Lecture Notes in Computer Science 10416

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/7408

J. Christopher Beck (Ed.)

Principles and Practice of Constraint Programming

23rd International Conference, CP 2017
Melbourne, VIC, Australia, August 28 – September 1, 2017
Proceedings

Springer

Editor
J. Christopher Beck
University of Toronto
Toronto, ON
Canada

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-319-66157-5 ISBN 978-3-319-66158-2 (eBook)
DOI 10.1007/978-3-319-66158-2

Library of Congress Control Number: 2017949522

LNCS Sublibrary: SL2 – Programming and Software Engineering

Printed on acid-free paper

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

Preface

This volume contains the proceedings of the 23rd International Conference on the Principles and Practice of Constraint Programming (CP 2017) held from August 28 to September 1, 2017 in Melbourne, Australia and colocated with the 20th International Conference on Theory and Applications of Satisfiability Testing (SAT 2017) and the 33rd International Conference on Logic Programming. Detailed information about the CP 2017 conference with links to the colocated conferences can be found at http://cp2017.a4cp.org.

The CP conference is the annual international conference on all aspects of computing with constraints including theory, algorithms, environments, languages, models, systems, and applications such as decision making, resource allocation, scheduling, configuration, and planning. In addition to the main technical track and long-standing applications track, and as a continuation of the effort of the CP community to reach out to other research fields that intersect with constraint programming, CP 2017 featured thematic tracks in Machine Learning and CP, Operations Research and CP, Satisfiability and CP, and Test and Verification and CP. Each track had its own Track Chair(s) and Program Committee to ensure that the papers would be peer reviewed by expert reviewers with specific knowledge of the intersecting area.

The conference received 115 submissions across all tracks, including eight submissions to the Journal and Sister Conferences Track. Each paper was assigned to a Senior Program Committee member or the appropriate Track Chair and to three Program Committee members from either the technical track Program Committee or the relevant thematic track Program Committee. All papers received at least three reviews, following which the authors had an opportunity to respond. Detailed discussions were held on each paper by the PC members, led by the SPC member, Track Chair, and Program Chair. The Senior Program Committee, including Track Chairs, met in Padova, Italy, on June 5, 2017 with participation both in person and via video link. Each paper was discussed by the SPC with the decisions taken by consensus. The Journal and Sister Conferences Track papers followed a separate process, led by the Track Chair, to evaluate the relevance and significance of submitted papers that had been previously published in journals or other conferences. The Journal and Sister Conferences Track Program Committee met in Cupar, Scotland, UK on June 18, 2017 to make the final decisions. The final outcome of these meetings was the acceptance of 46 papers across all technical and thematic tracks, resulting in an acceptance rate of approximately 44%, and the acceptance of all eight papers submitted to the Journal and Sister Conferences Track.

The Senior Program Committee awarded four best paper awards, generously supported by Springer.

– Best Paper Award: Grigori German, Olivier Briant, Hadrien Cambazard, and Vincent Jost, “Arc Consistency via Linear Programming”

- Distinguished Paper Award: Fahiem Bacchus, Antti Hyttinen, Matti Järvisalo, and Paul Saikko, "Reduced Cost Fixing in MaxSAT"
- Best Student Paper Award: Adrian Goldwaser and Andreas Schutt, "Optimal Torpedo Scheduling"
- Distinguished Student Paper Award: Guillaume Derval, Jean-Charles Regin, and Pierre Schaus, "Improved Filtering for the Bin-Packing with Cardinality Constraint"

The Program Chair and the Journal-Publication-Fast-Track Chair, Louis-Martin Rousseau, invited four papers from across the technical and thematic tracks to participate in the *Constraints* journal fast-track process to publish an extended version in the journal at the same time as the conference, while also presenting the work at the conference. Due to the tight editorial deadlines, one paper accepted this invitation and so appears in this volume as an abstract with the full paper in *Constraints*.

The conference program included five invited talks in coordination with SAT 2017 and ICLP 2017 by Agostino Dovier, Holger Hoos, Nina Narodytska, Enrico Pontelli, and Mark Wallace. The conference also shared the workshop program with the two colocated conferences, resulting in seven workshops overseen by the Joint Workshop Chairs: Charlotte Truchet, Enrico Pontelli, and Stefan Rümmele. The tutorial program, also chaired by Charlotte Truchet with support from the SAT 2017 and ICLP 2017 Program Chairs, consisted of four tutorials on CP, SAT, Mixed Integer Nonlinear Programming, and Machine Learning and Data Science. The Doctoral Program, jointly organized by CP 2017 and ICLP 2017 and chaired by Chris Mears and Neda Saeedloei, hosted 24 students from around the world. The students had an opportunity to present their work, meet one-on-one with a senior researcher mentor, and attend invited talks targeted to the experiences of a PhD student.

