Heuristiques pour les problèmes de la coloration des sommets d'un graphe et d'affectation de fréquences avec polarités

Nicolas Zufferey, ROSE, IMA, EPFL
(Résumé de la thèse)

Malgré le développement exponentiel de l'informatique et des calculateurs, de nombreux problèmes ne peuvent pas être résolus de manière exacte en un temps de calcul raisonnable. Il en va ainsi des deux problèmes d'affectation sous contraintes étudiés dans cette thèse, à savoir le problème de la coloration des sommets d'un graphe (PCSG) et le problème de l'affectation de fréquences avec polarités (PAFP), connus dans le milieu de l'optimisation combinatoire comme étant très difficiles. Dans ces deux problèmes, étant donné un ensemble d'objets, il s'agit d'attribuer un ou plusieurs nombres entiers à chaque objet en satisfaisant certaines contraintes, le tout en minimisant une fonction objectif. Afin de trouver une solution de qualité satisfaisante en un temps raisonnable à ces problèmes, l'utilisation de méthodes particulières, appelées heuristiques, est inévitable si l'on considère les instances les plus générales. Une heuristique peut se définir comme une méthode d'optimisation qui se contente de trouver en un temps raisonnable une solution de qualité satisfaisante pour un problème donné. Contrairement aux méthodes exactes, une heuristique ne garantit généralement pas l'obtention de la solution optimale. Malgré la relative simplicité des concepts qui caractérisent la plupart des heuristiques, leur adaptation efficace à un problème particulier nécessite un grand effort de modélisation et une bonne connaissance des propriétés du problème à traiter. L'un des objectifs de cette thèse consiste à donner quelques idées générales permettant d'adapter efficacement ce genre de méthodes. Ces principes seront illustrés en étudiant les deux problèmes susmentionnés.

Les heuristiques peuvent se classer en trois catégories. Les heuristiques constructives se contentent de construire pas à pas une seule solution. Elles se caractérisent par une grande rapidité mais leur performance est souvent décevante. Les heuristiques de recherche locale travaillent quant à elles sur une solution qu'elles tentent d'améliorer itérativement. Lors d'une itération, la solution courante est légèrement modifiée afin d'obtenir une solution voisine. Ces algorithmes obtiennent en général de bien meilleurs résultats que les procédures constructives, mais n'ont pas toujours une grande capacité à explorer des régions très différentes de l'espace des solutions. Enfin, les heuristiques évolutives agissent sur une population d'individus (des solutions ou des morceaux de solutions) qui coopèrent et s'adaptent individuellement. Elles ont la plupart du temps un fort potentiel pour trouver des solutions très différentes lors de leur application, mais manquent souvent d'agressivité, car malgré la diversité des solutions rencontrées, celles-ci ne sont pas toujours de grande qualité. Les heuristiques les plus puissantes connues à ce jour consistent à intégrer un algorithme de recherche locale (souvent un algorithme tabou) dans une méthode évolutive, afin de lui donner l'agressivité qui lui fait souvent défaut. Ces méthodes se nomment heuristiques évolutives hybrides.

Dans cette thèse, on expose dans un premier temps une revue de la littérature des heuristiques les plus connues pour chacune des trois familles de méthodes citées ci-dessus, puis on donne un aperçu des adaptations déjà réalisées pour le PCSG. À partir de ces travaux, on va tirer les premières leçons quant à l'adaptation pertinente d'une heuristique à un problème. Ensuite, on propose trois heuristiques très récentes pour résoudre le PCSG : l'algorithme des fourmis, la recherche à voisinages variables et la méthode à mémoire adaptative. La grande efficacité de nos méthodes à mémoire adaptative, qui sont des méthodes évolutives hybrides, renforce d'une part la thèse qui met sur le devant de la scène ce type d'heuristiques, et ouvre d'autre part une nouvelle voie de recherche, étant donné la simplicité des stratégies utilisées et l'absence d'une multitude de paramètres à ajuster. Ces deux points font en effet souvent défaut aux méthodes évolutives usuelles. Enfin, on propose une heuristique générale s'appuyant sur deux algorithmes tabou pour résoudre le PAFP. Ce problème peut être vu comme une extension du PCSG. Notre algorithme a été implémenté dans le cadre d'un concours international d'optimisation proposé par le Centre Électronique de l'Armée (France) par l'intermédiaire de la Société Française de Recherche Opérationnelle. Notre méthode nous a permis d'intégrer le groupe des six finalistes choisis parmi près de trente équipes concurrentes.

