distinct distances of a larger matrix D is obviously smaller than that of a smaller D . Thus, the control parameter for the fraction $\rho(n, b)$ of distinct distances must be in the form of $b \cdot f(n)$, where $f(n)$ is a rescaling function on the effect of the number of digits.

To find the scaling function $f(n)$, consider the number of distinct distances of matrix D for a given integer range R . The problem of finding the number of distinct distances is equivalent to the bin-ball problem as follows. We are given M balls and R bins, and asked to place the balls into the bins. Each ball is independently put into one of the bins with an equal probability. We are interested in the fraction of bins that are not empty after all the placements. Here, for asymmetric TSP $M = n^2 - n$ balls correspond to the total number of nondiagonal distances of matrix D , and R bins represent the possible integers to be selected from. Since each ball (distance) is thrown independently and uniformly into one of R bins (integers), the probability that one bin is not empty after throwing

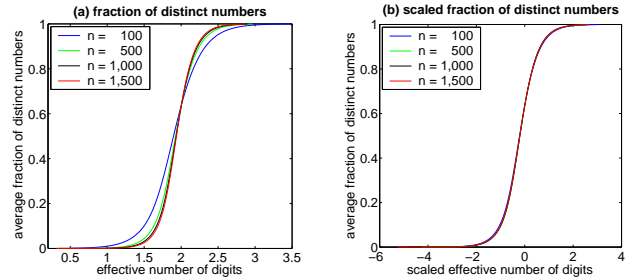


Figure 1: (a) Average fraction of distinct distances in matrix D , $\rho(n, b)$, controlled by the effective number of digits, $\beta = b \log_{10}^{-1}(n)$. (b) Average $\rho(n, b)$ after finite-size scaling. $\beta_0 = 2.00195 \pm 0.00025$ and $v = 6.4350 \pm 0.0003$.

M balls is $1 - (1 - 1/R)^M$. The expected number of occupied bins (distinct distances) is simply $R(1 - (1 - 1/R)^M)$. Thus, the expected fraction of distinct distances in matrix D is $E[\rho(n, b)] = R(1 - (1 - 1/R)^M) / M$. Note that if M or n is fixed, $E[\rho(n, b)] \rightarrow 1$ as $R \rightarrow \infty$, since in this case the expectation of the number of distinct distances approaches M . On the other hand, when R is fixed, $E[\rho(n, b)] \rightarrow 0$ when M or n goes to infinity, since all of a finite number of R bins will be occupied by an infinite number of balls in this case. We use decimal values for distances in this paper. Thus $R = 10^b$, where b is the number of digits for distances, and we can rewrite $E[\rho(n, b)]$ as $E[\rho(n, b)] = 10^b (1 - (1 - 1/10^b)^M) / M$.

It turns out that if we plot $\rho(n, b)$ against $b \log_{10}^{-1}(n)$, it will have relatively the same scale for different problem sizes n , as shown in Figure 1(a). This means that the scaling function for the effective number of digits is $f(n) = \log_{10}^{-1}(n)$. Function $b \log_{10}^{-1}(n)$ is thus the *effective number of digits* that controls the fraction of distinct distances in matrix D , which we denote as $\beta(n, b)$. To have the same effective number of digits β on two different problem sizes, say n_1 and n_2 with $n_1 < n_2$, the range R should be different. On these two problems, R needs to be n_1^β and n_2^β , respectively, giving $n_1^\beta < n_2^\beta$.

Note that R can also be represented by a number in other bases, such as binary. Which base to use will not affect the results quantitatively, but introduces a constant factor to the results. In fact, since $b = \log_{10}(R)$, the control parameter $\beta(n, b) = b \log_{10}^{-1}(n) = \log_n(R)$, which is independent of the base for intercity distances.

