Smoothing Funnels in the TSP with Genetic Algorithms

Doug Hains
dhains@cs.colostate.edu

Adele Howe
howe@cs.colostate.edu

Darrell Whitley
whitley@cs.colostate.edu

Colorado State University
Computer Science Department
Fort Collins, CO

ABSTRACT

The search space of the traveling salesman problem was recently shown to contain clusters of solutions known as funnels which are correlated in distance and evaluation within funnels, but the funnels themselves are separated by an non-trivial distance. These structures trap meta-heuristics such as Chained Lin-Kernighan, one of the best performing approximation algorithms. In this paper, we investigate how a hybrid Genetic Algorithm impacts the funnel structures. We examine two different recombination operators, generalized partition crossover and edge recombination. We find that both operators when used in a hybrid Genetic Algorithm produce populations which smooth out the funnel structure. The diverse population maintained by the hybrid GA locates the global optimum more often than Chained-LK.

Keywords

Traveling Salesman Problem, Genetic Algorithms, Local Search

1. INTRODUCTION

The search space for the Traveling Salesman Problem (TSP) has been shown in [5] to contain clusters of high quality solutions which form pockets or 'funnels' from which some approximation algorithms have a difficulty in escaping. Funnels are clusters of solutions which are correlated in distance and evaluations. Funnels should not be confused with the basins of attraction associated with local optima. In fact, funnels contain many local optima and consequently their basins of attraction. Figure 3 shows some of the local optima in the four funnels found in the ATT532 instance from TSPLIB (http://comopt.ifi.uni-heidelberg.de/software/TSPLIB95).
While the solutions within funnels are correlated in distance and evaluation, the funnels themselves are separated by a non-trivial distance. This makes it difficult for approximation algorithms to escape from funnels.

Modern approximation algorithms for the TSP are quite powerful and employ meta-heuristics, such as iterated local search, to escape local optima which traditionally foiled local search heuristics. Such algorithms have no problem driving the search deep into the search space to find high quality tours. However, funnels are more difficult to escape than local optima and prevent some approximation algorithms from finding the global optimum.

For example, Chained Lin-Kernighan (Chained LK) is one of the best perform approximation algorithms for the TSP [3,6]. When run for sufficiently long on the ATT532 instance from the TSPLIB repository, it finds the global optimum approximately 50% of the time. Of the times it does not find the global optimal, it locates tours with one of only three other evaluations. This is quite remarkable when realizing that for this particular instance, there are 531!/2 possible tours. The tours returned belong to one of four sets of tours which form plateaus found at the bottoms of the four funnels identified in [5] for this instance.

One reason Chained-LK becomes stuck in non-optimal funnels is that the funnels are separated from each other by a non-trivial *bond distance* from one another, where bond distance between two solutions v and w is defined as in [2],

$$d(v,w)=N-s_{v,w}$$

where N is the number of cities and s is the number of shared edges between v and w. The bond distance between the four funnel bottoms of ATT532 ranges from 36 to 86. Many meta-heuristics employ a type of operation known as a perturbation to modify a given solution in order to escape local optima. In the case of Chained-LK, the perturbation operator used is the double bridge move [6]. This operation is depicted in Figure 1 and produces a new tour which is at most a bond distance of 4 from the original tour. Once the search enters a funnel, there is little hope of perturbing a tour with a single double bridge far enough that the search will escape the funnel [5]. Rather, the search will find other local optima in the same funnel, as they are closer in distance, and eventually reach the funnel bottom from which there is little chance to escape.

The occurrence of funnels is clearly a function of the perturbation operator and the local search heuristic. We examine a hybrid genetic algorithm to determine if the genetic algorithm causes a breakdown in the funnel structure without a sacrifice in the efficiency of CLK. We look at two different recombination operators to explore the trade-off in being more disruptive in order to effect a larger perturbation to preserving good edges found by local search.

