

Dagstuhl Seminar on

The Travelling Salesman Problem

June 23–28, 2002

Organized by:

David S. Johnson (AT&T Labs)
Jan Karel Lenstra (TU Eindhoven)
Gerhard J. Woeginger (TU Twente)

Summary

The Traveling Salesman Problem belongs to the most basic, most important, and most investigated problems in optimization and theoretical computer science: A salesman has to visit each city from a given set exactly once. In doing this, he starts from his home city, and in the very end he has to return again to this home city. He wants to visit the cities in such an order that the total of the distances traveled in his tour becomes as small as possible, since this will save him time and gas. The Traveling Salesman Problem (TSP) consists in identifying this shortest tour through the cities. The TSP has many important applications in vehicle routing, VLSI design, production scheduling, cutting wallpaper, job sequencing, data clustering, curve reconstruction, etc. etc. etc., and research on the TSP has followed many different paths.

The Dagstuhl seminar on *The Travelling Salesman Problem* brought together researchers from Theoretical Computer Science, Operations Research, Mathematical Programming, Discrete Applied Mathematics, and Combinatorics who discussed new developments and new progress made on the TSP during the last 15 years. All in all, there were 42 participants with affiliations in Brazil, Canada, Denmark, France, Germany, Italy, Netherlands, Spain, Switzerland, United Kingdom, USA. In 31 talks the participants presented their latest results on implementational issues, computational challenges, case studies, theoretical questions, polyhedral links, polynomially solvable special cases, domination analysis, exponential neighborhoods, probabilistic analysis, parallel algorithms, online algorithms, approximation algorithms, local search approaches, and many more. The abstracts of these talks can be found in this report.

As a special event, on Tuesday afternoon we were all watching the semi-final match between South-Korea and Germany in the 2002 soccer world championship. Moreover, we moved the traditional hiking tour from Wednesday afternoon to Thursday afternoon, since on Wednesday afternoon we were all watching the other semi-final match between Brazil and Turkey. We enjoyed both soccer matches and also the hike a lot.

Due to the outstanding local organization and the pleasant atmosphere, this seminar was a most enjoyable and memorable event.

David S. Johnson
Jan Karel Lenstra
Gerhard J. Woeginger

Schedule of the week

Monday

9:00–12:00

David S. Johnson: The TSP-challenge

Neil Simonetti: Tour Improvement by Dynamic Programming

Chris Walshaw: Multilevel Refinement for Combinatorial Optimization Problems: A Case Study in the Travelling Salesman Problem

David Neto: Efficient Cluster Compensation for Lin-Kernighan Heuristics

16:00–18:00

Luciana Buriol: A Memetic Algorithm to Asymmetric Traveling Salesman Problem

Juan Jose Salazar: The Travelling Salesman Problem with Pickups and Deliveries

Alexander Barvinok: Distribution of Values in the Travelling Salesman Problem

Tuesday

9:00–12:00

Egon Balas: Polyhedral links between the TSP and related problems

Sándor Fekete: Traveling salesmen in the age of competition

Vladimir Deineko: On solvable TSP cases

13:30–15:30 Semi-final: South-Korea vs. Germany (Room S025)

16:00–18:00

Michelangelo Grigni: The TSP and graph minors

Anders Yeo: Domination analysis for the TSP

Gregory Gutin: When greedy-type heuristics fail

Chris Potts: Some exponential neighborhoods for the TSP

Wednesday**9:00–12:00**

David Applegate: Proofs and Linear Programming

Michael Steele: Some probability of the TSP and MST

Andrea Lodi: Polynomial time separation of simple comb inequalities

Walter Kern: The Traveling Scientist

13:30–15:30 Semi-final: Brazil vs. Turkey (Room S025)**16:00–18:00**

Dennis Naddef: Solving TSP to optimality

Harald Büsching: An augmented version of the LK-algorithm and some remarks about neighbourhood lists

André Rohe: Parallel TSP algorithms

Michelangelo Grigni: TSP and space-filling curves

Thursday**9:00–12:00**

Jan Karel Lenstra: Complexity classification of dial-a-ride problems

Leen Stougie: On-line TSP (I)

Willem de Paepe: On-line TSP (II)

Cor Hurkens: Bad (graphical and Euclidean) TSP instances

Martin Zachariasen: The TSP as a playground for new algorithmic ideas

Afternoon: Hike

Friday

9:00–12:00

George Nemhauser: Sports scheduling and the travelling tournament salesman

Lyle A. McGeoch: ATSP heuristics

Ralf Keuthen: Hybrid Local Search Techniques for the Asymmetric Travelling Salesman Problem

Catherine McGeoch: So you want to do experiments ...

