


EDWARD TSANG



Foundations of Constraint Satisfaction

Edited by Thom Fruehwirth



AI Fundamentals: Constraints Satisfaction Problems

Maria Simi



Constraints Propagation

LESSON 2

CONSTRAINT PROPAGATION – LOCAL CONSISTENCY – PROPERTIES.

Constraint propagation and related concepts

Constraint propagation

- Constraints are used to reduce the number of legal values for a variable, which in turn can reduce the legal values for another variable, and so on ...

Problem reduction techniques

- Techniques for transforming a CSP into equivalent problems which are hopefully easier to solve or recognizable as insoluble.

Enforcing local consistency

- The process of enforcing **local consistency** properties in a constraint graph causes inconsistent values to be eliminated
- Different types of local consistency have been studied

Problem reduction

Reducing a problem means removing from the constraints (legal assignments) those assignments which appear in no solution tuples.

Two CSP problems are **equivalent** if they have identical sets of variables and solutions.

A CSP problem \mathcal{P}_1 is **reduced** to a problem \mathcal{P}_2 when

1. \mathcal{P}_1 is equivalent to \mathcal{P}_2
2. Domains of variables in \mathcal{P}_2 are subsets wrt those in \mathcal{P}_1
3. The constraints in \mathcal{P}_2 are at least as restrictive than in \mathcal{P}_1

These conditions guarantee that a solution to \mathcal{P}_2 is also a solution to \mathcal{P}_1 , only **redundant** values and assignments are removed (no solution is lost).

The problem is easier to solve.

Problem reduction strategies

Problem reduction involves two possible tasks:

1. removing redundant values from the domains of the variables
2. tightening the constraints so that fewer compound labels satisfy them

Example: if $x < y$ is a constraint and $Dx = \{3, 4, 5\}$ and $Dy = \{1, 2, 4\}$ domains can be safely reduced to $\{3\}$ and $\{4\}$.

Constraints are seen as sets, then this means removing redundant compound labels from the constraints. If the domain of any variable or any constraint is reduced to an empty set, then one can conclude that the problem is insoluble.

Problem reduction is also called *consistency maintenance* since it relies on [re-]establishing local consistency properties.

Local consistency properties

- Node consistency
- Arc consistency
- Directional arc consistency
- Generalized arc consistency
- Path consistency
- K-consistency
- Forward Checking

All these operations do not change the set of the solutions, do not necessarily solve a problem but, used in conjunction with search, make the search more efficient by pruning the search tree.

Node consistency

A node is **consistent** if all the values in its domain satisfy unary constraints on the associated variable. In formula, given a unary constraint

$$C_i = \langle (x_i), R_i \rangle \quad D_i \subseteq R_i$$

A constraint network is *node-consistent* if all its nodes are consistent

Unary constraints can be easily satisfied by reducing the domains of variables.

$$D'_i \leftarrow D_i \cap R_i \quad D_i \leftarrow D'_i$$

The algorithm, called NC-1, is $O(d.n)$

Example: in the map coloring problem of Australia

- Suppose South Australia dislikes green: $(SA \neq \text{green})$ is a unary constraint.
- SA starts with domain {**red**, **green**, **blue**}, and we can make it *node-consistent* by eliminating **green**, leaving SA with the reduced domain {**red**, **blue**}

Arc consistency (for binary constraints)

A variable in a CSP is **arc-consistent** if every value in its domain satisfies the binary constraints on this variable.

X_i is **arc-consistent** with respect to another variable X_j if for every value in the current domain D_i there is some value in the domain D_j that satisfies the binary constraint on the arc (X_i, X_j) .

Example: $X = \{X, Y\}$ $D_X = D_Y = \{0, 1, 2, 3, 4, 5, 6, 7, 8, 9\}$

Constraint: $\langle (X, Y), \{(0, 0), (1, 1), (2, 4), (3, 9)\} \rangle$ i.e. $X = Y^2$ $X \rightarrow Y$

To make X arc-consistent with respect to Y , we reduce X 's domain to $\{0, 1, 2, 3\}$.

