

Traveling Salesman Heuristics and Embedding Dimension in the Hopfield Model

Andrew B. Kahng

Department of Computer Science and Engineering
Mailcode C-014
University of California, San Diego
La Jolla, CA 92093
kahng@cs.ucsd.edu

Abstract. *We discuss methods for improving performance of the Hopfield-Tank quadratic minimization approach to TSP. A wide range of geometric (e.g., convex-hull based), topological and cutting-plane heuristics are investigated. We also investigate performance on non-Euclidean and non-metrizable TSP instances using a new concept of embedding dimension. Implications concerning the nature of the Hopfield energy surface are discussed. We conclude that the Hopfield-Tank formulation is not as robust as might be hoped; however, it remains well-suited to many important applications.*

1. Introduction.

The oft-cited papers of Hopfield and Tank [11][23] are among the earliest to report successful application of neuro-computing to NP-complete [6] optimization problems. In particular, their work on the traveling salesman problem (TSP) has generated much interest due to the elegance of their formulation and the status of TSP as a canonical "intractable" problem.

The results in [11] have drawn attention for their irreproducibility; [25] and others have demonstrated that the Hopfield-Tank (H-T) TSP algorithm will often fail to produce a valid tour and is highly sensitive to parametrization and initial conditions. The obvious problem is that the terms of the H-T energy function separately attempt to enforce structure (i.e., a valid tour is exactly equivalent to a permutation matrix in the output) and cost minimization (i.e., the tour length should be minimum). When structure is enforced, output tours are of poorer quality; when we attempt to minimize cost, invalid outputs result. A host of papers, notably [4][9][19][21], are essentially devoted to scaling heuristics which are designed to make the H-T method yield correct tours.

It may be asked whether "fixing" the Hopfield-Tank approach to TSP is worthwhile. The fact that the network

must scale as $O(n^2)$ in problem size does not bode well for practical implementation. However, folklore has it that the Hopfield network will perform well on non-Euclidean or even non-metric (i.e., failing to satisfy the triangle inequality) TSP incidence matrices, while alternatives such as the elastic net [5] and the adaptive ring and its variants [1][12], despite linear scaling with problem size, are not as robust for such problems. This is the intuitive consequence of the "physical" and geometric heritage of the latter methods: after all, if we cannot even embed our cities in the plane, then seemingly there is little sense in using, e.g., a planar elastic net. It is well-known that the vast majority of practical TSP instances are non-planar or non-Euclidean, and so the model's purported robustness has been a motivating force behind continuing Hopfield-Tank TSP research. At the same time, the H-T TSP approach has been treated with kid gloves: it has not been judged by the same standards as other TSP heuristics. Among the references listed below are several which report only valid TSP solutions; others speak of heuristic "percentiles" in assessing solution quality (the ability to actually calculate TSP percentiles would imply the resolution of several open questions in combinatorial geometry and optimization).

With this in mind, the goal of our work is to explore the strengths and limitations of the Hopfield-Tank formulation, drawing on the rapidly-growing body of Hopfield TSP literature and additional intuition derived from the operations research, discrete optimization, and computer science theory literatures. (This is particularly appropriate as the Hopfield-Tank formulation is essentially an instance of classical constrained quadratic minimization.) We experiment with a number of heuristic H-T variants which stem from computational geometry (convex hull, clustering analysis), combinatorial geometry and discrete optimization (forbidden edge, cutting planes), etc. In contrast to virtually all previous work except [9][25], we evaluate performance not by the quality of an isolated "best" solution, but rather by the incidence of valid tours, the relationship between the average tourlength and the expected random tourlength, and so forth.

BACKGROUND

An instance of the traveling salesman problem is given by the following:

Definition: A TSP instance for n cities consists of a symmetric n by n matrix A with non-negative real off-diagonal entries a_{ij} , where a_{ij} denotes the distance between the i^{th} and j^{th} cities.