Controlled by the effective number of digits $b \log_{10}^{-1}(n)$, the fraction of distinct distances ρ has a phase transition of its own, also shown in Figure 1(a). The larger the problem, the sharper the transition, and there exists a crossover point among the transitions of problems with different sizes. We can examine the phase transitions more closely using finite-size scaling. Finite-size scaling [2] is a method that has been successfully applied to phase transitions in similar systems of different sizes. Based on finite-size scaling, around a critical parameter (temperature), problems of different sizes tend to be indistinguishable except for a change of scale given by a power law in a characteristic length. Thus, finite-size scaling can help to characterize a phase transition precisely around a critical point of the control parameter as the problem scales to infinity. For the problem at hand, the effective

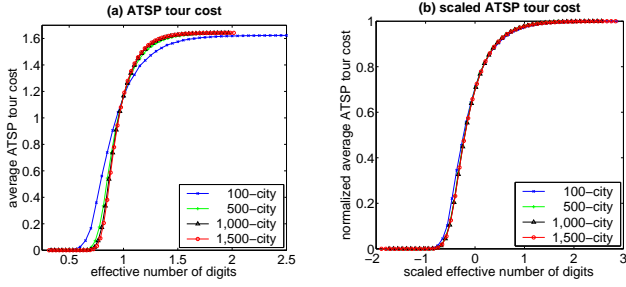


Figure 2: (a) Average optimal ATSP tour cost. (b) Scaled and normalized average optimal tour cost. $\beta_0 = 0.995 \pm 0.002$ and $v = 7.143 \pm 0.005$.

number of distinct distances $\beta = b \log_{10}^{-1}(n)$ and the problem size n play the roles of the temperature and the characteristic length, respectively. This means that as n goes to infinity, the control parameter becomes $(\beta - \beta_0)n^{1/v}$, where β_0 is the critical point and $n^{1/v}$ specifies the change of scale. Using $E[\rho(n, b)] = R(1 - (1 - 1/R)^M)/M$ and numerical methods, we obtained the critical value $\beta_0 = 2.00195 \pm 0.00025$ and rescaling component $v = 6.4350 \pm 0.0003$, where the error bounds represent the 95% confidence intervals. The rescaled phase transitions are shown in Figure 1(b).

Note that the number of digits used for intercity distances is nothing but a measurement of the precision of the distances. The larger the number of digits, the higher the precision becomes. This agrees with the common practice of using more effective digits to gain precision. Therefore, the phase transition of the control parameter is in turn determined by the precision of intercity distances.

4.2 Phase transitions

With the control parameter, the effective number of digits $\beta(n, b)$ for intercity distances, identified, we are now in a position to investigate possible phase transitions in (1) the ATSP cost, (2) the probability that an AP cost is equal to the corresponding ATSP cost, (3) the relative error or accuracy of the AP lower-bound function, and finally (4) the backbone of the ATSP. The answers to the first three problems provides a detailed picture on the accuracy of the AP cost function, and the answer to the last problem reveals the intrinsic constrainedness among the cities as the precision of distances changes. We examine these four problems in turn.

We generated uniformly random problem instances of 100-, 200-, 300- to 1,000-cities and 1,500-cities. Intercity distances are independently and uniformly chosen from $\{0, 1, 2, \dots, R-1\}$ for a given range R , which is controlled by the number of digits b . We varied b , with an increment of 0.1, from 1.0 to 6.0 for instances with up to 1,000-cities and from 1.0 to 6.5 for instances with 1,500-cities. For each combination of n and b , we generated 10,000 instances for problems (1)-(3) listed above, and 1,000 instances for problem (4) due to its high computational cost. To make the result figures readable, we only use the curves from 100-, 500-, 1,000- and 1,500-city problems.