The hybrid genetic algorithm employs local search as well as selection and recombination. As such, the population will consist of local optima which are input to the recombination operators. In this sense, the recombination operators perform the role of the perturbation operator used in iterated local search to escape local optima. We examine two different recombination operators, edge recombination [8] and generalized partition crossover (GPX) [9,10]. GPX uses only edges found in the parent tours to create offspring. Edge recombination attempts to use edges found in the parent tours, but will also introduce new edges into the offspring. Edge recombination is more disruptive than GPX and will cause a greater perturbation from the original local optima. We find that both operators create a sufficient perturbation to 'smooth' out the funnel structures. However , the disruptiveness of edge recombination prevents it find finding the global optimum as often as GPX, which finds the global optimum more often than Chained-LK.

2. Background

2.1 Chained Lin-Kernighan

Chained Lin-Kernighan is an iterated local search algorithm. The general structure of an iterated local search, given an initial solution s, is in the form of:

1. Let t = s.
2. Perturb t.
3. Apply local search to t.
4. If the evaluation of t is less than the evaluation of s,

let s = t.

5. Goto 1.

The particular implementation of Chained-LK used in this paper is from the latest version of the Concorde software package for the TSP (http://www.tsp.gatech.edu/concorde.html).

In step 2, the random double bridge move is used as the perturbation operator. Given a tour, four edges are randomly selected and replaced as shown in Figure 1.

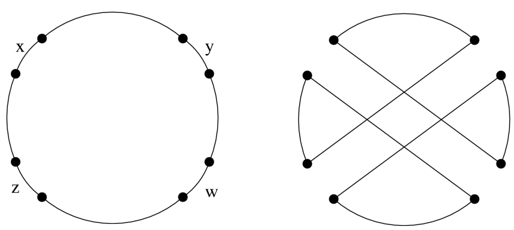

Figure 1: A double bridge move. Edges x, y, w and z are chosen randomly and replaced as shown on the right.

In step 3, Chained-LK uses the Lin-Kernighan local search heuristic [7] (LK-search). LK-search is perhaps the best performing local search heuristic known for the TSP [6]. Though the details are beyond the scope of this paper, it is a depth limited, back-tracking search that attempts to build sequences of edge replacements that result in an improved tour. The same implementation of LK-search is used in Chained LK and in the genetic algorithm described below. For purposes of this paper, it can be treated as a black box which accepts any solution and outputs a locally optimal solution. For more details on the heuristic, see [6,7].

2.2 The Hybrid Genetic Algorithm

In this paper, we use the general definition of a genetic algorithm (GA) as any heuristic algorithm which uses recombination and selection to optimize a population of solutions. A hybrid genetic algorithm additionally applies a local search heuristic at some stage to some or all of the solutions in the population. The general framework of our hybrid genetic algorithm is, given an initial population P of k solutions,

1. Recombine the solution from P which has the best evaluation with the other k-1 solutions in P.
2. Select k solutions from the solutions produced by recombination.
3. Apply local search to the k solutions.
4. Replace P with the k local optima produced in 3.
5. goto 1.

We will use two different recombination operators, edge recombination and GPX, in step 1 and examine the differences between the two in our results.

2.2.1 Edge Recombination

Edge recombination [8] takes two parent solutions and creates a single offspring. It does so by creating a list of n cities and the edges connected to each city from both parent solutions. It then randomly selects the initial city of one of the parent tours and heuristically chooses one of the edges which are connected to the initial city, removing any edges which connect to a visited city from the list. It follows this process and randomly selects an unvisited city when it encounters a city with no edges remaining on its list.

To create two offspring from the same two parent solutions, we create one offspring starting with the initial city from the first parent and the second offspring using the initial city from the second parent.

2.2.2 Generalized Partition Crossover

GPX [10] is a generalized version of partition crossover first presented in [9]. The operator works by taking two parent solutions represented as graphs and union together their edge sets to create a new graph G=(V,E) which contains n vertexes (cities) and the edges from both parents. An edge in E can either be a *common edge*, if the edge is found in both parents, or an *uncommon edge*, if the edge is found only in one of the parents.