Directeurs de thèse : Alain Hertz (alain.hertz@gerad.ca) et Dominique de Werra (dewerra.ima@epfl.ch)

Finding augmenting chains in extensions of claw-free graphs

A. Hertz (GERAD-Ecole polytechnique de Montréal, Canada),
V. Lozin (RUTCOR, Rutgers University, New Jersey, USA),
D. Schindl (ROSE, EPFL, Lausanne)

Résumé de la publication

Soit G un graphe simple et non orienté. Le problème du stable maximum est celui de trouver le plus grand sous-ensemble de sommets de G deux à deux non adjacents. Ce problème est NP-dur dans le cas général (si on ne se restreint pas à une sous-classe de graphes).

Trouver des chaînes augmentantes est la clé de la résolution en temps polynomial du problème de couplage maximum, qui est équivalent au problème du stable maximum dans les graphes de ligne. En 1980, G. J. Minty et N. Sbihi ont trouvé indépendamment une généralisation de ce résultat au cas des graphes ne contenant pas de griffe (graphe constitué d'un sommet central relié à trois sommets non reliés entre eux; en anglais 'claw') comme sous-graphe induit.

Dans la publication présentée ici, un algorithme polynomial permettant de détecter les chaînes augmentantes dans les graphes ne contenant pas de drapeau (graphe formé d'un cycle sur quatre sommets et d'un sommet pendant) ni de 'S12i' (graphe formé d'un sommet duquel partent un chemin de longueur 1, un de longueur 2 et un de longueur i) est présenté. Il est également montré que

ce résultat, combiné avec d'autres résultats sur la recherche de graphes augmentants, permet de conclure que le problème du stable maximum dans la classe (contenant celle des graphes sans griffe) des graphes sans drapeau ni S124 admet un algorithme polynomial, généralisant ainsi le résultat de Minty et Sbihi.

Ce papier de 8 pages a été soumis et récemment accepté pour publication dans Information Processing Letters, journal spécialisé dans les contributions intéressantes en informatique (au sens large) sous forme d'articles courts et concis.

Towards Adaptive Management Systems in Manufacturing: An Agent-Supported Approach

Jens Henoch, IFOR-ETHZ
(Abstract of the Thesis)

Owing to changing customer demands and expectations, a shift from build-to-forecast to build-to-order production for competitive reasons often becomes necessary in manufacturing. Management systems as a vital part of manufacturing systems have to master these new requirements, i.e. management systems have to cope with complexity and uncertainty. The motivation for this study is based on the observation that a rigid centralized approach to manufacturing planning and control is no longer appropriate. Therefore, this dissertation promotes an agent-based approach as a means of supporting operationally human-centered management systems towards an adaptive and flexible conduct, nowadays essential in manufacturing. The core of this thesis presents a modeling and simulation framework integrating concepts of the fields of management cybernetics, production logistics management, artificial intelligence, and object-oriented concurrent programming. As such, the work is at the crossroads of these fields. The framework features three modeling levels. The resulting models can be verified separately utilizing simulation. Physical layout considerations-such as the selection of machines to be utilized-are mapped on the physical level. Operational issues are addressed at the logical level, meaning that instructions are given on how to operate a physical system. Organizational structures and task responsibilities regarding the operations are modeled on the management level. The organizations being mapped on the management level of this framework are multi-agent systems providing means for management systems in their striving to improve their adaptiveness to the situational conditionalities of the shop-floor.