There is a phase transition in the ATSP tour cost, $ATSP(D)$, shown in Figure 2(a). The reported tour costs are obtained by dividing the integer tour costs from the DFBnB algorithm by $n \times (R - 1)$, where n is the number of cities and R

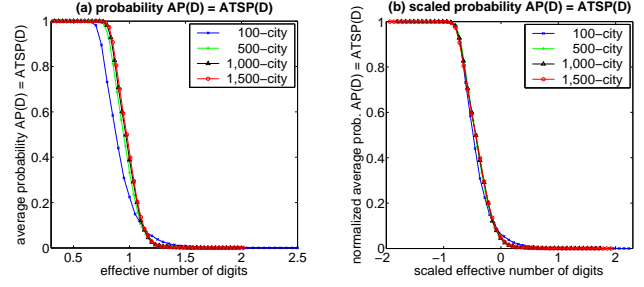


Figure 3: (a) Average probability that $AP(D) = ATSP(D)$. (b) Average probability after finite-size scaling. $\beta_0 = 1.169 \pm 0.02$, and $v = 8.929 \pm 0.004$.

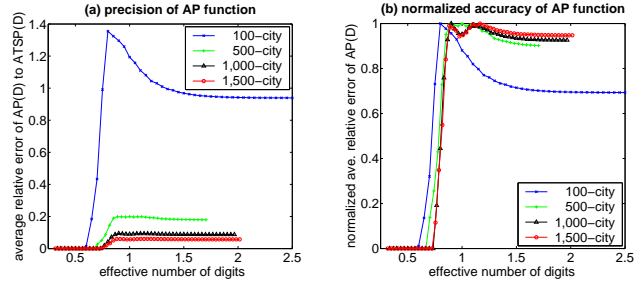


Figure 4: (a) Average accuracy of AP lower-bound function, measured by the error of AP cost relative to ATSP cost. (b) Normalized average accuracy.

the range of intercity costs. Equivalently, an intercity distance was virtually converted to a real value in $[0, 1]$. By doing this, we can verify the existing analytical result on the AP and ATSP costs. This will be discussed in detail in the next subsection.

As shown, the ATSP tour cost increases abruptly and dramatically as the effective number of digits increases, exhibiting phase transitions. The transitions become sharper as the problem becomes larger, and there exist crossover points among curves from different problem sizes. By finite-size scaling, we further determine the critical value of the control parameter at which the phase transitions occur. The scaled result is shown in Figure 2(b). It is worthwhile to mention that the AP cost (not shown) follows almost the same phase-transition pattern.

Our numerical results show that when the number of digits for intercity distances is very small, for example, less than 1.9 digits or $R \approx 80$ for $n = 1,500$, the AP and ATSP costs are equal to zero, meaning that these two costs are the same as well. Given a random distance matrix D , how likely is it that an AP cost will differ from the ATSP tour cost as the effective number of digits β increases? We address this question by examining the probability that an AP cost $AP(D)$ is equal to the corresponding ATSP cost $ATSP(D)$ as β increases. Figure 3(a) shows the results, averaged over the same set of instances for Figure 2. As shown, the probability that $AP(D) = ATSP(D)$ also experiences phase transitions. Figure 3(b) shows the phase transitions after finite-size scaling.

The results in Figure 3 also imply that the quality of the AP function degrades as the effective number of digits β increases. The degradation also follows a phase-transition process. This is verified by Figure 4, using the data from the same set of

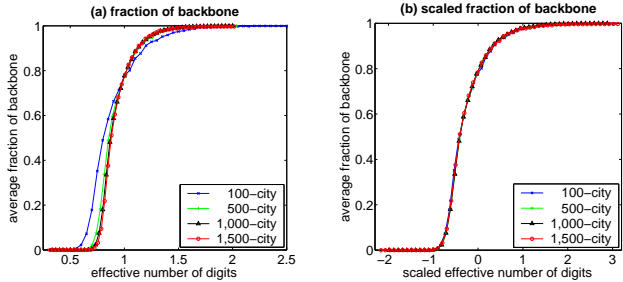


Figure 5: (a) Average fraction of backbone. (b) Scaled backbone fraction. $\beta_0 = 1.012 \pm 0.007$, $v = 6.452 \pm 0.002$.