All the paths in G of length greater than 1 between arbitrary vertexes v and w are replaced by a single edge from v to w to create a reduced graph G*=(V,E*). A partition on G* is separation of the V into two sets with an empty intersection and the union of which is V. The cost of the partition is the number of edges which are incident to vertexes in both sets. GPX identifies all partitions of cost 2 on G*.

For example, in Figure 2, two solutions to a 14 city TSP instance are represented as Hamiltonian circuits on a graph, one circuit is made of dashed edges the other of solid edges. Two partitions of cost two exist on the graph, labeled A and B.

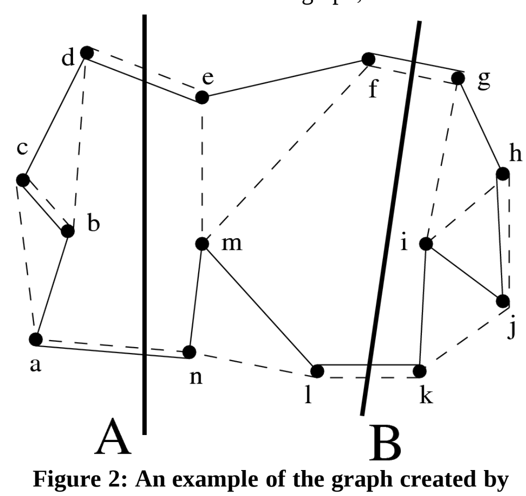

Figure 2: An example of the graph created by unioning two tours for GPX. One solution is represented by dashed lines, the other by solid lines. Two partitions of cost 2 exist depicted by the heavy lines labeled A and B.

It is proven in [9] that by taking the edges from one parent on one side of a partition of cost 2 and the edges from the second parent on the other side of the partition, a valid solution is formed which differ from either parent. In figure 2, the dashed edges could be taken from the left of partition A and the solid edges from the right. By inheriting the common edges across the partition, a solution different from either parents can be created. A second solution can be constructed by likewise taking the solid edges from the left of partition A and the dashed edges from the right.

In [10], it was proven that given g partitions on graph G* g^2-2 solutions can be created which are different than the parents. In our implementation of GPX we create two children from p partitions, one of which is guaranteed to be the best solution of all possible offspring. In the case that no partition of cost 2 exists in G*, the parent tours are perturbed using the random double bridge move.

It is important to note that GPX will create solutions containing only edges found in either of the parent solutions. Edge recombination, on the other hand, attempts to create solutions containing the edges found in either of the parent solutions but can also introduce new edges to the offspring which are found in neither parent solution. In this way, edge recombination is more disruptive than GPX and can create offspring with a greater bond distance from either of the parents.

2.2.3 Selection

In step 2, some method of selecting k solutions must be applied to the 2(k-1) offspring produced by recombination in order to prevent the population from growing exponentially. We developed a selection method known as diversity selection for use in the hybrid genetic algorithm. Solutions containing edges found in fewer members of the population are favored over those containing edges which are more frequently found in the population. The following formula is used to obtain the *diversity measure* for tour t,

$$m(t) = \sum_{e(j,k)} \frac{1}{M_{(j,k)}}$$

where e(j,k) is the edge from vertex j to k in tour t and M(j,k) is the number of times e(j,k) appears in the population. The solution with the minimum tour cost is selected along with the k-1 tours with highest diversity measures.

In step 3, the local search used is the same implementation of LK-search used in Chained-LK which can be found in the Concorde software package.

2.3 Funnels in the TSP

The existence of multiple funnels in the TSP were first reported in [5] when analyzing the search space under Chained-LK. Funnels are clusters of high quality solutions which are correlated in bond distance and evaluation to other solutions within the same funnel, but not to solutions in other funnels. Solutions in different funnels will have the same evaluation but will be separated by a non-trivial bond distance.