In order to effectively support management systems, multi-agent systems must follow the structural guidelines as proposed by cybernetics. Operationalized, the framework supports recursively structured multi-agent systems, in which the behaviors of planning, scheduling and execution can be

found in every multi-agent system. Thus, a behavior determines an agent's entire action repertoire, denoted also as capabilities.

The advanced level of adaptive conduct as a source for building eigen-variety is supported in a way that agents feature multiple levels of adaptation. Owing to the multi-level adaptation process, robustness can be gained by attenuating uncertainty. In order to maintain the system's cohesion, the adjustment of agents' variety is brought about by the introduction of constraint blocks to be respected by them. The information base of agents is embodied by their knowledge bases and internal models.

Additionally, a distinction has been made between decision-related agents (management agents) and decision-supporting agents (service agents). Multi-agent systems containing management agents are being developed on the management level and are subject to performance assessments regarding their organizational structure. The primary reason for the introduction of the management level is not the analysis of algorithmic competence, but rather to examine its utilization by decision-making agents.

The conceptual development of agents is technically backed up by an implementation based on an active object approach, which grants agents communication and process autonomy. The conceptual separation of the three modeling levels is technically supported as well. However, agents must interfere on the logical level. Proxies as special types of service agents can encapsulate objects on the logical level. As a result, the proxies are able to redirect the information flow to managerial agents, which was originally connecting the objects on the logical level. Thus, these agents are, among other things, able to allocate resources differently than originally planned.

Hedging Strategy and Electricity Contract Engineering

Gustaf Unger, IFOR-ETHZ
(Abstract of the Thesis)

This thesis studies risk management in the electricity market in general and the interaction between physical production and electricity contracts in particular. From a risk management point of view, a power portfolio differs substantially from a traditional financial portfolio. Electricity is non-storable, which together with the marginal production cost characteristics creates jumps in the spot price. The return of a power portfolio is hence typically heavy-tailed, and a risk measure, such as CVaR, that captures this heavy-taildness is needed. To be able to compare production and contracts on a unified basis, we identify the set of contracts that corresponds to each power plant. These contracts build up a replicating portfolio of the power plant. This engineering of contracts allows us to risk manage these often complex contracts, through production. Further, a producing electricity company can through a simple absence of arbitrage argument assess these contracts by studying the costs associated with the

corresponding power plant. Flexible production units, such as a gas turbine, relate to options whereas inflexible units, such as a nuclear plant, relate to futures.

The electricity market is heavily incomplete, why perfect hedges are not achievable for a number of contracts. Hence we introduce the concept of best hedge. The best hedge is found through an optimization, where risk, measured as CVaR, is minimized subject to a constraint on the expected profit. It turns out that this problem can be solved with linear programming, allowing us to handle problems of substantial size.

When a whole portfolio is considered we try to utilize our risk mandate at the best possible way. This leads us to the well-known problem in finance of portfolio optimization. However, this problem needs to be tailored for the electricity market because of the special characteristics of power portfolios. An optimal portfolio implies also an optimal dispatch of the production assets. We focus on the challenging hydro storage plant, which because of its flexible nature corresponds to a series of options. These options are however interdependent through the stored water in the reservoir. An exercise of an option, i. e. production, decreases the amount of stored water and may prohibit production at a later point in time. We develop a dynamic dispatch strategy, which takes this interdependence into account. The optimization of a portfolio consisting of a hydro storage plant and electricity contracts hence needs to derive the optimal portfolio of contracts and the optimal dispatch strategy, or with financial terms the optimal exercise conditions for the corresponding options. We solve the problem with linear programming by maximizing the expected profit over a specified time horizon under the constraint that CVaR of the portfolio may not exceed some threshold, typically determined by the risk preferences of the firm.

