Towards 40 years of constraint reasoning

Pedro Meseguer

Received: 1 September 2011 / Accepted: 25 October 2011 / Published online: 9 February 2012
© Springer-Verlag 2012

Abstract Research on constraints started in the early 1970s. We are approaching 40 years since the beginning of this successful field, and it is an opportunity to revise what has been reached. This paper is a personal view of the accomplishments in this field. We summarize the main achievements along three dimensions: constraint solving, modelling and programming. We devote special attention to constraint solving, covering popular topics such as search, inference (especially arc consistency), combination of search and inference, symmetry exploitation, global constraints and extensions to the classical model. For space reasons, several topics have been deliberately omitted.

Keywords Constraint reasoning · Constraint solving · Constraint modelling · Constraint programming

1 Introduction

In 1970, Richard Fikes published a paper—partially based on his recently presented PhD thesis—on how to solve some problems using constraint satisfaction [35]. A bit later, David Waltz, a MIT PhD student working in vision introduced some processing based on constraints to understand the sense of edges in 3D images [95]. In 1974, Ugo Montanari published a paper [71], setting the basic concepts of constraint processing. In 1977, Alan Mackworth presented several algorithms to achieve basic forms of local consistency [62], some of which are still in use. These publications can be seen as the

seminal works setting the basis of Constraint Reasoning, a new field at that time which has been very flourishing and successful until today. After the initial years, where these topics were mainly in academic workshops, this new field matured producing many publications in major AI conferences, having its own conferences and setting its own journals. In parallel, constraints were introduced in the field of logic programming, producing a new approach named constraint logic programming. Some companies started selling software based on constraints. The term *constraint programming* was coined, independently of the underlying programming style. Continuous improvements in constraints solving algorithms have caused substantial improvements in performance, while different extensions to the basic constraint model have significantly enlarged its expressivity. Now, we are closer to 40 years of this starting point, and it is meaningful to revise what has been done in this field. This is the purpose of this paper, which is necessarily short and leaves many topics untouched. The interested reader may consult some of the books existing on this topic [1,28,50,63,78,89] for a more detailed historical perspective, as well as for a much broader and deeper description of its different aspects.

This new field took constraints as its central element; a *constraint* on a set of variables permits some value combinations while forbids others (assuming a finite number of possible assignments for variables). The task of interest was (and remains) finding an assignment of values to variables satisfying all constraints. Many problems of combinatorial nature can be formulated in terms of a finite number of variables and constraints, abstracting details that are application dependent. This caused a lot of interest, because many real-world problems could benefit from this new technology (for instance, planning and scheduling, time-tabling, vehicle routing, configuration, and different applications in networking and bioinformatics, to name a few).

Partially supported by the Spanish project TIN2009-13591-C02-02 and Generalitat de Catalunya grant 2009-SGR-1434.

P. Meseguer (✉)
IIIA - CSIC, Campus Universitat Autònoma de Barcelona,
08193 Bellaterra, Spain
e-mail: pedro@iiia.csic.es

In this paper, we report the main accomplishments in Constraint Reasoning along three dimensions: constraint solving, constraint modelling and constraint programming. We devote substantially more space to constraint solving because it is the topic on which more work has been done by the research community working on constraints. We describe popular topics such as search (and in particular the well-known backtracking algorithm), inference (complete and incomplete, in this last class falls the arc-consistency idea), combinations of search and incomplete inference (forward checking and MAC algorithms), symmetry exploitation, global constraints and extensions to the classical constraint network model.

A personal view in these almost 40 years of history considers three consecutive periods, which do not have sharp boundaries but they are recognizable considering the dominant research topics in the community. They are:

- Initially, researchers were mainly interested in search procedures (including heuristics to improve search performance) in the context of binary constraints. Another main topic of interest was local consistency, mainly considered as preprocessing before search.
- Interests evolved towards more sophisticated search procedures, many of them maintaining some form of local consistency during search. Constraints were no longer assumed to be binary, and global constraints appeared. The role of inference methods was properly understood. Constraint modelling started to attract attention. Initial constraint programming environments were available.
- In the last decade, we have seen considerable performance improvements on symmetry exploitation, new global constraints and decompositions. Some extensions to the classical model have been proposed: continuous domains, soft constraints, distributed environments, etc. Connections with related fields, such as Operations Research or SAT, have been clarified. Constraint modelling and programming have reached some maturity.

The rest of this paper is organized as follows. First, we define the problem in Sect. 2, providing some classical examples. Second, we summarize the main solving approaches, grouped in the three main families of search, inference and hybrids (combining search and inference) in Sect. 3. Algorithmic improvements and extensions to the classical model are also described in Sect. 3. Third, we consider the issue of constraint modelling, i.e., given a real-world problem producing a constraint model for it in Sect. 4. Fourth, we describe different approaches to implement the constraint model into a constraint programming environment in Sect. 5. Finally, we conclude the paper in Sect. 6, mentioning those topics which have been deliberately omitted. Because space limitations, descriptions are necessarily short.

2 The problem

A *constraint network* is defined by a triplet $\langle \mathcal{X}, \mathcal{D}, \mathcal{C} \rangle$, where $\mathcal{X} = \{x_1, x_2, \dots, x_n\}$ is a set of n variables taking values in a collection of finite and discrete domains $\mathcal{D} = \{D_1, D_2, \dots, D_n\}$, such that x_1 takes value in D_1 , x_2 takes value in D_2 and so on, under a set of constraints $\mathcal{C} = \{C_1, C_2, \dots, C_e\}$. Given a constraint $C \in \mathcal{C}$, $var(C)$ is the sequence of variables involved in the constraint (also called its *scope*, or in some cases *schema*), while $rel(C)$ is the set of value tuples on $var(C)$ which are permitted by the constraint, $rel(C) \subset \prod_{x_i \in var(C)} D_i$. Often, $rel(C)$ is not given in extension, but providing a function that determines whether a tuple is permitted or not by C . The number of variables involved in C , $|var(C)|$, is the arity of C . If the arity of C is 1, 2, ... we say that C is a unary, binary, ... constraint. A constraint network is binary if all its constraints are unary or binary.

Often, it is implicitly assumed that there may exist only a single constraint among a particular set of variables. Constraint networks with this property are said to be normalized. In the following, we always assume normalized constraint networks. A unary constraint on x_i is written as C_i , and a binary constraint between x_i and x_j is written as C_{ij} .

Several operations can be done on a constraint network. One of these possible tasks is to assign a value to each variable from its domain such that *all* constraints are satisfied. This task is called *solving* the network, which is also referred to as the *constraint satisfaction problem* (CSP). Most of the algorithms presented here consider the task of solving a constraint network; their complexities are given in terms of n the number of variables, d the maximum domain size and e the number of constraints. This classical CSP definition assumes *hard* constraints: they are mandatory and any solution must satisfy all of them (in contrast with the *soft* constraint concept, see Sect. 3.5).

A particular case of CSP is the SAT problem, defined as finding a consistent assignment for a propositional formula expressed in clausal form satisfying the formula. A SAT instance can be seen as a CSP as follows: SAT variables correspond to CSP variables with domains $\{true, false\}$, and each clause corresponds to a constraint requiring that at least one literal (affirmed or negated variable) of the clause has to be assigned to *true*.

A constraint network $\langle \mathcal{X}, \mathcal{D}, \mathcal{C} \rangle$ can be visualized using a graph representation. There are three popular representations:

- Primal graph: each node represents a different variable, and there is an edge between two nodes if the corresponding variables appear together in the scope of some constraint.
- Dual graph: each node represents a different constraint, and there is an edge between two nodes if the

Fig. 1 *Left* the map formed by Norway, Sweden and Finland as instance of the map coloring problem. *Right* primal graph, dual graph and hypergraph of: (up) this instance formalized with binary constraints; (down) this instance formalized with a single *all-different* constraint

$$N \neq S \neq F \neq N$$

$$\text{all-different}(N, S, F)$$

corresponding constraints share a common variable in their scopes.

- **Hypergraph:** each node represents a different variable and each hyperedge represents a different constraint. The hyperedge corresponding to constraint c connects all variables in $\text{var}(c)$. When the constraint network contains unary and binary constraints only, the hypergraph coincides with the primal graph.

If, in addition to the constraint network $\langle \mathcal{X}, \mathcal{D}, \mathcal{C} \rangle$, there is an objective function on a subset of the variables, a popular task is to optimize that function subject to the constraints. This task is usually called the *constraint optimization problem* (also known by the acronym COP). About complexity, it is well-known that these problems are intractable in general: classical CSP is NP-complete (it is enough to realize that 3-SAT, a famous NP-complete problem, can be seen as a CSP), while COP is NP-hard (specifically, Δ_2 -complete).

As examples, we mention two classical ones: the *map coloring problem* and the *n-queens problem*. The map coloring problem consists of assigning colors from a given set to regions in a map, such that two adjacent regions have different colors. This problem can be easily expressed as a constraint network: variables are regions, values are colors, all variables have the same domain, constraints are $\neq$ between variables corresponding to adjacent regions. The map of Norway, Sweden and Finland is a nice example of map coloring. It appears in Fig. 1, with its corresponding graphs.

The *n-queens problem* consists of locating n queens in a $n \times n$ chessboard such that no queens are attacking each other. This problem admits several formalizations in terms of constraints. One of them considers variables representing queens: there is one variable per row (obviously, two queens cannot be located at the same row), the values are the available columns (initially, from 1 to n), under the constraints that two variables cannot be in the same column and in the same diagonal of the chessboard. This formalization appears in Fig. 2.

3 Constraint solving

Given a constraint network $\langle \mathcal{X}, \mathcal{D}, \mathcal{C} \rangle$, the task of interest finding an assignment of values to variables such that all constraints are satisfied (= solving the network). This task is known as the CSP. This can be done by several methods; we describe next those which we consider more relevant.

In this Section, several algorithms assume binary constraints. This is done for the easiness of presentation, nowadays constraint solving approaches have been generalized to handle constraints of any arity.

