Thursday 31 March

EvoCOP 1 - Real World Applications, 11:30-13:10

Chaired by Francisco Chicano

Particle Swarm Optimization for Multi-Objective Web Service Location Allocation,

Boxiong Tan, Yi Mei, Hui Ma, Mengjie Zhang

Automated Web service composition, which refers to the creation of a complex application from pre-existing building blocks (Web services), has been an active research topic in the past years. The advantage of having an automated composition system is that it allows users to create new applications simply by providing the required parameters, instead of having to manually assemble the services. Existing approaches to automated composition rely on planning techniques or evolutionary computing (EC) to modify and optimise composition solutions directly in their tree/graph form, a complex process that requires several constraints to be considered before each alteration. To improve the search efficiency and simplify the checking of constraints, this work proposes an indirect Particle Swarm Optimisation (PSO)-based approach. The key idea of the indirect approach is to optimise a service queue which is then decoded into a composition solution by using a planning algorithm. This approach is compared to a previously proposed graph-based direct representation method, and experiment results show that the indirect representation can lead to a greater (or equivalent) quality while requiring a lower execution time. The analysis conducted shows that this is due to the design of the algorithms used for building and evaluating the fitness of solutions.

Evaluating hyperheuristics and local search operators for periodic routing problems,

Yujie Chen, Philip Mourdjis, Fiona Polack, Peter Cowling, Stephen Remde

Meta-heuristics and hybrid heuristic approaches have been successfully applied to Periodic Vehicle Routing Problems (PVRPs). However, to be competitive, these methods require careful design of specific search strategies for each problem. By contrast, hyperheuristics use the performance of low level heuristics to automatically select and tailor search strategies. Hyperheuristics have been successfully applied to problem domains such as timetabling and production scheduling. In this study, we present a comprehensive analysis of hyperheuristic approaches to solving PVRPs. The performance of hyperheuristics is compared to published performance of state-of-the-art meta-heuristics.

A Hybrid Constructive Mat-heuristic Algorithm for the Heterogeneous Vehicle Routing Problem with Simultaneous Pickup and Delivery,

Baris Kececi, Fulya Altiparmak, Imdat Kara

In this paper, a variant of Vehicle Routing Problem, called Heterogeneous Vehicle Routing Problem with Simultaneous Pick-up and Delivery (HVRPSPD), is considered. The HVRPSPD can be defined as determining the routes and vehicle types on each route in such a way that the pickup and delivery demands of each customer must be performed with same vehicle, while minimizing the total cost. We propose a mathematical model for the problem and some valid inequalities for the model. Since the HVRPSPD is an NP-hard problem, the proposed mathematical model can be used to find the optimal solution for the small-size problems. Therefore we propose a hybrid mat-heuristic approach based on the formulation and Local Search to solve medium and large-size HVRPSPDs. A series of experiments is performed to evaluate the performance of proposed algorithm. Computational results show that hybrid mat-heuristic is computationally efficient to find good quality of initial solutions.

Modifying Colourings between Time-steps to Tackle Changes in Dynamic Random Graphs,

Bradley Hardy, Rhyd Lewis, Jonathan Thompson

Many real world operational research problems can be formulated as graph colouring problems. Algorithms for this problem usually operate under the assumption that the size and constraints of a problem are fixed, allowing us to model the problem using a static graph. For many problems however, this is not the case and it would be more appropriate to model such problems using dynamic graphs. In this paper we will explore whether feasible colourings for one graph at time-step *t* can be modified into a colouring for a similar graph at time-step *t+1* in some beneficial manner.

An Evolutionary Approach to the Full Optimization of the Traveling Thief Problem,

Nuno Lourenço, Francisco B. Pereira, Ernesto Costa

Real-World problems usually consist of several different small sub-problems interacting with each other. These interactions promote a relation of interdependence, where the quality of a solution to one sub-problem influences the quality of another partial solution. The Traveling Thief Problem (TTP) is a recent benchmark that results from the combination of the Traveling Salesman Problem (TSP) and the Knapsack Problem (KP). Thus far, existing approaches solve the TTP by fixing one of the components, usually the TSP, and then tackling the KP. We follow in a different direction and propose an Evolutionary Algorithm that addresses both sub-problems at the same time. Experimental results show that solving the TTP as whole creates conditions for

discovering solutions with enhanced quality, and that fixing one of the components might compromise the overall results.

EvoCOP 2 - Theoretical Studies, 14:15-15:55

Chaired by Gabriela Ochoa

Determining the Difficulty of Landscapes by PageRank Centrality in Local Optima Networks,

Sebastian Herrmann

The contribution of this study is twofold: First, we show that we can predict the performance of Iterated Local Search (ILS) in different landscapes with the help of Local Optima Networks (LONs) with escape edges. As a predictor, we use the PageRank Centrality of the global optimum. Escape edges can be extracted with lower effort than the edges used in a previous study. Second, we show that the PageRank vector of a LON can be used to predict the solution quality (average fitness) that can be achieved by ILS in different landscapes.