12:00 Concluding remarks

Abstracts

Projection, lifting, and links between the ATSP and related polytopes EGON BALAS

Projection and lifting play an important role in establishing connections between related combinatorial structures. Here we illustrate this through two examples. They are based on joint work with Matteo Fischetti (in the first case) and Maarten Oosten (in the second).

First, we look at the connection between the symmetric and asymmetric travelling salesman polytopes. Any (symmetric or asymmetric) facet inducing inequality for the ATS polytope on the complete digraph with n nodes gives rise to a (symmetric) facet inducing inequality for the STS polytope defined on a certain incomplete undirected graph with $2n$ vertices. The latter inequality can then be lifted into a facet inducing inequality for the STS polytope defined on the complete undirected graph with $2n$ vertices.

Second, we examine the connection between the ATS polytope on a digraph G and the cycle polytope (the convex hull of incidence vectors of directed cycles) of G . Any facet inducing inequality for the ATS polytope on G can be lifted into a facet inducing inequality for the cycle-and-loops polytope (the convex hull of incidence vectors of cycles and loops, where every node not in the cycle has a loop) defined on the digraph with loops G_L , by computing the coefficients of the loop variables. This lifting is largely sequence-independent and computationally inexpensive. Every facet inducing inequality for the cycle-and-loops polytope G_L can then be projected onto the subspace of the arc variables, where it gives rise to a facet inducing inequality for the cycle polytope of G .

The distribution of values in the Traveling Salesman Problem ALEXANDER BARVINOK

The objective function in the Traveling Salesman Problem can be viewed as a real-valued function f on the symmetric group S_n . We introduce the Hamming metric on S_n , so that the distance between two permutations of an n -element set is the number of elements on which the permutations disagree. Let \bar{f} be the average

value of f on S_n and let $f_0 = f - \bar{f}$. Let τ be an optimal permutation where the maximum of f_0 is attained (the minimum is treated in the same way). We show that for the symmetric TSP the average value of f_0 over the Hamming sphere centered at τ steadily increases as the sphere contracts to τ and that for the general asymmetric TSP this “bullseye” type behavior is “diluted” (in a certain rigorous sense). The distribution of values in a general Quadratic Assignment Problem can exhibit a very different behavior (“spikes”).

Joint work with Tamon Stephen.

An augmented version of the LK heuristic and some remarks on neighbourhood lists

HARALD BÜSCHING

The heuristic is more or less classified as a subtour ejection heuristic with refinements for example for doing changes in two different regions of the map in one step. Currently backtracking is used which results in slow performance but quite good quality.

Furthermore perfect neighbourhood lists give the chance to calculate connections which has to be part of real solution and which are definitely not part of it. The calculation can be done locally. They maybe some link of such techniques to the field of cutting plane algorithms which calculate the solution by Linear Programming.

A New Memetic Algorithm to Asymmetric Traveling Salesman Problem

LUCIANA BURIOL

We introduce a new memetic algorithm particularly designed to be effective with asymmetric instances of the traveling salesman problem (ATSP). The method incorporates a new local search engine and many other features that contribute to its effectiveness, such as: i) the topological organization of the population of agents as a complete ternary tree with thirteen nodes; ii) the hierarchical organization of the population in overlapping clusters leading to special selection and reproduction schemes; iii) efficient data structures. Computational experiments are conducted on all ATSP instances from the TSPLIB and ATSP challenge. The comparisons show that the results obtained by our method are competitive comparing with those obtained by other metaheuristics recently proposed for the ATSP.

Joint work with Paulo Morelato Franca and Pablo Moscato.