The program for the conference is the result of a substantial amount of work by many people to whom I am grateful. I would like to thank the authors for their submission of high-quality scientific work and the substantial efforts of the Program Committees and external reviewers, who jointly prepared 341 high-quality reviews. The Senior Program Committee and Track Chairs played a crucial role in managing the reviews and discussions, in writing meta-reviews and recommendations for each submission, and in making the final decisions. I would like to specifically acknowledge the efforts of the Track Chairs to attract new contributors to the conference: Yael Ben-Haim and Yehuda Naveh (Satisfiability and CP Track Chairs), David Bergman and Andre Cire (Operations Research and CP Track Chairs), Ken Brown (Application Track Chair), Arnaud Gotlieb and Nadjib Lazaar (Test and Verification and CP Track Chairs), Tias Guns and Michele Lombardi (Machine Learning and CP Track Chairs), Karen Petrie (Journal and Sister Conferences Track Chair), Enrico Pontelli (Biology and CP Track Chair), and Louis-Martin Rousseau (Journal-Publication-Fast-Track Chair).

Beyond the peer review process, there is a substantial team that made the program and conference possible. I would like to particularly thank: Peter Stuckey and Guido Tack (CP 2017 Conference Chairs), Christopher Mears and Neda Saeedloei (Doctoral Program Chairs), Charlotte Truchet (Tutorial Chair), Charlotte Truchet, Enrico Pontelli, and Stefan Rümmele (Joint CP/SAT/ICLP Workshop Chairs), Tommaso Urli (Publicity Chair), Maria Garcia de la Banda (ICLP 2017 Conference Chair), Serge

Gaspers and Toby Walsh (SAT 2017 Conference and Program Chairs), and Ricardo Rocha and Tran Cao Son (ICLP 2017 Program Chairs).

I would also like to thank the sponsors of the conference for their generous support. At the time of writing, these sponsors include: the Artificial Intelligence Journal Division (AIJD) of IJCAI, the Association for Constraint Programming, the Association for Logic Programming, the City of Melbourne, CompSustNet, Cosling, Cosytec, CSIRO Data61, the European Association for Artificial Intelligence, IBM, Monash University, Satalia, Springer, and the University of Melbourne.

July 2017 Chris Beck

Tutorials and Workshops

Tutorials

Introduction to Constraint Programming - If You Already Know SAT or Logic Programming

Guido Tack	Monash University, Australia

An Introduction to Satisfiability

Armin Biere	Johannes Kepler University Linz, Austria

Introduction to Machine Learning and Data Science

Tias Guns	Vrije Universiteit Brussel, Belgium

Mixed Integer Nonlinear Programming: An Introduction

Pietro Belotti	FICO, UK

Workshops

Pragmatics of Constraint Reasoning

Daniel Le Berre	Université d'Artois, France
Pierre Schaus	UCLouvain, Belgium

Workshop on Answer Set Programming and Its Applications

Kewen Wang	Griffith University, Australia
Yan Zhang	Western Sydney University, Australia

Workshop on Constraint Solvers in Testing, Verification, and Analysis

Zakaria Chihani	Atomic Energy Commission (CEA), France

Workshop on Logic and Search

David Mitchell	Simon Fraser University, Canada

Progress Towards the Holy Grail

Eugene Freuder	University College Cork, Ireland

International Workshop on Constraint Modeling and Reformulation

Özgür Akgün	University of St Andrews, UK

Colloquium on Implementation of Constraint Logic Programming Systems

Jose F. Morales IMDEA Software Institute, Spain
Nataliia Stulova IMDEA Software Institute, Spain

Conference Organization

Program Chair

J. Christopher Beck	University of Toronto, Canada

Conference Chairs

Peter Stuckey	University of Melbourne, Australia
Guido Tack	Monash University, Australia

Application Track Chair

Ken Brown	University College Cork, Ireland

Biology and CP Track Chair

Enrico Pontelli	New Mexico State University, USA

Machine Learning and CP Track Chairs

Tias Guns	Vrije Universiteit Brussel, Belgium
Michele Lombardi	DISI, University of Bologna, Italy

Operations Research and CP Track Chairs

David Bergman	University of Connecticut, USA
Andre Augusto Cire	University of Toronto Scarborough, Canada

Satisfiability and CP Track Chairs

Yael Ben-Haim	IBM Research, Israel
Yehuda Naveh	IBM Research, Israel

Test and Verification and CP Track Chairs

Nadjib Lazaar	LIRMM, France
Arnaud Gotlieb	Simula Research Laboratory, Norway

Journal and Sister Conferences Track Chair

Karen Petrie	University of Dundee, UK

Journal-Publication-Fast-Track Chair

Louis-Martin Rousseau	École Polytechnique de Montréal, Canada

Doctoral Program Chairs

Christopher Mears	Redbubble, Australia
Neda Saeedloei	University of Minnesota Duluth, USA

