

Principles of Constraint Programming

Krzysztof R. Apt

CWI, Amsterdam, The Netherlands



PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, UK
40 West 20th Street, New York, NY 10011-4211, USA
477 Williamstown Road, Port Melbourne, VIC 3207, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain
Dock House, The Waterfront, Cape Town 8001, South Africa

<http://www.cambridge.org>

© Cambridge University Press 2003

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2003

Printed in the United Kingdom at the University Press, Cambridge

Typeface Computer Modern 10/13pt *System* L^AT_EX 2

A catalogue record for this book is available from the British Library

ISBN 0 521 82583 0 hardback

Contents

<i>Acknowledgements</i>	<i>page</i> xi
1 <i>Introduction</i>	1
1.1 Basic characteristics of constraint programming	1
1.2 Applications of constraint programming	3
1.3 A very short history of the subject	5
1.4 Our approach	6
1.5 Organisation of the book	6
2 <i>Constraint satisfaction problems: examples</i>	8
2.1 Basic concepts	9
2.2 Constraint satisfaction problems on integers	11
2.3 Constraint satisfaction problems on reals	16
2.4 Boolean constraint satisfaction problems	19
2.5 Symbolic constraint satisfaction problems	21
2.6 Constrained optimization problems	43
2.7 Summary	47
2.8 Exercises	48
2.9 Bibliographic remarks	51
2.10 References	52
3 <i>Constraint programming in a nutshell</i>	54
3.1 Equivalence of CSPs	55
3.2 Basic framework for constraint programming	58
3.2.1 PREPROCESS	59
3.2.2 HAPPY	60
3.2.3 ATOMIC	61
3.2.4 SPLIT	61
3.2.5 PROCEED BY CASES	64
3.2.6 CONSTRAINT PROPAGATION	66

3.2.7	Constraint propagation algorithms	70
3.3	Example: Boolean constraints	71
3.4	Example: polynomial constraints on integer intervals	74
3.5	Summary	80
3.6	Bibliographic remarks	81
4	<i>Some complete constraint solvers</i>	82
4.1	A proof theoretical framework	83
4.1.1	Proof rules	84
4.1.2	Derivations	87
4.2	Term equations	92
4.2.1	Terms	93
4.2.2	Substitutions	94
4.2.3	Unifiers and mgus	95
4.2.4	Unification problem and solving of CSPs	98
4.2.5	The <i>UNIF</i> proof system	99
4.2.6	The MARTELLI–MONTANARI algorithm	103
4.3	Linear equations over reals	107
4.3.1	Linear expressions and linear equations	107
4.3.2	Substitutions, unifiers and mgus	110
4.3.3	Linear equations and CSPs	111
4.3.4	The <i>LIN</i> proof system	112
4.3.5	The GAUSS–JORDAN ELIMINATION algorithm	115
4.3.6	The GAUSSIAN ELIMINATION algorithm	118
4.4	Linear inequalities over reals	121
4.4.1	Syntax	121
4.4.2	Linear inequalities and CSPs	122
4.4.3	The <i>INEQ</i> proof system	123
4.4.4	The FOURIER–MOTZKIN ELIMINATION algorithm	124
4.5	Summary	131
4.6	Exercises	131
4.7	Bibliographic remarks	132
4.8	References	133
5	<i>Local consistency notions</i>	135
5.1	Node consistency	136
5.2	Arc consistency	138
5.3	Hyper-arc consistency	143
5.4	Directional arc consistency	144
5.5	Path consistency	147
5.6	Directional path consistency	155

5.7	k -consistency	157
5.8	Strong k -consistency	164
5.9	Relational consistency	166
5.10	Graphs and CSPs	170
5.11	Summary	175
5.12	Exercises	175
5.13	Bibliographic remarks	176
5.14	References	176
6	<i>Some incomplete constraint solvers</i>	178
6.1	A useful lemma	180
6.2	Equality and disequality constraints	181
6.3	Boolean constraints	184
6.3.1	Transformation rules	185
6.3.2	Domain reduction rules	186
6.3.3	Example: full adder circuit	188
6.3.4	A characterisation of the system <i>BOOL</i>	191
6.4	Linear constraints on integer intervals	192
6.4.1	Domain reduction rules for inequality constraints	194
6.4.2	Domain reduction rules for equality constraints	196
6.4.3	Rules for disequality constraints	199
6.4.4	Rules for strict inequality constraints	200
6.4.5	Shifting from intervals to finite domains	200
6.4.6	Example: the <i>SEND + MORE = MONEY</i> puzzle	201
6.4.7	Bounds consistency	202
6.4.8	A characterisation of the <i>LINEAR EQUALITY</i> rule	206
6.5	Arithmetic constraints on integer intervals	211
6.5.1	Domain reduction rules: first approach	211
6.5.2	Domain reduction rules: second approach	213
6.5.3	Domain reduction rules: third approach	217
6.5.4	Implementation of the third approach	221
6.5.5	Shifting from intervals to finite domains	223
6.6	Arithmetic constraints on reals	224
6.6.1	Interval arithmetic	226
6.6.2	Domain reduction rules	227
6.6.3	Implementation issues	233
6.6.4	Using floating-point intervals	236
6.6.5	Correctness and efficiency issues	238
6.7	Arithmetic equations over reals	242
6.8	Summary	245