The original formulation of Hopfield and Tank [11] uses n^2 output nodes in the form of a square matrix, where rows correspond to the individual cities and columns correspond to position in the tour. Thus, a valid tour will be in the form of a permutation matrix, i.e., exactly one city is in a given tour position, and every city is visited.

The Hopfield network finds a local minimum of the energy functional

$$E = \frac{A}{2} \sum_X \sum_i \sum_{j \neq i} v_{Xi} v_{Xj} + \frac{B}{2} \sum_i \sum_X \sum_{Y \neq X} v_{Xi} v_{Yi} + \frac{C}{2} (N - \sum_X \sum_i v_{Xi})^2 + \frac{D}{2} \sum_X \sum_{Y \neq X} \sum_i d_{XY} v_{Xi} (v_{Y,i+1} + v_{Y,i-1}) \quad (1)$$

The A term constrains each row to have one nonzero element; the B term acts similarly to constrain each column; the C term is a global scaling factor which ensures that exactly n nodes are active; and the D term represents tour distance. Thus, the final term is the only one which reflects cost; the remaining three terms enforce structure in the solution. It is sometimes useful to write down the implicit connection matrix, which is defined by

$$T_{Xi,Yj} = -A \delta_{XY} (1 - \delta_{ij}) - B \delta_{ij} (1 - \delta_{XY}) - C - D d_{XY} (\delta_{j,i+1} + \delta_{j,i-1}) \quad (2)$$

where δ_{ij} denotes Kronecker delta.

Many researchers have noticed that with the parameters values suggested in [11], tour quality is good *when a valid tour is found*, but unfortunately only a small percentage of outputs are valid tours. Our own investigations with these parameters indicate that approximately 85 percent of the attempts result in invalid output for random 10-city instances in the Euclidean square. A method which fails almost all the time, even when given relatively trivial tasks, will certainly never be very viable, and therefore much effort has been spent to increase the incidence of valid H-T tours without sacrificing tour quality.

Various groups have reported methods for forcing the Hopfield network to yield a valid permutation matrix output.

Tagliarini and Page [21][22] describe a normalization which enforces "k out of n" constraints in the output. For example, a valid TSP solution must satisfy a "1 out of n" constraint in each row and column of the output. Rewriting (1) as

$$E = -\frac{1}{2} \sum_i \sum_j T_{ij} v_i v_j - \sum_i V_i I_i$$

they show the following:

Fact: Given a subset of S neurons in the Hopfield network, if the $T_{ij} \equiv -M$ for all $i, j \in S$, M an arbitrary positive real, then setting $I_i \equiv Mk - 1$ will guarantee that exactly k neurons in S will be active at equilibrium. \square

This allows normalization of A , B , and C in (1) so that only the value of D changes ([24] achieves a similar reduction in parameterization). Tagliarini and Hanrahan [20] report that the incidence of valid tours does not vary strictly inversely with D , as might be expected. This implies an unexplained sensitivity to initial conditions. Furthermore, in order to completely enforce valid tours, they conclude empirically that D must scale *sublinearly* with problem size. In fact, Hegde et al. [9] show, again empirically, that the ratio D/C vanishes with increasing n if we wish to guarantee valid tours. Our own experiments confirm these observations. Since $D = 0$ will yield random tours as output, it seems that there is little hope that Hopfield TSP can be made to yield good tours on every input.

Van den Bout and Miller [24] propose a battery of heuristic fixes to the Hopfield-Tank TSP formulation; they argue that the revised energy function

$$E = \frac{B}{2} \sum_i \sum_X \sum_{Y \neq X} v_{Xi} v_{Yi} + \frac{1}{2} \sum_X \sum_{Y \neq X} \sum_i d_{XY} v_{Xi} (v_{Y,i+1} + v_{Y,i-1})$$

with $B = 4\sqrt{N}$ will guarantee that no more than one city can occupy any tour position. In addition, [24] proposes heuristic seeding of the tour, e.g., by finding a pair of close neighbors and assuming that they will be adjacent in the tour. This heuristic is a member of the class of symmetry-breaking heuristics discussed below. Finally, the authors propose renormalizations which essentially eliminate any asymmetry in the actual distribution of city locations. This is a subset of the class of weight-skewing techniques discussed below.