3.1 Search

Conceptually, the simplest way to solve a CSP is by searching the state space of the problem. This is done by guessing values for the problem variables. There are two main families of search methods: systematic search and local search. Both are described next.

3.1.1 Systematic search

Since the number of variables is finite, and the number of values in their domains is also finite, the state space—although it may be very large—is finite, so it can be traversed systematically until exhaustion. This is the rationale of systematic search for solving constraint networks. The search space is

Fig. 2 The 4-queens problem as CSP: variables represent *rows* and values are the available *columns* of the chessboard. The constraint graph is a clique of binary constraints

Fig. 3 Search tree of the 4-queens problem, following formalization of Fig. 2

represented as a tree, where each variable x_i is associated with a level in the tree, and each node at this tree level has d_i successors at the next level, one for each possible value of x_i ($d_i = |D_i|$). Each tree node corresponds to one assignment of variables (the one specified by the unique path from this node to the root); interior nodes correspond to partial assignments, while leaf nodes represent complete assignments. The search tree for the 4-queens problem appears in Fig. 3.

The simplest algorithm one could think of follows a *generate-and-test* schema. It generates all search tree leaves (or equivalently, all complete assignments involving all variables), testing the constraints on each of them. Any assignment satisfying all constraints is a solution. This algorithm is simple but extremely inefficient because it may test an exponential number of assignments which are inconsistent due to the same reason. For instance, consider the 4-queens problem. We know that any assignment with $x_1 = 1, x_2 = 1$ is inconsistent, because the queens in the first two rows are in the same column. However, this algorithm will generate all combinations of values for queens in the third and fourth rows, to finally conclude that there is no solution below.

The *backtracking* (BT) algorithm [15,48] solves some of the issues of generate-and-test. It performs a depth-first search traversal of the search tree, checking at each node whether all constraints among assigned variables are satisfied or not. If the answer is yes, BT continues its depth-first traversal. If the answer is no, BT prunes the subtree below the current node and continues depth-first in the rest of the tree. The justification for pruning the subtree rooted at the current node is simple: if at that node there is a constraint among assigned variables which is violated, it will remain violated in all the nodes of the subtree below the current node. There will not be any solution below the current node, so there is no point searching such subtree. The spatial and temporal complexities of BT are those of depth-first search: exponential in time but linear in space.

BT is a chronological backtracker, i.e., when it finds a dead-end it changes the last decision taken. Sometimes, this strategy is not effective because the last decision taken may

Fig. 4 A state of the 8-queens problem. There is no consistent value for x_6 but backtracking to x_5 does not release any consistent value for x_6 . Intelligent backtracking jumps directly to x_4

not be the culprit for the dead-end. As an example, consider the instance of the 8-queens problem depicted in Fig. 4. The last decision taken—the value assigned to x_5 —is not the culprit for not having a consistent value for variable x_6 . *Dependency-directed backtracking* [88] and the most general *intelligent backtracking* [17] solve this issue, but requiring exponential memory in its most general form. Keeping linear memory, several restricted forms of intelligent backtracking for constraint search have been developed under the names of *backjumping* [42], *graph-based-backjumping* [25], *conflict-directed-backjumping* [74] (CBJ). The idea of CBJ in the context of binary constraint networks is as follows: each variable keeps a conflict set, where for each value it records the first assigned variable which is incompatible with that value. When the search algorithm finds a dead-end, it backtracks to the closest variable in its conflict set (instead of backtracking to the closest variable, as BT will do). Upon backtracking, the conflict set of the new current variable is enlarged with the conflict set of the old current variable minus the new current variable. Intermediate variables between the jumping variable and the arrival variable are deassigned.

A further step in this line is the *dynamic backtracking* (DB) algorithm [45]. It is an intelligent backtracking algorithm which requires linear memory, while keeping some elaborated information about the culprits for constraint violations. When facing a dead-end, DB backtracks to the closest variable responsible for the conflict, keeping the values of the variables between the jumping variable and the arrival variable. This more sophisticated backtracking may have also some caveats regarding efficiency [4].

A last word for *variable and value ordering heuristics*: Associating statically one variable with each tree level, so this variable always appears at the same position when BT traverses one branch down, is more than needed to assure BT completeness. What is required is that each variable has to be considered in each branch, but the order in which the same variable appears may differ from branch to branch. This observation gives rise to *dynamic variable ordering* (also called DVO) heuristics, trying to select dynamically

the variable which shortens most the remaining search time. Solving a constraint network, most of the time is spent traversing branches which do not lead to any solution; in such case, a good strategy is to detect as soon as possible when there is no solution below the current node to perform pruning and jump to another branch. This is the basis of the *first-fail* principle: select the variable that is most likely to fail [49]. A simple implementation of this principle is selecting as the next variable the one with less consistent values with previous instantiations (called *minimum remaining values* [3] or simply *min-domain* heuristic [87]; often it is presented in connection with constraint propagation). Assuming the task of finding one solution, a similar reasoning can be done on value ordering as follows. When generating the successors of a node, i.e., the possible assignments for the variable associated with that node, any value ordering is legal. BT will stop after finding one solution that necessarily will contain all assigned variables. Since a value has to be chosen for the current variable, a reasonable heuristic is to start with those values which are most likely to succeed. BT starts searching successors from left, so values should be ordered left to right from most to less likely to succeed.

3.1.2 Local search

While systematic search procedures keep somehow memory of where they are in the search space, in order to perform an exhaustive traversal, local search procedures only keep information local to the current state. Standard local search procedures cannot guarantee that they will find a solution if there is one, or that a particular state will not be visited more than once. While lacking these theoretical guarantees, local search procedures usually scale up better than systematic ones; because of that, local search is sometimes preferred when the size of the problem to solve is large, often in the context of constraint optimization.

There are a wide set of search strategies: evolutionary algorithms [47], simulated annealing¹ [55], tabu search [46], just to name a few. Often they are named under the umbrella term of *metaheuristics* [16]. All of them consider three basic elements (here we consider states² s with all variables instantiated):

- Objective function: $F(s)$ it maps each state s to a cost given by $F(s)$; the solution is the global minimum of such function.

¹ In some communities, simulated annealing is considered as a global search strategy. Here it is included into the local search class as a non-systematic search method.

² Following the concept of state space in Artificial Intelligence, a *state* is a particular configuration of variables; in some cases, some variables might not be assigned (as in the interior nodes of the search tree of Fig. 3).

- Neighborhood: $N(s)$ set of states s' where you can go from the current state s in the next iteration. Typically, s' differs from s in the values of i variables (typically i is between 1 and 2).
- Selection criterion: given s , $F(s)$, $N(s)$, how to choose the next state s' .

A local search procedure considers the greedy minimization of function $F(s)$ following an iterative approach. At each iteration, after moving to the current state s , it computes the values of $F(s)$ and $N(s)$. It moves from state s to s' if $F(s) \geq F(s')$, starting a new iteration. However, this procedure may get trapped in local minima, where $F(s') \geq F(s)$, $\forall s' \in N(s)$. In some cases, this issue is tackled with restarts or random moves (since in the next state $F(s')$ may be worse than $F(s)$, these moves are called *worsening moves*). This combination is named *stochastic local search* (SLS) [52]. A SLS procedure ends when it finds a solution or when exhausting computational resources (such as a maximum number of iterations). There is a large number of combinations of SLS methods for CSP solving. Next, we summarize some of the most successful ideas in SLS:

- Min Conflicts Heuristic [68]: $F(s)$ counts the number of not satisfied constraints (conflicts) in s ; the next state s' tries to minimize the number of violated constraints, while changing the value of one variable only.
- Penalty Methods [24, 72, 92]: each constraint has a weight and $F(s)$ includes the sum of weights of violated constraints; weights are not constant but they change during search (for instance, in [72] when the system reaches a local minimum, the weights of violated constraints are incremented).
- GSAT [83]: this is a procedure to solve SAT instances, a particular case of CSP. Given a boolean formula in clausal form and a complete variable assignment (chosen initially at random), at each step GSAT chooses a boolean variable and flips its value (from *true* to *false* or vice versa) such that the number of unsatisfied clauses is minimized. If a solution has not been found, GSAT restarts every *maxSteps* flips. The whole process iterates a maximum of *maxTries*.
- Random Walk [65, 84]: instead of performing worsening moves only when the system is stuck in a local minimum, a worsening move can also occur in non-local minima with a small probability.
- Large Neighborhood Search [18]: when computing $N(s)$ some parts of the current state s (assumed feasible) are maintained while others are varied. Then, $N(s)$ includes states that differ in more than one or two values from s . Using consistency techniques, $N(s)$ is pruned, keeping its size manageable.

Local search appears as a promising strategy for solving constraint-based problems [51].

3.2 Inference

Processing a constraint network by *inference* means to perform some legal operations on constraints, such that original constraints are modified and a new constraint network naturally replaces the original one. The new network has exactly the same solutions as the original one. The interesting point is that the new constraint network is more explicit than the original one, and this allows traversing its search space more efficiently when looking for a solution. In this sense, we say that new network is *easier to solve* than the original.

Inference methods are divided into two large families: incomplete and complete inference. The main difference between them is that complete inference methods produce directly the solution after their execution, while this is not true for incomplete inference methods. The most relevant methods of both families are described next.

3.2.1 Incomplete inference

Incomplete inference is a umbrella term for a number of procedures working on constraints such that, when applied on any constraint network, although reducing the search space of the original network, they do not produce directly the solution which has to be found by subsequent search. Procedures manipulating and simplifying constraints have been developed since the very beginning of Constraint Reasoning,³ first as preprocessing and later combined with search.

Incomplete inference is also known as local consistency. An essential part of it is constraint propagation, when changes in a part of the problem (such as assigning a variable or removing values in a domain) are propagated to other parts of the problem because the action of constraints (typically, by enforcing some necessary condition for solutions). When the effect of local consistency is just removing values from domains, it is also called domain filtering.

ARC CONSISTENCY. A form of incomplete inference is arc consistency. Basically, it consists of analyzing each constraint in turn, to detect elementary values of variables that are not permitted by the constraint. These values are removed from its domain (because they will not be in any solution), and their removal is propagated through other constraints. This process continues until finding that no further changes are induced in the network. Then, we say that we have reached a *fix point*.