n	digits	AP cost	ATSP cost	relative AP error (%)
2100	8.6810	1.6435 ± 0.0008	1.6444 ± 0.0008	0.0402 ± 0.0009
2300	8.7555	1.6439 ± 0.0007	1.6445 ± 0.0008	0.0368 ± 0.0008
2500	8.8235	1.6441 ± 0.0007	1.6442 ± 0.0007	0.0336 ± 0.0008
2700	8.8861	1.6438 ± 0.0007	1.6444 ± 0.0007	0.0309 ± 0.0007
2900	8.9440	1.6439 ± 0.0007	1.6444 ± 0.0007	0.0286 ± 0.0007

Table 1: Results on AP cost, the ATSP cost and AP error relative to the ATSP cost, on random ATSP. Each datum is averaged of 10,000 instances. All error bounds represent 95 percent confidence intervals.

problem instances for the previous two figures.

We now turn to the backbone of the ATSP, which is the fraction of directed arcs that appear in all optimal solutions. The backbones also exhibit phase transitions as the effective number of digits for distances increases. The result is included in Figure 5. Interestingly, the phase-transition pattern of the backbone is almost identical to that of the fraction of distinct entities in the distance matrix, shown in Figure 1, and that of the ATSP tour cost, shown in Figure 2.

The fraction of backbone captures, in essence, the tightness of the constraints among the cities. As more intercity distances become distinct, the number of tours of unique lengths increases. Consequently, the number of optimal solutions decreases and the fraction of backbone grows inversely. When more arcs are part of the backbone, optimal solutions become more restricted and the number of optimal solutions decreases, making it more difficult to find an optimal solution.

4.3 Asymptotic AP precision

As a by-product of the phase-transition results, we now provide asymptotic values of the ATSP cost, the AP cost and its accuracy. We attempt to extend the previous theoretical results on the AP cost, which is known to be within $(1.51, 1.94)$ [8], and the observation that the accuracy of the AP lower bounds increases as the problem size increases [1].

We need to be cautious in selecting the number of digits for intercity distances. As discussed earlier, the same number of digits for distances gives rise to different effective numbers of digits on problems of different sizes. Therefore, the number of digits must be scaled properly to have the same effect on problems of different sizes when we exam an asymptotic feature.

Therefore, in our experiments, we fixed the scaled effective number of digits β to a constant. Based on the phase-transition of the control parameter in Figure 1, we took $\beta = (b \log_{10}^{-1}(n) - 2.002)n^{1/6.435}$ a constant of 2. $\beta = 2$ is suf-

ficiently large so that all distances are distinct, regardless of problem size, and the quantities to be examined do not change substantially after finite-size scaling. $\beta = 2$ is also relatively small so that we can experiment on large problems. In our implementation of the DFBnB algorithm, distances are integers of 4 bytes. Thus the number of digits must be less than 9.4 without causing an overflow in the worst case. Using $\beta = 2$, we can go up to roughly 2,500-city ATSPs in the worst case.

Table 1 shows the results, with up to 2,900 cities. The AP cost approaches to 1.6439 and the ATSP cost to 1.6444. The accuracy of AP function indeed improves as the problem size increases, with relative error reduced to about 0.0286% for 2,900-city problems.

5 Threshold Behavior of Subtour Elimination

The phase-transition results indicate that the ATSP becomes more constrained and difficult as the distance precision becomes higher. We now study how the DFBnB subtour elimination algorithm behaves. We separate this issue from the phase transitions studied before because we now consider the behavior of a particular algorithm, which may not be necessarily a feature of the underlying problem. Nevertheless, this is still an issue of its own interest because this algorithm is among the best known methods for the ATSP, and we hope that a better understanding of this algorithm can shed light on the typical case complexity of the problem.