Brandt et al. [3] give an alternative energy function wherein the first three terms of (1) are modified to

$$\frac{A}{2} \sum_i (\sum_j x_{ij} - 1)^2 + \frac{B}{2} \sum_j (\sum_i x_{ij} - 1)^2$$

Again, this ostensibly enforces permutation matrix structure for the output.

In [4], Clement et al. propose yet another scaling method which uses a heuristic estimate of tourlength to normalize parameters. Their claim is that the transformation

$$D' = D / \text{hypotenuse}$$

will allow valid solutions to result, where *hypotenuse* denotes the length of the diagonal of the bounding box of city locations.

Finally, Szu [19] adopts a rather different approach from any of the above work. He "repudiates" the work of Wilson and Pawley [25] with a Hopfield-Tank variant featuring binary output neurons and highly non-standard values for, e.g., global inhibition. However, the results presented seem to indicate only 9.1% incidence of valid 10-city tours; furthermore, the average successful tourlength in the data shown is much worse than can be achieved with the H-T parameters in [11].

Overall, it may be argued that the pre-processing analysis in, e.g., [24], requires a large amount of non-neural computation on a conventional processor. The associated CPU expenditure can very easily exceed the cost of a solution by heuristics such as nearest-neighbor or minimum spanning tree. This is obviously an important objection. However, as argued in [14], analog solutions to TSP are potentially useful for several reasons. In particular, they can provide fast heuristic initial solutions for hybrid neural-serial algorithms; they also give immediate lower bounds for problem classes such as matching and assignment. Furthermore, they can have small constant-factor error bounds, unlike, e.g., nearest-neighbor methods. Thus, we move on to discuss several classes of Hopfield TSP variants.

HOPFIELD TSP HEURISTICS

Initial experiments led to a number of straightforward assumptions and observations which constrained the scope of our work. We originally examined Hopfield-Tank TSP results for groups of 1000 randomly generated Euclidean instances of 10 and 20 cities with locations uniformly distributed in the unit square.

We noted many symptoms of numerical instability and sensitivity to initial conditions. For example, the final tour is highly dependent on the vector used for initial perturbation, or "symmetry breaking", of the network. We do not believe that this is due to the $2n$ -fold tour degeneracy; rather, the initial state seems to lock us into a particular region of configuration space. Heuristic fixes to this problem have many side effects, e.g., reducing the magnitude of the initial perturbation often leads to better-quality tours, but the number of false convergences increases.

Early experiments also allowed us to search for a useful performance metric. It was clear that the incidence of valid tours should be a part of any measure of success. Also, it was not necessary to evaluate 181,440 tours in order to find optimum 10-city tourlengths; known asymptotics and the law of large numbers suffice. If we denote the length of the optimal tour by T_{opt} , it is well known that for uniform distributions of N points in a Euclidean rectangle of area A , the ratio $\frac{T_{opt}}{\sqrt{NA}} \rightarrow \tau_2$, τ_2 a constant, as N becomes large [15]. Monte Carlo simulations and annealing results by Kirkpatrick and others give the estimate $\tau_2 = 0.749$ [2][15]. We can show by

equivalence of norms that such a constant τ_p must exist for any L_p norm (the Euclidean metric corresponds to L_2). Held and Karp use a one-tree approach which for random distributions yields an asymptotically good lower bound on Euclidean tour length of $0.708 \cdot \sqrt{NA}$ [15]. Upper bounds are provided by any TSP heuristic, e.g., Kernighan-Lin k -opts or the nearest-neighbor algorithm. Since asymptotics may not be close for such small values as $n = 10$, we evaluate the heuristic solution by using an average of Held-Karp and Kernighan-Lin results over a large number of instances. This procedure yields a value T^* that is *almost certainly* very close to the expected length of an optimal tour. A perfectly reasonable alternative is to use the expected value of a random tour as a benchmark; by the theory of matching lower bounds for TSP tourlength, this value grows at the same rate as T^* and has the added advantage of being trivial to calculate.