Workshop and Tutorial Chair

Charlotte Truchet	University of Nantes, France

Publicity Chair

Tommaso Urli	CSIRO/Data61 and the Australian National University, Australia

Senior Program Committee

Yael Ben-Haim	IBM Research, Israel
David Bergman	University of Connecticut, USA
Ken Brown	University College Cork, Ireland
Andre Augusto Cire	University of Toronto Scarborough, Canada
Sophie Demassey	CMA, MINES ParisTech, France
Bistra Dilkina	Georgia Institute of Technology, USA
Arnaud Gotlieb	Simula Research Laboratory, Norway
Tias Guns	Vrije Universiteit Brussel, Belgium
Nadjib Lazaar	LIRMM, France
Christophe Lecoutre	CRIL, Université d'Artois, France
Jimmy Lee	The Chinese University of Hong Kong, Hong Kong
Michele Lombardi	DISI, University of Bologna, Italy
Yehuda Naveh	IBM Research, Israel
Justyna Petke	University College London, UK
Karen Petrie	University of Dundee, UK
Enrico Pontelli	New Mexico State University, USA
Louis-Martin Rousseau	École Polytechnique de Montréal, Canada
Pierre Schaus	UCLouvain, Belgium
Andreas Schutt	Data61, CSIRO, and The University of Melbourne, Australia

Technical Track Program Committee

Carlos Ansótegui	Universitat de Lleida, Spain
Nicolas Beldiceanu	IMT Atlantique (LS2N), France

David Cohen	Royal Holloway, University of London, UK
Pierre Flener	Uppsala University, Sweden
Emmanuel Hebrard	LAAS-CNRS, Université de Toulouse, France
John Hooker	Carnegie Mellon University, USA
Marie-José Huguet	LAAS-CNRS, Université de Toulouse, France
Said Jabbour	CRIL CNRS, Université d'Artois, France
Joris Kinable	Carnegie Mellon University, USA
Zeynep Kiziltan	University of Bologna, Italy
Philippe Laborie	IBM, France
Chavalit Likitvivatanavong	Thailand
Boon Ping Lim	NICTA, Australia
Andrea Lodi	École Polytechnique de Montréal, Canada
Samir Loudni	GREYC, CNRS UMR 6072, Université de Caen Basse-Normandie, France
Ines Lynce	INESC-ID/IST, Universidade de Lisboa, Portugal
Arnaud Malapert	Université Côte d'Azur, CNRS, I3S, France
Ciaran McCreesh	University of Glasgow, UK
Kuldeep S. Meel	Rice University, USA
Peter Nightingale	University of St Andrews, UK
Justin Pearson	Uppsala University, Sweden
Gilles Pesant	École Polytechnique de Montréal, Canada
Thierry Petit	Worcester Polytechnic Institute, USA
Patrick Prosser	University of Glasgow, UK
Claude-Guy Quimper	Université Laval, Canada
Jean-Charles Regin	University Nice-Sophia Antipolis/I3S/CNRS, France
Emma Rollon	Technical University of Catalonia, Spain
Francesca Rossi	IBM Research and University of Padova, Italy
Mohamed Siala	Insight Centre for Data Analytics, University College Cork, Ireland
Michael Trick	Carnegie Mellon University, USA
Tommaso Urli	CSIRO Data61 and the Australian National University, Australia
Peter van Beek	University of Waterloo, Canada
Pascal Van Hentenryck	University of Michigan, USA
Willem-Jan Van Hoeve	Carnegie Mellon University, USA
Petr Vilím	IBM, Czech Republic
Christel Vrain	LIFO, University of Orléans, France
Mohamed Wahbi	Insight, University College Cork, Ireland
Roland Yap	National University of Singapore, Singapore
William Yeoh	New Mexico State University, USA
Alessandro Zanarini	ABB Corporate Research, Switzerland
Roie Zivan	Ben Gurion University of the Negev, Israel
Stanislav Živný	University of Oxford, UK

Application Track Program Committee

Carmen Gervet	Université de Montpellier, France
Philip Kilby	Data61 and the Australian National University, Australia
Deepak Mehta	United Technologies Research Centre Ireland, Ireland
Laurent Michel	University of Connecticut, USA
Laurent Perron	Google, France
Christian Schulte	KTH Royal Institute of Technology, Sweden
Paul Shaw	IBM, France
Helmut Simonis	Insight Centre for Data Analytics, University College Cork, Ireland
Sylvie Thiébaux	The Australian National University, Australia

