DIMACS

Workshop on Network Design: Connectivity & Facility Location

April 28-30, 1997

Organizers:

Professor Panos Pardalos
University of Florida

Professor Ding-Zhu Du
University of Minnesota

Contents

Nearly linear time approximation schemes for Euclidean TSP, Steiner Tree and other geometric problems

Sanjeev Arora
Princeton University
Dept. of Computer Science
Princeton, NJ 08544-2087
arora@CS.Princeton.edu

Abstract

Much effort has been invested in designing approximation algorithms for NP-hard geometric problems such as Euclidean Traveling Salesman and Minimum Steiner Tree.

Last year we presented new algorithms for many of these problems that are approximation schemes: for any fixed $c > 0$, the algorithm computes a solution of cost at most $1 + c$ times the optimal. The running time was $n^{O(1/c)}$ in the plane and slightly superpolynomial in higher dimensions. (Mitchell independently discovered a similar scheme for the planar case.)

This talk will describe a new approximation scheme that is simpler and much more efficient. For any fixed $c > 0$, it finds a $(1 + c)$-approximate solution in $n * \text{polylog}(n)$ time, which is nearly linear in n. The scheme also works in nearly linear time on instances in d dimensions (for fixed d).

Our algorithm generalizes to many other geometric problems, including the famous Minimum Steiner Tree problem, and the Euclidean Matching problem.

Approximating Dense Cases of Covering Problems

Marek Karpinski[1]
Department of Computer Science
University of Bonn
53117 Bonn
marek@cs.uni-bonn.de

and

Alexander Zelikovsky[2]
Computer Science Department
University of Virginia
alexz@cs.virginia.edu

Abstract

We study approximation complexity of dense cases of several covering problems: Set Cover, Steiner Tree and Vertex Cover.

1. Set Cover. Let $X = \{x_1, ..., x_k\}$ be a finite set and $P = \{p_1, ..., p_m\} \subseteq 2^X$ be a family of its subsets. The `Set Cover Problem` (SCP) asks for a minimum size sub-family M of P such that $X \subseteq \cup\{p|p \in M\}$. It is known that SCP cannot be approximated to within less than $\ln k - o(\ln k)$-factor unless $NP \subseteq DTIME[n^{loglogn}]$ and the greedy heuristic gives $1 + lnk$ approximation.

We consider ϵ-*dense* SCP where any element of X belongs to at least $\epsilon|P|$ sets for some $\epsilon < 1$. We show that the size of the greedy solution for ϵ-dense SCP is at most $\log_{1/(1-\epsilon)} k$. Therefore, the size of the optimal set cover is $O(\log k)$ and a simple $O(m^{O(\log k)})$-time exhaustive search chooses the optimal solution. So unless $NP \subseteq DTIME[n^{\log n}]$, the ϵ-dense SCP is not NP-complete.

On the other hand, the approximation complexity of ϵ-dense SCP is strictly less than in the general case. We show that for any $c > 0$ and $1 > \epsilon > 0$, there is a $c \ln k$-approximation algorithm for ϵ-dense SCP. The problem of finding constant factor approximation is still open.

Further densification leads to polynomial-time solvability of SCP. If each element of X is covered by at least $m - o(m^\delta)$ sets of P for some $\delta < 1$, then SCP can be solved exactly.

[1]Research partially supported by the International Computer Science Institute, Berkeley, California, by the DFG Grant KA 67314, by the ESPRIT BR Grants 7097 and EC-US 030

[2]Research partially supported by Volkswagen Stiftung and Packard Foundation

2. Steiner Tree. Given a connected graph $G = (V, E)$ and $S \subseteq V$. The `Steiner Tree Problem` (STP) asks for a minimum size tree within G which spans all terminals from S. A well-known minimum spanning tree heuristic gives 2-approximation for STP and the best up-today polynomial-time approximation guarantee is about 1.644. From the other side, STP is known to be MAX SNP - complete.

In an ϵ-*dense* instance of STP (for some $\epsilon < 1$) any terminal has at least ϵn neighbors outside S. Similarly to dense SCP, unless $NP \subseteq DTIME[n^{\log n}]$, the ϵ-dense STP is not NP-complete.

If the size of S is bounded by a constant, then we can find an exact solution in polynomial time. Otherwise, we apply to the dense STP a variant of Rayward-Smith heuristic with an approximation ratio at most $1+\log_{1/(1-\epsilon)}(|S|-\frac{1}{\epsilon})/(|S|-\frac{1}{\epsilon})$ Thus we show the existence of a polynomial-time approximation scheme (PTAS) for the ϵ-dense STP.

Similarly to SCP, we define δ-*superdense* STP to be the case of STP where any terminal has at least $n - o(n^\delta)$ neighbors outside S. The δ-superdense STP can be solved exactly in polynomial time.

3. Vertex Cover Problem. Given a graph $G = (V, E)$, `Vertex Cover Problem` (VCP) asks for a minimum size vertex set $C \subseteq V$ such that at least one end of any edge belongs to C.

The graph G is ϵ-*dense* if any vertex has at least $\epsilon|V|$ neighbors for some $\epsilon > 0$. We show that as in general case ϵ-dense VCP is MAX SNP-hard. From the other side, the best known approximation algorithm in the general case has performance ratio 2 and we find a better algorithm for ϵ-dense VCP with a performance ratio at most $\frac{2}{1+\epsilon}$.