Figure 6(a) shows the average complexity of the DFBnB algorithm, measured by the number of calls to the AP function. The result is averaged over the same problem instances for each data point as used for the phase transitions in the previous section. Note that the number of AP calls increases significantly from small problems to large ones using the same effective number of digits for distances. Thus, we normalize the result in such a way that for a given problem size, the minimal and maximal AP calls among all problem instances of the same size are mapped to zero and one, respectively, and the other AP calls are proportionally adjusted to a ratio between 0 and 1. This allows us to compare the results from different problem sizes on one figure. The curves in Figure 6(a) follow a pattern similar to that of the phase transitions in the previous section. The complexity of the subtour elimination algorithm increases with the effective number of digits, and exhibits a threshold behavior similar to phase transitions. Indeed, we can use finite-size scaling to capture the behavior as the problem size grows, as illustrated in Figure 6(b). The results in Figure 6 and the results in the previous section indicate that the complexity of the subtour elimination algorithm goes hand-in-hand with the constrainedness of the problem.

Similar results have been reported in [26]. The results of this subsection extend that in [26] to different sizes of problems and by applying finite-size scaling to capture the threshold behavior as problem size increases.

6 Conclusions and Discussions

Our main contributions are twofold. First, we provided strong positive evidences to the long-standing question of whether the Traveling Salesman Problem (TSP) has phase transitions [17]. We studied this issue on the more general asymmetric TSP (ATSP). We empirically showed that many properties, including the ATSP tour cost and the fraction of backbone, have

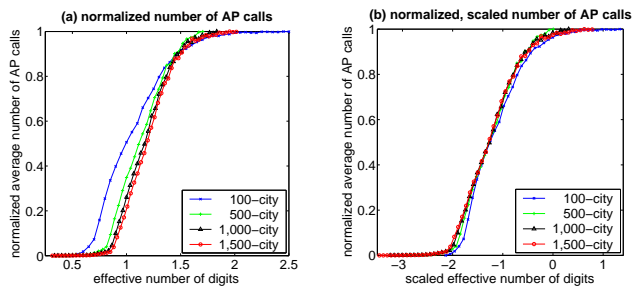


Figure 6: (a) Normalized average number of AP calls of DF-BnB subtour elimination. (b) Scaled average number of AP calls. $\beta_0 = 1.702 \pm 0.045$ and $v = 8.015 \pm 0.032$.

two characteristically different values, and the transitions between them are rather abrupt and dramatic, displaying phase-transition phenomena. The control parameter of the phase transitions is the effective number of digits representing the intercity distances, which in essence measures the precision of distances. Our results revealed the connection between distance precision and phase transition properties in the ATSP. Distance precision is the control parameter for various phase transitions of the ATSP. We believe that the concept of precision determining problem properties such as phase transitions is rather universal and may very well be applicable to other problems, including various scheduling and planning problems.

Second, our phase transitions results provide a practical guidance to how to generate difficult random ATSP problem instances and which instances to use to compare the asymptotic performance of two algorithms. A common practice in comparing algorithms when using a random ensemble is to generate problems with a fixed distance precision. Our results imply that the correct way is to use problem instances of different sizes that have the same or similar features such as the same fraction of backbones. This, in turn, requires to increase the number of digits for intercity distances as the problem size grows. We believe that this guidance is general and can be applied to other optimization problems.