It turns out that a simultaneous optimization of the dispatch and the contracts is needed, since the dispatch depends on the volume risk in the entered contracts. A main result is the high value related to the operational flexibility of the hydro storage plant. By studying the dual of our linear portfolio optimization problem, we can actually quantify this value. In a performed case study it is shown that this value of flexibility can be substantial. Any valuation that does not take this operational flexibility into account may hence underestimate flexible power plants.

A Graph Theoretical Approach for Reconstruction and Generation of Oriented Matroids

Lukas Finschi, IFOR-ETHZ
(Abstract of the Thesis)

This thesis studies the reconstruction and generation of oriented matroids. Oriented matroids are a combinatorial abstraction of discrete geometric objects such as point configurations or hyperplane arrangements. Both problems, reconstruction and generation, address fundamental questions of representing

and constructing (classes of) oriented matroids. The representations which are discussed in this thesis are based on graphs that are defined by the oriented matroids, namely tope graphs and cocircuit graphs. The first part of this thesis studies properties of these graphs and the question as to what extent oriented matroids are determined by these graphs. In the second part, these graph representations are used for the design of generation methods which produce complete lists of oriented matroids of given number of elements and given rank. These generation methods are used in the third part for the construction of a catalog of oriented matroids and of complete listings of the combinatorial types of point configurations and hyperplane arrangements.

The reconstruction problem is the problem of whether an oriented matroid can be reconstructed from some representation of it, which is here the tope graph and the cocircuit graph. It is known that tope graphs determine oriented matroids up to isomorphism. However, there is no simple graph theoretical characterization of tope graphs of oriented matroids. We strengthen the known properties of tope graphs and prove that for every element f the topes that are not bounded by f induce a connected subgraph in the tope graph. This property is later used for the design of generation methods that are based on tope graphs.

On the contrary to the tope graph case, it is known that cocircuit graphs do not determine isomorphism classes of oriented matroids. However, if every vertex is labeled by its supporting hyperplane, oriented matroids can be reconstructed up to reorientation. We present a simple algorithm which gives a constructive proof for this result. Furthermore, we extend the known results and show that the isomorphism class of a uniform oriented matroid is determined by its cocircuit graph. In addition, we present polynomial algorithms which provide a constructive proof to this result, and it is shown that the correctness of the input of the algorithms can be verified in polynomial time.

The generation problem asks for methods for listing all oriented matroids of given cardinality of the ground set and given rank. The known generation methods have been designed primarily for uniform oriented matroids in rank 3 or 4. Our methods are based on tope graph and cocircuit graph representations and generate all isomorphism classes of oriented matroids, including non-uniform ones in arbitrary rank. The generation approach incrementally extends oriented matroids by adding single elements. These single element extensions are studied in terms of localizations of graphs, which are signatures on the vertex sets that characterize single element extensions.

The first two generation methods are based on tope graphs. These methods make use of the properties of tope graphs studied earlier in this thesis, especially of the new connectedness property. The first method is a reverse search method for the generation of generalized localizations in the tope graph. In the second method graph automorphisms are used to reduce the amount of isomorphic single element extensions. Furthermore we discuss techniques which reduce multiple extension of the same oriented matroid from different minors.

Two algorithms based on cocircuit graph representations are designed similarly to those based on tope graphs. However, all these first four generation methods lack efficiency, and a reason for this is that they do not use a good characterization of localizations. Due to a result of Las Vergnas, localizations of cocircuit graphs can be characterized by sign patterns on the coline cycles in the cocircuit graph. This allows us to design a fifth method which is efficient in practice. This method is a backtracking algorithm which enumerates all sign patterns of coline cycles that are feasible in terms of the characterization. It turns out that the method is similar to a method of Bokowski and Guedes de Oliveira for the uniform case. Our method is more general as it is capable to handle all oriented matroids in arbitrary rank, including non-uniform oriented matroids. Furthermore it uses an efficient data structure and a new dynamic ordering in the backtrack procedure.