If the degree of any vertex in G is at least $n - o(n^\delta)$, then VCP has PTAS.

A chunking based genetic algorithm for the Steiner tree problem in graphs

Stefan Voß **Kai Gutenschwager**

Technische Universität Braunschweig
Institut für Wirtschaftswissenschafte
Abt-Jerusalem-Straße 7
D - 38106 Braunschweig, Germany
stefan.voss@tu-bs.den

Abstract

Given an edge-weighted graph the *Steiner tree problem in graphs* (SP) is to determine a minimum cost subgraph spanning a set of specified vertices. In order to achieve this minimum subgraph additional vertices, so-called Steiner vertices, may be included. Real world problems arising in the layout of connection structures in networks as, e.g., in topological network design and locational analysis may often be decomposed into a number of well known optimization problems including SP. According to its complexity one is interested in developing efficient heuristics. In this paper we provide new ideas on a hybrid approach combining genetic algorithms and chunking.

Genetic algorithms are a class of adaptive search procedures based on principles derived from the dynamics of natural population genetics. SP may be represented by a bit string indicating Steiner vertices as either included or excluded from the solution. The fitness of any solution is computed by a minimum spanning tree algorithm restricted to the graph induced by the specified vertices plus all included Steiner vertices.

As a hybrid concept we investigate the combination of genetic algorithms with the concept of chunking. Chunking may be referred to as a basic learning concept used to keep certain properties in mind, but also to structure complex systems. When faced with a complex structure, one strategy is to form groups of solution attributes, i.e., to find higher level relationships in the search space. Chunks, in the context of genetic algorithms, define schemata. Here chunks will not be generated randomly. The general idea of the chunking approach is to speed up the process of evolution by combining chunks of attributes that seem to be superior and to define schemata being responsible for the solution quality. This is done by additionally analyzing each solution to form a few interesting chunks, which will be separated with only a small probability by the crossover operator.

Successful numerical investigations are reported on a set of classical SP benchmark instances from the OR library.

On the Power-p and Bottleneck Steiner Tree Problems

Piotr Berman
Dept. of Computer Science & Engineering
Penn State
`berman@cse.psu.edu`

and

Alexander Zelikovsky[1]
Computer Science Department
University of Virginia.
`alexz@cs.virginia.edu`

Abstract

Many VLSI routing applications, such as routing to minimize estimates of electrical delay as well as the facility location problem involve computation of Steiner trees in which the weight of an edge is a nonlinear function of its length. We consider the following problem. Given a complete graph $G = (V, E, cost)$ with nonnegative edge costs, and a set of *terminls* S, a tree T within G is called a *Steiner tree* if T spanns S and any *Steiner point* of T (i.e. non-terminal vertex) has degree at least 3. The *power-p Steiner tree problem* (p-STP) seeks for a Steiner tree that minimizes the sum of the edge costs each raised to the p power. We also consider the *bottleneck Steiner tree problem* (BSTP) which asks for a Steiner tree that minimizes the maximum cost of an edge.

Although the both problems have been considered in literature, no complexity results have been established. We show that the both problems are NP-hard and, moreover, the bottleneck Steiner trees cannot be approximated with the factor less than 2, unless P = NP.

It is easy to see that, similarly to the standard Steiner tree problem, if Steiner points are not allowed, then the minimum spanning tree is the best Steiner tree for the both p-STP and BSTP. So we study the *Steiner ratio*, i.e. the quality of approximation of optimal Steiner trees with the minimum spanning trees.

We prove the conjecture of Ganley & Salowe (1994) that the Steiner ratio for the p-STP is constant-bounded for any p. In particular, for $p = 2$, we show that the minimum spanning tree is at most 23.3 times worse than the

[1]Research partially supported by Volkswagen Stiftung and Packard Foundation.

optimum. We do not think that this bound is tight, since the worst instance of 2-STP that we construct has the Steiner ratio 17.2.

We also prove that the bottleneck Steiner ratio is no more than $0.5 \log_2 |S|$. It is easy to show that this bound is tight.

Shortest networks on spheres

J. H. Rubinstein
University of Melbourne
Dept. of Mathematics
rubin@maths.mu.OZ.AU

M. Brazil
University of Melbourne
Dept. of Mathematics

D. Thomas
University of Melbourne
Dept. of Electrical Engineering
d.thomas@mullian.ee.mu.oz.au

J. Weng
University of Melbourne
Dept. of Mathematics

and

N. Wormald
University of Melbourne
Dept. of Mathematics

Abstract

We discuss various methods for finding the shortest network or Steiner tree connecting up a given set of points on the surface of a sphere. The edges of such a tree will be arcs of great circles and can be parametrised by their normals. It turns out that there is a set of equations for these normals which has a nice geometric interpretation.

An additional problem with finding such networks is that as well as many topologies for the network, there can be many networks with a given topology, unlike the situation in the plane. We discuss an interesting connection with Morse theory (calculus of variations) which gives some bounds on the growth of such numbers.

Some simple descent methods for finding such networks will also be discussed. A solution of the ratio conjecture on spheres will also be given.