Let us consider a binary constraint network, where constraints correspond to arcs in the primal constraint graph

(C_{ij} connects x_i and x_j). Instead of considering the whole constraint network, we consider each subnetwork of 2 variables connected by a constraint. Analyzing such constraint, it is possible to detect if some values will not be in a solution. Then, these values are removed and the effect of their removal is propagated to other domains by the rest of constraints. If a variable becomes without values (empty domain), the whole network has no solution; in this case, the “no-solution” condition has been found inspecting subnetworks of size 2 only. Otherwise, if some values have been removed, the search space of the instance becomes smaller and search can be done more efficiently.

More formally, constraint C_{ij} is directionally arc consistent (from x_i to x_j) iff for any value $a \in D_i$ there exists a value $b \in D_j$ such that $(a, b) \in \text{rel}(C_{ij})$. C_{ij} is arc consistent iff it is directionally arc consistent in both directions. A constraint network is arc consistent iff all its constraints are arc consistent [28].

If for a particular $a' \in D_i$ there is no value in D_j consistent with it, a' can be removed from D_i because it will not be in any solution (the assignment $x_i = a'$ will violate at least one constraint, C_{ij}). The removal of a' is propagated by all other constraints involving x_i . Let us consider C_{ik} . If it was previously made arc consistent, a' removal may cause C_{ik} to become directionally arc-inconsistent, from x_k to x_i , while in the other direction it remains arc consistent. So C_{ik} has to be rechecked from x_k to x_i . An example of arc-consistency processing appears in Fig. 5.

A number of algorithms have been proposed for enforcing arc consistency on a binary constraint network. Among them, we mention AC-3 [62], the optimal AC-4 [70] and AC-2001 [13]. Arc consistency can be achieved in a polynomial number of steps (optimal temporal complexity $O(ed^2)$, e is the number of binary constraints, d is the largest domain size).

Originally, arc consistency was developed for binary constraint networks, when the hypergraph coincides with the primal graph. When considering constraints of any arity, the same concept is called *generalized arc consistency* (GAC). A constraint C is GAC iff for any $x_p \in \text{var}(C)$, a value $a \in D_p$ appears in some permitted value tuple $t \in \text{rel}(C)$ as value for x_p . A constraint network is GAC if all its constraints are GAC [28].

Less demanding than arc consistency is *bound consistency* (BC), applicable to totally ordered domains. Only the lower/upper bounds of each domain are checked in each constraint using a weaker form of consistency that assumes the domains of other variables compact and without holes, perfectly defined by their lower/upper bounds (there are some variants on the definition of bound consistency, see [9]).

Arc consistency is one of the topics on which more work has been done in the context of constraint processing. For a comprehensive study, see [9].

³ The early work of Waltz in vision [95] considers arc consistency on line edges.

Fig. 5 Example of arc-consistency processing. *Left* clique of three variables, connected by binary constraints; permitted value pairs are connected by a link (*microstructure*). *Right* crossed values are elimi-

nated from their corresponding domains because they are arc inconsistent; the problem solution is the assignment of value c to each variable

PATH-CONSISTENCY. Assuming again binary constraint networks, let us consider subnetworks of three variables, say x_i, x_j, x_k . If a pair of consistent values between x_i and x_j cannot be consistently extended including x_k , the initial pair is considered inconsistent and recorded in a modified C_{ij} . If at some point a constraint becomes empty (= no pair of values is permitted), the network has no solution.

Formally, a pair of variables (x_i, x_j) is path-consistent iff for any pair of consistent values $(a, b) \in \text{rel}(C_{ij})$ and any third variable x_k , there exists $c \in D_k$ such that $(a, c) \in \text{rel}(C_{ik})$ and $(b, c) \in \text{rel}(C_{jk})$. A constraint network is path consistent iff any pair of variables is path-consistent [28].

As in the arc-consistency case, if constraints are given in extension path consistency can be achieved in polynomial time (the optimal complexity is $O(n^3 d^3)$ [70]). However, the complexity is substantially higher than in the arc-consistency case, so usually it does not pay-off when compared with the simpler but cheaper arc-consistency property (especially, when local consistencies are enforced at each node visited by a search procedure, as described in Sect. 3.3). In addition, path-consistency modifies the constraints, which is more difficult to handle than the easier task of modifying the domains, the only effect of arc consistency. Because of all these, path consistency remains unattractive in practice, and it is usually not available in most of the existing constraint solvers.

K-CONSISTENCY. K -consistency is a generalization of the above concepts. A binary constraint network is k -consistent iff given a subset of $k - 1$ variables consistently assigned, for any k th variable it is possible to find a value consistent with the values of the $k - 1$ variables previously assigned [36]. Then, arc consistency is 2-consistency, path-consistency is 3-consistency, etc. Optimal k -consistency can be achieved in time exponential in k [22]. A constraint network is k -strong consistent, when it is p -consistent for $1 \leq p \leq k$. An interesting result is the following: given a constraint network where variables are ordered such that the width of the order (= the maximum number of edges that go back from a

node in the primal constraint graph) is lower than k , a constraint network that is strong k -consistent is backtrack-free, i.e., BT will find a solution without performing backtracking [37].

This result tempts us to compute the width of the primal graph (and its associated ordering), processing the network afterwards to achieve the required level of strong consistency. However, the procedure for strong consistency adds edges to the primal graph, so the width of the selected ordering may change after executing this procedure (except when the primal graph is a tree, then it can be solved backtrack-free after processing it by directional arc consistency).

Interestingly, the *adaptive consistency* approach [29] was a step forward along this line. First, it was able to adjust the level of consistency required for each variable of the ordered primal graph to assure a backtrack-free search. And second, it introduced the concept of induced width, as the width of the ordered graph after adaptive consistency processing. But all this is better explained as bucket elimination in the next paragraphs.

3.2.2 Complete inference

Complete inference procedures, when applied to any constraint network, produce directly a solution so no subsequent backtrack search is needed. Let us assume that all constraints are given in extension. A *tuple* t with scope S is an ordered set of values assigned to each variable in $S \subseteq \mathcal{X}$. The *projection* $t[T]$ of a tuple t on $T \subset S$ is a new tuple which only includes the values assigned to the variables in T . We define the join $\bowtie$ of two constraints C and C' , written $C \bowtie C'$, as a new constraint with scope $\text{var}(C) \cup \text{var}(C')$ such that $t \in \text{rel}(C \bowtie C')$ if and only if $t[\text{var}(C)] \in \text{rel}(C) \wedge t[\text{var}(C')] \in \text{rel}(C')$. The join of constraints is associative and commutative.

By definition, we observe that the tuples satisfying constraints C and C' satisfy $C \bowtie C'$, and vice versa. Then, a direct way to compute all solutions of a constraint network

is computing the join of all its original constraints, $\bowtie_{C \in \mathcal{C}} C$. Tuples satisfying $\bowtie_{C \in \mathcal{C}} C$ are the solutions of the original network, and there are no other solutions.

Since solving a constraint network is NP-complete, it causes no surprise that complete inference has temporal complexity exponential in the *induced width*, w^* , a parameter of the primal constraint graph that basically measures its cyclicity.

BUCKET ELIMINATION. As mentioned before, the join of all the original constraints will produce a single constraint whose permitted tuples will be the solutions of the problem. However, this is more than needed to solve the problem by inference. If we do the joins of constraints following some particular order, we can save space and time. To properly describe the bucket elimination algorithm, we still need to define a new operation: given a constraint C , projecting a variable $x \in \text{var}(C)$ out of C , written $C[-x]$, is a new constraint with scope $\text{var}(C) - \{x\}$ whose allowed tuples are those of C eliminating the values of x and removing duplicates.

Before describing the bucket elimination algorithm [27], we introduce some concepts from [37]. A *ordered graph* (G, o) is simply a graph G with a total order o of its nodes. Given (G, o) , the *width of a node* x is the number of arcs that connect x with other nodes prior x in the order o ; the *width of an order* o is the maximum width of its nodes; the *width of a graph* is the minimum width of all its possible orderings. Given (G, o) , where G is a primal constraint graph, we define the *induced graph* G' along the ordering o [29] as G enlarged with some extra arcs computed as follows: nodes are processed from last to first along o ; let x be a node, all its parents (nodes connected with x and before x in o) are connected by a clique of arcs. Given the induced graph G' , previous definitions of width associated with node, order or graph also apply here. In particular, the parameter called induced width w^* is the width of the induced graph G' . Computing the minimum induced width of a graph is NP-complete [28].

The bucket elimination algorithm works as follows: given an ordered graph (G, o) of a constraint network, the bucket of each variable is the set of constraints which mention that variable as the last variable in their scope following o . Buckets are processed following o from last to first. If x is the current variable, all constraints in bucket B_x are combined together into a single constraint C_x which summarizes the effect of B_x . Then, since variable x is not further required, it is projected out from C_x , obtaining the constraint R_x which replaces B_x in the set of constraints $\mathcal{C}$. This new constraint R_x —which does not mention x —is properly located in the bucket of the last variable of its scope, and it will be processed in the future.

Once all buckets have been processed, one knows if the original problem has solutions. To recover solutions, variables are processed along o —from first to last—assigning variable x a value of a tuple allowed by C_x and consistent

with all previous variables in o . It is easy to see that such value always exists, and the process is backtrack-free. The spatial and temporal complexities of bucket elimination are exponential in the induced width of the constraint graph [27].

MINI-BUCKETS. It may happen that the system has not enough memory to compute $C_i \leftarrow \bowtie_{C \in B_i} C$ (or you want to be faster than computing this exponential-time step). To alleviate these exponential bounds, the *mini-bucket* approximation comes into play. The idea is to partition bucket B_i in subsets such that the join of all constraints in each subset produces a new constraint whose arity is upper bounded by r , a parameter that establishes the maximum memory available to keep the new constraint (at most d^r tuples) [26]. Mini-bucket elimination is an approximation of the exact bucket elimination algorithm, it is not guaranteed to find a true solution (although the final assignment will be closer to a solution). The spatial and temporal complexities of mini-bucket elimination are exponential in r . Mini-buckets can also be used to generate heuristics for optimization tasks [54].