Further observations, conventions and assumptions include the following:

- The constants A and B are symmetric in the formulation, so we force them to be equal (we arbitrarily set $A = B = 500.0$; scalings of C and D are thus the only ones allowed. It can be seen that for some choices of C , invalid tours will result even when $D = 0$. For each Hopfield-Tank variant, we found a heuristic *operating point* as follows: with $A = B = 500$ and $D = 0$, find the integer value of C for which a maximal percentage of valid tours results (or take the median of all such values). Then, the values of A , B and C define the operating point of the algorithm, and we parametrize results by the value of D .
- Each data point gives average results for a set of 1000 randomly generated planar Euclidean TSP instances, consisting of 10 or 20 cities with locations uniformly distributed in the unit square. In the figures below, we plot percentage of valid tours and a normalized ratio of average tourlength and expected random tourlength (the latter number differs by a constant factor from the ratio of average tourlength and T^*). The normalizing factor is \sqrt{n} , i.e., the rate of growth of expected tourlength.
- There are two basic ways of "skewing" the input: (i) modifications of the TSP distance matrix which are then propagated to the T_{ij} matrix, and (ii) directly modifying the T_{ij} values (note that $2n$ -fold degeneracy exists for this modification). We observe that the latter technique is much more successful in general; below we speculate as to why this is so. There is a sharp drop in solution quality if the magnitude of change in the T_{ij} entries exceeds C .
- Even when A , B and C are held constant, D does not necessarily vary with solution quality (average tourlength) or inversely with the percentage of valid tours. These relationships seem to be highly complicated (see Figure 1) and are further evidence of instability.

The following classes of heuristic variants were evaluated.

Nearest-Neighbor Skewing. We can see empirically that in an optimum TSP tour, a given city is usually connected to one of its close neighbors. It is therefore natural to use a sparse incidence matrix, where only elements which are among the k lowest in their respective rows and columns are retained. Hopfield and Tank in their original paper proposed exactly this heuristic enhancement with $k = 4$. Heuristic *NN 1* gives a smooth implementation of T_{ij} skewing, i.e., a pair of nearby cities will have greater encouragement to be adjacent in the tour than would two cities which are further apart. *NN 2* perturbs all T_{ij} values by a constant.

Cutting-Plane Skewing. The dual of nearest-neighbor skewing consists of forbidding those edges that are unlikely to occur in the optimal tour. For example, it is improbable that a city will be connected to its most distant neighbor, and so we can arbitrarily inhibit the formation of such a link. In heuristic *CP 1*, we simply inhibit all links corresponding to the four most expensive adjacencies of each city. Heuristic *CP 2* is considerably more sophisticated, and uses what are essentially cutting-plane techniques [13] to determine forbidden edges. Both *CP 1* and *CP 2* are "discontinuous"; there is no variation of the incremental inhibition with distance or other parameters.

Convex Hull Analysis. In tracking the convergence of the H-T output nodes, one may observe that the first $n - 1$ cities are usually fixed in a sensible manner; the last city is then interpolated into the tour at the remaining feasible position. An analysis of 1000 runs yielded 998 cases where the last city to converge (call it C) was part of the convex hull of tour locations. Furthermore, we have found that C is almost always either the city furthest from the remaining cities (i.e., maximal sum of distances to two nearest neighbors) or the one for which the remaining cities subtended the greatest angle (because it is an extreme point of the set of locations, there is a line l passing through C such that the remaining cities lie on one side of l). Heuristic *CH 1* reduces the inhibitions for T_{ij} entries corresponding to the greatest subtended angle of any extreme point; heuristic *CH 2* reduces the inhibitions for entries corresponding to the nearest neighbors of the most isolated city. Note that *CH 2* is analogous to the renormalization technique proposed in [24].