A New Model of Generalised Steiner Networks[1]

J. F. Weng
Department of Electrical Engineering
University of Melbourne
Victoria 3052, Australia
`weng@maths.mu.OZ.AU`

Abstract

Suppose T is a 2-dimensional network with a Steiner topology. If any edge e joining 2 Steiner points s and s' satisfies that the other edges of s, s' are, in pairs, parallelly lying on two sides of e, then T is called a generalised Steiner network. As an example, we show that the drive shafts for an underground mine form a network of this type since the shafts have to have a maximum gradient. In this paper we give an algorithm that can compute generalised Steiner networks with given Steiner topologies in linear time.

Using Hadwiger Numbers in Network Design

Dietmar Cieslik
University of Greifswald
Germany
breier@rz.uni-greifswald.de

Abstract

Minimum spanning trees (MST) are classical examples of combinatorial optimization problems. Several efficient algorithms to construct such trees are known. Solutions of many other problems hinge on the construction of an MST as an intermediary step. Often problems do not look for an ordinary MST but for a special tree or a generalization, see for instance the MST with degree constraints or a Steiner Minimal Tree (SMT). In the most cases now we find the problems as hard in the sense of design and computational complexity. If we look for such trees in metric spaces equiped with a geometry then there are constraints in the combinatorial structure are given which influence the complexity.

Let $M_d(B)$ be the d-dimensional linear space with norm derived from the unit ball B. The Hadwiger number $z_d(B)$ is the maximum number of nonoverlapping translates of B which can be brought into contact with B. This number is an upper bound for the degree of any vertex in a MST, in a SMT and in some other relatives of an MST. We prove this fact and determine the Hadwiger number for some special spaces more exactly. Additionally, we give the best-known bounds for the maximum SMT degree for the two-dimensional spaces.

As a consequence it will be show that

1. The BDMST Problem read as follows: Given a finite set of points in a metric space and an integer $\beta > 1$. Find a tree interconnecting N with minimal length in which each vertex has degree at most β. The problem is polynomially solvable in $M_d(B)$ if $\beta \geq z_d(B)$.
2. We introduce a k-SMT in the way that we allow at most k additional points in the tree interconnecting a finite set of points, where $k \geq 0$ is a given integer. Then such tree can be found by a procedure in polynomially bounded time. The relative defect going from a $(k-1)$-SMT to a k-SMT tends to zero, when k runs to infinity.

The Steiner Tree Problem for Terminals
on the Boundary of a Rectilinear Polygon

Siu-Wing Cheng
Department of Computer Science
HKUST, Clear Water Bay
Hong Kong
scheng@cs.ust.hk

Abstract

We study the problem of computing an optimal Steiner tree interconnecting a set of terminals lying on the boundary of a simple rectilinear polygon. The optimal solution desired is required to lie inside the given polygon. It is shown in [2] that the problem finds application in homotopic routing in VLSI design.

If the solution is required to be a subgraph of a given underlying planar graph G, then the problems becomes the well-known Steiner tree problem for planar graphs. Note that the terminals lie on the exterior face of G. Polynomial-time algorithms have been obtained for general planar graphs. In our problem, G is actually a grid graph. Kaufmann *et al* [2] exploit this fact and the rectilinear nature of the problem to improve the running time to $O(\min\{n^4 \log n, mn^2\})$, where m is the number of vertices in G. The value of m is at least n and m can be very large even when the geometry of the problem is not complicated. The requirement of an input underlying graph is not necessary. Richards and Salowe [3] and Cheng and Tang [1] solve the problem for rectilinear convex polygons in $O(k^4 n)$ and $O(k^2 n)$ time respectively, where k is the number of sides of the polygon. Note that k can range from $O(1)$ to $\Theta(n)$.

In this paper, we present an $O(k^3 n)$-time algorithm for any simple rectilinear polygon. (The time complexity analysis assumes that $k = O(n)$.) Our main contributions are proving some nice structural properties for a selective set of optimal Steiner trees, and developing a dynamic programming algorithm that exploits these properties. The structural properties are extensions of the results in [3] to simple rectilinear polygons. The dynamic programming strategy is an extension as well as a simplification of the result in [1].

Bibliography

[1] S.W. Cheng and C.K. Tang, A fast algorithm for computing optimal rectilinear Steiner tree for extremal point sets. *Proceedings of the Sixth Annual International Symposium on Algorithms and Computation (ISAAC)*, 1995, pp. 322–331.

[2] M. Kaufmann, S. Gao, and K. Thulasiraman, On Steiner minimal trees in grid graphs and its application to VLSI routing. *Proceedings of International Symposium on Algorithms and Computation (ISAAC)*, 1994, pp. 351–359.

[3] D. Richard and J. Salowe, A linear-time algorithm to construct a rectilinear Steiner minimal tree for k-extremal point sets. *Algorithmica*, 7 (1992), pp. 247–276.

Approximation results for the optimum cost chromatic partition problem

Klaus Jansen
Fachbereich IV - Mathematik
Universität Trier
54 286 Trier, Germany
`jansen@dm3.uni-trier.de`

Abstract

In this talk, we give an overview about the optimum cost chromatic partition (OCCP) problem for several graph classes. The OCCP problem is the problem of coloring the vertices of a graph such that adjacent vertices get different colors and that the total coloring costs are minimum. Alternatively, the OCCP problem can be formulated as follows: Given a graph $G = (V, E)$ with n vertices and a sequence of coloring costs $(k_1, \ldots, k_n)$, find a partition into independent sets $U_1, \ldots, U_s$ with $s \leq n$ such that the costs $\sum_{c=1}^{s} k_c \cdot |U_c|$ are minimum.