6.9	Exercises	245
6.10	Bibliographic remarks	248
6.11	References	251
7	<i>Constraint propagation algorithms</i>	254
7.1	Generic iteration algorithms	256
7.1.1	Iterations	256
7.1.2	Algorithms for arbitrary partial orderings	261
7.1.3	Algorithms for cartesian products of partial orderings	264
7.2	From partial orderings to CSPs	268
7.3	A node consistency algorithm	269
7.4	An arc consistency algorithm	271
7.5	A hyper-arc consistency algorithm	273
7.6	A directional arc consistency algorithm	275
7.7	A path consistency algorithm	277
7.8	A directional path consistency algorithm	281
7.9	A k -consistency algorithm	283
7.10	A relational consistency algorithm	286
7.11	Implementations of incomplete constraint solvers	287
7.12	Summary	290
7.13	Exercises	291
7.14	Bibliographic remarks	295
7.15	References	297
8	<i>Search</i>	299
8.1	Search trees	301
8.2	Labeling trees	303
8.2.1	Complete labeling trees	304
8.2.2	Reduced labeling trees	308
8.2.3	<i>prop</i> labeling trees	310
8.3	An example: <i>SEND + MORE = MONEY</i>	313
8.4	Instances of <i>prop</i> labeling trees	315
8.4.1	Forward checking	315
8.4.2	Partial look ahead	319
8.4.3	Maintaining arc consistency (MAC)	321
8.5	Search algorithms for the labeling trees	324
8.5.1	Backtrack-free search	325
8.5.2	Backtrack-free search with constraint propagation	327
8.5.3	Backtracking	329
8.5.4	Backtracking with constraint propagation	330
8.6	Instances of backtracking with constraint propagation	332

8.6.1	Forward checking	332
8.6.2	Partial look ahead	333
8.6.3	Maintaining arc consistency (MAC)	334
8.6.4	Searching for all solutions	335
8.7	Search algorithms for finite constrained optimization problems	335
8.7.1	Branch and bound	337
8.7.2	Branch and bound with constraint propagation	339
8.7.3	Branch and bound with constraint propagation and cost constraint	339
8.8	Heuristics for search algorithms	341
8.8.1	Variable selection	341
8.8.2	Value selection	343
8.9	An abstract branch and bound algorithm	344
8.10	Summary	347
8.11	Exercises	347
8.12	Bibliographic remarks	348
8.13	References	349
9	<i>Issues in constraint programming</i>	351
9.1	Modeling	352
9.1.1	Choosing the right variables	352
9.1.2	Choosing the right constraints	353
9.1.3	Choosing the right representation	356
9.1.4	Global constraints	358
9.2	Constraint programming languages	359
9.2.1	Constraint logic programming	360
9.2.2	ILOG solver	362
9.2.3	Generation of constraints	363
9.3	Constraint propagation	364
9.4	Constraint solvers	366
9.4.1	Building constraint solvers	366
9.4.2	Incrementality	367
9.4.3	Simplification of constraints	368
9.5	Search	369
9.5.1	Search in modeling languages	369
9.5.2	Depth-first search: backtracking and branch and bound	370
9.5.3	Breadth-first search and limited discrepancy search	371
9.5.4	Local search	372
9.5.5	Search in constraint programming languages	375
9.5.6	Biology-inspired approaches	378

9.6	Over-constrained problems	379
9.6.1	Partial, weighted and fuzzy CSPs	380
9.6.2	Constraint hierarchies	381
9.6.3	Generalisations	383
9.6.4	Reified constraints	383
9.7	Summary	384
9.8	Bibliographic remarks	384
	<i>Bibliography</i>	387
	<i>Author index</i>	401
	<i>Subject index</i>	404

Introduction

1.1	Basic characteristics of constraint programming	1
1.2	Applications of constraint programming	3
1.3	A very short history of the subject	5
1.4	Our approach	6
1.5	Organisation of the book	6

1.1 Basic characteristics of constraint programming

THIS BOOK IS about *constraint programming*, an alternative approach to programming which relies on a combination of techniques that deal with *reasoning* and *computing*. It has been successfully applied in a number of fields including molecular biology, electrical engineering, operations research and numerical analysis. The central notion is that of a constraint. Informally, a *constraint* on a sequence of variables is a relation on their domains. It can be viewed as a requirement that states which combinations of values from the variable domains are admitted. In turn, a *constraint satisfaction problem* consists of a finite set of constraints, each on a subsequence of a given sequence of variables.