Biology and CP Track Program Committee

Nicos Angelopoulos	Wellcome Sanger Institute, UK
Alexander Bockmayr	Freie Universität Berlin, Germany
Simon De Givry	INRA - MIAT, France
Agostino Dovier	Università di Udine, Italy
Ferdinando Fioretto	University of Michigan, USA

Machine Learning and CP Track Program Committee

Christian Bessiere	CNRS, University of Montpellier, France
Bruno Cremilleux	Université de Caen, France
Thi-Bich-Hanh Dao	University of Orléans, France
Georgiana Ifrim	University College Dublin, Ireland
Michela Milano	DISI Università di Bologna, Italy
Siegfried Nijssen	Université Catholique de Louvain, Belgium

Operations Research and CP Track Program Committee

Serdar Kadioglu	Oracle Corporation, USA
Nick Sahinidis	Carnegie Mellon University, USA
Olivia Smith	IBM Research, Australia
Christine Solnon	LIRIS CNRS UMR 5205/INSA Lyon, France
Pascal Van Hentenryck	University of Michigan, USA
Tallys Yunes	University of Miami, USA

Satisfiability and CP Track Program Committee

Alan Frisch	University of York, UK
George Katsirelos	MIAT, INRA, France
Ian Miguel	University of St Andrews, UK
Nina Narodytska	Samsung Research America, USA

Steve Prestwich	Insight Centre for Data Analytics, Ireland
Ofer Strichman	Technion, Israel

Test and Verification and CP Track Program Committee

Sébastien Bardin	CEA LIST, France
Mats Carlsson	RISE SICS, Sweden
Roberto Castañeda Lozano	SICS, Sweden
Catherine Dubois	ENSIIE, Samovar, France
Vijay Ganesh	University of Waterloo, Canada
Anastasia Paparrizou	CRIL-CNRS, Université d'Artois, France
Marie Pelleau	LIP6, UPMC, France
Andreas Podelski	University of Freiburg, Germany
Michel Rueher	Université Côte d'Azur, CNRS, France
Pascal Van Hentenryck	University of Michigan, USA
Lebbah Yahia	University of Oran 1, Algeria

Journal and Sister Conferences Track Program Committee

Özgür Akgün	University of St Andrews, UK
Munsee Chang	University of St Andrews, UK
Ian Gent	University of St Andrews, UK
Christopher Jefferson	University of St Andrews, UK
Peter Nightingale	University of St Andrews, UK

Additional Reviewers

Rui Abreu
Özgür Akgün
Suguman Bansal
Johannes Gerhardus Benade
Clément Carbonnel
Mats Carlsson
Supratik Chakraborty
Kenil Cheng
Christel Christel
Martin Cooper
Thi-Bich-Hanh Dao
Daniel J. Fremont
Luca Di Gaspero
Chrysanthos Gounaris
Alban Grastien
Peter Jeavons
Christopher Jefferson
Ryo Kimura
Philippe Laborie
Javier Larrosa
Nadjib Lazaar
Olivier Lhomme
Barnaby Martin
Jacopo Mauro
Ciaran McCreesh
Jean-Noël Monette
Christian Muise
Cemalettin Ozturk
Alexandre Papadopoulos
Guillaume Perez
Yash Puranik
Ashish Sabharwal
Lakhdar Sais
Paul Scott

Thiago Serra
Gilles Simonin
Carsten Sinz
Friedrich Slivovsky
Atena Tabakhi
James Trimble
Gilles Trombettoni
Matt Valeriote
Christoph M. Wintersteiger
Ghiles Ziat
Roie Zivan
Ed Zulkoski

Journal Fast Track (Abstract)

Improved Filtering for the Bin-Packing with Cardinality Constraint

Guillaume Derval[1], Jean-Charles Régin[2], and Pierre Schaus[1]

[1] UCLouvain, Belgium
{guillaume.derval,pierre.schaus}@uclouvain.be
[2] University of Nice Sophia-Antipolis, France
jcregin@gmail.com

Previous research [2, 3] shows that a cardinality reasoning can improve the pruning of the bin-packing constraint, even when cardinalities are not involved in the original model. Our contribution is two-fold.

We first introduce a new algorithm, called BPCFlow, that filters both load and cardinality bounds on the bins, using a flow reasoning similar to the one used for the Global Cardinality Constraint.

Moreover, we detect impossible assignments of items by combining the load and cardinality of the bins using a new reasoning method called "too-big/too-small". This new method attempts to construct for each bin with load and cardinality bounds $[\underline{L}, \overline{L}]$ and $[\underline{C}, \overline{C}]$ a maximum-weighted set of $\overline{C} - 1$ items. Once this set is constructed, we detect that items with weight $w < \underline{L} - \sum_{i \in S} w_i$ cannot be assigned to the current bin. Similar arguments can be used to detect a maximum weight. The "too-big/too-small" reasoning is then adapted to the existing propagators, namely SimpleBPC [3], Pelsser's method [2] and BPCFlow.