The OCCP problem introduced by Supowit corresponds to a VLSI layout problem. A net is a group of two terminals (e.g. points on a circle or points on two opposite parallel lines) that need to be electrically connected. A routing segment is a line that connects two terminals of a net. Given a netlist and an associated cost value for each layer, the objective is to partition the net list and place them on different layers in such a way that routing segments do not cross on the same layer and that the total costs are minimum. Therefore, this layout problem is equivalent to the OCCP problem (restricted e.g. to circle graphs or to permutation graphs).

Another application is given by Kroon, Sen, Deng and Roy. The OCCP problem for interval graphs is equivalent to the Fixed Interval Scheduling Problem (FISP) with machine dependent processing costs. In this scheduling problem each job $j \in J$ must be executed during a given time interval (s_j, f_j). We assume that a sufficient number of machines is available and that each job must be executed by one of the machines. If job j is executed by machine m, then the associated processing costs are k_m. The objective is to find a feasible schedule for all jobs with minimum total processing costs.

We prove several approximation results for the OCCP problem restricted to bipartite, chordal, comparability, interval, permutation, split and unimodular graphs. We prove that there exists no polynomial approximation algorithm

with ratio $O(|V|^{0.5-\epsilon})$ for the OCCP problem restricted to bipartite and interval graphs, unless $P = NP$. Furthermore, we propose approximation algorithms with ratio $O(|V|^{0.5})$ for bipartite, interval and unimodular graphs. Finally, we prove that there exists no polynomial approximation algorithm with ratio $O(|V|^{1-\epsilon})$ for the OCCP problem restricted to split, chordal, permutation and comparability graphs, unless $P = NP$.

Constrained Spanning Tree Problems:
Fast Approximate Methods and Parallel Computation

Narsingh Deo *and* **Nishit Kumar**
Department of Computer Science
University of Central Florida
Orlando, FL 32816-2362
{deo, kumar}@cs.ucf.edu

Abstract

In most real-life network optimization situations requiring a minimum spanning tree (MST), it is not enough for the spanning tree to satisfy just the minimum-weight condition. The MST must also satisfy some additional constraint; such as a bound on the degree, or a bound on the diameter of the tree. This additional constraint often makes the problem NP-hard. For instance (in a given connected, edge-weighted, undirected graph), the problem of finding an MST in which the degree of every node is 2 or less is nothing but the problem of computing a shortest Hamiltonian path — an NP-hard problem. Here, we discuss 35 such constrained-spanning-tree problems culled from various sources in the literature.

Since most of these constrained MST problems are NP-hard, good approximate algorithms are needed. It appears that the approximate algorithms for these problems can be designed using two distinct strategies. In one method we first obtain an MST ignoring the constraint, and then iteratively refine this MST by edge-replacements till the constraint is satisfied. In each iteration of this *iterative refinement* heuristic, those edges which violate the constraints are penalized (i.e., their weights are increased) by means of a carefully designed penalty-function to discourage them from appearing in subsequent iterations. In the second method, on the other hand, one constructs a spanning tree once and for all in a modified greedy fashion, employing problem-specific heuristics in the selection of the next edge (or edges) to be added during the tree construction stage itself. We show how both these approaches are used to design simple and fast algorithms, which lend themselves to efficient parallel implementation, specially on massively-parallel SIMD machines.

In this paper, we specifically discuss formulations, parallel implementations, empirical results, and our experience with both methods on three problems: *Degree-Constrained MST*, *Diameter-Constrained MST*, and *Minimum-Length Fundamental Cycle Set*. The iterative refinement method is applied for solving

the first and the second problems, whereas the approach of constructing the spanning tree once-and-for-all is employed for the second and the third problems. Extensive empirical results are presented for large graphs using the 8192-processor, MasPar MP-1 — an SIMD parallel computer. For instance, our parallel program required under half a minute to compute a spanning tree with a degree bound of 3 for a benchmark problem from the TSPLIB (Traveling Salesman Problem Library) on 5934 cities, i.e., a graph with about 18 million edges. The weight of the constrained spanning tree obtained was no worse than 0.23% above that of the optimal solution. We also discuss various implementational issues for the MasPar in computing an unconstrained MST, which must be fast for any effective use of iterative refinement. Also presented here is an average-case analysis, for the convergence of iterative refinement, using the *Degree-Constrained MST* as an example. This analysis is for the space of all $n \times n$ symmetric weight matrices (representing every edge-weighted, undirected graph) and uses a well-known result from the theory of random graphs.

Approximation algorithms for facility location problems

David B. Shmoys
Cornell University
shmoys@cs.cornell.edu

Abstract

We shall present new approximation algorithms for several facility location problems. In each facility location problem that we shall discuss, there is a set of locations at which we may build a facility (such as a warehouse), where the cost of building at location i is f_i; furthermore, there is a set of client locations (such as stores) that require to be serviced by a facility, and if a client at location j is assigned to a facility at location i, a cost of c_{ij} is incurred. The objective is to determine a set of locations at which to open facilities, so as to minimize the total facility and assignment costs. In the uncapacitated case, each facility can service an unlimited number of clients, whereas in the capacitated case, each facility can serve at most u clients.

