International Centre for Scientific Culture “E. Majorana”
School of Mathematics “G. Stampacchia”

“GRAPH THEORY, ALGORITHMS AND APPLICATIONS”

ERICE, SEPTEMBER 8-16,2008

PROGRAMME AND ABSTRACTS

TUESDAY 9
9:00am - 12:30 pm
PAOLO TOTH, University of Bologna, IT
Exact Algorithms for the Asymmetric Travelling Salesman Problem
2:30 – 6:30 pm
R.E. TARJAN, Princeton University and HP Labs, Princeton, NJ, US
Graph algorithms

WEDNESDAY 10
9:00am - 12:30 pm
K. MEHLHORN, Max-Planck Institut für Informatik, Saarbrücken, DE
Cycles in Graphs Theory, Algorithms, Applications, and Implementation
2:30 – 6:30 pm
A. GOLDBERG, Microsoft Research, Silicon Valley, CA, US
Efficient Point-to-Point Shortest Path Algorithms WITH Preprocessing

THURSDAY 11
9:00am - 12:30 pm
G.F. ITALIANO, University "Tor Vergata", Rome, IT
Dynamic graph algorithms
2:30 – 6:30 pm
B. GOLDEN, University of Maryland, Baltimore, MD, US
Metaheuristics for the vehicle routing problem and its variants

SATURDAY 13
9:00am - 12:30 pm
P. PARDALOS, University of Florida, Gainesville, FL, US
Cliques, Quasi-Cliques and Clique Partitions in Graphs
2:30 – 6:30 pm
G. RINALDI, CNR, Rome, IT
Max-Cut and extensions: seeking for optimal solutions

SUNDAY 14
9:00am - 12:30 pm
M.G. SPERANZA, University of Brescia, IT
The split delivery vehicle routing problem

TUESDAY 9
9:00am - 12:30 pm

PAOLO TOTH, University of Bologna, IT

Exact Algorithms for the Asymmetric Travelling Salesman Problem

The Travelling Salesman Problem (TSP) is one of the best-known combinatorial optimization and graph theory problems. In this talk, the "asymmetric" version of the problem, called ATSP in the following, is considered. ATSP can be defined as follows. Given a complete directed graph G, where each arc has an associated cost, a Hamiltonian circuit of G is a circuit passing through each vertex of G exactly once. The ATSP is to find a Hamiltonian circuit of G whose total cost (given by the sum of the costs of its arcs) is a minimum. Such a problem is known to be NP-hard in the strong sense, and finds many practical applications (e.g. in vehicle routing, scheduling, wiring, FMS, and sequencing problems). ATSP can be formulated as an Integer Linear Programming Problem.

The talk will review the most effective exact algorithms proposed for finding the optimal solution of ATSP. These algorithms are based on the branch-and-bound and the branch-and-cut approaches.

Several lower bounding procedures for ATSP can be designed, by considering different substructures of the problem, each associated with a subset of constraints defining a well-structured relaxed problem (Min-sum Assignment Problem, rooted Shortest Spanning Arborescence and Anti-Arborescence Problems, …), whose solution gives a valid lower bound for ATSP. Tighter lower bounds can be obtained by means of more sophisticated bounding procedures based on the "restricted Lagrangian approach" or on the "additive approach". Effective branch-and-bound algorithms based on the previously mentioned bounding procedures will be described.

In the branch-and-cut algorithms, at each node of the branch-decision tree a lower bound is obtained by solving a Linear Programming (LP) relaxation of the corresponding ATSP, containing only a subset of the constraints. The LP relaxation is iteratively tightened by adding valid inequalities (cuts) that are violated by the current LP optimal solution. These cuts are identified by solving, for each family of valid inequalities for ATSP, the corresponding separation problem by using exact or heuristic separation procedures. Effective branch-and-cut algorithms based on this approach will be described.

Extensive computational results comparing the performance of the previously described algorithms on random and real-world instances from the literature will be reported.