On polynomially solvable cases

VLADIMIR DEINEKO

We consider the bipartite traveling salesman problem (BTSP) with a symmetric distance matrix. The BTSP is defined as the problem of finding the shortest tour on the set of $n = 2k$ points with k black and k white points that alternates black and white points. The BTSP is NP-hard and, therefore, finding polynomially solvable cases is one of the possible branches of research. We show that the technique developed for the analysis of the TSP can successfully be used to characterise some solvable cases of the BTSP as well. In particular, we present a polynomial time algorithm that decides whether there exists a renumbering of the cities such that the resulting distance matrix allows to write the optimal tour implicitly without further analysis of input data. The results presented generalise some previously published solvable cases of the BTSP known also as the shoelace problem.

Joint work with Gerhard Woeginger.

Traveling Salesmen in the Age of Competition

SÁNDOR FEKETE

We propose the “Competing Salesmen Problem” (CSP), a 2-player competitive version of the classical Traveling Salesman Problem. This problem arises when we are considering two competing salesmen instead of just one. The concern for a shortest tour is replaced by the necessity to reach any of the customers before the opponent does. In particular, we consider the situation where players are taking turns, moving one edge at a time within a graph $G = (V, E)$. The set of customers is given by a subset $V_C \subseteq V$ of the vertices. At any given time, both players know of their opponent’s position. A player wins if he is able to reach a majority of the vertices in V_C before the opponent does.

We give a number of positive results for special cases of the problem, in particular, for the case where G is a bipartite graph. We also point out some of the difficulties: Even if both players start at the same vertex, the starting player may lose, and there may be draws. We also show that the problem is PSPACE-complete, even on bipartite graphs with both players starting at distance two from each other.

Joint work with Aviezri Fraenkel, Rudolf Fleischer, and Matthias Schmitt.

Graph Minors and the TSP

MICHELANGELO GRIGNI

We study approximation schemes for the metric TSP where the metric is shortest paths in an edge-weighted graph. In general this problem is MaxSNP-hard, but $1 + \varepsilon$ approximation schemes are known for certain restricted classes graphs. In increasing generality, we consider planar graphs, bounded genus graphs, and any graph family with a fixed forbidden minor. The running times of these schemes are either a polynomial or quasi-polynomial, with an exponent depending on ε . Their performance depends on finding a light $(1 + \varepsilon)$ -spanner in the given graph. We define a “detour gap” number which is useful in our analysis of spanners, and we present a conjecture which (if true) would greatly improve our known results in the case of a forbidden minor.

Joint work with Papa Sissokho.

When greedy-type heuristics fail

GREGORY GUTIN

We introduce and study greedy-type heuristics for combinatorial optimization (CO) problems. We prove that, for several CO problems, greedy-type algorithms will produce the unique worst possible solution in the worst case.

Joint work with A. Vainstein and A. Yeo.

Bad Instances of the Graphical and Euclidean TSP

COR HURKENS

We describe instances of the graphical and euclidean TSP for which classical heuristics as Farthest Insertion and Nearest Neighbor may yield bad results. A simple graphical instance of the TSP with all cities lying along a cycle can trick Farthest Insertion into delivering a tour of length $2 - \mathcal{O}(\varepsilon^{-1})$ times optimal. More complicated examples can be constructed yielding performance ratios as bad as 6.5. Also for the Euclidean case, Farthest Insertion may yield tours with length more than twice the optimum. A relatively simple example is given with a performance ratio arbitrarily close to 2.43. The final example demonstrates the well-known $\mathcal{O}(\log n)$ worst case ratio for Nearest Neighbor. In contrast with the examples from literature, this class of instances of the graphical TSP is suitable for class room use. The underlying graph is a set of 2^N triangles the bases of which form a cycle. All edges have length one.

Hybrid Local Search Techniques for the Asymmetric Travelling Salesman Problem

RALF KEUTHEN

The talk presents a new class of heuristics which embed an exact algorithm within the framework of a local search heuristic, inspired by a practical problem arising in electronics manufacture. The basic idea of this heuristic is to break the original problem into small subproblems having similar properties to the original problem. These subproblems are then solved using time intensive heuristic approaches or exact algorithms and the solution is re-embedded into the original problem. Here we develop our embedded search heuristic, *HyperOpt*, and use the asymmetric TSP to test its performance in comparison to other local search based approaches. We introduce an interesting hybrid of *HyperOpt* and 3-opt for asymmetric TSPs which proves more efficient than *HyperOpt* or 3-opt alone. We report extensive computational results to investigate the performance of our heuristic approach for single and iterative local searches.

