# Lecture Notes in Computer Science 10848

*Commenced Publication in 1973*
Founding and Former Series Editors:
Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

More information about this series at http://www.springer.com/series/7407

Willem-Jan van Hoeve (Ed.)

# Integration of Constraint Programming, Artificial Intelligence, and Operations Research

15th International Conference, CPAIOR 2018
Delft, The Netherlands, June 26–29, 2018
Proceedings

Springer

*Editor*
Willem-Jan van Hoeve
Carnegie Mellon University
Pittsburgh, PA
USA

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-319-93030-5 ISBN 978-3-319-93031-2 (eBook)
https://doi.org/10.1007/978-3-319-93031-2

Library of Congress Control Number: 2018944423

LNCS Sublibrary: SL1 – Theoretical Computer Science and General Issues



Printed on acid-free paper

This Springer imprint is published by the registered company Springer International Publishing AG
part of Springer Nature
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

# Preface

This volume contains the papers that were presented at the 15th International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research (CPAIOR 2018), held in Delft, The Netherlands, June 26–29, 2018. It was co-located with the 28th International Conference on Automated Planning and Scheduling (ICAPS 2018).

The conference received a total of 111 submissions, including 96 regular paper and 15 extended abstract submissions. The regular papers reflect original unpublished work, whereas the extended abstracts contain either original unpublished work or a summary of work that was published elsewhere. Each regular paper was reviewed by at least three Program Committee members, which was followed by an author response period and a general discussion by the Program Committee. The extended abstracts were reviewed for appropriateness for the conference. At the end of the reviewing period, 47 regular papers were accepted for presentation during the conference and publication in this volume, and nine abstracts were accepted for presentation at the conference. Three papers were published directly in the journal *Constraints* via a fast-track review process. The abstracts of these papers can be found in this volume. The EasyChair system was used to handle the submissions, reviews, discussion, and proceedings preparation.

In addition to the regular papers and extended abstracts, three invited talks were given, by Michela Milano (University of Bologna; joint invited talk with ICAPS), Thorsten Koch (Zuse Institute Berlin and Technische Universität Berlin), and Paul Shaw (IBM). The abstracts of the invited talks can also be found in this volume.

The conference program included a Master Class on the topic “Data Science Meets Combinatorial Optimization,” with the following invited talks:

- Siegfried Nijssen (Université catholique de Louvain): Introduction to Machine Learning and Data Mining
- Tias Guns (Vrije Universiteit Brussel): Data Mining Using Constraint Programming
- Kate Smith-Miles (University of Melbourne): Instance Spaces for Objective Assessment of Algorithms and Benchmark Test Suites
- Bistra Dilkina (University of Southern California): Machine Learning for Branch and Bound
- Elias Khalil (Georgia Institute of Technology): Learning Combinatorial Optimization Algorithms over Graphs
- Barry O’Sullivan (University College Cork): Recent Applications of Data Science in Optimization and Constraint Programming

The organization of this conference would not have been possible without the help of many individuals. First, I would like to thank the Program Committee members and external reviewers for their hard work. Several Program Committee members deserve additional thanks because of their help with timely reviewing of fast-track papers, shepherding regular papers, or overseeing the discussion of papers for which I had a conflict of interest. I am also particularly thankful to David Bergman (Master Class Chair), Bistra Dilkina (Publicity Chair), and Joris Kinable (Sponsorship Chair) for their help in organizing this conference. Special thanks is reserved for the conference chair, Mathijs de Weerdt, who also acted as the liaison with Delft University and the organization of ICAPS. His support was instrumental in making this event a success.

Lastly, I want to thank all sponsors for their generous contributions. At the time of writing, these include: the *Artificial Intelligence* journal, Decision Brain, SAS, Springer, Delft University of Technology, the Association for Constraint Programming (ACP), AIMMS, Gurobi, GAMS, Pipple, the European Association for Artificial Intelligence (EurAI), and Cosling.

April 2018 Willem-Jan van Hoeve

# Organization

## Program Chair

| | |
|---|---|
| Willem-Jan van Hoeve | Carnegie Mellon University, USA |

## Conference Chairs

| | |
|---|---|
| Mathijs de Weerdt | Delft University of Technology, The Netherlands |
| Willem-Jan van Hoeve | Carnegie Mellon University, USA |

## Master Class Chair

| | |
|---|---|
| David Bergman | University of Connecticut, USA |

## Publicity Chair

| | |
|---|---|
| Bistra Dilkina | University of Southern California, USA |

## Sponsorship Chair

| | |
|---|---|
| Joris Kinable | Eindhoven University of Technology, The Netherlands |