We then experiment our four new algorithms on Balanced Academic Curriculum Problem and Tank Allocation Problem instances.

BPCFlow is shown to be indeed stronger than previously existing filtering, and more computationally intensive. We show that the new filtering is useful on a small number of hard instances, while being too expensive for general use.

Our results show the introduced "too-big/too-small" filtering can most of the time drastically reduce the size of the search tree and the computation time. This method is profitable in 88% of the tested instances.

This work is published in the Constraints journal [1].

References

1. Derval, G., Régin, J.C., Schaus, P.: Improved filtering for the bin-packing with cardinality constraint. In: Constraints. Springer (2017)
2. Pelsser, F., Schaus, P., Régin, J.C.: Revisiting the cardinality reasoning for binpacking Constraint. In: Schulte, C. (eds.) Principles and Practice of Constraint Programming, CP 2013. LNCS, vol 8124, pp. 578–586. Springer, Heidelberg (2013)

3. Schaus, P., Régin, J.C., Van Schaeren, R., Dullaert, W., Raa, B.: Cardinality reasoning for bin-packing constraint: application to a tank allocation problem. In: Milano, M. (eds) Principles and Practice of Constraint Programming. LNCS, vol 7514, pp. 815–822. Springer, Heidelberg (2012)

Journal and Sister Conference Tracks (Abstracts)

Ranking Constraints

Christian Bessiere[1], Emmanuel Hebrard[2], George Katsirelos[3], Zeynep Kiziltan[4], and Toby Walsh[5]

[1] LIRMM, CNRS, Université de Montpellier
[2] LAAS-CNRS, Université de Toulouse
[3] MIAT, INRA
[4] University of Bologna
[5] University of New south Wales

Abstract. In many problems we want to reason about the ranking of items. For example, in information retrieval, when aggregating several search results, we may have ties and consequently rank orders. (e.g. [2, 3]). As a second example, we may wish to construct an overall ranking of tennis player based on pairwise comparisons between players. One principled method for constructing a ranking is the Kemeny distance [5] as this is the unique scheme that is neutral, consistent, and Condorcet. Unfortunately, determining this ranking is NP-hard, and remains so when we permit ties in the input or output [4]. As a third example, tasks in a scheduling problem may run in parallel, resulting in a ranking. In a ranking, unlike a permutation, we can have ties. Thus, 12225 is a ranking whilst 12345 is a permutation. To reason about permutations, we have efficient and effective global constraints. Regin [7] proposed an $O(n^4)$ GAC propagator for permutations. For BC, there is an even faster $O(n \log n)$ propagator [6]. Every constraint toolkit now provides propagators for permutation constraints. Surprisingly, ranking constraints are not yet supported. In [1], we tackle this weakness by proposing a global ranking constraint. We show that simple decompositions of this constraint hurt pruning. We then show that GAC can be achieved in polynomial time and we propose an $O(n^3 \log n)$ algorithm for achieving RC as well as an efficient quadratic algorithm offering a better tradeoff.

References

1. Bessiere, C., Hebrard, E., Katsirelos, G., Kiziltan, Z., Walsh. T.: Ranking constraints. In: Proceedings of the of IJCAI, pp. 705–711 (2016)
2. Brancotte, B., Yang, B., Blin, G., Denise S. Cohen-Boulakia, Hamel, S.: Rank aggregation with ties: experiments and analysis. In: Proceedings of the VLDB Endowment (PVLDB) (2015)
3. Fagin, R., Kumar, R., Mahdian, M., Sivakumar, D., Vee, E.: Comparing and aggregating rankings with ties. In: Proceedings of the PODS, pp. 47–58. ACM (2004)
4. Hemaspaandra, E., Spakowski, H., Vogel, J.: The complexity of kemeny elections. Theoret. Comput. Sci. **349**(3), 382–391 (2005)
5. Kemeny, J.G.: Mathematics without numbers. Daedalus **88**(4), 577–591 (1959)

6. Ortiz, A., Quimper, C.-G., Tromp, J., van Beek, P.: A fast and simple algorithm for bounds consistency of the all different constraint. In: Proceedings of the of IJCAI, pp. 245–250 (2003)
7. Régin. J.-C., A filtering algorithm for constraints of difference in CSPs. In: Proceedings of the AAAI, pp. 362–367 (1994)

Modeling with Metaconstraints and Semantic Typing of Variables

André Ciré[1], J.N. Hooker[2], and Tallys Yunes[3]

[1] University of Toronto
[2] Carnegie Mellon University
[3] University of Miami

Research in hybrid optimization shows that a combination of constraint programming and optimization technologies can significantly speed up computation. A key element of hybridization is the use of high-level metaconstraints in the problem formulation, which generalize the global constraints that are characteristic of constraint programming models. Metaconstraints aid solution by communicating problem structure to the solver.