To solve a given problem by means of constraint programming we first formulate it as a constraint satisfaction problem. To this end we

- introduce some variables ranging over specific domains and constraints over these variables;
- choose some language in which the constraints are expressed (usually a small subset of first-order logic).

This part of the problem solving is called *modeling*. In general, more than

one representation of a problem as a constraint satisfaction problem exists. Then to solve the chosen representation we use either

- domain specific methods,

or

- general methods,

or a combination of both.

The *domain specific methods* are usually provided in the form of implementations special purpose algorithms. Typical examples are:

- a program that solves systems of linear equations,
- a package for linear programming,
- an implementation of the unification algorithm, a cornerstone of automated theorem proving.

In turn, the *general methods* are concerned with the ways of reducing the search space and with specific *search methods*. The algorithms that deal with the search space reduction are usually called *constraint propagation algorithms*, though several other names have been often used. These algorithms maintain equivalence while simplifying the considered problem. They achieve various forms of *local consistency* that attempt to approximate the notion of (global) consistency. The (top down) search methods combine various forms of constraint propagation with the customary back-track and branch and bound search.

The definition of constraint programming is so general that it embodies such diverse areas as Linear Algebra, Global Optimization, Linear and Integer Programming, etc. Therefore we should stress one essential point. If domain specific methods are available they should be applied *instead* of the general methods. For example, when dealing with systems of linear equations, the well-known linear algebra algorithms are readily available and it does not make sense to apply to these equations the general methods.

In fact, one of the aims of constraint programming is to look for efficient domain specific methods that can be used instead of the general methods and to incorporate them in a seamless way into a general framework. Such a framework usually supports

- domain specific methods by means of specialised packages, often called *constraint solvers*,
- general methods by means of various built-ins that in particular ensure or facilitate the use of the appropriate constraint propagation algorithms and support various search methods.

Once we represent a problem as a constraint satisfaction problem we need to solve it. In practice we are interested in:

- determining whether the chosen representation has a solution (is consistent),
- finding a solution, respectively, all solutions,
- finding an optimal solution, respectively, all optimal solutions w.r.t. some quality measure.

After this short preview we can formulate the following basic characteristics of constraint programming:

Two Phases Approach: The programming process consists of two phases: a generation of a problem representation by means of constraints and a solution of it. In practice, both phases consist of several smaller steps that can be interleaved.

Flexibility: The representation of a problem by means of constraints is very flexible because the constraints can be added, removed or modified. This flexibility is inherited by constraint programming.

Presence of Built-ins: To support this approach to programming several built-in methods are available. They deal with specific constraint solvers, constraint propagation algorithms and search methods.

An additional aspect brought in by constraint programming is that modeling by means of constraints leads to a representation of a problem by means of relations. This bears some resemblance to database systems, for instance relational databases. In fact, constraints are also studied in the context of database systems. They are useful in situations where some information, for instance the definition of a region of a map, needs to be provided implicitly, by means of constraints on reals.

The difference is that in the context of database systems the task consists of efficiently querying the considered relations, independently on whether they are defined explicitly (for instance by means of tables) or implicitly (for example by means of recursion or inequalities). In contrast, in constraint programming the considered relations are usually defined implicitly and the task consists of solving them or determining that no solution exists. This leads to different methods and different techniques.

1.2 Applications of constraint programming

Problems that can be best solved by means of constraint programming are usually those that can be naturally formulated in terms of requirements,

general properties, or laws, and for which domain specific methods lead to overly complex formalisations. Constraint programming has already been successfully applied in numerous domains including:

- interactive graphic systems (to express geometric coherence in the case of scene analysis),
- operations research problems (various optimization problems, in particular scheduling problems),
- molecular biology (DNA sequencing, construction of 3D models of proteins),
- business applications (option trading),
- electrical engineering (location of faults in the circuits, computing the circuit layouts, testing and verification of the design),
- numerical computation (solving polynomial constraints with guaranteed precision),
- natural language processing (construction of efficient parsers),
- computer algebra (solving and/or simplifying equations over various algebraic structures).

More recent applications of constraints involve generation of coherent music radio programs, software engineering applications (design recovery and code optimization) and selection and scheduling of observations performed by satellites. Also, constraint programming proved itself a viable approach to tackle certain computationally intractable problems.

While an account of most of these applications cannot be fit into an introductory book, like this one, an interested reader can easily study the research papers on the above topics, after having acquainted himself/herself with the methods explained in this book.