It is possible to use search as an alternative to the mini-buckets approximation. Given an ordering o of variables to be processed by bucket elimination, and r an upper bound on the arity of the new constraints, this approach works as follows: if processing the bucket of a variable along o produces a new constraint of arity $\leq r$, then the standard bucket elimination algorithm is used; otherwise, a search algorithm is used, branching on the possible values for this variable [59].

3.3 Hybrid algorithms

Incomplete inference (also called local consistency) procedures can be used either (i) as preprocessing, before search starts, to simplify the original instance, or (ii) inside the search process, typically at each visited node. In the second case, we talk about *hybrid algorithms* because they combine search (typically systematic search, in its simplest form the BT algorithm) with some form of incomplete inference. This combination often offers significant savings if compared with pure search, and it has become a kind of standard when solving constraint networks. The benefits of this combination are due to the reduction that local consistency causes in the search space of subproblems found by search, causing to visit less nodes. On the other hand, some extra work is done when enforcing the local consistency at each visited node (because of that, local consistencies of low complexity are preferred). It has been observed experimentally that the best trade-off between benefits and extra work occurs for arc consistency (or lower local consistencies). Because of that, we will concentrate our presentation on hybrid algorithms to the combination of BT with some forms of arc consistency.

Although hybrid algorithms often offer better performance, their worst case temporal complexity does not improve over the BT one, which is exponential. An interesting work

Fig. 6 FC search tree (left) and MAC search tree (right) on the 4-queens problem. Observe the more powerful MAC inference in the left part of the search tree. Light cells are forbidden, dark cells are free

compares the relative performance of some hybrid algorithms [57].

3.3.1 Forward checking

Originally proposed for solving binary constraint networks, the *forward checking* (FC) algorithm [49], follows a BT schema with the following addition: when a variable is assigned, constraints connecting this variable with unassigned ones are made arc consistent. This may cause deletions in the domains of these variables; if one of these domains becomes empty, there is no solution below the current node so the subtree below it can be pruned, saving search effort. Upon backtracking, when the value assigned to a variable changes, the deletions caused by enforcing arc consistency with the old value for the variable that has to be undone. The search tree developed by FC in the 4-queens problem appears in Fig. 6.

In FC, it is worth noting that when a variable is assigned a value, no checking for consistency is done on already assigned variables because all values remaining in the domains of unassigned variables are already consistent (because of FC behavior). The FC algorithm has been extended to deal with non-binary constraints in [12].

3.3.2 Maintaining arc consistency

The *Maintaining arc consistency* (MAC) algorithm [79] represents a step further in the application of arc consistency inside a systematic search schema. The main idea of this MAC is that arc consistency must be maintained at all times in the current subproblem considering all its constraints (not only constraints connecting the current variable with unassigned ones, as FC). In addition, the subproblem where the current variable is still unassigned must be arc consistent. This is done as follows. Let x be the current variable to be assigned. When this variable takes value a , arc consistency

is enforced in the remaining subproblem. If some domain becomes empty, there is no solution below the current node (observe that this step is stronger than the corresponding step in FC, because arc consistency is enforced on more constraints), so the subtree below it can be pruned causing search savings. Otherwise, a recursive call is done indicating which is the next variable. If no solution can be found under the assignment $x \leftarrow a$, then a is removed from D_x and arc consistency is enforced on the new subproblem. This step is new with respect to FC. As in the previous algorithm, upon backtracking when the value assigned to a variable changes, the deletions caused by enforcing arc consistency with the old value for that variable have to be undone. The search tree developed by MAC in the 4-queens problem appears in Fig. 6.

An original view of the MAC algorithm is that of traversing a binary tree where left branches correspond to assignments of variables ($x = a$) and right branches correspond to rebuttal of the same values (adding the constraint $x \neq a$). An example of such binary tree on the 4-queens problem appears in Fig. 7.

Fig. 7 MAC binary tree on the 4-queens problem. Left branches are assignments, right branches are value rebuttals (only rebuttal $x_1 \neq 1$ appears). Light cells are forbidden, dark cells are free

3.4 Algorithmic enhancements

3.4.1 Symmetry exploitation

In many constraint networks there are symmetries as transformations involving variables and values that leave the problem invariant. For instance, in the 4-queens problem following Fig. 2 formulation, if we exchange x_1 with x_4 and x_2 with x_3 , we obtain exactly the same problem (this is called *variable symmetry*). Alternatively, if we exchange columns 1 and 4, and columns 2 and 3, again the problem remains the same (this is called *value symmetry*). The importance of symmetry comes from this simple observation: if a state is solution, all its symmetric states are also solutions (and the same occurs with non-solutions). Therefore, it is enough to visit one state of all states related by a symmetry (ignoring the symmetry these states look different, but considering the symmetry these states belong to the same equivalence class—equivalence relation induced by the symmetry—and it is enough to visit one representative of such class). In some cases, this causes enormous reductions in the size of the space to explore, producing very substantial gains in search efficiency. Because of that, symmetries have received some attention in constraint solving.

Variable and value symmetries have been known for a long time, but a general definition of symmetry in Constraint Reasoning has been only recently provided. In [21], authors propose (1) *solution symmetry*: a permutation of variables and values that leaves the set of solutions invariant, (2) *constraint symmetry*: a permutation of variables and values that leaves the set of constraints invariant (so basically the problem does not change). Both types of symmetry allow variable/value symmetries as special cases. Interestingly, constraint symmetry is included in solution symmetry; in fact, there are many more solution symmetries than constraint symmetries. Programmers detect usually constraint symmetries only because without the set of solutions, it is difficult to identify solution symmetries that are not constraint symmetries [21].

Perhaps the simplest way to deal with symmetries is by model reformulation, looking for a model with a lower number of symmetries. Unfortunately, this technique is really case dependent, and requires a deep knowledge of the problem at hand (see [43] for some examples). When the problem presents a number of indistinguishable variables, typically we numbered them, introducing artificially symmetries that do not exist in the original problem. A way to solve this is to use set variables (whose value is a set instead of a single value). Set variables break some variable symmetries and represent more naturally set solutions.

An approach to deal with symmetries is to add some constraints before search starts to break symmetries. For instance, imagine that your problem has ten variables which

are fully symmetric, $x_1, x_2, \dots, x_{10}$ (for any pair of variables, you can freely exchange them). A way to deal with this symmetry is to add constraints requiring a lexicographical order among them: $x_1 \leq x_2 \leq x_3 \leq \dots \leq x_{10}$. This approach is called *lex-leader* [23]. This strategy requires a static variable ordering compatible with the lex-leader ordering. In some cases, a matrix of variables appears in constraint programming; if rows are interchangeable and columns are interchangeable, lex-leader constraints can be imposed in rows and columns, producing what is called *double-lex* approach [56]. Also lexicographic ordering over multisets has been investigated.

Alternatively, one can break symmetries by adding some constraints dynamically during search. One approach following this line is Symmetry Breaking During Search (SBDS) [44]. The basic idea is, upon backtracking, to add constraints that prevent visit states that are symmetrical to those already visited. Another approach—somehow related—is called Symmetry Breaking by Dominance Detection (SBDD) [33]. The basic idea is, when going to visit a new node of the search tree, check if this node is symmetrically equivalent to an already visited one. To do so, it is not needed to store all visited nodes; it is enough to store the roots of fully explored subtrees. Unlike adding symmetry breaking constraints before search, both SBDS and SBDD can be used with DVO heuristics and they can deal with symmetries involving values.

In some cases, the number of symmetries could be large (even exponential). Adding static or dynamically new constraints to break these symmetries could be counterproductive at some point. Checking extra constraints causes some overhead, to be compared with the benefits of symmetry breaking. In practice, it is commonly assumed that the best efficiency is achieved breaking a only part of the symmetries present in the problem. For a more in-depth presentation, see [43].

3.4.2 Global constraints

MOTIVATING EXAMPLE. A good way to introduce global constraints is with an example. Let us consider the Sudoku puzzle: it is a grid of 9×9 cells, which contains 9 non-intersecting subgrids of size 3×3 . The goal is to assign to each cell a digit from 1 to 9, such that the following constraints are satisfied:

- there is no repeated digit in each row of the 9×9 grid,
- there is no repeated digit in each column of the 9×9 grid,
- there is no repeated digit in each 3×3 subgrid.

While finding a legal assignment for an empty Sudoku is trivial, the problem usually considered is the Sudoku completion: when some cells have already a digit assigned and

Fig. 8 The Sudoku completion problem: assign a digit from 1, . . . , 9 to each *empty cell* such that the assignment of the whole grid forms a legal Sudoku (as explained in the *text*)

7			5	8	3			6
	6			1	4	5		
	5	2			6		8	3
3			2			9	5	8
5				7	8		6	
6	4	8		1		3		
	6		8		2	5		
			3	1	5			7
2	1	5	6				3	

the task is to find an assignment of digits to the other cells such that the whole assignment forms a legal Sudoku. An example of Sudoku completion appears in Fig. 8.

A natural way to model this problem as a CSP is to consider each cell as a variable. Empty cells have a domain $\{1, 2, 3, 4, 5, 6, 7, 8, 9\}$, while cells with a digit assigned have a domain formed by that digit, under the above mentioned constraints. It is not difficult to see that Sudoku constraints can be expressed as cliques of binary inequality constraints among the variables involved. Thus, there are 36 binary inequalities per row, 36 binary inequalities per column and 36 binary inequalities per 3×3 subgrid.

Alternatively, we can consider the *all-different* constraint on a set of variables *SetVar*, requiring all variables of *SetVar* to be assigned a different value. Using this constraint, we can model the Sudoku completion problem with one *all-different* constraint per row, one per column and one per 3×3 subgrid.