## Program Committee

| | |
|---|---|
| Tobias Achterberg | Gurobi, Germany |
| Chris Beck | University of Toronto, Canada |
| Nicolas Beldiceanu | IMT Atlantique, France |
| David Bergman | University of Connecticut, USA |
| Timo Berthold | Fair Isaac Germany GmbH, Germany |
| Natashia Boland | Georgia Institute of Technology, USA |
| Andre Augusto Cire | University of Toronto, Canada |
| Mathijs de Weerdt | Delft University of Technology, The Netherlands |
| Bistra Dilkina | University of Southern California, USA |
| Ambros Gleixner | Zuse Institute Berlin, Germany |
| Carla Gomes | Cornell University, USA |
| Tias Guns | Vrije Universiteit Brussel, Belgium |
| Matti Järvisalo | University of Helsinki, Finland |
| Serdar Kadioglu | Oracle Corporation, USA |
| Joris Kinable | Eindhoven University of Technology, The Netherlands |
| Philippe Laborie | IBM, France |
| Jeff Linderoth | University of Wisconsin-Madison, USA |
| Andrea Lodi | Polytechnique Montréal, Canada |

| | |
|---|---|
| Michele Lombardi | University of Bologna, Italy |
| Laurent Michel | University of Connecticut, USA |
| Michela Milano | University of Bologna, Italy |
| Nina Narodytska | VMware Research, USA |
| Siegfried Nijssen | Université catholique de Louvain, Belgium |
| Laurent Perron | Google France, France |
| Gilles Pesant | Polytechnique Montréal, Canada |
| Jean-Charles Régin | Université Nice-Sophia Antipolis, France |
| Andrea Rendl | University of Klagenfurt, Austria |
| Louis-Martin Rousseau | Polytechnique Montréal, Canada |
| Ashish Sabharwal | Allen Institute for Artificial Intelligence (AI2), USA |
| Domenico Salvagnin | University of Padua, Italy |
| Pierre Schaus | Université catholique de Louvain, Belgium |
| Andreas Schutt | CSIRO and The University of Melbourne, Australia |
| Peter Stuckey | CSIRO and The University of Melbourne, Australia |
| Michael Trick | Carnegie Mellon University, USA |
| Charlotte Truchet | Université de Nantes, France |
| Pascal Van Hentenryck | University of Michigan, USA |
| Willem-Jan van Hoeve | Carnegie Mellon University, USA |
| Alessandro Zanarini | ABB Corporate Research, Switzerland |

## Additional Reviewers

Babaki, Behrouz
Bai, Junwen
Belov, Gleb
Björck, Johan
Bliem, Bernhard
Booth, Kyle E. C.
Bridi, Thomas
Cherkaoui, Rachid
Cohen, Eldan
Csizmadia, Zsolt
Davarnia, Danial
De Filippo, Allegra
de Una, Diego
Dey, Santanu
Duc Vu, Minh
Emadikhiav, Mohsen
Fages, Jean-Guillaume
Farias, Vivek
Fischetti, Matteo
Flener, Pierre
Gagnon, Samuel
Gavanelli, Marco
Gottwald, Robert Lion
Grenouilleau, Florian
Gualandi, Stefano
Hendel, Gregor
Hojny, Christopher
Huang, Teng
Khalil, Elias
Kimura, Ryo
Kiziltan, Zeynep
Legrain, Antoine
Lin, Qin
Lippi, Marco
Lo Bianco, Giovanni
Lozano Sanchez, Leonardo
Marinakis, Adamantios
McCreesh, Ciaran
Mehrani, Saharnaz
Miltenberger, Matthias
Müller, Benjamin
Nonato, Maddalena

Olivier, Philippe
Pecin, Diego
Perez, Guillaume
Pralet, Cédric
Prouvost, Antoine
Raghunathan, Arvind
Rehfeldt, Daniel
Roli, Andrea
Römer, Michael
Serra, Thiago
Serrano, Felipe
Shi, Qinru
Tanneau, Mathieu
Tesch, Alexander
Tubertini, Paolo
Turner, Mark
Urli, Tommaso
Van Cauwelaert, Sascha
van den Bogaerdt, Pim
Vigerske, Stefan
Witzig, Jakob

## Extended Abstracts

The following extended abstracts were accepted for presentation at the conference:

- Magnus Björk, Pawel Pietrzak and Andriy Svynaryov: Modelling Real-World Strict Seniority Bidding Problems in Airline Crew Rostering
- Emir Demirović, Nicolas Schwind, Tenda Okimoto and Katsumi Inoue: Recoverable Team Formation: Building Teams Resilient to Change
- Andreas Ernst, Dhananjay Thiruvady, Davaatseren Baatar, Angus Kenny, Mohan Krishnamoorthy and Gaurav Singh: Mining, Matheuristics, and Merge-Search
- Alexandre Gondran and Laurent Moalic: Finding the Chromatic Number by Counting k-Colorings with a Randomized Heuristic
- Elias Khalil and Bistra Dilkina: Training Binary Neural Networks with Combinatorial Algorithms
- Varun Khandelwal: Solving Real-World Optimization Problems Using Artificial Intelligence
- Shiang-Tai Liu: Objective Bounds of Quadratic Programming with Interval Coefficients and Equality Constraints
- Günther Raidl, Elina Rönnberg, Matthias Horn and Johannes Maschler: An A*-Based Algorithm to Derive Relaxed Decision Diagrams for a Prize-Collecting Sequencing Problem
- Mark Wallace and Aldeida Aleti: Using CP to Prove Local Search Is Effective

# Abstracts of Invited Talks

# Same, Same, but Different: A Mostly Discrete Tour Through Optimization

Thorsten Koch[1,2]

[1] Zuse Institute Berlin, Takustr 7, 14195, Berlin, Germany
[2] Technische Universität Berlin, Str. des 17. Juni 136, 10623, Berlin, Germany
koch@zib.de
http://www.zib.de/koch

**Abstract.** This talk will give a short tour through selected topics in mathematical optimization. Though these topics are quite diverse, they also have a lot in common.

The tour will start at mixed-integer non-linear optimization (MINLP), proceed to mixed-integer optimization (MILP), it will then make short detour to linear programming (LP) and exact solutions, then proceed to algorithms, software, modelling, and parallel computing, jumping to gas networks as an application, from there visit Steiner tree problems, and finally arrive back at MILP.

On route, we will take the opportunity to point out a few challenges and open problems.

# Empirical Model Learning: Boosting Optimization Through Machine Learning

Michela Milano

DISI, University of Bologna, Bologna, Italy
michela.milano@unibo.it

**Abstract.** One of the biggest challenges in the design of decision support and optimization tools for complex, real-world, systems is coming up with a good combinatorial model. The traditional way to craft a combinatorial model is through interaction with domain experts: this approach provides model components (objective functions, constraints), but with limited accuracy guarantees. Often enough, accurate predictive models (e.g. simulators) can be devised, but they are too complex or too slow to be employed in combinatorial optimization.

In this talk, we propose a methodology called Empirical Model Learning (EML) that relies on Machine Learning for obtaining decision model components that link decision variables and observables, using data either extracted from a predictive model or harvested from a real system. We show how to ground EML on a case study of thermal-aware workload allocation and scheduling. We show how to encapsulate different machine learning models in a number of optimization techniques.

We demonstrate the effectiveness of the EML approach by comparing our results with those obtained using expert-designed models.

# Ten Years of CP Optimizer

Paul Shaw

IBM. 1681, route des Dolines, 06560 Valbonne, France
paul.shaw@fr.ibm.com

**Abstract.** CP Optimizer is the IBM constraint solving engine and part of CPLEX Optimization Studio. This talk takes a look at both the motivation and history of CP Optimizer, and the ten year journey from its beginnings until today.

At selected points, I will delve into the operation of different features of the engine, and the motivation behind them, together with how performance improvements in the automatic search were achieved.

From more recent history, I will concentrate on important developments such as the CP Optimizer file format, presolve, explanations for insolubility and backtrack, and lower bounds on the objective function.

# Abstracts of Fast-Track Journal Papers

# Online Over Time Processing of Combinatorial Problems

Robinson Duque[1], Alejandro Arbelaez[2], and Juan Francisco Díaz[1]

[1] Universidad del Valle, Cali, Colombia
{robinson.duque,juanfco.diazg}@correounivalle.edu.co
[2] Cork Institute of Technology, Cork, Ireland
alejandro.arbelaez@cit.ie

In an online environment, jobs arrive over time and there is no information in advance about how many jobs are going to be processed and what their processing times are going to be. We study the online scheduling of Boolean Satisfiability (SAT) and Mixed Integer Programming (MIP) instances that are well-known NP-complete problems. Typical online machine scheduling approaches assume that jobs are completed at some point to minimize functions related to completion time (e.g., makespan, minimum lateness, total weighted tardiness, etc).

In this work, we formalize and present an online over time problem where arriving instances are subject to waiting time constraints. To formalize our problem, we presented an extension of the Graham notation ($\alpha|\beta|\gamma$) that allowed us to represent the necessary constraints. We also proposed an approach for online scheduling of combinatorial problems that consisted of three parts. Namely, training/testing models for processing time estimations; implementation of a hybrid scheduling policy using SJF and MIP; and usage of instance interruption heuristics to mitigate the impact of inaccurate predictions.