The growing importance of this area can be witnessed by the fact that there are now annual conferences and workshops on constraint programming and its applications that consistently attract more than one hundred (occasionally two hundred) participants. Further, in 1996 an (unfortunately expensive) journal called ‘Constraints’ was launched. Also, several special issues of computer science journals devoted to the subject of constraints have appeared. But the field is still young and only a couple of books on this subject have appeared so far. This led us to writing this book.

1.3 A very short history of the subject

Before we engage in our presentation of constraint programming, let us briefly summarise the history of this subject. It will allow us to better understand the direction the field is heading.

The concept of a constraint was used already in 1963 in an early work of I. Sutherland on an interactive drawing system SKETCHPAD. In the seventies various experimental languages were proposed that used the notion of constraints and relied on the concept of constraint solving.

The concept of a constraint satisfaction problem was also formulated in the seventies by researchers in the artificial intelligence (AI). They also identified the main notions of local consistency and the algorithms that allow us to achieve them. Independently, various search methods were defined. Some of them, like backtracking can be traced back to the nineteenth century, while others, like branch and bound, were defined in the context of combinatorial optimization. The contribution of constraint programming was to identify various new forms of search that combine the known techniques with various constraint propagation algorithms. Some specific combinations were already studied in the area of combinatorial optimization.

In the eighties the first constraint programming languages of importance were proposed and implemented. The most significant were the languages based on the logic programming paradigm. This led to a development of *constraint logic programming*, an extension of logic programming by the notion of constraints. The programming view that emerged led to an identification of *constraint store* as a central concept. Constraint propagation and various forms of search are usually available in these languages in the form of built-ins.

In the late eighties and the nineties a form of synthesis between these two developments took place. The researchers found various new applications of constraint programming, most notably in the fields of operations research and numerical analysis. The progress was often achieved by identifying important new types of constraints and new constraint propagation algorithms. One also realised that further progress may depend on a combination of techniques from AI, operations research, computer algebra and mathematical logic. This turned constraint programming into an interesting hybrid area, in which theoretical work is often driven by applications and in turn applications lead to new challenges concerning implementations of constraint programming.

1.4 Our approach

In our presentation of the basic concepts and techniques of constraint programming we strive at a streamlined presentation in which we clarify the nature of these techniques and their interrelationship. To this end we organised the presentation around a number of simple principles.

Principle 1: Constraint programming is about a formulation of the problem as a constraint satisfaction problem and about solving it by means of domain specific or general methods.

This explains our focus on the constraint satisfaction problems and constraint solvers.

Principle 2: Many constraint solvers can be naturally explained using a rule-based framework. The constraint solver consists then of a set of rules that specify its behaviour and a scheduler. This viewpoint stresses the connections between rule-based programming and constraint programming.

This explain our decision to specify the constraint solvers by means of proof rules that transform constraint satisfaction problems.

Principle 3: The constraint propagation algorithms can be naturally explained as instances of simple generic iteration algorithms.

This view allows us to clarify the nature of the constraint propagation algorithms. Also, it provides us with a natural method for implementing the discussed constraint solvers, since a rule scheduler is just another instance of a generic iteration algorithm.

Principle 4: (Top down) search techniques can be conceptually viewed as traversal algorithms of the search trees.

This explains why we organised the chapter on search around the slogan:

Search Algorithm = Search Tree + Traversal Algorithm,

and why we explained the resulting algorithms in the form of successive reformulations.

1.5 Organisation of the book

The above explained principles lead to a natural organisation of the material. Here is a short preview of the remaining chapters. In **Chapter 2** we discuss several examples of constraint satisfaction problems. We stress there that in many situations several natural representations are possible. In **Chapter 3** we introduce a general framework that allows us to explain the basics of constraints programming. We identify there natural ingredients of this

framework. This makes it easier to understand the subject of the subsequent chapters.

Then, in **Chapter 4**, we provide three well-known examples of complete constraint solvers. They deal, respectively, with solving equations over terms, linear equations over reals and linear inequalities over reals. In turn, in **Chapter 5** we introduce several notions of local consistency and characterise them in the form of proof rules. These notions allow us to study in **Chapter 6** in more detail a number of incomplete constraint solvers that involve Boolean constraints and linear and arithmetic constraints on integers and reals.

In **Chapter 7** we study the constraint propagation algorithms that allow us achieve the forms of local consistency discussed in Chapter 5. The characterisation of these notions in the form of proof rules allows us to provide a uniform presentation of these algorithms as instances of simple generic iteration algorithms. Next, in **Chapter 8**, we discuss various (top down) search algorithms. We present them in such a way that one can see how these algorithms are related to each other. Finally, in **Chapter 9**, we provide a short overview of the research directions in constraint programming.

Those interested in using this book for teaching may find it helpful to use the transparencies that can be downloaded from the following website: <http://www.cwi.nl/~apt/pcp>.