If we also make Y arc-consistent with respect to X , then Y 's domain becomes $\{0, 1, 4, 9\}$ and the whole **edge** is arc-consistent.

A relational algebra view

Consider variable x_i with associated domain D_i . We further assume a constraint between x_i and x_j , expressed by relation $R_{i,j}$.

Arc $x_i \rightarrow x_j$ is arc consistent iff $D_i \subseteq \pi_i(R_{i,j} \bowtie D_j)$

Where \bowtie and π are the join and projection operator of relational algebra. The operation is a *left semijoin* (\ltimes)

Arc $x_i \rightarrow x_j$ can be made *arc consistent* by computing:

$$D_i' \leftarrow D_i \cap \pi_i(R_{i,j} \bowtie D_j) \quad D_i \leftarrow D_i'$$

Example:

$$\begin{aligned} \pi_x(R_{x,y} \bowtie D_x) &= \pi_x(\{(0, 0), (1, 1), (2, 4), (3, 9)\} \bowtie \{0, 1, 2, 3, 4, 5, 6, 7, 8, 9\}) \\ &= \pi_x(\{(0, 0), (1, 1), (2, 4), (3, 9)\}) = \{0, 1, 2, 3\} \end{aligned}$$

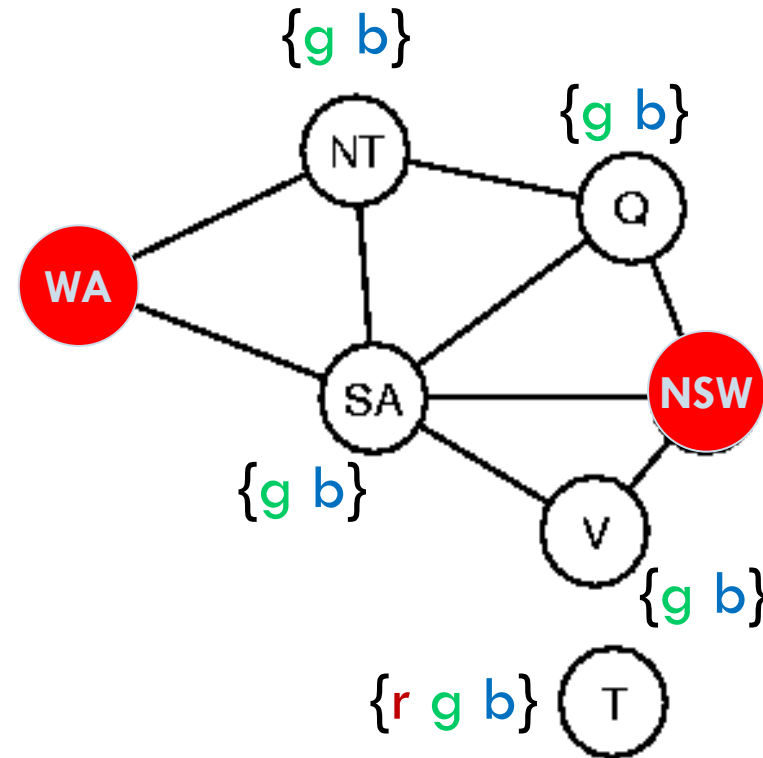
$$D_x' \leftarrow D_x \cap \{0, 1, 2, 3\} = \{0, 1, 2, 3\}$$

Arc consistent but no solutions

Arc consistency does not guarantee a solution.

In this case all the arcs are consistent but there is no solution

Impossible to color three fully connected nodes with two colors



Algorithms for arc consistency (AC-3)

The most popular algorithm for arc consistency is called AC-3 [Mackworth, 1977]. AC-3(*csp*) maintains a queue of arcs to consider; initially all the arcs in *csp*. An **edge** produces two arcs.