To identify funnels, Chained-LK is run on an instance until no improving tour is found for 10,000 consecutive iterations of the iterated local search loop described previously. Using this procedure, we identified four funnels in ATT532 (a pseudo-random instance from the TSPLIB with 532 cities) four in U574, (a PCB drilling problem from the TSPLIB with 574 cities) and two funnels in RAND500 (a Euclidean random instance with 500 cities [5]).

A random walk process was used to find local optima in each funnel in [5]. The process begins by initiating random walk from the solutions found when identifying funnels. At each step of the random walk, steepest descent 2-opt [3,6] is applied until a local optimum was found. If the local optima had the same evaluation as the funnel bottom, the random walk continued from the last step in the walk prior to 2-opt being applied. If the local optima was different from the funnel bottom, the solution was saved and a new random walk was started.

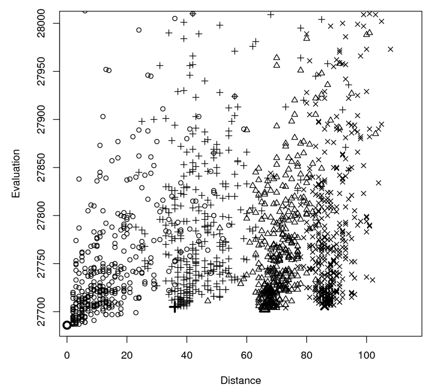

Figure 3: Distance versus evaluation of local optima (2-opt) found in ATT532.

We repeated this procedure until 1,000 local optima were found on the ATT532 instance and plotted the evaluation of the optima versus the distance to the global optimum in Figure 3. Each point represents a local optima and the shape of each point in Figure 3 corresponds to the funnel bottom where the random walk which led to the local optima started. The funnel bottoms are shown with larger and bolder points. The clusters of tours are clear at lower evaluations and tend to spread out as the evaluations grow. What we have found is that at some point in the search, Chained-LK enters a funnel. It is then able to drive the search to the bottom of the funnel but is unable to escape to different funnels.

3. Funnels and Genetic Algorithms

We have previously shown that a hybrid genetic algorithm using the GPX recombination operator is capable of finding higher quality solutions more quickly than Chained-LK when given the same number of LK-search calls and also finds the global optimum more often [10].

We now examine the impact of a hybrid genetic algorithm on the funnel structure. We use the hybrid genetic algorithm outlined above and consider two different algorithms, one using edge recombination in step 1 and the other using GPX. The algorithms are identical with the exception of the recombination operator. Edge recombination attempts to use the edges found in the parent tours but, unlike GPX, the edges in the offspring are not guaranteed to be found in the parents. Edge recombination can inject random edges into the offspring, where GPX will always produce offspring containing only edges found in the parent solutions. Thus, edge recombination may provide more of the 'jump' required to escape funnels than GPX but will also cause more disruption which can prevent a focused search. By examining the results of both operators, we will explore this trade-off between larger perturbations and a focused search.