2:30 – 6:30 pm

R.E. TARJAN, Princeton University and HP Labs, Princeton, NJ, US

Graph algorithms

WEDNESDAY 10
9:00am - 12:30 pm

K. MEHLHORN, Max-Planck Institut für Informatik, Saarbrücken, DE
Cycles in Graphs Theory, Algorithms, Applications, and Implementation

The cycles of a graph convey important information about the graph. A cycle basis concise description of the set of all cycles. It is a family of cycles that spans all cycles of graph. In an undirected graph, a cycle is simply a set of edges with respect to which every vertex has even degree. We view cycles as vectors indexed by edges. The entry for an edge is one the edge belongs to the cycle and is zero otherwise. Addition of cycles corresponds to vector addition modulo 2 (= symmetric difference of the underlying edge sets). In this way, the cycles of a graph form a vector space over the field F2 of two elements and a cycle basis is simply basis of this vector space. In other words, a cycle basis is a minimal family of cycles that every other cycle can be formed as the symmetric difference of some cycles in the basis. In a planar map, the face cycles form a cycle basis. The notion for directed graphs is slightly more involved. The weight of a cycle is either the number of edges in the cycle (in unweighted graphs) or the sum of the weights of the edges in the cycle (in weighted graphs). A minimum cycle basis is basis of total minimum weight. The analysis of the cycle space has applications various fields, e.g., electrical engineering [Kir47], mechanical engineering [CHR76], biology chemistry [Gle01], surface reconstruction [GKM+07], and periodic timetabling [Lie06]. Some of these applications require bases with special properties [LR07].I will discuss some of applications. The first polynomial time algorithms for constructing minimum cycle bases in undirected graphs were discovered by Horton [Hor87] and de Pina [dP95]. Faster realizations of the latter approach are discussed in the papers [BGdV04, KMMP04, MM]. Both approaches can generalized to directed graphs [LR05, KM05, HKM06, Kav05]. The running time of the best algorithm for undirected graphs is O(m2n+mn2logn). For directed graphs, Monte Carlo algorithms achieve essentially the same running time. However, the best deterministic algorithms have a running time that is higher by a multiplicative factor of m. I conjecture that improvements are possible. Since the running times of the exact algorithms for minimum cycle basis are fairly high degree polynomials, approximation algorithms are a sensible alternative. They attempt to construct bases of nearly minimal weight. Approximation algorithms are discussed in [KMM07]. They combine construction techniques for minimum bases with the construction of sparse spanning subgraphs [ADD+93, TZ01]. A cycle basis of a directed graph is integral if every directed cycle of the graph is an integer linear combination of the cycles in the basis and not just a rational linear combination. Integral cycle basis are required for the application to periodic timetabling. The complexity status of finding minimal integral cycle basis is open. Construction and approximation algorithms are described in [Lie03, Lie06, Kav, ELR07].

2:30 – 6:30 pm

A. GOLDBERG, Microsoft Research, Silicon Valley, CA, US
Efficient Point-to-Point Shortest Path Algorithms WITH Preprocessing

We study the point-to-point shortest path problem in a setting where preprocessing is allowed. The two main techniques we address are ALT and REACH. ALT is A* search with lower bounds based on landmark distances and triangle inequality. The REACH approach precomputes locality values on vertices and uses them to prune the search. We also discuss related topics, such as adding shortcuts ("virtual highways") to improve algorithm efficiency.

The resulting algorithms are quite practical for our motivating application, computing driving directions, both for server and for portable device applications. The ideas behind our algorithms are elegant and may have other applications as well.

THURSDAY 11
9:00am - 12:30 pm

G.F. ITALIANO, University "Tor Vergata", Rome, IT

Dynamic graph algorithms

In many applications of graph algorithms, including communication networks, graphics, assembly planning, and VLSI design, graphs are subject to discrete changes, such as additions or deletions of edges or vertices. In the last decades there has been a growing interest in such dynamically changing graphs, and a whole body of algorithms and data structures for dynamic graphs has been discovered. This talk is intended as an overview of this field.

In a typical dynamic graph problem one would like to answer queries on graphs that are undergoing a sequence of updates, for instance, insertions and deletions of edges and vertices. The goal of a dynamic graph algorithm is to update efficiently the solution of a problem after dynamic changes, rather than having to recompute it from scratch each time. Given their powerful versatility, it is not surprising that dynamic algorithms and dynamic data structures are often more difficult to design and analyze than their static counterparts.

2:30 – 6:30 pm

B. GOLDEN, University of Maryland, Baltimore, MD, US

Metaheuristics for the vehicle routing problem and its variants

The Consistent Vehicle Routing Problem

In the small-package shipping industry (as in other industries), companies try to differentiate themselves by providing high levels of customer service. This can be accomplished in several ways including on-line tracking of packages, ensuring on-time delivery, and offering residential pick ups. Some companies want their drivers to develop relationships with customers on a route and would like the same drivers to visit the same customers at roughly the same time on each day that the customers need service. These service requirements together with traditional constraints on vehicle capacity and route length define a variant of the classical capacitated vehicle routing problem (VRP) that we call the Consistent VRP (ConVRP).