These models have been the subject of extensive research over the past thirty-five years. It is well-known that the natural integer programming formulations have linear relaxations that yield extremely good lower bounds. Since these problems are NP-hard, it is unlikely that there exist efficient algorithms to find optimal solutions. Instead, we consider algorithms that are guaranteed to find solutions within a given factor of the optimum. We shall consider the case in which the assignment costs are symmetric and satisfy the triangle inequality. For the uncapacitated facility location, we give a polynomial-time algorithm that finds a solution within a factor of 3.16 of the optimal. This is the first constant performance guarantee known for this problem. The algorithm is based on rounding the optimal solution to the linear programming relaxation, and hence we also prove a bound on the quality of this lower bound. We first give a simple analysis, and the show that a randomized variant of this technique (which can then be derandomized) yields an improved performance guarantee. We shall also briefly discuss extensions to the capacitated case.

This is joint work with Éva Tardos (Cornell) and Karen Aardal (Utrecht).

Star, Grid, and Ring Topologies in Facility Location and Network Design

Wu-ji Li
ATT
111 Madison Avenue
Morristown, New Jersey 07960
wujili@attmail.com

and

J. MacGregor Smith
Department of MIE
University of Massachusetts
Amherst Massachusetts 01003
jmsmith@ecs.umass.edu

Abstract

Star, grid, and ring networks and their combinations constitute the fundamental topologies for most manufacturing, telecommunication and many other network design applications. In this paper, we present experimental results using a quadratic assignment algorithm for designing and evaluating these network topologies where random traffic flows occur. Two stochastic models are presented, one based on an $M/G/\infty$ congestion model while the other is based on an $M/G/1/PS$ model.

Massive Dynamic Decomposition Methods for Dynamic Network Flows

Warren B. Powell **Zhi-Long Chen**
Greg Godfrey **Joel Shapiro**
CASTLE Laboratory
Department of Civil Engineering & Operations Research
Princeton University
Princeton, NJ 08544
{powell, zhichen, joel}@dragon.princeton.edu, ggodfrey@princeton.edu

Abstract

We address the problem of finding integer flows that optimize discrete, dynamic resource scheduling problems. We propose a metastrategy that decomposes large problems into potentially thousands of smaller subproblems that are easily solved, producing optimal integer solutions in a small amount of time. We show how these subproblems can be linked in a simple way, producing near-optimal integer solutions for ultra-large scale problems. The approximations are particularly amenable to stochastic problems. We demonstrate the performance of the procedure on a dynamic crew scheduling problem involving over 6,000 drivers covering 35,000 tasks.

Greedy and prohibition-based diversification heuristics for Graph Partitioning

Roberto Battiti and Alan Bertossi
Dipartimento di Matematica
Università di Trento,
38050 Povo (TN), Italy
{battiti,bertossi}@science.unitn.it

Abstract

The balanced k-way partitioning problem, where a division of a graph into k disjoint subsets with minimum cut size is searched for, is paradigmatic in VLSI and network design problems. Successful heuristic approaches start from Kernighan-Lin (KL) variable depth search, and include now Simulated Annealing, Tabu Search (TS), Greedy Randomized procedures, and Genetic Algorithms.

In the paper we present a review of greedy and search procedures, with particular emphasis on the relationship between the KL and TS methods. In addition, two new algorithms are introduced and experimentally evaluated: a very fast and effective greedy procedure and a prohibition-based (TS) technique with state-of-the-art performance.

A Frequency Assignment Problem
in Cellular Phone Networks

Andreas Eisenblaetter *and* **Martin Groetschel**
Konrad-Zuse-Zentrum
Takustraße 7
D-14195 Berlin-Dahlem
{eisenblaetter, groetschel}@zib.de

Abstract

In cellular phone networks, stationary antennas work as base stations and supply access for cellular phones to the network. Frequencies have to be assigned to base station in such a way that the (expected) interference between neighboring base stations is minimized subject to several technical side constraints. The feasibily problem can be regarded as a generalization of list- and T-colorings.

We will present a mathematical model and give a survey of several heuristics which were successfully applied to various large real-world instances. As a starting point for lower bound computations, a 0/1-LP formulation is given and preliminary results are shown.

Dimensioning of survivable telecommunication networks

Roland Wessaely *and* **Martin Groetschel**
Konrad-Zuse-Zentrum
fuer Informationstechnik Berlin
Takustraße 7
D-14195 Berlin, Germany
`wessaely@zib.de`

Abstract

Designing low cost networks that survive certain failure situations belongs to one of the prime tasks in the telecommunications industry. In this talk we describe a mathematical model combining several aspects of survivability that are elsewhere treated in a hierarchical fashion. We present mathematical investigations of this integrated model, a cutting plane algorithm, as well as several heuristics for its solution. Moreover, we report computational results with real world data.

Strong Formulations for Network Design Problems with Connectivity Requirements

S. Raghavan
U. S. WEST Advanced Technologies, Inc.
`raghava@uswest.com`

and

T. L. Magnanti
Massachusetts Institute of Technology
`magnanti@MIT.EDU`

Abstract

The network design problem with connectivity requirements (NDC) is an important combinatorial optimization problem, with many applications in telecommunications, transportation and VLSI design. It generalizes several celebrated combinatorial optimization problems like the minimum spanning tree, the Steiner tree, and the survivable network design problem.