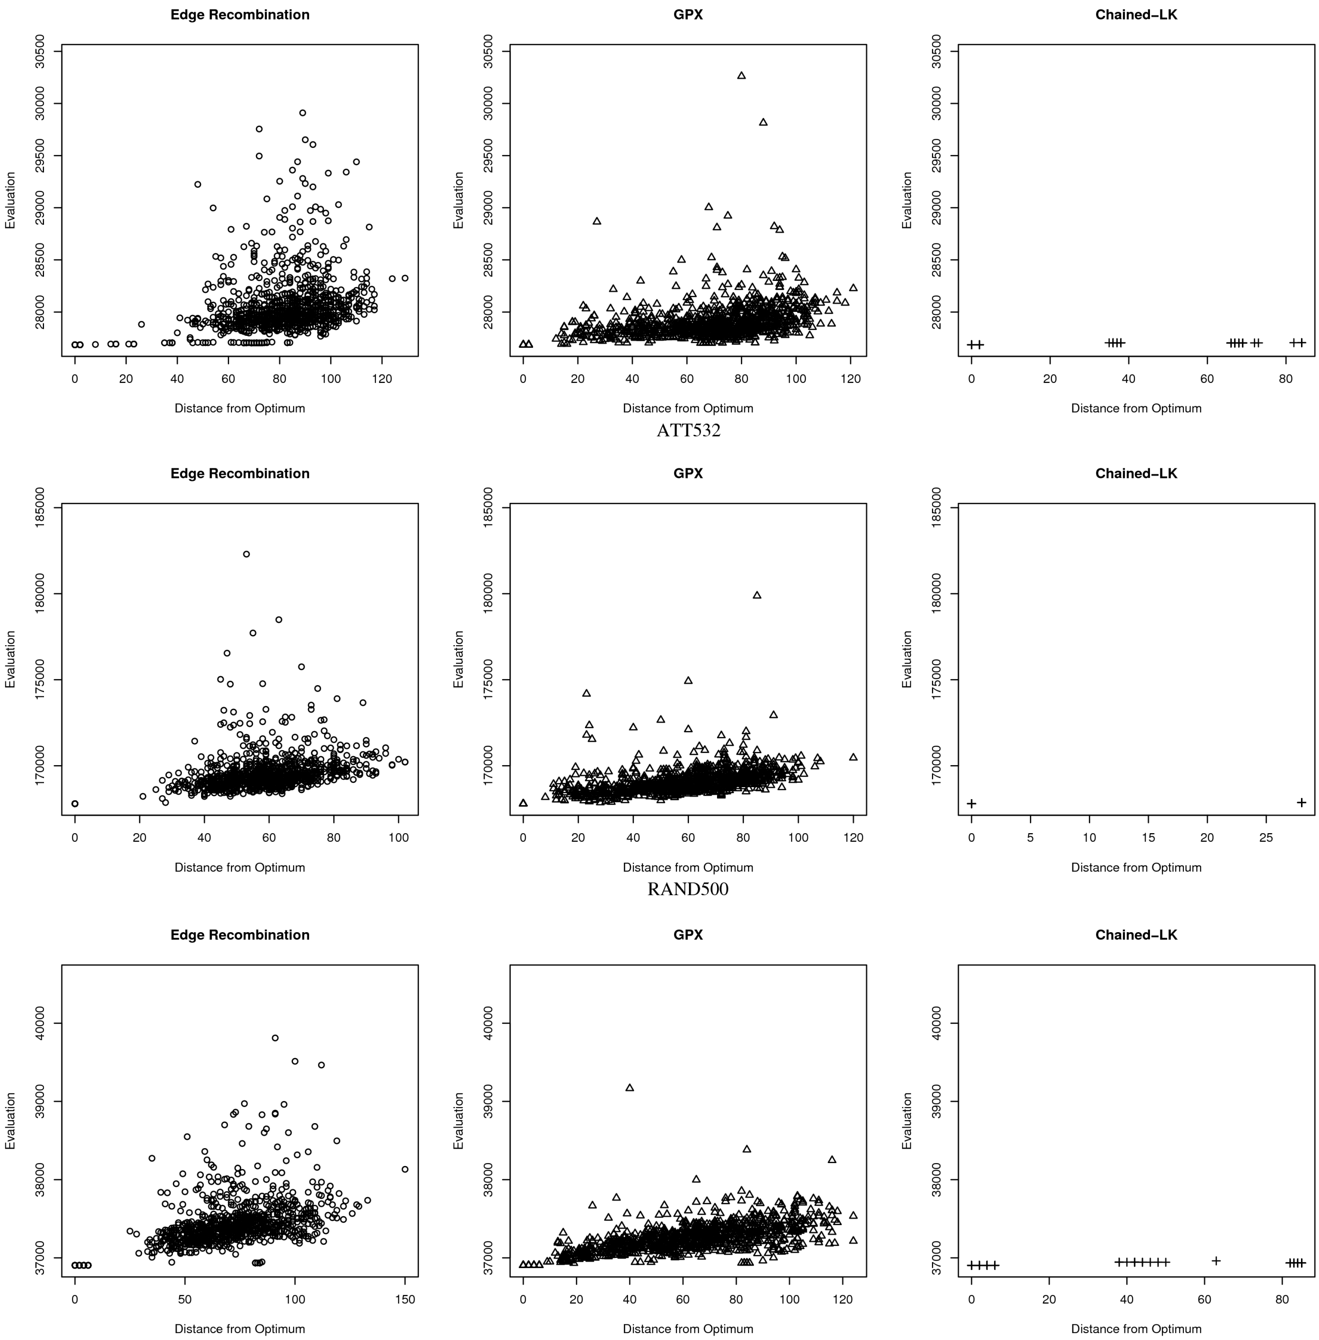

Figure 4: Plots of 1,000 solutions found by a hybrid GA using edge recombination, a hybrid GA using GPX and Chained-LK for instances ATT532 (top row), RAND500 (middle row) and U574 (bottom row). Although each algorithm produced 1,000 tours, Chained-LK produced only a handful of unique tours in the 1,000 trials.

To examine if the genetic algorithms are susceptible to the funnels in the same way as Chained-LK, we ran Chained-LK, the hybrid genetic algorithm using edge recombination and the hybrid genetic algorithm using GPX on ATT532, U574 and RAND500. Each algorithm was allowed to make 10,000 LK-search calls before terminating. The GAs both used a population of 10 solutions, therefore they ran for 1,000 generations. Chained-LK maintains a single tour and ran for 10,000 iterations. We did this so that 1,000 solutions were produced by each algorithm and plotted each tour using its distance from the global optimum versus evaluation in Figure 4.

In the Chained-LK results (right hand column of Figure 4), the 1,000 tours have a small number of unique evaluations and are therefore plotted on top of one another. The ATT532 and U574 results (top and bottom rows of Figure 4) show that the tours fall into four groups which are associated with plateaus on the bottom of each of the four funnels identified for these instances in [5]. In RAND500 (the middle row of Figure 4), only two unique evaluations are produced by Chained-LK and again these solutions correspond to the bottom of the two funnels found in RAND500 [5].

The results of the hybrid GAs show that the majority of tours do not fall into a small handful of evaluations but rather form a single cluster which contains a far greater number of unique tours than the Chained-LK results. The more disruptive nature of edge recombination seems to produce tours with higher evaluations than the hybrid genetic algorithm using GPX. However, both genetic algorithms effectively 'smooth' out the the distinct groups of similar tours produced by Chained-LK.