A Property Preserving Method for Extending a Single-objective Problem Instance to Multiple Objectives with Specific Correlations,

Ruby L. V. Moritz, Enrico Reich, Matthias Bernt, Martin Middendorf

A method is proposed to generate multi-objective optimization problem instances from a corresponding single-objective instance. The user of the method can specify the correlations between the generated the objectives. Different from existing instance generation methods the new method allows to keep certain properties of the original single-objective instance. In particular, we consider optimization problems where the objective is defined by a matrix, e.g., a distance matrix for the Traveling Salesperson problem (TSP) or a flow matrix for the Quadratic Assignment problem. It is shown that the method creates new distance matrices with specific correlations between each other and also have the same average distance and variance of distances as the distance matrix of the original instance. This property is important, e.g., when the influence of correlations between the objectives on the behavior of metaheuristics for the multi-objective TSP are investigated. Some properties of the new method are shown theoretically. In an empirical analysis the new method is compared with instance generation methods from the literature.

Efficient Hill Climber for Multi-Objective Pseudo-Boolean Optimization,

Francisco Chicano, Darrell Whitley, Renato Tinós

Local search algorithms and iterated local search algorithms are a basic technique. Local search can be a stand-alone search method, but it can also be hybridized with

evolutionary algorithms. Recently, it has been shown that it is possible to identify improving moves in Hamming neighborhoods for k-bounded pseudo-Boolean optimization problems in constant time. This means that local search does not need to enumerate neighborhoods to find improving moves. It also means that evolutionary algorithms do not need to use random mutation as a operator, except perhaps as a way to escape local optima. In this paper, we show how improving moves can be identified in constant time for multiobjective problems that are expressed as k-bounded pseudo-Boolean functions. In particular, multiobjective forms of NK Landscapes and Mk Landscapes are considered.

Evolutionary Algorithms for Finding Short Addition Chains: Going the Distance,

Stjepan Picek, Carlos A. Coello Coello, Domagoj Jakobovic, Nele Mentens

The problem of finding the shortest addition chain for a given exponent is of great relevance in cryptography, but is also very difficult to solve since it is an NP-hard problem. In this paper, we propose a genetic algorithm with a novel representation of solutions and new crossover and mutation operators to minimize the length of the addition chains corresponding to a given exponent. We also develop a repair strategy that significantly enhances the performance of our approach. The results are compared with respect to those generated by other metaheuristics for instances of moderate size, but we also investigate values up to $2^{127} - 3$. For those instances, we were unable to find any results produced by other metaheuristics for comparison, and three additional strategies were adopted in this case to serve as benchmarks. Our results indicate that the proposed approach is a very promising alternative to deal with this problem.

Limits to Learning in Reinforcement Learning Hyper-heuristics,

Fawaz Alanazi, Per Kristian Lehre

Learning mechanisms in selection hyper-heuristics are used to identify the most appropriate subset of heuristics when solving a given problem. Several experimental studies have used additive reinforcement learning mechanisms, however, these are inconclusive with regard to the performance of selection hyper-heuristics with these learning mechanisms. This paper points out limitations to learning with additive reinforcement learning mechanisms. Our theoretical results show that if the probability of improving the candidate solution in each point of the search process is less than 1/2 which is a mild assumption, then additive reinforcement learning mechanisms perform asymptotically similar to the simple random mechanism which chooses heuristics uniformly at random. In addition, frequently used adaptation schemes can affect the memory of reinforcement learning mechanisms negatively. We also conducted experiments on two well-known combinatorial optimisation problems, bin-packing and flow-shop, and the obtained results confirm the theoretical findings. This study

suggests that alternatives to the additive updates in reinforcement learning mechanisms should be considered.

EvoCOP 3 - Methodologies, 16:15-17:45

Chaired by Christian Blum

Experimental Evaluation of Two Approaches to Optimal Recombination for Permutation Problems,

Anton V. Eremeev, Julia V. Kovalenko

We consider two approaches to formulation and solving of optimal recombination problems arising as supplementary problems in genetic algorithms for the Asymmetric Travelling Salesman Problem and the Makespan Minimization Problem on a Single Machine. All four optimal recombination problems under consideration are NP-hard but relatively fast exponential-time algorithms are known for solving them. The experimental evaluation carried out in this paper shows that the two approaches to optimal recombination are competitive with each other.