Symmetry Breaking. Several researchers claim that the Hopfield-Tank TSP formulation is hampered by a high "degeneracy" of (good) solutions. Under such a paradigm, the network fails because it cannot decide on any one of several equivalent good tours; partially satisfying each of these good solutions results in an invalid output. Though there is no real evidence for the notion of such symmetric local minima, we tried two heuristics: (i) *SB 1* simply forces city 1 to be in position 1, and (ii) *SB 2* forces city 1 to be in position 1 and its nearest neighbor to be in position 2.

The eight heuristics were all (with the possible exception of *CH 1*) relatively plausible at the start of experimentation; they were all motivated by accepted results in the optimization and neural network literatures. But, as seen in the results below, some of them significantly worsen the performance of H-T TSP. There is little in the literature to explain such a phenomenon, and so we find it necessary to introduce a new characterization of "difficult" input TSP matrices for the Hopfield network.

Non- Δ TSP INSTANCES and EMBEDDING DIMENSION

As noted above, the bulk of neural network TSP literature deals with planar Euclidean TSP instances, and some heuristics are dependent on this aspect of the input data. However, suppose we have an instance where

$$d(a,c) > d(a,b) + d(b,c)$$

for some choice of cities a , b and c . We say that such an input is non-metrizable, and we note that such inputs can be created from Euclidean instances by executing heuristics listed above. Non-metrizable instances comprise the vast majority of those encountered in real-world situations; in such cases, geometric concepts such as convex hull or area maximization become nearly meaningless. Sometimes, though, a middle ground exists.

Consider the following example. Imagine that travel costs are 100 dollars from A to B , 200 dollars from B to C , and 800 dollars from A to C . Clearly the triangle inequality does not hold, but this does not mean that all is lost: in real life, a poor salesman who must travel from A to C will take the detour through B even if he has already visited B .

Thus in this situation, as with certain practical applications, one can smooth the TSP incidence matrix, e.g., by applying Floyd's shortest-path ("triple operation") algorithm:

```

for i = 1, n {
  for j = 1, n {
    for k = 1, n {
      d(i,j) = max [d(i,j), d(i,k) + d(k,j)]
    }
  }
}

```

The resulting distance matrix will satisfy the triangle inequality. For TSP instances that satisfy the triangle inequality (i.e., the class of Δ -TSP instances, in the terminology of [6]), we have investigated a new hierarchy defined by the *embedding dimension*. (These results are the subject of current study.) We digress briefly to characterize embedding dimension.

Consider a symmetric n by n TSP distance matrix A . There are

$$\binom{n}{2} = \frac{n(n-1)}{2}$$

inter-city distances A_{ij} that are of interest. If each city's location in R^k is given by a k -tuple of reals, then we have kn unknowns. (Alternatively, there are $k(n-1)$ unknowns, if we assume one city to be at the origin.) Because there are only $\frac{n(n-1)}{2}$ equations, we have an underspecified system when the dimension k is greater than $n/2$, and thus a solution (i.e., an embedding in R^k) exists.

[A simple inductive argument can also be used. Consider the distance a_{12} : it can be realized by placing the first and second cities a_{12} units apart on the line. Now, we can solve for the location of city 3 using distances a_{13} and a_{23} , and so forth. In general, the first k cities will define a hyperplane in R^k , and the $(k+1)^{st}$ city can therefore be embedded in R^k . This is obviously a poorer bound on k than that derived above.]

It is NP-hard to find the minimum embedding dimension of a given symmetric TSP incidence matrix, and the concept is little-studied. (Graph-theoretic variants of this problem have been studied, e.g., by Graham and Winkler [8] and Sherlekar and JaJa [18].)

As embedding dimension increases, no clear trend in solution quality or frequency is discernible; this is perhaps due to the limited problem sizes we are studying. (We conjecture that the probability of H-T TSP finding a valid tour will in general decrease as embedding dimension increases.)