It is clear that when a variable x is assigned a value v , that value has to be removed from the domains of the other variables constrained with x . In addition, the number of different available values (cardinality of the union of the domains) for the 9 constrained variables (corresponding to the same row, column or 3×3 subgrid) should be 9; otherwise, there is no way to satisfy the corresponding Sudoku constraint. The first reasoning is achieved when enforcing GAC in both formalizations, clique of binary inequalities and *all-different* constraints. The second reasoning is only achieved when enforcing GAC on *all-different* constraints; a binary inequality does not include all the variables involved in a Sudoku constraint, so enforcing GAC on it cannot account for the number of their different available values.

Why *all-different* has more pruning power than the clique of binary constraints? The reason is based on the set of variables viewed by each constraint. The *all-different* constraint considers 9 variables simultaneously, so it is able to make better deductions than when considering pairs of variables in isolation. Having a more *global* view of the problem allows for better deductions, which cannot be achieved with a more local view (which in occasions could be myopic). An example of this appears in Fig. 9: a map coloring instance with 3 variables constrained to have different values but only two available values.

It is worth noting that both formalizations for the Sudoku completion problem are correct and obtain true solutions.

Fig. 9 Map coloring instance, with three variables but only two available colors formulated with an *alldifferent* constraint (left) and its binary decomposition into a clique of binary inequalities (right). While GAC on the binary decomposition deduces nothing, GAC on *alldifferent* detects that there is no solution

However, assuming a solving strategy based on search, one is more efficient than the other because it is able to detect (and therefore prune) sooner than the other branches not leading to any solution.

GLOBAL CONSTRAINTS. A *global constraint* captures the relation among a non-fixed number of variables with specific semantics [90]. C is a class of constraints defined by a Boolean function f_C whose arity is not fixed. Constraints $C(x_{i_1}, \dots, x_{i_k})$, $C(x_{j_1}, \dots, x_{j_q})$ with different arities can be defined by the same Boolean function. For instance, *alldifferent*(x_1, x_2, x_3) and *alldifferent*(x_4, x_5, x_6, x_7) are two instances of the *alldifferent* global constraint, where $f_{\text{alldifferent}}(T)$ returns true iff $x_i \neq x_j, \forall x_i, x_j \in T$.

Examples of global constraints are:

- *alldifferent*: different values have to be assigned to a set of variables;
- *atmost/atleast*: an upper/lower bound on the number of times a value may appear assigned in a set of variables;
- *cardinality* (also written *gcc*): it establishes lower and upper bounds to the number of times that several values may appear assigned in a set of variables;
- *sum*: the addition of the numerical values of a set of variables should be equal to some parameter;
- *circuit*: it is satisfied when a set of variables are assigned nodes of a directed graph such that they form a Hamiltonian circuit;
- *lexicographic*: considering two vectors of variables, it is satisfied when the two vectors are lexicographically ordered after assignment.

A global constraint is defined by a boolean function, no matter the arity of the constraint. This permits to develop GAC enforcing programs for this particular global constraint in which the number of variables involved in the constraint is an input parameter. Exploiting the semantic of this global constraint, those GAC enforcing programs can be made specific for it, with lower complexity than the corresponding to a generic GAC propagator. As an example, consider the *alldifferent* constraint on k variables; a generic GAC propagator requires to pass through all the tuples so it has a com-

plexity $O(d^k)$, while the specialized GAC propagator [77] can do it in $O(k^2 d^2)$. The exploitation of this fact has been crucial for the success of the constraint technology in practice (see Sect. 5).

Some global constraints are *semantically decomposable* in a polynomial number of simpler constraints of fixed arity, in the sense that the set of solutions of the global constraint is exactly equal to the set of solutions of a polynomial number of simpler constraints on the same set of variables [10]. An example of this is the *alldifferent* constraint: *alldifferent*(x_1, x_2, x_3) is semantically equivalent to $x_1 \neq x_2, x_1 \neq x_3$ and $x_2 \neq x_3$. However, there are constraints that cannot be decomposed into simpler constraints without adding a polynomial number of extra variables. An example of this is the *atleast* constraint. Allowing extra variables, any constraint is decomposable in smaller constraints.

A global constraint is *operationally decomposable* into smaller constraints of fixed arity, if enforcing GAC on the global constraint removes the same values as enforcing GAC on the set of smaller constraints. A global constraint may be semantically decomposable into a set of smaller constraints but not operationally decomposable in such set, in the sense that GAC on the global constraint is more powerful (removes more values) than GAC on the decomposition. For example, *alldifferent* is decomposable in a clique of binary inequalities; however, GAC on *alldifferent* prunes more than binary AC on the set of inequalities (see Fig. 9). So it is interesting to find decompositions which are as powerful in GAC propagation as the initial global constraint. Some constraints can do that on the same set of variables as the original constraint (e.g., the *lex* constraint). Others require a polynomial number of constraints on extra variables (for instance, the *atleast* constraint). And others cannot achieve the same propagation in any polynomial number of simpler constraints of fixed arity (e.g., the *alldifferent* constraint) [11].

Regarding GAC propagation, there are several classes of global constraints:

- those for which GAC propagation is NP-hard (the *sum* constraint),
- those for which GAC propagation is polynomial and
 - there is no equivalent decomposition, even allowing a polynomial number of extra variables (the *alldifferent* constraint),
 - there is an equivalent decomposition (the *lex* constraint on the same variables, the *atleast* constraint adding a polynomial number of extra variables).

The interest in decomposing global constraints comes from the fact that currently there are hundreds of reported global constraints (the current version of the global constraints catalog collects more than 350 global constraints,

with a description of more than 2,800 pages [6]). In such situation, it is of great interest finding if there is a short set of basic constraints such that any global constraint could be seen as a combination of such basic constraints (this would facilitate enormously the implementation of CP solvers). Although some successful work in this direction has been done, some global constraints escape the intended decomposition on basic constraints.

3.5 Extensions

3.5.1 Continuous constraints

The defined model of constraint network is extended to continuous domains when variable domains are intervals over the real numbers, and constraints are equations on these variables. Again, a solution is an assignment of values to variables satisfying the constraints.

This topic falls in the numerical analysis field. A common approach is to replace numerical analysis by interval analysis, where floating-point numbers are replaced by intervals. As in the discrete case, it is of great interest to discard parts of the search space that will not contain any solution. The concept of arc consistency from discrete domains has been successfully extended to the continuous case [34]. Given the computational difficulties to compute arc consistent intervals, a simpler form of local consistency called *hull* consistency (the version of bound consistency for continuous domains) is pursued. In fact, an approximation called *box* consistency, is achieved in practice [7].

Solving algorithms usually follow a branch-and-reduce schema. The idea is to obtain a set of interval boxes (Cartesian products of variable intervals) covering all solutions of the problem. The algorithm chooses (often using heuristics) a box and splits it, following a branching procedure. Resulting boxes are reduced by enforcing some kind of local consistency. The process iterates until boxes are small enough [8].

In general, solutions are found applying numerical methods (Newton method, etc.). The role of constraints is more focused to remove parts of the space without solutions, so numerical methods could be applied more efficiently. Hybrid approaches have been proposed, trying to combine different strategies. For a comprehensive description see [8].

3.5.2 Soft constraints

A constraint is *soft* if an assignment violating it can still be considered a possible solution. Classical constraints are *hard*: they must necessarily be satisfied by any solution. Typically, hard constraints represent physical properties that must necessarily be satisfied, while soft constraints represent priorities or preferences, which might be violated without invalidating a solution.

There are different frameworks for dealing with soft constraints. To unify their representation, two generic frameworks have been proposed: valued CSP [80] and semiring-based CSP [14], the latter being more general than the former (essentially, they differ in the required ordering for the intermediate satisfaction values: total order for valued CSP, partial order for semiring-based CSP).

Assuming the weighted CSP model where constraints are replaced by cost functions (they return a cost that one has to pay if the instantiating tuple is in the solution), the goal is to find the assignment that minimizes the addition of costs. Solving a weighted CSP becomes an optimization problem, typically solved by branch-and-bound methods. Different lower bounds have been proposed for branch-and-bound pruning [67].

Several soft local consistencies have been proposed. Interestingly, the arc consistency in the classical hard case splits in several levels in the soft case [60]. Maintaining soft local consistency during branch-and-bound search allows to compute a lower bound. Some soft versions of global hard constraints have also been successfully used for solving soft CSP instances [91]. Inference methods in CSP solving (such as bucket elimination, or cluster-tree elimination) can also be successfully used for computing the global optimum/assignment of a weighted CSP [28].

3.5.3 Distributed constraints

A CSP is distributed (DisCSP) if its parts (variables and constraints) are distributed among different agents and cannot be grouped into a single agent, often because of privacy reasons. In this setting, each agent knows a part of the problem, but no agent knows the whole problem. The solution satisfying all constraints is reached through message passing.

The first complete algorithm for DisCSP solving was Asynchronous Backtracking (ABT) [97]. It is a distributed search algorithm, able to perform when the problem has no solution; when a solution is found the network remains silent, which is detected by external specialized algorithms. Initially designed with a static variable ordering, ABT enhanced with DVO has been proposed [98]. Among other approaches, we mention the AAS and AFC algorithms [66,85]. In the context of local search (complete but with exponential memory), there is the Asynchronous Weak Commitment algorithm [97].

Replacing constraints by cost functions (as in the weighted CSP model), the so-called distributed constraint optimization problem (DCOP) has received a lot of attention lately. Now, the solution minimizes the addition of costs. The first complete algorithm for optimal DCOP solving was ADOPT [69], a distributed algorithm following a best-first strategy. Replacing best-first by depth-first with branch-and-bound, we obtain the BnB-ADOPT [96] which reaches better per-

formance. These are distributed search algorithms, exchanging short but many messages (in the worst case, exponentially many). Alternatively, the DPOP algorithm [73] performs distributed inference, requiring a linear number of messages but they can be exponentially large.

4 Constraint modelling

Modelling is the process of passing from a real-world problem into a constraint network. In principle, it is not very difficult to develop a model of a real-world problem. What could be difficult is to produce a model that solves the problem in a reasonable time. There is wide evidence that the modelling process has a dramatic effect in how easy the problem can be solved in practice [39,76]. Despite some high level languages for constraint specification, modelling resists automatization and remains an art (sometimes a *black art*) in which constraint programmers spend their time and energy looking for efficient models of real-world problems.