Joint work with Edmund Burke and Peter Cowling.

Computer-Aided Complexity Classification of Dial-a-Ride Problems

JAN KAREL LENSTRA

In dial-a-ride problems, items have to be transported from a source to a destination. The characteristics of the servers involved as well as the specific requirements of the rides may vary. Problems are defined on some metric space, and the goal is to find a feasible solution that minimizes a certain objective function. The structure of these problems allows for a notation similar to the standard notation for scheduling and queueing problems. We introduce such a notation and show how a class of 7,930 dial-a-ride problem types arises from this approach. In examining their computational complexity, we define a partial ordering on the problem class and incorporate it in the computer program *DaRClass*. As input *DaRClass*. uses lists of problems whose complexity is known. The output is a classification of all problems into one of three complexity classes: polynomially solvable, NP-hard, or open. For a selection of the problems that form the input for *DaRClass*, we exhibit a proof of polynomial solvability or NP-hardness. Note that not all of these problems are for real and only exist because they have been generated by a computer program.

Joint work with Willem de Paepe, Jiri Sgall, Rene Sitters, and Leen Stougie.

Polynomial-Time Separation of Simple Comb Inequalities

ANDREA LODI

The *comb* inequalities are a well-known class of facet-inducing inequalities for the Travelling Salesman Problem, defined in terms of certain vertex sets called the *handle* and the *teeth*. We say that a comb inequality is *simple* if the following holds for each tooth: either the intersection of the tooth with the handle has cardinality one, or the part of the tooth outside the handle has cardinality one, or both. The simple comb inequalities generalize the classical *2-matching* inequalities of Edmonds, and also the so-called *Chvátal comb* inequalities.

In 1982, Padberg and Rao gave a polynomial-time algorithm for *separating* the 2-matching inequalities — i.e., for testing if a given fractional solution to an LP relaxation violates a 2-matching inequality. We extend this significantly by giving a polynomial-time algorithm for separating the simple comb inequalities. The key is a result on $\{0, \frac{1}{2}\}$ -cuts due to Caprara and Fischetti.

Joint work with Adam Letchford.

Heuristics for the Asymmetric Traveling Salesman Problem

LYLE A. MCGEOCH

We discuss two heuristics for the Asymmetric Traveling Salesman Problem and report the results of experiments on these and other algorithms. The Kanellakis-Papadimitriou algorithm is a local improvement heuristic that is similar to Lin-Kernighan and uses focussed search to find a set of tour edges that can be replaced to yield a shorter tour. An iterated version of this algorithm can be obtained by applying a random 4-opt move whenever a local optimum is found. Zhang's algorithm works by solving assignment problems and patching the results. It uses truncated branch-and-bound to identify edges to include or exclude in the final tour.

Zhang's algorithm and iterated Kanellakis-Papadimitriou are both effective heuristics, although their relative performance varies widely for different kinds of instances. Our experiments suggest that Zhang's algorithm performs best on instances that have small gap between the assignment problem and Held-Karp lower bounds.

Joint work with Jill Cirasella, David Johnson, and Weixiong Zhang. It was presented at ALENEX 2001.

Solving TSP to optimality

DENIS NADDEF

A few crucial points in the solution of very large instances to optimality are addressed: 1) Primal bounding, 2) Separation and 3) Branching. A new class of facet inducing inequalities for the traveling salesman polytope is also described in the talk.

Sports Scheduling and The Traveling Tournament Problem

GEORGE NEMHAUSER

The *Traveling Tournament Problem* is a sports timetabling problem that abstracts the important issues in creating timetables where team travel is an important issue. Instances of this problem seem to be very difficult to solve even for very small cases. Given the practical importance of solving instances similar to these, this makes this problem an interesting challenge for combinatorial optimization techniques. We introduce this problem, give some interesting classes of instances and give some base computational results.

Joint work with Kelly Easton and Mike Trick.