As it is sometimes impossible to "smooth" an incidence matrix, we must examine how Hopfield TSP performs for random or other non-metrizable input. Preliminary results for non-metrizable distance matrices are very interesting, and are reported below. As there is no established measure of "nonmetric-ness" for an incidence matrix A , we use the value $P(A)$ as a parameter, where $P(A)$ denotes the probability that three given cities (a "triple") determine distances which violate the triangle inequality. In particular, it seems that Hopfield-Tank TSP is not as robust as previously believed.

RESULTS AND CONCLUSIONS

Sample results for the four classes of TSP variants are presented in Figures 2-5. Figure 6 shows some results for the standard Hopfield-Tank implementation. As noted earlier, we are interested in (i) the percentage of final configurations which correspond to valid tours; and (ii) the average valid tourlength as a percentage of expected random tourlength. Figure 7 gives an example of how the frequency of valid tours varies inversely with $P(A)$ for non-metrizable instances.

The Hopfield-Tank TSP algorithm is extremely sensitive to initial perturbative values, scaling of parameters, and even convergence criteria. Such characteristics have been noted by previous workers and ascribed to frustration of the energy configuration. Heuristic fixes proposed in the literature are disappointing: the sparse matrix and forbidden-edge methods proposed by Hopfield and Tank are unsuccessful unless they are "smoothed", and symmetry-breaking also fails to improve

performance. Of course, our work is by no means unsuccessful. For example, the NN1 heuristic was able to obtain better-quality solutions than those reported by Hopfield and Tank, while at the same time succeeding on almost 92% of its inputs. The NN1 heuristic was also by construction the least disruptive of a smooth energy configuration.

The most interesting aspect of our results concerns embedding dimension and metrization of the TSP incidence matrix. In implementing the various heuristics, we observe that as the distance matrix accumulates a large number of non- Δ triples, the incidence of invalid tours rises dramatically. The same good solutions are still feasible in the revised distance matrix (and they remain in the same "percentile" of all solutions), but somehow they become harder, not easier, to find. We conclude that although Hopfield-Tank remains a pioneering vehicle for TSP solution, its robustness for random or non-metric inputs may not be as great as previously supposed.

FUTURE WORK

The erratic performance of Hopfield-Tank TSP algorithm is partly due to the discrete nature of structure terms in the energy functional. We are experimenting with a dual method of enforcing the permutation matrix structural constraint. Essentially, we define a less strict penalty term which will theoretically allow trajectories outside the hypercube, but which does not do so in practice. Convergence to permutation matrix structure is a consequence of the Cauchy-Schwartz inequality.

The discovery that NN1 is a good heuristic may confirm the intuition in [14] concerning space-filling curves and Hopfield TSP formulations. It is easy to see that the NN1 heuristic can be extended to yield a TSP algorithm which optimizes by the generalized curvature metric of [14]. But will the algorithm fail for non-metrizable inputs?

With regard to embedding dimension, we ask if discrepancy and matching lower bounds on optimal TSP value can help us predict performance ratios for Hopfield TSP in higher dimensions. In particular, does the error grow with n at the same rate for all dimensions ≥ 3 ?

Furthermore, we note that the purely structural constraint afforded by such work as [21] is sufficient to allow solution of such problems as n -queens or graph k -coloring. In addition, important mathematical questions might well be resolved through neural computation, particularly those concerning the theory of block designs, Ramsey theory, and the characterization of finite projective planes of various orders. We consider such mathematical problems to be the realm where Hopfield networks excel. We note in passing that the dichotomy between structure and cost (similar to that of feasibility/optimality in mathematical programming) resembles the relationship between recognition and enumeration. This too is an area for investigation. Finally, we mention fixed- or bounded-dimension linear and quadratic program-

ming as ideal areas of practical application which have not been fully explored in the literature.

Many problems in CAD or network flows, such as routing through congestions, can be formulated as small-dimension linear programs with (combinatorially) many columns. Is it possible to design a Hopfield network that will perform what is essentially column-generation in, e.g., revised simplex techniques? Similarities in, e.g., facet structure of the associated optimization polytopes suggest that a network which solves TSP can be used for a large class of problems.