Modeling with metaconstraints, however, raises a fundamental issue of variable management that must be addressed before its full potential can be realized. The solver frequently creates auxiliary variables as it relaxes and/or reformulates metaconstraints. Variables created for different constraints may actually have the same meaning, or they may relate in some more complicated way to each other and to variables in the original model. The solver must recognize these relationships among variables if it is to generate the necessary channeling constraints and formulate a tight overall continuous relaxation of the problem.

We address this problem systematically with a semantic typing scheme that reveals relationships among variables while allowing simpler, self-documenting models. We view a model as organized around user-defined, multiplace predicates that denote relations akin to those that occur in a relational database. A variable declaration is viewed as a database query that has the effect of assigning a semantic type to the variable. Relationships between variables are then deduced from their semantic types.

We develop this idea for a wide variety of constraint types, including systems of all-different constraints, employee scheduling constraints, general scheduling constraints with interval variables, sequencing problems with side constraints, disjunctions of linear systems, and constraints with piecewise linear functions. We develop three very general classes of channeling constraints that can be automatically inferred and are based on such relational database operations as projection. Finally, we discuss the advantages of semantic typing for error detection and model management.

This is an extended abstract of the full paper, which appears in *INFORMS Journal on Computing* **28** (2016) 1–13.

MaxSAT-Based Large Neighborhood Search for High School Timetabling

Emir Demirović and Nysret Musliu

Institute of Information Systems, Databases and Artificial Intelligence Group, Vienna University of Technology, Vienna, Austria
{demirovic,musliu}@dbai.tuwien.ac.at

Extended Abstract

The problem of high school timetabling (HSTT) is to coordinate resources (e.g. rooms, teachers, students) with times in order to fulfill certain goals (e.g. scheduling lectures). It is a well known and widespread problem, as every high school requires some form of timetabling. Unfortunately, HSTT is hard to solve and just finding a feasible solution for simple variants of HSTT has been proven to be NP-complete. When solving hard combinatorial problems such as HSTT, there are two solving paradigms that are used often: local search algorithms, which usually find fast local optimal solutions, but cannot guarantee the optimality, and complete algorithms, which provide optimal results by exhaustively enumerating all solutions over longer periods of time.

In this paper [1], we aim to obtain the best of both worlds by combining the two strategies. More precisely, we develop a new anytime algorithm for HSTT which combines local search with a novel maxSAT-based large neighborhood search. A local search algorithm is used to drive an initial solution into a local optimum and then more powerful large neighborhood search (LNS) techniques based on maxSAT are used to further improve the solution. During the course of the algorithm, the solution is iteratively *destroyed*, by using one of the two neighborhood vectors, and *repaired* by maxSAT. The size of the neighborhood vectors is increased with time until the complete search space is explored, allowing the algorithm to prove optimality if given enough computational time.

The computational results demonstrate that we outperform the state-of-the-art solvers on numerous benchmarks and provide four new upper bounds. To the best of our knowledge, this is the first time maxSAT is used within a large neighborhood search scheme. In addition, we experiment with several variants to show the importance of each component of the algorithm. Furthermore, our algorithm is more efficient than a pure maxSAT-based approach for the given computational setting (20 min runtime).

Reference

1. Demirovic, E., Musliu, N.: Maxsat-based large neighborhood search for high school timetabling. Comput. OR **78**, 172–180 (2017)

Android Database Attacks Revisited

Behnaz Hassanshahi and Roland H.C. Yap

School of Computing, National University of Singapore, Singapore
b.hassanshahi@u.nus.edu
ryap@comp.nus.edu.sg

Many Android apps (applications) employ databases for managing sensitive data. In [1], we systematically study attacks targeting databases in benign Android apps and also study a new class of database vulnerabilities, which we call *private database* vulnerabilities.

We propose an analysis framework, extending the framework in [2], to find Android database vulnerabilities which are confirmed with a proof-of-concept (POC) exploit, i.e. zero-day. Our analysis combines static dataflow analysis, symbolic execution and constraint solving and finally dynamic testing to certify the exploit. In order, to generate a POC malware, our analysis uses an SMT solver to solve the path constraints in the program which together with the Android manifest is used to generate parameters for API calls which may exploit the app database vulnerabilities. Dynamic testing on the generated POC malware confirms whether or not the malware exploits the app database vulnerabilities, if not, alternative malware are generated.