While it appears from these results that Chained-LK consistently finds lower evaluation solutions than the GAs, this does not imply better results. Table 1 reports the number of times the global optimum was found over 100 trials for each instance. In all cases, the hybrid GA using GPX was able to find the global optimum more often than Chained-LK and in two cases the hybrid GA using edge recombination was able to find the global optimum more often than Chained-LK. Interestingly, both hybrid GAs were able to find the global optimum in all trials for the RAND500 instance. This instance has only two funnels, and inspecting the distribution for Chained LK in Figure 4, we see that these funnels are separated by a bond distance of approximately 30. Both recombination operators must provide enough of a jump that neither become stuck in a non-optimal funnel.

Table 1: Number of Global Optimal solutions found over 100 trials when each method was allowed 1000 LK-search calls.

Instance	Edge recombination	GPX	Chained LK
ATT532	27	61	44
RAND500	100	100	85
U574	79	86	56

We see in Figure 4 that the majority of tours maintained by both genetic algorithms do not separate into funnels but rather form a single, diverse cluster. Because of this diversity, a much wider range of tours are maintained in the population. This enables the recombination operators to escape funnels and locate the global optimum more often than Chained-LK.

We also want to determine whether the final tours in the populations of genetic algorithms have a different distribution in terms of what funnels they belong to. We know that Chained-LK will drive the search to one of a few funnel bottoms for each instance. The tours in the final population of the GAs are already highly optimized and will be used as the initial tours to Chained-LK. We will run Chained-LK on each initial tour until it becomes stuck in a funnel bottom. We will then say the initial tour belongs to the funnel which contains the funnel bottom found by Chained-LK.