Efficient cluster compensation for Lin-Kernighan heuristics

DAVID NETO

We describe cluster compensation, a technique for adjusting the Lin-Kernighan heuristic for the Traveling Salesman Problem so that it is less prone to wasting effort in the presence of clustered inputs. We outline the essential features of the Lin-Kernighan heuristic, and argue how those features combine to force the heuristic to waste effort on clustered inputs. We then define cluster compensation, proposing it as a way to help Lin-Kernighan better estimate future costs in its gain function. We indicate how cluster compensation may be performed with low overhead. We provide experimental results over several classes of inputs. They show that cluster compensation has very little impact on the quality of the tours produced. For geometric instances, cluster compensation reduces running time between 23 and 68 percent. For random distance matrix instances, cluster compensation does not reduce running time enough to cover its initialization costs, increasing overall running time by 14 percent on average. Typically, cluster compensation saves time when computing a Minimum Spanning Tree is fast in comparison to the optimization phase of Lin-Kernighan. Furthermore, cluster compensation provides greater benefit for more sharply clustered instances.

Work done under the supervision of Derek Corneil.

On-line dial-a-ride problems under a restricted information model

WILLEM DE PAEPE

In on-line dial-a-ride problems, servers are traveling in some metric space to serve requests for rides which are presented over time. Each ride is characterized by two points in the metric space, a *source*, the starting point of the ride, and a *destination*, the end point of the ride. Usually it is assumed that at the release of such a request complete information about the ride is known. We diverge from this by assuming that at the release of such a ride only information about the source is given. At visiting the source, the information about the destination will be made available to the servers. For many practical problems, our model is closer to reality. However, we feel that the lack of information is often a *choice*, rather than inherent to the problem: additional information *can* be obtained, but this requires investments in information systems. In this paper we give mathematical evidence that for the problem under study it pays to invest.

Joint work with M. Lipmann, X. Lu, R.A. Sitters, L. Stougie.

Exponential Neighborhoods for the Traveling Salesman Problem

CHRIS N. POTTS

Exponential neighborhoods for the traveling salesman that can be searched in polynomial time by dynamic programming are discussed. In particular, the neighborhoods are obtained by disjoint reversals or nested reversals of sections of the tour, or a combination of the two. For different combinations of disjoint reversals and nested reversals, we provide the time complexity of the dynamic program for searching the neighborhood and give the neighborhood size. Included in our classification are the dynasearch neighborhoods that are formed from a combination of disjoint 2-opt or 3-opt moves.

Joint work with Richard Congram.

Parallel TSP algorithms

ANDRE ROHE

We describe several ideas to parallelize TSP algorithms. The main focus is on very large ($\geq 10,000,000$ nodes) problems. The main ideas of the algorithms can be described as follows:

- A theorem (similar to ideas of Karp and Arora) for efficiently computing lower bounds for very large problems in parallel (even on small machines).
- An efficient use of the branch and cut approach in parallel
- A cluster based approach to generate good tours from scratch for large problems.
- Direct Lin-Kernighan parallelization schemes
- A 'window-based' algorithm to postoptimize tours based on any given local opt algorithm.

Solving some Travelling Salesman Problems with Pickups and Deliveries

JUAN JOSE SALAZAR GONZALEZ

Let us consider a depot and a set of customers located in a region. The distance between each pair of locations is considered symmetrical and known in advance. Each customer is associated with a demand which can be positive or negative. If the demand is positive then the customer is named *delivery customer* and when it is negative then the customer is named *pickup customer*. A capacitated vehicle, originally located at the depot, must visit each customer exactly once with the right load of the product to satisfy all the demands, and finally return to the depot. The aim of the problem is to find the route for the vehicle with minimum travel distance.

We present a new mathematical model and develop a branch-and-cut approach using valid inequalities from the well-known Travelling Salesman Problem, Set Packing Problem and Capacitated Vehicle Routing Problem. We also address how to solve other interesting variants of the problem considering two commodities and a full/empty load requirement from the depot. Several computational results on instances from literature shows the good performance of the exact algorithm on instances with up to 100 nodes. We also present a heuristic algorithm to solve instances with up to 500 nodes.

Joint work with Hipolito Hernandez Perez.

Tour Improvement by Dynamic Programming

NEIL SIMONETTI

Consider the following restricted (symmetric or asymmetric) traveling salesman problem (TSP): Given an initial ordering of the cities and a positive integer k , find a minimum cost tour with the constraint that if city i precedes city j by k or more places in the initial ordering, then city i must precede city j in the final ordering.