In order to analyse how apps use databases, it is necessary to accurately handle URI objects and libraries which use them. We build accurate models for URI objects connecting them to appropriate constraints. Simple URI methods can be directly translated to SMT formulas while more complex URI methods are modelled using Symbolic Finite Transducers together with the SMT solver.

We evaluate our analysis on popular Android apps, successfully finding many database vulnerabilities. Surprisingly, our analyzer finds new ways to exploit previously reported and fixed vulnerabilities. We also propose a fine-grained protection mechanism which extends the Android manifest to protect against database attacks.

References

1. Hassanshahi, B., Yap, R.H.C.: Android database attacks revisited. In: ACM Asia Conference on Computer and Communications Security (ASIACCS), pp. 625–639, ACM (2017)
2. Hassanshahi, B., Jia, Y., Yap, R.H.C., Saxena, P., Liang, Z.: Web-to-application injection attacks on android: characterization and detection. In: 20th European Symposium on Research in Computer Security, LNCS, vol. 9327, pp. 577–598. Springer, Cham (2015)

This is a summary of paper [1].

Hybrid Optimization Methods for Time-Dependent Sequencing Problems (Abstract)

Joris Kinable[1,2], Andre A. Cire[3], and Willem-Jan van Hoeve[2]

[1] Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213, USA
jkinable@cs.cmu.edu

[2] Tepper School of Business, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213, USA
acire@utsc.utoronto.ca

[3] Department of Management, University of Toronto Scarborough, 1265 Military Trail, Toronto, ON M1C 1A4, Canada
vanhoeve@andrew.cmu.edu

Abstract. A large number of practical problems in manufacturing, transportation, and distribution require the sequencing of activities over time. Often activities in a sequencing problem are subject to operational constraints and optimization criteria involving *setup times*, i.e., the minimum time that must elapse between two consecutive activities in a sequence. A setup time typically models the time to change jobs in an assembly line or the travel time between two cities in traveling salesman problems. In classical sequencing problems, the setup time is only defined between pairs of activities. However, in many practical applications the setup time is also a function of the order of the activities in the sequence. Such *position-dependent* setup times are useful in modeling different states of a resource throughout a schedule, for example when the internal components of a machine degrade after performing a number of tasks.

In this paper, we introduce a novel optimization method for sequencing problems with position-dependent setup times. Our proposed method relies on a hybrid approach where a constraint programming model is enhanced with two distinct relaxations: A discrete relaxation based on multivalued decision diagrams, and a continuous relaxation based on linear programming, which are combined via the method of additive bounding. The relaxations are used to generate bounds and enhance constraint propagation. We conduct experiments on three variants of the time-dependent traveling salesman problem: the first considers no side constraints, the second considers time window constraints, and the third considers precedence constraints between pairs of activities. The experiments indicate that our techniques substantially outperform general-purpose methods based on mixed-integer linear programming and constraint programming models.

This paper appeared as "Joris Kinable, Andre A. Cire, and Willem-Jan van Hoeve. Hybrid Optimization Methods for Time-Dependent Sequencing Problems. *European Journal of Operational Research* 259(3):887–897, 2017".

Learning Rate Based Branching Heuristic for SAT Solvers

Jia Hui Liang, Vijay Ganesh, Pascal Poupart, and Krzysztof Czarnecki

University of Waterloo, Waterloo, Canada

Abstract. In this paper, we propose a framework for viewing solver branching heuristics as optimization algorithms where the objective is to maximize the *learning rate*, defined as *the propensity for variables to generate learnt clauses*. By viewing online variable selection in SAT solvers as an optimization problem, we can leverage a wide variety of optimization algorithms, especially from machine learning, to design effective branching heuristics. In particular, we model the variable selection optimization problem as an online multi-armed bandit, a special-case of *reinforcement learning*, to learn branching variables such that the learning rate of the solver is maximized. We develop a branching heuristic that we call *learning rate branching* or LRB, based on a well-known multi-armed bandit algorithm called *exponential recency weighted average* and implement it as part of MiniSat and CryptoMiniSat. We upgrade the LRB technique with two additional novel ideas to improve the learning rate by accounting for *reason side rate* and exploiting *locality*. The resulting LRB branching heuristic is shown to be faster than the VSIDS and conflict history-based (CHB) branching heuristics on 1975 application and hard combinatorial instances from 2009 to 2014 SAT Competitions. We also show that CryptoMiniSat with LRB solves more instances than the one with VSIDS. These experiments show that LRB improves on state-of-the-art. The original version of this paper appeared in the SAT 2016 proceedings [1].