The histograms counting the number of tours belonging to each funnel are shown in Figure 5. The counts are sorted from lowest evaluation to highest, with the leftmost bin being the global optimum.

We can see that in ATT532 the distribution is skewed more towards the optimal solution when using GPX. There was no significant difference in the distributions for RAND500 between the three methods. For U574, we see that Chained-LK has a stronger tendency to find tours located in the funnel with a bottom evaluation of 36935 than the other two methods. In all cases, the hybrid genetic algorithm using GPX has a number of tours in the optimal funnel greater than or equal to the other two methods.

The hybrid genetic algorithm using GPX is able to skew the distributions more towards the global optimum than edge recombination. This results in a higher number of globally optimal solutions found when using GPX over edge recombination. This is most likely a result of the disruptive nature of edge recombination which may prevent the search from drilling down to funnel bottoms.

4. Future Work

We see that the hybrid GA using GPX skews the distribution of solutions towards the optimal funnel more so than Chained-LK which results in a larger number of optimal tours found. However, simply running the GA for a number of generations may not take full advantage of the power of a population. Given the diverse populations of low cost solutions created by the hybrid GA, an alternative strategy is to use the tours found by the GA in a method known as Tour merging [3].

Tour merging in the TSP refers to merging together multiple solutions to create a significantly reduced search space than

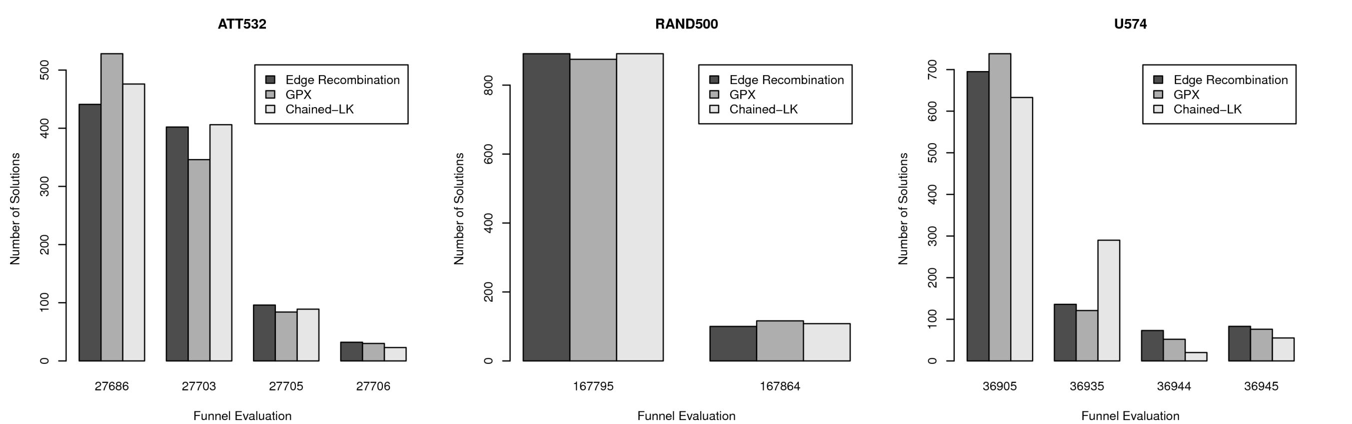

Figure 5: Distributions of tours found by the three algorithms among the identified funnels for each instance.

found in the initial problem. When represented as a graph, a TSP instance is a complete graph on n vertexes. When merging k tours, we union the edges found in each tour onto a single graph containing n cities and at most k*n edges. This creates a much smaller space which needs to be searched and, if all the optimal edges are represented, an exact algorithm working on the reduced problem will find the same optimal solution as would be found when searching the original complete graph.

To determine if we capture the optimal edges in the populations of the hybrid GAs, we looked at the populations after only five generations of each genetic algorithm and counted the number of optimal edges which were found in the population. We repeated this for 100 trials and the average of the counts are shown in Table 2.