Formally it is stated as follows.

Network Design Problem with Connectivity Requirements: We are given an *undirected* graph $G = (N, E)$, with node set N and edge set E, and a cost vector $\mathbf{c} \in \mathcal{R}_+^{|E|}$ on the edges E. We are also given a symmetric $|N| \times |N|$ matrix $\mathbf{R} = [r_{ij}]$ that prescribes the number of edge–disjoint paths needed between each pair of nodes: r_{ij} represents the required number of edge-disjoint paths between nodes i and j.

We wish to select a set of edges that satisfy these requirements at minimum cost, as measured by the sum of costs of edges we choose.

We consider the unitary NDC, a restriction of the NDC problem where the solution network must be connected, and present a new directed formulation for the unitary NDC problem that is stronger than a natural undirected formulation. We then project out several classes of valid inequalities–partition inequalities, odd-hole inequalities, and combinatorial design inequalities–that generalize known classes of valid inequalities for the Steiner tree problem to the unitary NDC problem.

Next we consider the Steiner forest problem and show how to strengthen its' formulation based on a directizing procedure. We then generalize this approach to the non-unitary version of the NDC problem. To our knowledge, the model we present in this paper is the first directed model for the Steiner forest problem in the literature.

Approximation Algorithms for
Dynamic K-Center Problems

Samir Khuller
Univ of Maryland at College Park
`samir@cs.umd.edu`

Abstract

Facility location problems arise in a variety of contexts. Location of emergency facilities, information servers in a network etc., are all situations where facility location problems arise. One of the basic facility location problem is called the K-center problem, and asks for a placement of K centers in a graph, so that each vertex is guaranteed to have a center close to it. Specifically, we wish to minimize the maximum distance from each vertex to its closest center. This problem is NP-hard, and several optimal approximation algorithms are known for it.

We will consider generalizations of this problem to the case when the underlying metric is not static, but changes dynamically in a predictable way. For example, travel time is a function of traffic and may vary periodically over a day. We would like to choose the facilities, so that each vertex is always close to a center, at any point of time during the day.

(This are joint work with Bhatia, Guha and Sussmann.)

A Case Study of Derandomization Methods for Combinatorial Approximation Algorithms

José D. P. Rolim
rolim@cui.unige.ch
and
Luca Trevisan
trevisan@cui.unige.ch

Abstract

Randomized algorithms are often used to solve or to approximate optimization problems related to network design. Theoretical and practical concern about the effective availability of a source of truly random bits, as well as the desire of improved reliability, motivate the search for deterministic versions of such algorithms.

While there is no general recipe to come up with the derandomization of a randomized algorithm, a small set of techniques seem to underlie most of the known results. In general, the way randomization is used in the algorithm calls for the proper derandomization method. The method *of conditional espectation* (and its generalizations) is suitable when we have a set of 0/1 values $x_1, \ldots, x_n$ that are chosen independently with probability $p_1, \ldots, p_n$, and the objective function is a linear combination of such values. Randomized rounding falls in this case. Another classical use of randomness is to *sample* a small set of points from some larger space. Several mtehods are known to sample using few random bits. In several cases, the number of random bits can be made logarithmic and thus a trivial derandomization is possible. In some cases, randomization is used to construct combinatorial structures via the probabilistic method. A slight change of the algorithm can make the structure simple enough to be constructed in deterministic polynomial time.

The above three scenarios (that are themselves a random sample out of a somewhat larger variety) are explained via some case studies.

Randomized rounding for routing, and packing and covering ILP Randomized rounding was introduced in Raghavan and Thompson (1987), Raghavan (1988). It allows to approximate routing problems as well as problems that can be formulated as integer linear programs in packing or covering form. The approximation guarantee is proved using Chernoff bounds. Derandomization involves the use of *pessimistic estimators*, a generalization of the method of conditional espectation. Srinivasan (1996) has improved Raghavan's bound for set cover, and has provided pessimistic estimators to derandomize his algorithm.

Non-independent randomized rounding for the assignment problem Arora, Frieze and Kaplan (1996) develop a version of randomized rounding for assignment problems (where the solution is a one-one function between two sets of n objects). Applications include the linear arrangment problem in dense graphs. Randomization is used in two phases of their algorithm: an *oversampling* phase and a *merging* phase. Chernoff bounds are used in the proof of correctness. Derandomization involves randomness-efficient sampling and merging.

Multicommodity flow Linial, London and Rabinovich (1994) (and independently Aumann and Rabani (1994)) show that the multicommodity flow problem with k commodities can be approximated within a factor $O(\log k)$. The algorithm uses a mapping of a certain metric (obtained by solving a linear programming relaxation of the problem) into a ℓ_1 metric. Randomization is used to generate the mapping. Linial, London and Rabinovich (1995) have successively found a deterministic way of constructing the mapping (even if it requires more dimensions than the original one). Garg (1995) has independently found a deterministic algorithm for the special case of sparsest cut.

Open questions We remark that for some important network optimization problems, only randomized approximation algorithms are known. One example is the all-terminal network reliability problem (Karger (1995)), a $\#P$-hard counting problem.