Here practical experience is a grade. In a report on learned lessons on crosswords [5], authors indicate how the model, the search algorithm and the heuristics interact, and none of these decisions can be made totally independent of the others. Nevertheless, in the following, we will assume that the solution is found by a combination of search and incomplete inference (interleaving search with constraint propagation is the standard solving approach), ignoring particular algorithms and heuristics (otherwise it would be very difficult to deal with all the possible combinations of these elements).

Perhaps a way to start when assessing two constraint models is to answer two questions (i) what is the size of the search space? and (ii) which model prunes more? Let us consider three different models of the n -queens problem:

1. One boolean variable per cell: n^2 variables, equal domains $\{0, 1\}$, constraints: different rows, columns, diagonals;
2. One variable per row: n variables, equal domains $\{1, \dots, n\}$, constraints: different columns, diagonals,
3. One variable per row + *alldifferent*: n variables, equal domains $\{1, \dots, n\}$, constraints: one *alldifferent* ($x_1, \dots, x_n$), different diagonals.

The search space size is: (1) 2^{n^2} , (2) and (3) n^n , which is far smaller than 2^{n^2} . Considering constraints, model (2) includes by construction the constraint “no two queens in the same row”, so it is preferred to model (1). Model (3) captures in a single global constraint that “no two queens in the same column”, which prunes more than the clique of binary inequalities when enforcing GAC, so (3) is preferred to (2). Therefore, simple modelling criteria are preferring smaller

Fig. 10 A state of the 6-queens problem, from [86]

search spaces and constraints with higher pruning power. As a rule of thumb, we prefer models using as few constraints as possible, as long as these constraints have efficient propagation algorithms.

Existing CP solvers maintain a different level of local consistency in each constraints, looking for a reasonable trade-off between the cost of enforcing such local consistency and the search reduction it produces. Typically, AC is enforced on binary constraints, GAC is enforced on global constraints and BC is enforced on arithmetic constraints, although some variants exist. Complex constraints may have little propagation (details depend on the particular solver used).

The first election in a model is what is called a *viewpoint* [61] that basically selects the set of variables and their domains. From this selection, constraints must assure that model solutions correspond to problem solutions. We prefer those viewpoints which allow for an easy and concise way to express the required constraints.

Given two constraints C_1 and C_2 to be satisfied simultaneously, a new constraint can be made as the *conjunction* of these constraints $C_1 \wedge C_2$. It will be tighter than any individual constraint of the conjunction. It may cause more pruning (however, this would not necessarily speed up solving in current solvers). An example of this effect appears in Fig. 10, where value 3 for x_5 would be deleted after AC if between any two queens there is one constraint “different columns and different diagonals” but it will remain if between any two queens there are two constraints “different columns” and “different diagonals”.

The use of *global constraints* has been shown very useful in practice. On one hand, global constraints have a more global view of the problem, so typically they can filter more (=remove more values) than their binary decompositions (when such decomposition are possible) when GAC is enforced on them. On the other hand, exploiting their semantics, global constraint propagators can achieve GAC more efficiently than using generic propagators. Current CP solvers provide a library of global constraints, which permit programmers to use them without implementing from scratch.

A constraint is *reified* when there is a 0/1 variable associated with that constraint such that the variable takes the value 1 when the constraint is satisfied and value 0 otherwise.

Some CP solvers offer reified constraints, which can be used to model disjunctions: if x_1 and x_2 are associated with constraints C_1 and C_2 , the expression $x_1 + x_2 \geq 1$ indicated that at least one constraint has to be satisfied.

Having the minimum set of constraints that allows for correct solutions is not always correlated with good performance. *Redundant* constraints, in the sense that they are implied from the basic constraints defining the problem, are useful when they discard some assignments that would be permitted by that minimum set of constraints. Assignments discarded by redundant constraints would also be found unfeasible by the basic constraints, but after some search that could be saved if redundant constraints are used [30]. The interaction between redundant constraints and variable ordering heuristics is relevant, to avoid that the above-mentioned assignments are never found during search. Again, a deep knowledge of the problem to solve is advisable.

Redundancy can also occur on variables. *Auxiliary* variables are variables which are not strictly needed to model the problem, but their introduction allow constraints to propagate better (or to express better the problem). Although they might be considered not as important as original problem variables, they are variables in its own right and can be used to guide search.

If the chosen viewpoint does not satisfy our expectations, one can look for another viewpoint in a *reformulation* process. Classical examples of reformulation are the *dual* and the *hidden variable* transformations [2], passing from non-binary formulations into binary ones (useful in the early days). An interesting class is the so-called *permutation* problems, which have as many variables as values and solutions are permutations of such values (although not every permutation is a solution). The n -queens problem is an example of this class. Permutation problems have two standard viewpoints, that exchange the roles of variables and values (in the n -queens problem, exchanging the roles of rows—variables—by columns—values). In some cases, constraints are more easily expressed in one viewpoint than in the other. Finally, an extreme case of reformulation is to use a SAT encoding of the considered problem, to be solved by a SAT solver [93].

A powerful mechanism in modelling is to connect several (generally two) views on a problem (viewpoints), in such a way that the pruning done in one view is propagated to another view by the so-called *channeling constraints*. These constraints connect the viewpoints, making available in a viewpoint what has been deduced in the other and vice versa. Viewpoints are mutually redundant, in the sense that one would suffice to achieve a correct solution, but combining them a solution can be found faster [20]. Interestingly, when two viewpoints are combined, some constraints of one (or both) viewpoints can be dropped without affecting correctness and improving performance [86]. This approach

of combining viewpoints can be extended to more than two [31].

The presence of *symmetry* can also be detected and exploited in modelling. Sometimes symmetry is introduced in the modelling phase, differentiating among indistinguishable variables (i.e., in the n -queens problem, we may have a variable for the 1st queen, another for the 2nd queen, ... which is not necessary because all what we know is a set of n -queens). Reformulation and the use of set variables are ways to avoid such traps. Nevertheless, there are problems with an inherent degree of symmetry, so to avoid it some of the methods described in the Sect. 3.4 should be used. In particular, if we add symmetry breaking constraints before search, we can use them to generate redundant constraints [41].

In many cases, there is a function f to optimize (say minimize) under the set of constraints $\mathcal{C}$. If f returns integers, a naive approach is to solve a sequence of CSPs with a constraint requiring f lower than or equal to a parameter which decreases in steps of one ($f \leq k$, $f \leq k - 1$, $f \leq k - 2$ and so on). Given the first instance of the sequence without solution, the minimum of f and the solution are those of the previous instance in the sequence. Alternatively, these problems can be solved by a branch-and-bound algorithm, which each time a new solution is found with $f = k$, a new constraint is added requiring $f < k$ (so only solutions with lower values of f are considered).

Some work has been done on generating automatically several models from problem specifications, even linking them by channeling constraints [86]. The identification of constraint patterns [94], as structures that repeatedly appear in constraint networks, seems a promising modelling direction. Once these patterns are well understood, they can be successfully reused through applications.

All these ideas have been successfully used in some applications, but no idea is universally successful [40]. The number of variables, the number of constraints and the quality of their propagation seem to have a crucial impact in the amount of search needed to solve a problem. However, we should have in mind that reducing search does not always reduce run-time (because of the different overheads of the solver). A deep knowledge of the problem to solve seems to be the best advice, to be able to generate different viewpoints, redundant constraints and auxiliary variables. Some of these elements could be finally used in the problem model if they really enhance efficiency [86].

5 Constraint programming

Programming is the process of passing from a constraint network, modelling a real-world problem into its effective implementation in some of the existing CP environments.

Obviously, you may start from scratch, programming everything from the very beginning, but nowadays this is not needed. Existing CP environments offer a number of facilities to declare variables, domains and constraints to represent your problem, to specify how search should be done and the required level of propagation (often per constraint). In addition, a library of constraints is available, many of them global offering specialized propagators already implemented. The existence of several CP environments has been crucial for the success of constraint technology. In the following, we describe some aspects of existing CP environments for finite domains (i.e., we do not consider systems that work with numerical constraints and interval domains).

It is worth noting that constraint programming follows a declarative approach [38, 75]. Typically, you declare the model of your problem in terms of variables, domains and constraints. Once everything is declared, you ask for solving, but the actual solving process is not programmed by the user. Solving is done by the CP solver, which in most cases is used as a black box that returns a solution or failure. Often, a function to be optimized can be added.

A CP environment provides facilities to define variables, domains and constraints. Variables and domains are internally represented in a rather direct way: variables are associated with domains, which are represented as sets of integers. Regarding constraints, by default they are not explicitly represented as collections of permitted tuples (to avoid memory problems and to maintain the constraint semantics), but by a collection of propagators (also known as filters or narrowing operators). A *propagator* is a function that maps domains into domains and it must be *contracting* (only remove values from the domain) and *monotonic* (if one domain is smaller than another one, this relation remains after the action of the propagator). A collection of propagators *implements* a constraint C iff they are *correct* (they do not remove assignments permitted by C) and they are *checking* C (when all variables involved in C are instantiated, the propagator must check if the tuple is permitted by C or not, reporting failure in this last case). The iterated application of propagators of the constraint set computes a reduction of domains which is the greatest mutual fixpoint of all propagators (this iterated application includes the execution of propagators for which the domains of its variables have been modified by the action of other propagators).

With this structure, the underlying search procedure (since propagators implement incomplete inference, some search is needed to achieve a solution) can be seen as a two-step procedure: (i) the application of a set of propagators until finding a fixpoint, and (ii) if no empty domain is detected, a variable is chosen on which a number of excluding alternatives are considered; each alternative is posted as a new constraint in turn, and their corresponding set of propagators is executed, applying again point (i). Step (i) is known as propagation, while

step (ii) is branching (or search). The state where branching is performed is known as *choice point*. Instead of developing a search tree from variable assignments, this approach develops the search tree by adding constraints on variables (of which assigning a value is a particular case). In this sense, this approach generalizes the classical search tree of Fig. 3. After reaching a fixpoint, it is of crucial importance to repeat the execution of a propagator only when the domains of the corresponding variables have changed. This is monitored by *events*, which signal changes in the domains and trigger the corresponding propagator. A number of improvements and simplifications have been suggested inside this strategy that combines search and propagation. They are clearly described in [81].