Unlike other approaches, we attempt to maximize the number of solved instances using single and multiple machine configurations. Our empirical evaluation with well-known SAT and MIP instances, suggest that our interruption heuristics can improve generic ordering policies to solve up to 21.6x and 12.2x more SAT and MIP instances. Additionally, our hybrid approach observed results that are close to a semi clairvoyant policy (SCP) featuring perfect estimations. We observed that with very limited data to train the models our approach reports scenarios with up to 90% of solved instances with respect to the SCP.

Finally, we experimented using models that were trained with different feature families and observed an interesting trade-off between the quality of the predictions and the computational cost to calculate such features. For instance, *Trivial* features are basically free to compute but they have impact on the quality of the models. On the other hand, *Cheap* features offer an interesting trade-off between prediction quality and computational cost. This abstract refers to the full paper [1].

## Reference

1. Duque, R., Arbelaez, A., Díaz, J.F.: Online over time processing of combinatorial problems. In: Constraints Journal Fast Track of CPAIOR (2018)

# Deep Neural Networks as 0-1 Mixed Integer Linear Programs: A Feasibility Study

Matteo Fischetti[1] and Jason Jo[2]

[1] Department of Information Engineering (DEI), University of Padova
matteo.fischetti@unipd.it
[2] Montreal Institute for Learning Algorithms (MILA) and Institute for Data Valorization (IVADO), Montreal
jason.jo.research@gmail.com

**Abstract.** Deep Neural Networks (DNNs) are very popular these days, and are the subject of a very intense investigation. A DNN is made by layers of internal units (or neurons), each of which computes an affine combination of the output of the units in the previous layer, applies a nonlinear operator, and outputs the corresponding value (also known as activation). A commonly-used nonlinear operator is the so-called rectified linear unit (ReLU), whose output is just the maximum between its input value and zero. In this (and other similar cases like max pooling, where the max operation involves more than one input value), for fixed parameters one can model the DNN as a 0-1 Mixed Integer Linear Program (0-1 MILP) where the continuous variables correspond to the output values of each unit, and a binary variable is associated with each ReLU to model its yes/no nature. In this paper we discuss the peculiarity of this kind of 0-1 MILP models, and describe an effective bound-tightening technique intended to ease its solution. We also present possible applications of the 0-1 MILP model arising in feature visualization and in the construction of adversarial examples. Computational results are reported, aimed at investigating (on small DNNs) the computational performance of a state-of-the-art MILP solver when applied to a known test case, namely, hand-written digit recognition.

# Intruder Alert! Optimization Models for Solving the Mobile Robot Graph-Clear Problem

Michael Morin[1,2], Margarita P. Castro[1], Kyle E. C. Booth[1], Tony T. Tran[1], Chang Liu[1], and J. Christopher Beck[1]

[1] Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, ON, Canada
{mmorin,mpcastro,kbooth,tran,cliu,jcb}@mie.utoronto.ca
[2] Department of Operations and Decision Support Systems, Université Laval, Québec, QC, Canada
michael.morin@osd.ulaval.ca

We investigate optimization-based approaches and heuristic methods for the *graph-clear problem* (GCP), an $\mathcal{NP}$-Hard variant of the pursuit-evasion problem. The goal is to find a schedule that minimizes the total number of robots needed to “clear” possible intruders from a facility, represented as a graph. The team of robots can use *sweep* actions to remove intruders from *contaminated* nodes and *block* actions to prevent intruders from traveling between nodes. A solution to the GCP is a schedule of sweep and block actions that detects all potential intruders in the facility while minimizing the number of robots required. Solutions such that cleared vertices at each time step form a connected subgraph are termed *contiguous*, while those that prevent recontamination and, therefore, the need to sweep a node more than once, are called *progressive*.

We prove, via a counter-example, that enforcing contiguity may remove all optimal solutions and, conversely, that preventing two special forms of recontamination does not remove all optimal solutions. However, the completeness for the general case of progressive solutions remains open.

We then present *mixed-integer linear programming* (MILP) and *constraint programming* (CP) approaches, as well as new heuristic variants for solving the GCP and compare them to previously proposed heuristics. This is the first time that MILP and CP have been applied to the problem. Our experimental results indicate that our heuristic modifications improve upon the heuristics in the literature, that constraint programming finds better solutions than the heuristics in run-times reasonable for the application, and that mixed-integer linear programming is the superior approach for proving optimality. Nonetheless, for larger problem instances, the optimality gap for CP and MILP remains very large, indicating the need for future research and improvement. Given the performance of CP and MILP compared to the heuristic approaches, coupled with the appeal of the model-and-solve framework, we conclude that they are currently the most suitable approaches for the graph-clear problem.

---

Margarita P. Castro, Kyle E. C. Booth— Equally contributing authors.

# Contents