The generation methods are used for the construction of a catalog of oriented matroids. This catalog is organized using basis orientations of oriented matroids. We discuss some properties of the catalog and a method to generate the catalog. The catalog of oriented matroids can be used to find complete listings of combinatorial types of point configurations and hyperplane arrangements. We study these listing problems and discuss solution methods. Furthermore we show by an example the potential of these complete listings in resolving geometric conjectures. The listings of oriented matroids, point configurations, and hyperplane arrangements can be accessed via the Internet on http://www.om.math.ethz.ch.

A Software Framework for Developing Distributed Cooperative Decision Support Systems

Alexandre Gachet, DIUF, Université de Fribourg
(Abstract of the Thesis)

In this thesis, we described the design and implemention of a computer-based solution for a well-defined multiparticipant decision making process evolving in distributed, highly decentralized environments. First, we showed that a software framework would be an appropriate computer-based solution to reach that goal and we explained why software frameworks should be included in the classification of DSS development tools. Next, we described our new model, which aims at turning a DSS into a natural extension of the decision makers' decision making capabilities. We identified two requirements to reach that goal: (1) the DSS should mimic multiparticipant decision making behaviors and (2) the DSS should always be as close as possible to the decision maker, ready to be used anywhere, at any time.

Mimicking multiparticipant decision making behaviors is not a trivial task and finds its justification in the fact that decision makers today seldom decide alone. Decision making today is mostly a multiparticipant process. Consequently, the

shift from single-user DSS to multiparticipant DSS needed to be broken down into two separate tasks: (1) design of a new distributed architecture for distributed DSS and (2) support for cooperation between the various actors of the distributed DSS. In addition, ensuring that a DSS will always be as close as possible to the decision maker, ready to be used anywhere, at any time, raised interesting considerations in a distributed environment. We showed that centralized systems are clearly inappropriate for such situations, and decided to design a new decentralized architecture. We based this distributed, decentralized architecture on a federalist model of cooperation, which allows DSS users to join in dynamic, self-forming, self-managed, and self-healing federations. This model frees the DSS users from many technical constraints.

We chose the Jini distributed technology to implement this federalist model as a software framework. Among other advantages, the use of the Jini technology allows DSS designers and DSS users to stay focused on the DSS problem space, and not on implementation details of the underlying distributed service. As for the requirement to support cooperation between the various actors of the distributed DSS, we extended an existing DSS model and made it suitable for distributed environments. Central to this model is the concept of distributed decision support objects. DDSO are reusable components, whose main purposes are to simplify the task of the decision assistants and decision makers, and to increase their efficiency. The possibility offered to the decision makers to reuse an existing piece of knowledge (for example, the specification of a decision situation, the goal of a decision task, a problem formulation, etc.), possibly in a slightly modified and/or recombined form, represents the real novelty of this approach.

We chose the JavaSpaces service (a core Jini service) to manage DDSO in the federalist model of cooperation mentioned above. The JavaSpaces service is a coordination tool for gluing processes together into a distributed application. It provides a fundamentally different programmating model that views an application as a collection of processes cooperating via the flow of objects into and out of one or more spaces. A space is a shared, network-accessible repository for objects.

The Jini technology provided a distributed infrastructure allowing DSS users to build dynamic, distributed communities. However, we showed that an infrastructure is not enough to deal with the specifics of the DSS architecture defined in this thesis. A new coordination layer was still needed to bring all the pieces of the framework together. To define this new abstraction layer, we introduced a new concept called the virtual twin of the DSS user. From a conceptual point of view, a virtual twin is the alter ego of a DSS user, managing her interests on the network. In that way, coordination processes are managed by a special abstraction layer containing all the virtual twins of the DSS users, which free the user from technical constraints.

This software framework for developing distributed cooperative decision support systems has been used to develop a prototype DSS for the Swiss government. This DSS will be used in the food supply sector.