AC-3 pops off an arc(X_i, X_j) from the queue and makes X_i arc-consistent with respect to X_j .

1. If this step leaves D_i unchanged, the algorithm just moves on to the next arc.
2. If D_i is made smaller, then we add to the queue all arcs (X_k, X_j) where X_k is a neighbor of X_i .
3. If D_i becomes empty, then we conclude that the whole CSP has no consistent solution.

When there are no more arcs to consider, we are left with a CSP that is equivalent to the original CSP, but simpler.

AC-3: AIMA pseudo-code

function AC-3(*csp*) **returns** false if an inconsistency is found and true otherwise

inputs: *csp*, a binary CSP with components (X, D, C)

local variables: *queue*, a queue of arcs, initially all the arcs in *csp*

while *queue* is not empty **do**

$(X_i, X_j) \leftarrow \text{REMOVE-FIRST}(\textit{queue})$

if REVISE(*csp*, X_i, X_j) **then**

if size of $D_i = 0$ **then return** false

for each X_k **in** $X_i.\text{NEIGHBORS} - \{X_j\}$ **do**

 add (X_k, X_i) to *queue*

return true

function REVISE(*csp*, X_i, X_j) **returns** true iff we revise the domain of X_i

revised \leftarrow false

for each x **in** D_i **do**

if no value y in D_j allows (x,y) to satisfy the constraint between X_i and X_j **then**


 delete x from D_i

revised \leftarrow true

return *revised*

Arc consistency: an example

- Variables $A \{1, 2, 3, 4\}$ $B \{1, 2, 3, 4\}$ $C \{1, 2, 3, 4\}$
- Constraints $A < B$; $A > C$



QUEUE	ARC	ARC DOMAIN
$\{(A, B), (B, A), (A, C), (C, A)\}$		
$\{(B, A), (A, C), (C, A)\}$	(A, B)	$A = \{1, 2, 3, 4\}$
$\{(A, C), (C, A)\}$	(B, A)	$B = \{1, 2, 3, 4\}$
$\{(C, A)\}$	(A, C)	$A = \{1, 2, 3\}$
$\{(B, A), (C, A)\}$	<i>add (B, A) for checking</i>	
$\{(C, A)\}$	(B, A)	$B = \{2, 3, 4\}$
$\{\}$	(C, A)	$C = \{1, 2, 3, 4\}$
At the end: $A = \{2, 3\}$ $B = \{3, 4\}$	$C = \{1, 2\}$	

Complexity of AC-3

Assume a CSP with n variables, each with domain size at most d , and with c binary constraints (arcs).

Each arc (X_k, X_i) can be inserted in the queue only d times because X_i has at most d values to delete.

Checking consistency of an arc can be done in $O(d^2)$ time, so we get $O(c d^3)$ total worst-case time

Complexity: $O(c d^3)$... polynomial time

The algorithm AC-4 is an improved version of AC-3, based on the notion of support, that doesn't need to consider all the incoming arcs. Some more information must be kept. $O(c d^2)$.

Directional Arc Consistency

Directional Arc Consistency (DAC) is defined under **total ordering of the variables**.

A CSP is **directional arc consistent** (DAC) under an ordering of the variables if and only if for every label $\langle x, a \rangle$ which satisfies the constraints on x , there exists a compatible label $\langle y, b \rangle$ for every variable y , **which is after x** according to the ordering.

In the algorithm for establishing DAC (DAC-1), each arc is examined exactly once, by proceeding from the last in the ordering, so the complexity is $O(cd^2)$.

We will see later an use of this property.


Warning: AC cannot always be achieved by running DAC-1 in both directions.

DAC in both directions weaker than AC

After running DAC with orderings ACB and BCA, the only effect is to delete 1 from the C domain.

However the resulting graph is not arc consistent.

Arc BC is non consistent: the value 1 should be deleted from the domain of B to make the arc BC consistent.