In the execution of the propagation step, propagators are in a list to be executed. Usually, this list is ordered placing low complexity propagators first (although the importance of the events triggering propagators is also considered). Anyway, propagator starvation (a propagator in the list waiting an unbounded amount of time to be executed) should be avoided. Although propagators have been defined as functions, in practice they have some internal, private state that has to be stored between executions. This state is used to achieve extra functionalities (for instance, incrementality) that improve performance.

Among propagators, we mention indexicals (especially for historical reasons). An *indexical* is a propagator for a simple constraint of the form *variable is in range*. Finite domain arithmetic constraints are compiled in terms of a set of indexicals. For instance $x \leq y$ is compiled into the indexicals $x \text{ in } -\infty \dots \min(y)$ and $y \text{ in } \min(x) \dots \infty$.

The common problem of domain representation has two popular formats: range sequences and bit vectors. A range sequence is a sequence of intervals such that the domain is exactly the union of such intervals. A bit vector is formed by a set of consecutive words in memory, such that the bit in position i is set to 1 iff the i th value is in the domain. Large domains are usually represented by range sequences.

A function that CP environments repeat many times is restoration, especially applied to domains. Different techniques have been tried, such as copy, trailing and recomputation. While copy is the simplest strategy, it has the problem of memory usage. Trailing records the changes done, so they can be undone if needed. Recomputation considers repeating the execution to achieve the desired state. Obviously, this last option requires almost no extra memory, but it may last longer than the others. Nowadays, trailing is the prevailing strategy in most CP environments.

Current CP environments are divided into two main classes:

- Autonomous systems: are environments for specific programming languages which include the elements for finite

domain constraint solving. Examples of such autonomous systems are Sicstus Prolog [19] and COMET [32].

- Library systems: are libraries which provide users of existing programming languages, such as C++ or Java, with the facilities to declare and work with variables, domains and constraints. Examples of library systems are Choco [58], Gecode [82] and IBM ILOG CPLEX CP Optimization Studio [53].

We also mention MiniZinc [64], an autonomous high level CP environment. MiniZinc programs are first compiled into another CP language called FlatZinc. From this language, they can be interfaced with different CP backend solvers. Here, we only mention some CP environments, but there are many more. For a comprehensive list, see [81].

6 Summary

Trying to summarize all the work done in Constraint Reasoning in the last 40 years in a few pages is simply impossible. We have presented in an informal—but hopefully precise—way the most relevant achievements in this field, according to the view and formation of the author.

Most of the work described in this paper follows an algorithmic approach for solving constraint networks. We have presented search as the basic solving strategy. In systematic search, the simple BT algorithm enhanced with some intelligent backtracking techniques appears as the kernel element for constraint solving. Also the basic blocks of local search have been described. Inference appears as an alternative way to solve constraint networks. Incomplete inference (also known as constraint propagation), specifically GAC or weaker forms of local consistency, has shown idoneous to be hybridized with systematic search strategies to achieve what is today the standard way of solving constraint networks which could be summarized in the expression *search + propagation*.

The inclusion of global constraints has been a huge step forward in achieving solutions to practical problems in reasonable time. Also symmetry exploitation has been crucial to avoid useless work in many problems showing some kind of symmetry. Relaxing some of the assumptions of the initial constraint network model, it has been extended in several dimensions generating extensions such as continuous constraints, soft constraints and distributed constraints.

Modelling, as the task of translating a real-world problem into a constraint network that could be solved in a reasonable time, appears as a difficult, hand-made task that resists automatization. Currently, a number of good practices and rules of thumb are known but a systematic way to attack modelling remains to be done. In the last part of the paper, we have described some of the elements of CP environments.

There is a wide variety of them, indicating that the engineering dimension of Constraint Reasoning is in good health.

However, there are a number of other topics in Constraint Reasoning that we have not talked about, which deserve some mention. Restarts in search procedures have received some attention, and different branching strategies to develop search trees have been studied. Domain consistencies, a class of local consistencies removing values from domains (wider than arc consistency), have been shown useful in a number of cases. Complexity has been a recurrent topic, considering both the problem topology and the internal form of constraints, looking for tractability islands. In many cases, there are close connections with Operation Research (especially when optimization is involved). A substantial amount of work has been devoted to constraint languages, from the early days of logic programming to current CP environments. Another classical topic has been the study of dynamic CSPs, when one has to solve a sequence of CSP instances, each differing little from the previous one. Further extensions to the classical problem such as unbounded domains (open CSPs) and variables with structure (for instance set variables) have been analyzed. Lately, the introduction of universal quantifiers in the definition of a solution of a constraint network has been studied.

Finally, we have not talked about applications. Constraints fit quite well in the field of temporal reasoning. Planning and scheduling has been, since the very beginning, a natural field for constraint applications. In addition, there are relevant constraint applications in vehicle routing, networking and bioinformatics.

Acknowledgments The author sincerely thanks Christian Bessiere, Christian Bliek, Javier Larrosa, Carme Torras and Mateu Villaret who read previous versions of this document providing many useful comments which largely improved the paper.

References

1. Apt, K.R.: Principles of Constraint Programming. Cambridge University Press (2003)
2. Bacchus, F., van Beek, P.: On the conversion between non-binary and binary constraint satisfaction problems. In: Proceedings of 15th AAAI, pp. 311–318 (1998)
3. Bacchus, F., van Run, P.: Dynamic variable ordering in csp. In: Proceedings of 1st CP, pp. 258–275 (1995)
4. Baker, A.B.: The hazards of fancy backtracking. In: Proceedings of 12th AAAI, pp. 288–293 (1994)
5. Beacham, A., Chen, X., Sillito, J., van Beek, P.: Constraint programming lessons learned from crossword puzzles. In: Proceedings of 14th Canadian Conference on AI, pp. 78–87 (2001)
6. Beldiceanu, N., Carlsson, M., Rampon, J.-X.: Global Constraint Catalog. Working version of SICS technical report T2010L07 (2011)
7. Benhamou, F., Goualart, F., Granvilliers, L., Puget, J.-F.: Revising hull and box consistency. In: Proceedings of 16th ICLP, pp. 124–138 (1999)
8. Benhamou, F., Granvilliers, L.: Continuous and interval Constraints. In: Handbook of Constraint Programming. Elsevier (2006)
9. Bessiere, C.: Constraint propagation. In: Handbook of Constraint Programming. Elsevier (2006)
10. Bessiere, C., Van Hentenryck, P.: To be or not to be ... a global constraint. In: Proceedings of 9th CP, pp. 789–794 (2003)
11. Bessiere, C., Katsirelos, G., Narodyska, N., Quimper, C.G., Walsh, T.: Decompositions of alldifferent, global cardinality and related constraints. In: Proceedings of 21th IJCAI, pp. 419–424 (2009)
12. Bessiere, C., Meseguer, P., Freuder, E., Larrosa, J.: On forward checking for non-binary constraint satisfaction. *Artif. Intell.* **141**, 205–224 (2002)
13. Bessiere, C., Regin, J.C.: Refining the basic arc consistency algorithm. In: Proceedings of 17th IJCAI, pp. 309–315 (2001)
14. Bistareli, S., Montanari, U., Rossi, F.: Constraint solving over semirings. In: Proceedings of 14th IJCAI, pp. 624–630 (1995)
15. Bitner, J.R., Reingold, E.M.: Backtrack programming techniques. *Commun. ACM* **18**, 651–656 (1975)
16. Blum, C., Roli, A.: Metaheuristics in combinatorial optimization: Overview and conceptual comparison. *ACM Comput. Surveys* **35**, 268–308 (2003)
17. Bruynooghe, M.: Solving combinatorial optimization problems by intelligent backtracking. *Inf. Process. Lett.* **11**, 36–39 (1981)
18. Carchae, T., Beck, J.: Principles for the design of large neighborhood search. *J. Math. Model. Algorithms* **8**, 245–270 (2009)
19. Carlsson, M., et al.: SICStus Prolog User's Manual (release 4.2.0). Swedish Institute of Computer Science (2011)
20. Cheng, B.M.W., Choi, K.M.F., Lee, J.H.M., Wu, J.C.K.: Increasing constraint propagation by redundant modelling: an experience report. *Constraints* **4**, 167–192 (1999)
21. Cohen, D.A., Jeavons, P., Jefferson, C., Petrie, K.E., Smith, B.: Symmetry definitions for constraint satisfaction problems. In: Proceedings of 11th CP, pp. 17–31 (2005)
22. Cooper, M.C.: An optimal k-consistency algorithm. *Artif. Intell.* **41**, 89–95 (1990)
23. Crawford, J., Ginsberg, M.L., Luks, E., Roy, A.: Symmetry-breaking predicates for search problems. In: Proceedings of 5th KR, pp. 148–159 (1996)
24. Davenport, A., Tsang, E., Wang, C., Zhu, K.: Genet: A connectionist architecture for solving constraint satisfaction problems by iterative improvement. In: Proceedings of 12th AAAI, pp. 325–330 (1994)
25. Dechter, R.: Enhancement schemes for constraint processing: back-jumping, learning and cutset decomposition. *Artif. Intell.* **41**, 273–312 (1990)
26. Dechter, R.: Mini-bucket: a general scheme of generating approximations in automated reasoning. In: Proceedings of 15th IJCAI, pp. 1297–1302 (1997)
27. Dechter, R.: Bucket elimination: a unifying framework for reasoning. *Artif. Intell.* **113**, 41–85 (1999)
28. Dechter, R.: Constraint Processing. Morgan Kaufmann. (2003)
29. Dechter, R., Pearl, J.: Network-based heuristics for constraint satisfaction problems. *Artif. Intell.* **34**, 1–38 (1987)
30. Dincbas, M., Simonis, H., van Hentenryck, P.: Solving the car-sequencing problem in constraint logic programming. In: Proceedings of 8th ECAI, pp. 290–295 (1988)
31. Dotu, I., del Val, A., Cebrian, M.: Redundant modeling for the quasigroup completion problem. In: Proceedings of 9th CP, pp. 288–302 (2003)
32. Dynadec.: Comet Tutorial (version 2.1). Dynamic Decision Technologies Inc. (2010)
33. Fahle, T., Schamberger, S., Sellmann, M.: Symmetry breaking. In: Proceedings of 7th CP, pp. 93–107 (2001)
34. Faltings, B.: Arc consistency for continuous variables. *Artif. Intell.* **65**, 363–376 (1994)