References

- [1] E. Balas and P. Toth. Branch and bound methods. In *The Traveling Salesman Problem*, pages 361–401. John Wiley & Sons, Essex, England, 1985.
- [2] M. N. Barber. Finite-size scaling. In *Phase Transitions and Critical Phenomena*, volume 8, pages 145–266. 1983.
- [3] M. Bellmore and J. C. Malone. Pathology of traveling-salesman subtour-elimination algorithms. *Operations Research*, 19:278–307, 1971.
- [4] C. Borgs, J. T. Chayes, and B. Pittel. Phase transition and finite-size scaling for the integer partitioning problem. *Random Structures and Algorithms*, 19:247–288, 2001.
- [5] G. Carpaneto, M. Dell’Amico, and P. Toth. Exact solution of large-scale, asymmetric Traveling Salesman Problems. *ACM Trans. on Mathematical Software*, 21:394–409, 1995.
- [6] G. Carpaneto and P. Toth. Some new branching and bounding criteria for the asymmetric traveling salesman problem. *Management Science*, 26:736–743, 1980.
- [7] P. Cheeseman, B. Kanefsky, and W. M. Taylor. Where the really hard problems are. In *Proc. IJCAI-91*, pages 331–337.
- [8] D. Coppersmith and G. B. Sorkin. Constructive bounds and exact expectations for the random assignment problem. *Random Structures and Algorithms*, 15:113–144, 1999.
- [9] A. Frieze, R. M. Karp, and B. Reed. When is the assignment bound asymptotically tight for the asymmetric traveling-salesman problem? In *Proc. of Integer Programming and Combinatorial Optimization*, pages 453–461, 1992.
- [10] A. Frieze and G. B. Sorkin. The probabilistic relationship between the assignment and asymmetric traveling salesman problems. In *Proc. of SODA-01*, pages 652–660, 2001.
- [11] I. Gent and T. Walsh. The TSP phase transition. *Artificial Intelligence*, 88:349–358, 1996.
- [12] C. P. Gomes, T. Hogg, T. Walsh, and W. Zhang. IJCAI-2001 tutorial: Phase transitions and structure in combinatorial problems. <http://www.cs.wustl.edu/~zhang/links/ijcai-phase-transitions.html>.
- [13] G. Gutin and A. P. Punnen, editors. *The Traveling Salesman Problem and Its Variations*. Kluwer Academic Publishers, 2002.
- [14] D. S. Johnson, G. Gutin, L. A. McGeoch, A. Yeo, W. Zhang, and A. Zverovich. Experimental analysis of heuristics for the ATSP. In G. Gutin and A. Punnen, editors, *The Traveling Salesman Problem and its Variations*, pages 445–488. 2002.
- [15] R. M. Karp and J. Pearl. Searching for an optimal path in a tree with random costs. *Artificial Intelligence*, 21:99–117, 1983.
- [16] R. M. Karp and J. M. Steele. Probabilistic analysis of heuristics. In *The Traveling Salesman Problem*, pages 181–205. 1985.
- [17] S. Kirkpatrick and G. Toulouse. Configuration space analysis of traveling salesman problems. *Journale de Physique*, 46:1277–1292, 1985.
- [18] E. L. Lawler, J. K. Lenstra, A. H. G. Rinnooy Kan, and D. B. Shmoys, editors. *The Traveling Salesman Problem*. John Wiley & Sons, Essex, England, 1985.
- [19] S. Martello and P. Toth. Linear assignment problems. *Annals of Discrete Mathematics*, 31:259–282, 1987.
- [20] C. J. H. McDiarmid. Probabilistic analysis of tree search. In *Disorder in Physical Systems*, pages 249–260. 1990.
- [21] D. Mitchell, B. Selman, and H. Levesque. Hard and easy distributions of SAT problems. In *Proc. AAAI-92*, 1992.
- [22] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky. Determining computational complexity from characteristic ‘phase transitions’. *Nature*, 400:133–137, 1999.
- [23] J. Slaney and T. Walsh. Backbones in optimization and approximation. In *Proc. IJCAI-01*, pages 254–259, 2001.
- [24] W. Zhang. Phase transitions and backbones of 3-SAT and Maximum 3-SAT. In *Proc. CP-2001*, pages 153–167, 2001.
- [25] W. Zhang and R. E. Korf. Performance of linear-space search algorithms. *Artificial Intelligence*, 79:241–292, 1995.
- [26] W. Zhang and R. E. Korf. A study of complexity transitions on the asymmetric Traveling Salesman Problem. *Artificial Intelligence*, 81:223–239, 1996.