Generalized arc consistency

An extension of the notion of arc consistency to handle n -ary rather than just binary constraints (also called *hyper-arc* consistency).

A variable X_i is **generalized arc consistent** with respect to a n -ary constraint if for every value v in the domain of X_i there exists a tuple of values that is a member of the constraint, has all its values taken from the domains of the corresponding variables, and has its X_i component equal to v .

For example, if all variables have the domain $\{0, 1, 2, 3\}$, then to make the variable X consistent with the ternary constraint $X < Y < Z$, we would have to eliminate 2 and 3 from the domain of X because the constraint cannot be satisfied when X is 2 or 3.

Path consistency [Montanari]

Arc consistency tightens down the domains (unary constraints) using the arcs (binary constraints).

Path consistency is a stronger notion: it tightens the binary constraints by using implicit constraints that are inferred by looking at triples of variables.

A path of length 2 between variables $\{X_i, X_j\}$ is **path-consistent** with respect to a third variable X_m if, for every consistent assignment $\{X_i = a, X_j = b\}$, there is an assignment to X_m that satisfies the constraints on $\{X_i, X_m\}$ and $\{X_m, X_j\}$.

In relational algebra:

$$R_{i,j} \subseteq \pi_{i,j}(R_{i,m} \bowtie D_m \bowtie R_{m,j})$$

Path consistency algorithm and properties

To achieve path consistency:

$$R_{i,j}' \leftarrow R_{i,j} \cap \pi_{i,j}(R_{i,m} \bowtie D_m \bowtie R_{m,j})$$

$$R_{i,j} \leftarrow R_{i,j}'$$

The algorithm is called PC-2.

If all path of length 2 are made consistent, then all path of any length are consistent [Montanari 1974], so longer path need not be considered.

This is called path consistency because one can think of it as looking at a path from X_i to X_j with X_m in the middle.

Path consistency: example

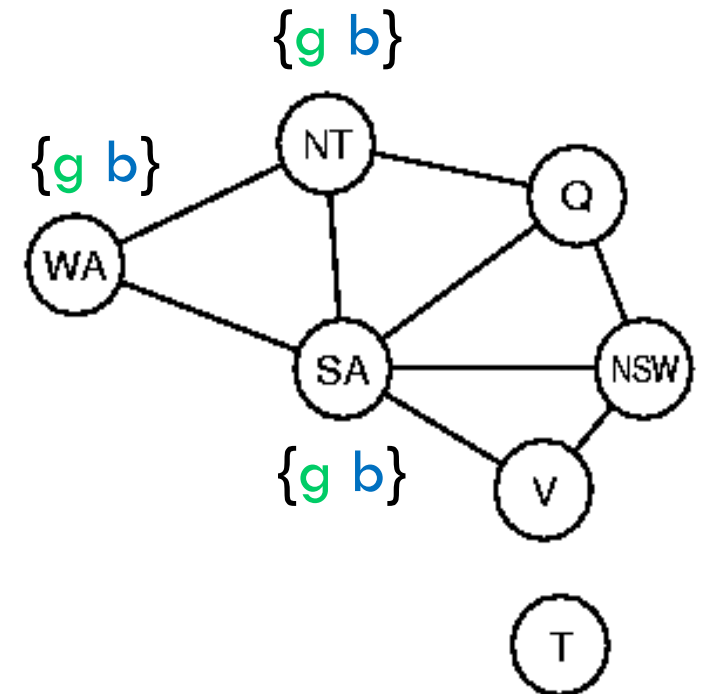
Coloring the Australia map with two colors is impossible, but arc-consistency is not able to discover it.

If we try to make the set $\{WA, SA\}$ *path consistent* with respect to NT.

The consistent assignments for WA and SA are only two:

1. $\{WA = \text{green}, SA = \text{blue}\}$
2. $\{WA = \text{blue}, SA = \text{green}\}$

Neither of them is compatible with $NT = \text{green}$ nor $NT = \text{blue}$, so the domain of WA and SA become empty and we can conclude that there are no solutions.



k -consistency

Stronger forms of propagation can be defined with the notion of **k -consistency**, a generalization of the other properties.