A linear algorithm for the "pos/neg-weighted" 1-median problem on a cactus

Rainer E. Burkard
Technische Universität Graz
Institut für Mathematik
Steyrergasse 30
A-8010 Graz
Austria
burkard@opt.math.tu-graz.ac.at

and

Jakob Krarup
DIKU
Department of Computer Science
University of Copenhagen
DK-2100 Copenhagen
Denmark
krarup@diku.dk

Abstract

The 1-median problem on a network asks for a vertex minimizing the sum of the weighted shortest path distances from itself to all other vertices, each associated with a certain positive weight. We allow for *negative* weights as well and devise an exact algorithm for the resulting "pos/neg-weighted" problem defined on a cactus. The algorithm is shown to run in linear time.

Network Parallel Routing and Maximum Partition Matching

Jianer Chen[1]
Department of Computer Science
Texas A&M University
College Station, TX 77843-3112
`chen@cs.tamu.edu`

Abstract

The star network has received considerable attention recently by researchers as a graph model for interconnection network. It has been shown that it is an attractive alternative to the widely used hypercube model. In this talk, we study the problem of parallel routing, i.e., constructing maximum number of node-disjoint shortest paths between two given nodes, on star networks. We first reduce the parallel routing problem on star networks to a combinatorial problem called Partition Matching. An efficient algorithm is derived for the Maximum Partition Matching problem, which induces an optimal algorithm that constructs the maximum number of node-disjoint shortest paths between two given nodes in the star network. We will show how short node-disjoint paths can be constructed efficiently from a node to a set of nodes in the star network. Our results significantly improve previous results for the problems.

This is a joint work with C. C. Chen of Department of Computer Science and Engineering, Tatung Institute of Technology, Taiwan.

[1]Supported in part by the National Science Foundation under Grant CCR-9110824.

Fault-tolerant Routing and Multicasting in Butterfly Networks

Feng Cao **D.Z. Du**
Department of Computer Science
University of Minnesota
Minneapolis, MN 55455
cao@augustus-230.cs.umn.edu du@cs.umn.edu

and

F.D. Hsu
Graduate Program in Computer Science
Fordham University
LL813, 113 West 60th Street
New York, NY 10023
Hsu@murray.fordham.edu

Abstract

In this paper, we study one of important architectures—Butterfly networks. Butterfly networks are widely used in switching networks, local area networks and parellel computing. The connectivity properties and robust routing and multicasting are two research issues in this study. We determine the fault diameters in Butterfly networks, which shows their connectivity properties in the faulty situations. We show how to construct a path between two nodes in the faulty Butterfly networks. We also give polynomial-time algorithms for generating the containers, and the multicasting paths.

Service-Constrained Network Design Problems

R. Ravi
Operations Research
Graduate School of Industrial Administration
Carnegie Mellon University
ravi@andrew.cmu.edu

Abstract

Several practical instances of network design problems often require the network to satisfy multiple constraints. In particular, we consider problems that involve a combination of connectivity and location constraints. They have the following form: find a low-cost network, under one cost function, that *services* every node in the graph, under another cost function, (i.e., every node of the graph is within a prespecified distance from the network). Such problems find applications in optical network design and the efficient maintenance of distributed databases.

We formulate these problems as bicriteria network design problems, and present approximation algorithms for a class of service-constrained network design problems. We also present lower bounds on the approximability of these problems that demonstrate that our performance ratios are close to best possible.

Joint work with Madhav Marathe and Ravi Sundaram

Genetic Algorithms and Multi-Coordination in Network Design

Robert R. Meyer
Computer Sciences Department
University of Wisconsin-Madison
`rrm@lowfat.cs.wisc.edu`

Abstract

We solve structured problems of enormous size by employing a very high-level view that at each iteration utilizes decomposition into subproblems of manageable size followed by a coordination phase that utilizes genetic algorithms followed by heuristic completion and search phases to combine subproblem solutions (which serve as building blocks in a genetic algorithm context) to achieve feasible solutions of the original problem. This procedure is iterated by appropriately modifying subproblems on the basis of optimality conditions for the original problem, and then repeating the subproblem solution and coordination steps. In large-scale optimization, the subproblems and coordination problems may require significant computing time. Thus, on a cluster of powerful workstations, this approach is particularly appropriate because the granularity of these tasks is relatively large compared to the amount of communication required to pass solutions between processors. High-level representations are also advantageous from a genetic algorithms viewpoint in that they allow building blocks to be more easily identified. Finally, in addition to solving subproblems in parallel, the methods that we consider also solve the coordinating problems in parallel (multi-coordination), thereby avoiding the serial bottleneck that a single coordination task would impose. We consider a variety of large-scale models, including multi-commodity flows, fixed-charge models, and network design problems.

A Dual-Ascent Algorithm for
Low-Connectivity Network Design

S. Raghavan
U. S. WEST Advanced Technologies, Inc.
raghava@uswest.com

and

T. L. Magnanti
Massachusetts Institute of Technology
magnanti@MIT.EDU

Abstract

The network design problem with low connectivity requirements (NDLC) arises frequently at local telephone companies as a subproblem in the design of low cost survivable networks. Formally, we can state the NDLC problem as follows. Given a graph $G = (N, E)$ and a connectivity requirement $r_i \in (0, 1, 2)$ for each node, design a network at minimum cost that contains $r_{st} = \min(r_s, r_t)$ edge-disjoint paths between nodes s and t. The typical design procedure (for such networks) is iterative and interactive, and so a fast (approximate) solution method is desired.