ACKNOWLEDGEMENTS

The author would like to thank G. Hueter (UCSD Department of Physics and HNC, Inc.) and R. Hecht-Nielsen (UCSD Department of ECE and HNC, Inc.) for calling attention to this problem, as well as for continued support and encouragement. Professor E. Page, G. Tagliarini and K. Hanrahan, all of Clemson University, provided helpful discussions of their work. The author is supported in part by NSF Grant MIP-8700767 and a grant from the State of California MICRO Program.

REFERENCES

- [1] B. Angeniol, G. de la Croix Vaubois and J. Y. le Texier, "Self-Organizing Feature Maps and the Travelling Salesman Problem", *Neural Networks* 1 (1988), pp. 289-293.
- [2] E. Bonomi and J.-L. Lutton, "The N-City Travelling Salesman Problem: Statistical Mechanics and the Metropolis Algorithm", *SIAM Review* 26 (1984), pp. 551-568.
- [3] R. D. Brandt, Y. Wang, A. J. Laub and S. K. Mitra, "Alternative Networks for Solving the Traveling Salesman Problem and the List-Matching Problem", *Proc. International Conf. on Neural Networks* II, July 1988, pp. 333-340.
- [4] W. I. Clement, R. M. Inigo and E.S. McVey, "Synaptic Strengths for Neural Simulation of the Traveling Salesman Problem", *Proc. Applications of Artificial Intelligence SPIE* 927, 1988, pp. 373-380.
- [5] R. Durbin and D. Willshaw, "An Analogue Approach to the Travelling Salesman Problem Using an Elastic Net Method", *Nature* 326(6114), April 1987, pp. 689-691.
- [6] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Theory of NP-Completeness*, New York, Freeman, 1979.
- [7] R. L. Graham, "An Efficient Algorithm for Determining the Convex Hull of a Finite Planar Set", *Inf. Proc. Letters* 10 (1980), pp. 132-136.
- [8] R. L. Graham and P. M. Winkler, "On Isometric Embeddings of Graphs", *Trans. Amer. Math. Soc.* 288(2) (1985), pp. 527-536.
- [9] S. U. Hegde, J. L. Sweet and W. B. Levy, "Determination of Parameters in a Hopfield/Tank Computational Network", *Proc. International Conf. on Neural Networks* II, July 1988, pp. 291-298.
- [10] F. Hoffmann and K. Kriegel, "Embedding Rectilinear Graphs in Linear Time", *Inf. Proc. Letters* 29 (1988), pp. 75-79.
- [11] J. J. Hopfield and D. W. Tank, "'Neural' Computation of Decisions in Optimization Problems", *Biological Cybernetics* 52, 1985, pp. 141-152.
- [12] G. Hueter, "Solution of the Traveling Salesman Problem With an Adaptive Ring", *Proc. International Conf. on Neural Networks* I, 1988, pp. 85-92.
- [13] R. Jonker and T. Volgenant, "Identification of Non-Optimal Arcs for the Traveling Salesman Problem", *Operations Research Letters* 1(3), July 1982, pp. 85-88.
- [14] A. B. Kahng, *draft*, February 1989.
- [15] E. Lawler, J. K. Lenstra, A. Rinnooy Kan, D. Shmoys (eds.), *The Traveling Salesman Problem*, New York, Prentice-Hall, 1986.
- [16] G. S. Lueker, "Optimization Problems on Graphs with Independent Random Edge Weights", *SIAM J. Computing* 10(2), May 1981, pp. 338-351.
- [17] E. W. Page and G. A. Tagliarini, "Algorithm Development for Neural Networks", *Proc. Symp. on Innovative Science and Technology SPIE* 880, January 1988, pp. 11-19.
- [18] D. D. Sherlekar and J. JaJa, "Layouts of Graphs of Arbitrary Degree", *Proc. 25th Allerton Conference on Combinatorics and Optimization*, 1987, pp. 355-364.
- [19] H. Szu, "Fast TSP Algorithm Based on Binary Neuron Output and Analog Neuron Input Using the Zero-Diagonal Interconnect Matrix and Necessary and Sufficient Constraints of the Permutation Matrix", *Proc. International Conf. on Neural Networks* II, July 1988, pp. 259-266.
- [20] G. A. Tagliarini and K. Hanrahan, *private communication*, December 1988.
- [21] G. A. Tagliarini and E. W. Page, "Solving Constraint Satisfaction Problems with Neural Networks", *Proc. International Conf. on Neural Networks* III, June 1987, pp. 741-747.
- [22] G. A. Tagliarini and E. W. Page, "A Neural-Network Solution to the Concentrator Assignment Problem", *Proc. IEEE Conf. on Neural Information Processing Systems*, November 1987, pp. 775-782.
- [23] D. W. Tank and J. J. Hopfield, "Simple 'Neural' Optimization Networks: An A/D Converter, Signal Decision Circuit, and a Linear Programming Circuit", *IEEE Trans. on Circuits and Systems* CAS-33(5), 1986, pp. 533-541.
- [24] D. E. Van den Bout and T. K. Miller, "A Traveling Salesman Objective Function That Works", *Proc. International Conf. on Neural Networks* II, July 1988, pp. 299-303.