35. Fikes, R.E.: Ref-arf: a system for solving problems stated as procedures. *Artif. Intell.* **1**, 27–120 (1970)
36. Freuder, E.C.: Synthesizing constraint expressions. *Commun. ACM* **21**, 958–965 (1978)
37. Freuder, E.C.: A sufficient condition for backtrack-free search. *J. ACM* **29**, 24–32 (1982)
38. Freuder, E.C.: In pursuit of the holy grail. *Constraints* **2**, 57–61 (1997)
39. Freuder, E.C.: Modeling: the final frontier. In: *Proceedings of PAC-LP*, pp. 15–21 (1999)
40. Frisch, A., Jefferson, C., Martinez, B., Miguel, I.: The rules of constraint modelling. In: *Proceedings of 19th IJCAI*, pp. 311–318 (2005)
41. Frisch, A., Jefferson, C., Miguel, I.: Symmetry-breaking as a prelude for implied constraints: a constraint modelling pattern. In: *Proceedings of 16th ECAI*, pp. 171–175 (2004)
42. Gaschnig, J.: A general backtracking algorithm that eliminates most redundant tests. In: *Proceedings of 5th IJCAI*, p. 457 (1977)
43. Gent, I., Petrie, K.E., Puget, J.F.: Symmetry in constraint programming. In: *Handbook of Constraint Programming*. Elsevier (2006)
44. Gent, I., Smith, B.: Symmetry breaking in constraint programming. In: *Proceedings of 14th ECAI*, pp. 599–603 (2000)
45. Ginsberg, M.L.: Dynamic backtracking. *J. AI Res.* **1**, 25–46 (1993)
46. Glover, F., Laguna, M.: *Tabu Search*. Kluwer, Dordrecht (1997)
47. Goldberg, D.E.: *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley (1989)
48. Golomb, S., Baumert, L.: Backtrack programming. *J. ACM* **12**, 516–524 (1965)
49. Haralick, R.M., Elliot, G.L.: Increasing tree search efficiency for constraint satisfaction problems. *Artif. Intell.* **14**, 263–313 (1980)
50. Van Hentenryck, P.: *Constraint Satisfaction in Logic Programming*. The MIT Press (1989)
51. Van Hentenryck, P., Michell, L.: *Constraint-Based Local Search*. MIT Press (2005)
52. Hoos, H., Stutzle, T.: *Stochastic Local Search: Foundations and Applications*. Elsevier/Morgan Kaufman (2004)
53. IBM. IBM ILOG CPLEX Optimization Studio documentation. IBM (2010)
54. Kask, K., Dechter, R.: A general scheme for automatic search heuristics from specification dependencies. *Artif. Intell.* **129**, 91–131 (2001)
55. Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P.: Optimization by simulated annealing. *Science* **220**, 671–680 (1983)
56. Kiziltan, Z.: Symmetry breaking ordering constraints. PhD thesis, Uppsala University, Department of Information Science (2004)
57. Kondrak, G., van Beek, P.: A theoretical evaluation of selected backtracking algorithms. *Artif. Intell.* **89**, 365–387 (1997)
58. Laburthe, F., Jussien, N.: *Choco Solver Documentation*. Ecole de Mines de Nantes (2011)
59. Larrosa, J., Dechter, R.: Boosting search with variable elimination in constraint optimization and constraint satisfaction problems. *Constraints* **8**, 303–326 (2003)
60. Larrosa, J., Schiex, T.: Solving weighted csp by maintaining arc consistency. *Artif. Intell.* **159**, 1–26 (2004)
61. Law, Y.C., Lee, J.H.M.: Model induction: a new source of csp model redundancy. In: *Proceedings of 15th AAAI*, pp. 54–60 (2002)
62. Mackworth, A.: Consistency in networks of relations. *Artif. Intell.* **8**, 99–118 (1977)
63. Marriot, K., Stuckey, P.J.: *Programming with Constraints*. The MIT Press (1998)
64. Marriot, K., Stuckey, P.J., De Koninck, L., Samulowitz, H.: *An introduction to MiniZinc (version 1.4)*. University of Melbourne and NICTA (2011)
65. McAllister, D., Selman, B., Kautz, H.: Evidence for invariants in local search. In: *Proceedings of 14th AAAI*, pp. 321–326 (1997)
66. Meisels, A., Zivan, R.: Asynchronous forward-checking for dis-csps. *Constraints* **12**, 131–150 (2007)
67. Meseguer, P., Rossi, F., Schiex T.: Soft constraints. In: *Handbook of Constraint Programming*. Elsevier (2006)
68. Minton, S., Johnston, M., Philips, A., Laird, P.: Minimizing conflicts: a heuristic repair method for constraint satisfaction and scheduling problems. *Artif. Intell.* **58**, 161–205 (1992)
69. Modi, P.J., Shen, W.M., Tambe, M., Yokoo, M.: Adopt: asynchronous distributed constraint optimization with quality guarantees. *Artif. Intell.* **161**, 149–180 (2005)
70. Mohr, R., Henderson, T.C.: Arc and path consistency revisited. *Artif. Intell.* **28**, 225–233 (1986)
71. Montanari, U.: Networks of constraints: fundamental properties and applications to picture processing. *Inf. Sci.* **7**, 95–132 (1974)
72. Morris, P.: The breakout method for escaping from local minima. In: *Proceedings of 11th AAAI*, pp. 40–45 (1993)
73. Petcu, A., Faltings, B.: A scalable method for multiagent constraint optimization. In: *Proceedings of 19th IJCAI*, pp. 266–271 (2005)
74. Prosser, P.: Domain filtering can degrade intelligent backtracking search. In: *Proceedings of 13th IJCAI*, pp. 262–267 (1993)
75. Puget, J.-F.: Constraint programming: a great AI success. In: *Proceedings of 13th ECAI*, pp. 698–675 (1998)
76. Puget, J.-F.: Constraint programming next challenge: simplicity of use. In: *Proceedings of 10th CP*, pp. 5–8 (2004)
77. Regin, J.C.: A filtering algorithm for constraints of difference in csps. In: *Proceedings of 12th AAAI*, pp. 362–367 (1994)
78. Rossi, F., Van Beek, P., Walsh, T.: *Handbook of Constraint Programming*. Elsevier (2006)
79. Sabin, D., Freuder, E.: Contradicting conventional wisdom in constraint satisfaction. In: *Proceedings of 11th ECAI*, pp. 125–129 (1994)
80. Schiex, T., Fargier, H., Verfaillie, G.: Values constraint satisfaction problems: hard and easy problems. In: *Proceedings of 14th IJCAI*, pp. 631–637 (1995)
81. Schulte, C., Carlsson, M.: Finite constraint programming systems. In: *Handbook of Constraint Programming*. Elsevier (2006)
82. Schulte, C., Tack, G., Lagerkvist, M.: *Modelling and Programming with Gecode (version 3.7.1)* (2011)
83. Selman, B., Johnston, M., Philips, A., Laird, P.: A new method for solving hard satisfiability problems. In: *Proceedings of 10th AAAI*, pp. 440–446 (1992)
84. Selman, B., Kautz, H., Cohen, B.: Noise strategies for improving local search. In: *Proceedings of 12th AAAI*, pp. 337–343 (1994)
85. Silaghi, M.C., Sam-Haroud, D., Faltings, B.: Asynchronous consistency and maintenance in distributed constraint satisfaction. *Artif. Intell.* **161**, 25–53 (2005)
86. Smith, B.: *Modelling*. In: *Handbook of Constraint Programming*. Elsevier (2006)
87. Smith, B., Grant, S.: Trying harder to fail first. In: *Proceedings of 13th ECAI*, pp. 249–253 (1998)
88. Stallman, R.M., Sussman, G.J.: Forward reasoning and dependency-directed backtracking in a system for computer-aided circuit analysis. *Artif. Intell.* **9**, 135–196 (1977)
89. Tsang, E.P.K.: *Foundations of Constraint Satisfaction*. Academic Press (1993)
90. van Hoewe, W.J., Katriel, I.: Global constraints. In: *Handbook of Constraint Programming*. Elsevier (2006)
91. van Hoewe, W.J., Pesant, G., Rousseau, L.-M.: On global warming: flow-based soft global constraints. *J. Heuristics* **12**, 347–373 (2006)
92. Voudouris, C., Tsang, E.: Guided local search. In: *Handbook of Metaheuristics*. Kluwer (2003)
93. Walsh, T.: Sat $\vee$ csp. In: *Proceedings of 6th CP*, pp. 441–456 (2000)
94. Walsh, T.: Constraint patterns. In: *Proceedings of 9th CP*, pp. 53–64 (2003)

95. Waltz, D.: Understanding line drawings of scenes with shadows. In: *The Psychology of Computer Vision*, pp. 19–91. McGraw-Hill (1975)
96. Yeoh, W., Felner, A., Koenig, S.: BnB-ADOPT: asynchronous branch-and-bound dcop algorithm. *J. AI Res.* **38**, 85–133 (2010)
97. Yokoo, M., Durfee, E., Ishida, T., Kuwabara, K.: The distributed constraint satisfaction problem: formalization and algorithms. *IEEE Trans. Knowl. Data Eng.* **10**, 673–685 (1998)
98. Zivan, R., Meisels, A.: Dynamic ordering for asynchronous backtracking on discsps. *Constraints* **11**, 179–197 (2006)