Reference

1. Liang, J.H., Ganesh, V., Poupart, P., Czarnecki, K.: Learning rate based branching heuristic for SAT Solvers. In: Proceedings of the 19th International Conference on Theory and Applications of Satisfiability Testing, SAT 2016, Bordeaux, France, 5–8 July 2016, pp. 123–140 (2016)

Three Generalizations of the FOCUS Constraint

Nina Narodytska[1], Thierry Petit[2,3], Mohamed Siala[4], and Toby Walsh[5]

[1] Samsung Research America, Mountain View, USA
nina.n@samsung.com
[2] School of Business, Worcester Polytechnic Institute, Worcester, USA
tpetit@wpi.edu
[3] LINA-CNRS, Mines-Nantes, Inria, Nantes, France
thierry.petit@mines-nantes.fr
[4] Insight Centre for Data Analytics, Department of Computer Science, University College Cork, Ireland
mohamed.siala@insight-centre.org
[5] UNSW, Data61 and TU Berlin, Sydney, NSW 2052, Australia
toby.walsh@data61.csiro.au

Abstract. The Focus constraint expresses the notion that solutions are concentrated. In practice, this constraint suffers from the rigidity of its semantics. To tackle this issue, we propose three generalizations of the Focus constraint. We provide for each one a complete filtering algorithm. Moreover, we propose ILP and CSP decompositions.

This work is published in [1, 2].

References

1. Narodytska, N., Petit, T., Siala, M., Walsh, T.: Three generalizations of the FOCUS constraint. In: Proceedings of the 23rd International Joint Conference on Artificial Intelligence, IJCAI 2013, Beijing, China, 3–9 August 2013, pp. 630–636 (2013)
2. Narodytska, N., Petit, T., Siala, M., Walsh, T.: Three generalizations of the FOCUS constraint. Constraints **21**(4), 495–532 (2016)

Conditions Beyond Treewidth for Tightness of Higher-Order LP Relaxations

Mark Rowland, Aldo Pacchiano, and Adrian Weller

UC Berkeley
pacchiano@berkeley.edu

We examine Boolean binary weighted constraint satisfaction problems without hard constraints, and explore conditions under which it is possible to solve the problem exactly in polynomial time [2]. We are interested in the problem of finding a configuration of variables $x = (x_1, \ldots, x_n) \in \{0, 1\}^n$ that maximizes a score function, defined by unary and pairwise rational terms $f(x) = \sum_{i=1}^{n} \psi_i(x_i) + \sum_{(i,j)\in E} \psi_{ij}(x_i, x_j)$. In the machine learning community, this is typically known as *MAP* (or *MPE*) *inference*.

In this work, we consider a popular approach which first expresses the MAP problem as an integer linear program (ILP) then relaxes this to a linear program (LP). If the LP optimum is achieved at an integral point we say the LP is tight. If the LP is performed over the marginal polytope, which enforces global consistency [1], then the LP will always be tight but exponentially many constraints are required. Sherali and Adams introduced a series of successively tighter relaxations of the marginal polytope: for any integer r, $\mathbb{L}_r$ enforces consistency over all clusters of variables of size $\leq r$. $\mathbb{L}_r$ is solvable in polynomial time and tight for graphs of treewidth $r - 1$ [1].

Most past work has focused on characterizing conditions for $\mathbb{L}_2$ and $\mathbb{L}_3$ tightness [3, 4]. Here we significantly improve on the result for $\mathbb{L}_3$ of [4], and provide important new results for when $LP + \mathbb{L}_4$ is tight, employing an interesting geometric perspective. The main result is to show that the relationship which holds between forbidden minors characterizing treewidth and $\mathbb{L}_r$ tightness for $r = 2$ and $r = 3$ breaks down for $r = 4$, hence demonstrating that treewidth is not precisely the right condition for analyzing tightness of higher-order LP relaxation.

References

1. Wainwright, M., Jordan, M.: Treewidth-based conditions for exactness of the Sherali-Adams and Lasserre relaxations. Technical report, University of California, Berkeley, 671:4 (2004)
2. Weller, A., Tang, K., Sontag, D., Jebara, T.: Understanding the Bethe approximation: when and how can it go wrong? In: Uncertainty in Artificial Intelligence (UAI) (2014)

This is a summary of the paper "M. Rowland, A. Pacchiano, A.Weller. Conditions Beyond Treewidth for Tightness of Higher-order LP Relaxations. *AISTATS 2017*"

3. Weller, A.: Characterizing tightness of LP relaxations by forbidding signed minors. In: Uncertainty in Artificial Intelligence (UAI) (2016)
4. Weller, A., Rowland, M., Sontag, D.: Tightness of LP relaxations for almost balanced models. In: Artificial Intelligence and Statistics (AISTATS) (2016)

Contents

Technical Track

Application Track