

1

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 explains 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 to 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>.

2

*Constraint satisfaction problems:
examples*

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

THE AIM OF this chapter is to discuss various examples of constraint satisfaction problems (CSPs ² in short). The notion of a CSP is very general, so it is not surprising that these examples cover a wide range of topics. We limit ourselves here to the examples of CSPs that are simple to explain and that illustrate the use of general methods of constraint programming. In particular, we included here some perennial puzzles, since, as it has been recognised for some time, they form an excellent vehicle to explain certain principles of constraint programming.

As already mentioned in Chapter 1 the representation of a problem as a CSP is usually called *modeling*. The selected examples clarify a number of aspects of modeling. First, as we shall see, some of the problems can be formalised as a CSP in a straightforward way. For other problems the appropriate representation as a CSP is by no means straightforward and relies on a non-trivial ‘background’ theory that ensures correctness of the

² For those knowledgeable in other areas of computer science: constraint satisfaction problems have nothing to do with Communicating Sequential Processes, a programming notation for distributed processes introduced by C.A.R. Hoare and also abbreviated to CSP.

adopted representation. Also for several problems, more than one natural representation exists.

When presenting the CSPs it is useful to classify them according to some criterion. In general, the techniques used to solve CSPs depend both on the domains over which they are defined and on the syntax of the used constraints. In most examples we use some simple language to define the constraints. Later, in Chapters 4 and 6, we shall be more precise and shall discuss in detail specific languages in which the constraints will be defined. But now it is too early to appreciate the role played by the syntax. So we rather classify the CSPs according to the domains over which they are defined. This explains the structure of this chapter.

First, we formalise in Section 2.1 the notion of a constraint and of a CSP. Then, in Section 2.2 we introduce some well-known problems and puzzles that can be naturally formalised as CSPs with integer domains. In Section 2.3 we consider examples of problems the formalisation of which leads to CSPs with variables ranging over reals. In turn, in Section 2.4 we consider *Boolean CSPs*. These are CSPs in which the variables range over the integer domain $[0..1]$ or, equivalently, $\{\text{false}, \text{true}\}$ and in which the constraints are expressed by means of Boolean expressions.

An important class of CSPs are the ones in which the variables range over non-numeric domains. We call them *symbolic CSPs*. They are considered in Section 2.5. In case we are interested in finding an optimal solution to a CSP we associate with each solution an objective function that we want to minimise or maximise. This leads to a modification of a CSP that we call a *constrained optimization problem*. They are considered in Section 2.6.

2.1 Basic concepts

As explained in the previous chapter constraint satisfaction problems, or CSPs, are a fundamental concept in constraint programming. To proceed we need to define them formally. The precise definition is completely straightforward. First we introduce the notion of a constraint.

Consider a finite sequence of variables $Y := y_1, \dots, y_k$ where $k > 0$, with respective domains $D_1, \dots, D_k$ associated with them. So each variable y_i ranges over the domain D_i . By a *constraint* C on Y we mean a subset of $D_1 \times \dots \times D_k$. When $k = 1$ we say that the constraint is *unary* and when $k = 2$ that the constraint is *binary*. By a *constraint satisfaction problem*, or a *CSP*, we mean a finite sequence of variables $X := x_1, \dots, x_n$ with respective domains $D_1, \dots, D_n$, together with a finite set $\mathcal{C}$ of constraints, each on a subsequence of X . We write such a CSP as $\langle \mathcal{C} ; \mathcal{DE} \rangle$, where

10 **Constraint satisfaction problems: examples**

$\mathcal{DE} := x_1 \in D_1, \dots, x_n \in D_n$ and call each construct of the form $x \in D$ a **domain expression**. To simplify the notation we omit the ‘{ }’ brackets when presenting specific sets of constraints $\mathcal{C}$.

We now define the crucial notion of a solution to a CSP. Intuitively, a solution to a CSP is a sequence of legal values for all of its variables such that all its constraints are satisfied. More precisely, consider a CSP $\langle \mathcal{C} ; \mathcal{DE} \rangle$ with $\mathcal{DE} := x_1 \in D_1, \dots, x_n \in D_n$. We say that an n -tuple $(d_1, \dots, d_n) \in D_1 \times \dots \times D_n$ **satisfies** a constraint $C \in \mathcal{C}$ on the variables $x_{i_1}, \dots, x_{i_m}$ if

$$(d_{i_1}, \dots, d_{i_m}) \in C.$$

Then we say that an n -tuple $(d_1, \dots, d_n) \in D_1 \times \dots \times D_n$ is a **solution** to $\langle \mathcal{C} ; \mathcal{DE} \rangle$ if it satisfies every constraint $C \in \mathcal{C}$. If a CSP has a solution, we say that it is **consistent** and otherwise we say that it is **inconsistent**.

Note that in the definition of a constraint and of a CSP no syntax was assumed. In practice, of course, one needs to define the constraints and the domain expressions. In what follows we assume that they are defined in some specific, further unspecified, language. In this representation it is implicit that each constraint is a subset of the Cartesian product of the associated variable domains. For example, if we consider the CSP $\langle x < y ; x \in [0..10], y \in [5..10] \rangle$, then we view the constraint $x < y$ as the set of all pairs (a, b) with $a \in [0..10]$ and $b \in [5..10]$ such that $a < b$.

Let us illustrate these concepts by a simple example. Consider the sequence of four variables x, y, z, u ranging over natural numbers and the following three constraints on them: $x^3 + y^3 + z^3 + u^3 = 100$, $x < u$, $x + y = z$. According to the above notation we write this CSP as

$$\langle x^3 + y^3 + z^3 + u^3 = 100, x < u, x + y = z ; x \in \mathcal{N}, y \in \mathcal{N}, z \in \mathcal{N}, u \in \mathcal{N} \rangle,$$

where $\mathcal{N}$ denotes the set of natural numbers.

Then the sequence $(1, 2, 3, 4)$ is a solution to this CSP since this sequence satisfies all constraints. Indeed, we have $1^3 + 2^3 + 3^3 + 4^3 = 100$, $1 < 4$ and $1 + 2 = 3$.

Finally, let us clarify one simple matter. When defining constraints and CSPs we refer to the sequences (respectively subsequences) of variables and *not* to the sets (respectively subsets) of variables. Namely, given a CSP each of its constraints is defined on a *subsequence* and not on a *subset* of its variables. In particular, the above constraint $x < y$ is defined on the subsequence x, y of the sequence x, y, z, u .

Also, the sequence z, y is not a subsequence x, y, z, u , so if we add to the above CSP the constraint $z = y + 2$ we cannot consider it as a constraint on z, y . But we can view it of course as a constraint on y, z and, if we wish,