This restricted TSP can be solved in time linear in n , but exponential in k . This model has proven to be very successful on some time-window constrained TSPs. This algorithm can also be used as an improvement heuristic to solve unrestricted TSPs. It searches an exponential neighborhood $\Omega((k/e)^n)$, has a neighborhood that has a small overlap with interchange heuristics, and has a neighborhood mostly populated with high-quality tours, assuming the initial ordering is a high-quality tour.

Joint work with Egon Balas.

The on-line travelling salesman

LEEN STOUGIE

In the On-Line Travelling Salesman Problem requests for visiting cities (points in a metric space) arrive on-line while a salesman is travelling. The salesman travels at no more than unit speed. The objective is to find a route for the salesman which finishes as early as possible. Two versions of the problem have been defined, differing in the requirement that the salesman has to return to his starting point or not.

Typically, algorithms for on-line optimization problems are subjected to competitive analysis. The competitive ratio is the worst-case ratio between the objective value produced by the algorithm and the optimal objective value. The lack of information at the outset of the problem usually obstructs any algorithm to find the optimal solution value for every instance of the problem, independent of the amount of computation time required by the algorithm. Thus, for almost all on-line optimization problems there are lower bounds on the competitive ratio of any algorithm. We call an algorithm then best possible if its competitive ratio is equal to the lower bound.

Lower bounds can be derived in a two-person game setting, in which one of the players is the algorithm and the other one an adversary giving the instance of the problem knowing the algorithm. In the on-line travelling salesman problem the most obvious adversary is unreasonably potent. We devise more fair adversary models.

In this lecture we give an overview of the results known on competitive analysis in the on-line travelling salesman problem. We also present some open questions concerning this problem.

Multilevel combinatorial refinement: A case study in the TSP

CHRIS WALSHAW

We consider the multilevel paradigm and its potential to aid the solution of combinatorial optimisation problems with specific reference to the TSP. The multilevel paradigm involves recursive coarsening to create a hierarchy of approximations to the original problem. An initial solution is found (sometimes for the original problem, sometimes at the coarsest level) and then iteratively refined at each level, coarsest to finest. As a general solution strategy the multilevel procedure has been in use for many years (most notably in the form of multigrid schemes) and has been applied to many problem areas. However, with the exception of graph partitioning, multilevel techniques have not been widely applied to combinatorial problems. In this talk we address the issue of multilevel refinement for such problems and in particular the TSP and, with the aid of examples and results, make a case for its use as a meta-heuristic. The results provide compelling evidence that, although the multilevel framework cannot be considered as a panacea for combinatorial problems, it can provide a valuable addition to the combinatorial optimisation toolkit. We also give a possible explanation for the underlying process and extract some generic guidelines for its future use.

Domination analysis of ATSP algorithms

ANDERS YEO

The asymmetric travelling salesman problem (ATSP), is the problem of finding a cheapest Hamilton cycle (called a tour), in a complete digraph, with weights on the arcs. The domination number of a tour, in a given instance of the ATSP, is the number of tours which have cost greater than or equal to the tour we are looking at. As there are $(n-1)!$ tours in an instance of the ATSP (with n vertices), a solution with domination number $(n-1)!$ would be optimum. The domination number of an algorithm for the ATSP, is the minimum possible domination number of any solution it could produce. So if the domination number of an ATSP algorithm is $(n-1)!$ then the algorithm always finds a optimum tour. If the domination number is one it may find the unique worst tour.

We present some polynomial algorithms with domination number $(n-2)!$, some of which are already well studied algorithms. We also mention a polynomial algorithm with domination number $(n-2)!/2$, but which is based on some unpublished results, and which is not a practical algorithm. We furthermore present some of the proof-techniques and open problems and conjectures.

TSP as a Playground for New Algorithmic Ideas

MARTIN ZACHARIASEN

The TSP has for decades been the prototype problem in combinatorial optimization. Ideas developed originally for the TSP have been transferred to many other problems with great success. As an example we summarize recent developments in the solution of Steiner trees in the plane. Using LP-based techniques similar to those for the TSP, problem instances with several thousand terminals have been solved to optimality. We conclude the talk with a discussion of the role of the TSP as a playground for new general heuristic algorithms, and point out that special care is needed — both when using TSP as an experimental test case and when judging the outcome of these experiments.

Joint work with David Warme and Pawel Winter.