A CSP is *k -consistent* if, for any set of $k - 1$ variables and for any consistent assignment to those variables, a consistent value can always be assigned to any k^{th} variable.

1-consistency says that, given the empty set, we can make any set of one variable consistent: this is what we called node consistency.

2-consistency is the same as arc consistency. For binary constraint networks

3-consistency is the same as path consistency.

Strong k -consistency

A CSP is **strongly** k -consistent if it is k -consistent and is also $(k - 1)$ -consistent, $(k - 2)$ -consistent, . . . all the way down to 1-consistent.

Now suppose we have a CSP with k nodes and make it strongly k -consistent. We can then solve the problem as follows: First, we choose a consistent value for x_1 . We are then guaranteed to be able to choose a value for x_2 because the graph is 2-consistent, for x_3 because it is 3-consistent, and so on.

For each variable x_i , we need only search through the d values in the domain to find a value consistent with x_1, \dots, x_{i-1} . We are guaranteed to find a solution in time $O(n^2d)$.

BUT: Any algorithm for establishing k -consistency must take time exponential in k in the worst case. Worse, k -consistency also requires space that is exponential in k .

Forward checking (FC)

A very weak, local and quick form of consistency checking which is triggered during the search process.

When you assign a value v to a variable X in the process of searching for a consistent assignment, check the neighbors variables and exclude values that are not compatible with v from their domains.

Specialized global constraints [AIMA]

A **global constraint** is one involving an arbitrary number of variables (but not necessarily all variables).

The *Alldiff* constraint says that all the variables involved must have distinct values and is very common (Crypto-arithmetic, Sudoku).

Treating *Alldiff* with special algorithms can be much more efficient.

Solutions for *AllDiffs*

Simple form of inconsistency detection for *Alldiff*:

if m variables are involved in the constraint, and if they have n possible distinct values altogether, and $m > n$, then the constraint cannot be satisfied.

Remove any variable in the constraint that has a singleton domain. Delete that variable's value from the domains of the remaining variables. Repeat as long as there are singleton variables. If at any point an empty domain is produced or there are more variables than domain values left, then an inconsistency has been detected.

Example: in the map coloring problem, the assignment {WA=red , NSW =red} and AC-3 do not detect the inconsistency on the variables NT, Q, SA. The *Alldiff* constraint is instead effective.

Resource constraints

The resource constraint is often called ***Atmost constraint***

Example: $Atmost(10, P_1, P_2, P_3, P_4)$, meaning that 10 is the maximum of personnel units to be assigned to 4 tasks. This constraint may be checked by summing the minimum requirement for personnel for each task.

Domains can be represented by upper and lower bounds and managed by **bounds propagation**.

Example: air scheduling with two flights F_1 and F_2

$D_1 = [0, 165]$ and $D_2 = [0, 385]$ with additional constraint $F_1 + F_2 = 420$

By propagating bounds constraints, we reduce the domains to

$D_1 = [35, 165]$ and $D_2 = [255, 385]$

Conclusions

- ✓ We have looked at problem reduction techniques which work by enforcing local consistency properties of different strength and complexity.
- ✓ These are properties that make the problem simpler: the more effort you put, the simpler the problem becomes.
- ✓ These techniques will be used in connection with search algorithms which is the topic of the next lecture.

Your turn

- ✓ Look at the Python implementation of AC-3 and use it to reduce a CSP problem
- ✓ Study and present AC-4
- ✓ Compare AC with Path Consistency (cost vs effect).

References

Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach* (3rd edition). Pearson Education 2010 [Cap 6 – CSP]

Handbook of Constraint Programming, Edited by F. Rossi, P. van Beek and T. Walsh. Elsevier 2006.

Edward Tsang, *Foundations of Constraints Satisfaction* [Cap 3]