Table 2: The mean number of optimal edges found in the population of both hybrid GAs after five generations.

Genetic Algorithm	ATT532	RAND500	U574
GPX	532+/- 0	500+/-0	573.93 +/- 0.1
Edge Recomb.	532+/-0	500+/-0	574 +/- 0

Table 2 shows us that using either of the recombination methods will yield a population containing the optimal edges in only five generations with a high probability (In fact, edge recombination captured the globally optimal edges in all trials). This costs only 60 LK-search calls to produce 10 high quality tours containing the optimal edges. The populations could then be merged to form a reduced problem containing only the edges found in the population and can be passed off to an exact algorithm, such as the dynamic programming algorithm described in [3].

5. Conclusions

We have shown that using a hybrid genetic algorithm can smooth out the funnel structures which are found in Chained-LK. The selection method maintains a diverse population and does not allow the search to converge on a single tour. This allows the recombination operators to escape from funnels that would normally trap a search like Chained-LK.

We explored two different recombination operators and found that GPX performs better than edge recombination. This is likely due to the more disruptive nature of edge recombination because of its ability to introduce new edges into the population. GPX preserves the sub-tours found on either side of the partitions which provide a basis for the recombination. This allows the operator to maintain some of the work done by the local search heuristic while still providing a larger perturbation than the double bridge move.

Our results motivate future work for an alternative to using the the hybrid genetic algorithms to directly attempt to find the global optimum. The algorithms may be best suited as the basis for a tour merging algorithm. We have shown that the optimal edges are found in the populations of the hybrid genetic algorithms using both operators the majority of trials when run for only a small number of generations. This quickly generates a handful of tours which contain the necessary information to find the global optimum by an exact algorithm but with a greatly reduced search space than found in the complete instance.

6. References

[1] Applegate, D., Cook, W., Roe, A.: Chained Lin-Kernighan for Large Traveling Salesman Problems. *INFORMS Journal on Computing 15, 1,* (2003), 82-92.

[2] Boese, K.D., Kahng, A.B., Muddu, S.: A New Adaptive Multi-start Technique for Combinatorial Global Optimizations. *Operations Research Letters, 16,* (1994), 101-113.

[3] Cook, W., Seymour, P.: Tour Merging via Branch Decomposition. *INFORMS Journal on Computing 14, 3,* (2003), 233-248.

[4] Croes, G.: A Method for Solving Traveling-Salesman Problems. *Operations Research*, 1958, 791-812.

[5] Hains, D., Whitley, D., Howe A.: Revisiting the Big Valley Search Space Structure in the TSP. *Journal of Operations Research Society*, (2010).

[6] Johnson, D.S., McGeoch, L.A.: The Traveling Salesman Problem: A Case Study in Local Optimization. *Local Search in Combinatorial Optimization,* Arts, E.H.L., Lenstra, J., eds., John Wiley and Sons Ltd., (1997), 215-310.

[7] Lin, S., Kernighan, B.: An Effective Heuristic for the Traveling-Salesman Problem. *Operations Research,* (1973), 498-516.

[8] Whitley, D., Starkweather, T., Shaner, D.: The Traveling Salesman and Sequence Scheduling: Quality Solutions Using Genetic Edge Recombination. *The Handbook of Genetic Algorithms*, Davis, L., ed., Van Nostrand Reinhold, (1991), 350-372.

[9] Whitley, D., Hains, D., Howe, A.: Tunneling Between Optima: Partition Crossover for the Traveling Salesman Problem. *Proceedings of the 11th Annual Conference on Genetic and Evolutionary Computation,* ACM, (2009), 915-922.

[10] Whitley, D., Hains, D., Howe, A.: A Hybrid Genetic Algorithm for the Traveling Salesman Problem using Generalized Partition Crossover. *Proceedings of the 11th International Conference on Parallel Problem Solving from Nature*, Springer, (2010).