In this talk we first describe a (strong) flow-based formulation for the NDLC problem. We then develop a dual-ascent solution procedure that builds upon Wong's algorithm for the Steiner tree problem. This solution method provides both a feasible solution, and a lower bound, thereby providing individual quality guarantees for each solution. We describe a dual-ascent algorithm that is based on a strong (directed) formulation for the problem.

Our computational results indicate that this solution method is effective and robust, rapidly solving problems with up to 300 nodes and 3000 edges to within 4 percent of optimality.

Nonlinear and Mixed-Integer Optimization in Process Network Systems

C. S. Adjiman **C. A. Schweiger**

and

C. A. Floudas
Department of Chemical Engineering
Princeton University
Princeton, NJ 08544
{claire, carl, floudas}@titan.princeton.edu

Abstract

Process Synthesis seeks to systematically develop process flowsheets that convert raw materials into desired products. In the optimization approach to process synthesis, one of the key aspects is the use of a process superstructure which represents all alternative structures of interest. This superstructure is a network of interconnected units that represents the alternative process flowsheets. This paper presents the mathematical modelling of the superstructure which has a mixed set of binary and continuous variables and results in a mixed-integer optimization model. Due to the nonlinearity of chemical models, these problems are generally classified as Mixed-Integer Nonlinear Programming (MINLP) problems.

A number of local optimization algorithms, developed for the solution of this class of problems, are presented in this paper: Generalized Benders Decomposition (GBD), Outer Approximation (OA), Generalized Cross Decomposition (GCD), Branch and Bound (BB), Extended Cutting Plane (ECP), and Feasibility Approach (FA). Some recent developments for the global optimization of nonconvex MINLPs are then introduced. In particular, two branch-and-bound approaches are discussed: the Special structure Mixed Integer Nonlinear αBB (SMIN-αBB), where the binary variables should participate linearly or in mixed-bilinear terms, and the General structure Mixed Integer Nonlinear αBB (GMIN-αBB), where the continuous relaxation of the binary variables must lead to a twice-differentiable problem. Both algorithms are based on the αBB global optimization algorithm for nonconvex continuous problems. Finally, an overview of optimization strategies for problems involving black-box models is given.

Once the theoretical issues behind local and global optimization algorithms for MINLPs have been exposed, attention is directed to their practical use.

The framework MINOPT is discussed as a computational tool for the solution of process synthesis problems. It is an implementation of a number of local optimization algorithms for the solution of MINLPs. In order to quantify the performance of the different methods, a process synthesis example of heat exchanger network synthesis is presented. The formulation as an MINLP and the solution through local and global algorithms are discussed. Results for a variety of process network problems are also presented.

A GRASP For The MultiTarget MultiSensor Tracking Problem

R.A. Murphey
`murphey@eglin.af.mil`

P.M. Pardalos
`pardalos@deming.ise.ufl.edu`

and

L. Pitsoulis
`leonidas@deming.ise.ufl.edu`

Abstract

A Greedy Randomized Adaptive Search Procedure (GRASP) is presented for computing approximate solutions to the MultiTarget MultiSensor Tracking (MTMST) problem. The MTMST problem may be described as follows: Given a large number of sensors, each with some error, reporting a large number of closely spaced objects at discrete intervals in time, determine trajectory estimates for any targets that may be present. At any single interval in time, the MTMST is shown to be a Multi Dimensional Assignment Problem (MAP). Since the objective is to select a target hypothesis and partition of the measurements that is "most likely" to occur, a likelihood cost function and partitioning constraint set are developed. Filtering techniques are established which promise to dramatically decrease the cardinality of the feasible partition space, enabling a GRASP heuristic to be applied effectively. Computational results on test problems will be presented using a recursive algorithm coded in FORTRAN.

Network Improvement Problems

Madhav Marathe
Los Alamos National Laboratory
`madhav@c3serve.c3.lanl.gov`

Abstract

We study *Budget Constrained Optimal Network Upgrading Problems*. Such problems aim at finding optimal strategies for improving a network under some cost measure subject to certain budget constraints. Given a edge weighted graph $G(V, E)$, in the *Edge Based Upgrading Model*, it is assumed that each edge e of the given network has an associated function $c(e)$ that specifies the cost of upgrading the edge by a given amount. A reduction strategy specifies for each edge e the amount by which the length $l(e)$ is to be reduced. In the *Node Based Upgrading Model* a node v can be upgraded at an expense of $cost(v)$. Such an upgrade reduces the cost of each edge incident on v by a fixed factor ϱ, where $0 < \varrho < 1$. For a given budget B, the goal is to find an improvement strategy such that the total cost of reduction is at most the given budget B and the cost of a subgraph (e.g. minimum spanning tree) under the modified edge lengths is the best over all possible strategies which obey the budget constraint.

After giving a brief overview of the various network upgrade models, I will summarize and describe a number of results (hardness as well as easiness) obtained by our group.

This represents joint work with S. O. Krumke, H. Noltemeier, (Department of Computer Science, University of Würzburg), R. Ravi (GSIA, Carnegie Mellon University), and S. S. Ravi (Department of Computer Science, University at Albany).