In this paper, we formulate the problem as a mixed integer program (MIP) and develop an algorithm to solve the ConVRP that is based on the record-to-record travel algorithm. We compare the performance of our algorithm to the optimal MIP solutions for a set of small problems and then apply our algorithm to five simulated data sets with 1,000 customers and a real-world data set with more than 3,700 customers. The solutions produced by our algorithm on all problem sets do a very good job of meeting customer service objectives with routes that have a low total travel time.

The Balanced Billing Cycle Vehicle Routing Problem

Utility companies typically send their meter readers out each day of the billing cycle in order to determine each customer's usage for the period. Customer churn requires the utility company to periodically remove some customer locations from its meter-reading routes. On the other hand, the addition of new customers and locations requires the utility company to add new stops to the existing routes. A utility that does not adjust its meter-reading routes over time can find itself with inefficient routes and, subsequently, higher meter-reading costs. Furthermore, the utility can end up with certain billing days that require substantially larger meter-reading resources than others. However, remedying this problem is not as simple as it may initially seem. Certain regulatory and customer service considerations can prevent the utility from shifting a customer's billing day by more than a few days in either direction. Thus, the problem of reducing the meter-reading costs and balancing the workload can become quite difficult. We describe this Balanced Billing Cycle Vehicle Routing Problem in more detail and develop an algorithm for providing solutions to a slightly simplified version of the problem. Our algorithm uses a combination of heuristics and integer programming via a three-stage algorithm. We discuss the performance of our procedure on a real-world data set.

The Split Delivery Vehicle Routing Problem: Applications, Algorithms, Test Problems, and Computational Results

In the split delivery vehicle routing problem (SDVRP), a customer's demand can be split among several vehicles. In this article, we review applications of the SDVRP including the routing of helicopters in the North Sea and solution methods such as integer programming and tabu search. We develop a new heuristic that combines a mixed integer program and a record-to-record travel algorithm. Our heuristic produces high-quality solutions to six benchmark problems that have 50-199 customers and generally performs much better than tabu search. On five other problems for which lower bounds exist, our heuristic obtains solutions within 5.85%, on average. Finally, we generate 21 new test problems that have 8-288 customers. A near-optimal solution can be visually estimated for each problem. We apply our heuristic to these new problems and report our computational results. In addition, we discuss several extensions.

SATURDAY 13

9:00am - 12:30 pm

P. PARDALOS, University of Florida, Gainesville, FL, US

Cliques, Quasi-Cliques and Clique Partitions in Graphs

During the last decade, many problems in social, biological, and financial networks require finding cliques, or quasi-cliques.
Cliques or clique partitions have also been used as clustering or classification data mining tools in data sets represented by networks. These networks can be very large and often massive and therefore external (or semi-external) memory algorithms are needed. We discuss applications, algorithms and computational approaches for all types of problems.

2:30 – 6:30 pm

G. RINALDI, CNR, Rome, IT

Max-Cut and extensions: seeking for optimal solutions

Several methods for finding the exact optimal solution of max-cut and of some extensions of this problem will be discussed

SUNDAY 14

9:00am - 12:30 pm

M.G. SPERANZA, University of Brescia, IT

The split delivery vehicle routing problem

In the vehicle routing problem (VRP) the objective is to construct a minimum cost set of routes serving all customers where the demand of each customer is not greater than the vehicle capacity and where each customer is visited exactly once. In the split delivery vehicle routing problem (SDVRP) the restriction that each customer is visited once is removed. Moreover, the demand of each customer can be greater than the capacity of the vehicles. The SDVRP is NP-hard, even under restricted conditions on the costs, when all vehicles have a capacity greater than two, while it is solvable in polynomial time when the vehicles have a maximum capacity of two.
The cost saving that can be obtained by allowing split deliveries can be up to 50% of the cost of the optimal solution of the VRP. The variant of the VRP in which the demand of a customer may be greater than the vehicle capacity, but where each customer has to be visited a minimum number of times, will also be considered. The cost saving that can be obtained by allowing more than the minimum number of required visits can be again up to 50%. Simple heuristics that serve the customers with demands greater than the vehicle capacity by full load out-and-back trips until the demands become less than the vehicle capacity may be quite far from the optimal solution. The VRP and the SDVRP solutions will be also compared through a simulation study.
Three heuristic methods have been proposed for the solution of the SDVRP: The local search by Dror and Trudeau, a simple and effective tabu search algorithm and a sophisticated heuristic that, using the information collected during the tabu search, builds promising routes and solves MILP models to decide which routes to use and how to serve the customers through those routes. The heuristics will be compared on a set of benchmark instances.