[25] G. V. Wilson and G. S. Pawley, "On the Stability of the TSP Problem Algorithm of Hopfield and Tank", *Biological Cybernetics* 58, 1988, pp. 63-70.




Figure 1. Typical graphs of valid tour frequency and normalized tour length. Both both values are expressed as percentages. Normalized tour length denotes percentage of expected random tourlength for Euclidean 10-city instances. All figures plot statistics for groups of 1000 random instances with city locations uniformly distributed in the unit square. We did not differentiate failure modes as in [9][25].




Figure 2. Performance of NN1 (x) and NN2 (o). [Note: Results in Figures 2-6 are for the best variant found in each of the respective classes. For example, NN1 changes the T_{XIJ} matrix by using a constant inhibitory skewing, while excitatory skewing varies with cume of neighbor rankings, ranging from a factor of 1.0 to 1.7. In all classes, convergence criteria, search for permutation matrix structure at intermediate stages of calculation, etc. were exploited to maximize incidence of valid tours.] All results are for C code executed on a Sun-4 running UNIX.




Figure 3. Performance of CP1 (x) and CP2 (o). Results are fairly good, but clearly this method is not truly dual to the NN approach.




Figure 4. Performance of *CH 1* (x) and *CH 2* (o). Convex hull requires $O(n \log n)$ calculation as in [7]. Results may be unimpressive because skewing is "discontinuous" and is not confined to the distance matrix. Preliminary experiments with reducing the relevant distances d_{ij} show slightly higher valid tour incidence.




Figure 6. Performance of Hopfield-Tank TSP. Parameters follow [11]. As with all the heuristic variants, we see that even very small values of D make results much better than random.




Figure 5. Performance of *SB 1* (x) and *SB 2* (o). Results are significantly different from those of [25]; this may be due to parametrization of the C value by the method described above.

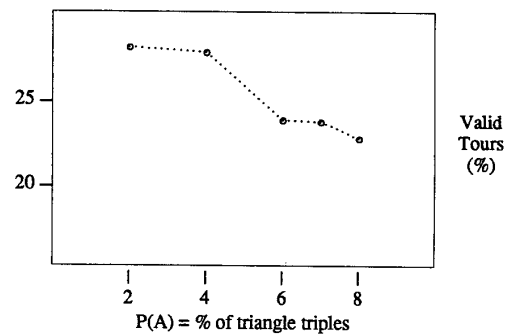


Figure 7. Sample graph of valid tour frequency versus $P(A)$ for Hopfield-Tank variant. Probability of a triple satisfying the triangle inequality was explicitly evaluated for each instance and rounded to the nearest percent; data points therefore do not represent thousands of trials. We ask if there is a better way to measure non-embeddability of the TSP adjacency matrix.