Sim-EDA: A Multipopulation Estimation of Distribution Algorithm Based on Problem Similarity,

Krzysztof Michalak

In this paper a new estimation of distribution algorithm Sim-EDA is presented. This algorithm combines a multipopulation approach with distribution modelling. The proposed approach is to tackle several similar instances of the same optimization problem at once. Each subpopulation is assigned to a different instance and a migration mechanism is used for transferring information between the subpopulations. The migration process can be performed using one of the proposed strategies: two based on similarity between problem instances and one which migrates specimens between subpopulations with uniform probability. Similarity of problem instances is expressed numerically and the value of the similarity function is used for determining how likely a specimen is to migrate between two populations. The Sim-EDA algorithm is a general framework which can be used with various EDAs. The presented algorithm has been tested on several instances of the Max-Cut and TSP problems using three different migration strategies and without migration. The results obtained in the experiments confirm, that the performance of the algorithm is improved when information is transferred between subpopulations assigned to similar instances of the problem. The migration strategy which transfers specimens between the most similar problem instances consistently produces better results than the algorithm without migration.

Solving the Quadratic Assignment Problem with Cooperative Parallel Extremal Optimization,

Danny Munera, Daniel Diaz, Salvador Abreu

Several real-life applications can be stated in terms of the Quadratic Assignment Problem. Finding an optimal assignment is computationally very difficult, for many useful instances. We address this problem using a local search technique, based on Extremal Optimization and present experimental evidence that this approach is competitive. Moreover, cooperative parallel versions of our solver improve performance so much that large and hard instances can be solved quickly.

Construct, Merge, Solve & Adapt: Application to the Repetition-Free Longest Common Subsequence Problem,
Christian Blum, Maria J. Blesa

In this paper we present the application of a recently proposed, general, algorithm for combinatorial optimization to the repetition-free longest common subsequence problem. The applied algorithm, which is labelled CONSTRUCT, MERGE, SOLVE & ADAPT, generates sub-instances based on merging the solution components found in randomly constructed solutions. These sub-instances are subsequently solved by means of an exact solver. Moreover, the considered sub-instances are dynamically changing due to adding new solution components at each iteration, and removing existing solution components on the basis of indicators about their usefulness. The results of applying this algorithm to the repetition-free longest common subsequence problem show that the algorithm generally outperforms competing approaches from the literature. Moreover, they show that the algorithm is competitive with CPLEX for small and medium size problem instances, whereas it outperforms CPLEX for larger problem instances.

Friday 1 April

EvoCOP 4 – Best Paper Candidates, 11:30-13:00

Chaired by Bin Hu

Hyperplane Elimination for Quickly Enumerating Local Optima,
Brian W. Goldman, William F. Punch
Examining the properties of local optima is a common method for understanding combinatorial-problem landscapes. Unfortunately, exhaustive algorithms for finding local optima are limited to very small problem sizes. We propose a method for exploiting problem structure to skip hyperplanes that cannot contain local optima, allowing runtime to scale with the number of local optima instead of with the landscape size. We prove optimality for linear functions and Concatenated Traps, and we provide empirical evidence of optimality on NKq Landscapes and Ising Spin Glasses. We further refine this method to find solutions that cannot be improved by flipping *r* or fewer bits, which counterintuitively can reduce total runtime. While previous methods were limited to landscapes with at most 2^{34} binary strings, hyperplane elimination can enumerate the same problems with 2^{77} binary strings, and find all 4-bit local optima of problems with 2^{200} binary strings.

Particle Swarm Optimisation with Sequence-Like Indirect Representation for Web Service Composition,

Alexandre Sawczuk da Silva, Yi Mei, Hui Ma, Mengjie Zhang

Automated Web service composition, which refers to the creation of a complex application from pre-existing building blocks (Web services), has been an active research topic in the past years. The advantage of having an automated composition system is that it allows users to create new applications simply by providing the required parameters, instead of having to manually assemble the services. Existing approaches to automated composition rely on planning techniques or evolutionary computing (EC) to modify and optimise composition solutions directly in their tree/graph form, a complex process that requires several constraints to be considered before each alteration. To improve the search efficiency and simplify the checking of constraints, this work proposes an indirect Particle Swarm Optimisation (PSO)-based approach. The key idea of the indirect approach is to optimise a service queue which is then decoded into a composition solution by using a planning algorithm. This approach is compared to a previously proposed graph-based direct representation method, and experiment results show that the indirect representation can lead to a greater (or equivalent) quality while requiring a lower execution time. The analysis conducted shows that this is due to the design of the algorithms used for building and evaluating the fitness of solutions.

Deconstructing the Big Valley Search Space Hypothesis,

Gabriela Ochoa, Nadarajen Veerapen

The big valley hypothesis suggests that, in combinatorial optimisation, local optima of good quality are clustered and surround the global optimum. We show here that the idea of a single valley does not always hold. Instead the big valley seems to de-construct into several valleys, also called ‘funnels’ in theoretical chemistry. We use the local optima networks model and propose an effective procedure for extracting the network data. We conduct a detailed study on four selected TSP instances of moderate size and observe that the big valley decomposes into a number of sub-valleys of different sizes and fitness distributions. Sometimes the global optimum is located in the largest valley, which suggests an easy to search landscape, but this is not generally the case. The global optimum might be located in a small valley, which offers a clear and visual explanation of the increased search difficulty in these cases. Our study opens up new possibilities for analysing and visualising combinatorial landscapes as complex networks.