Foundations of Artificial Intelligence

D4. Constraint Satisfaction Problems: Arc Consistency

Malte Helmert

University of Basel

April 9, 2025

Constraint Satisfaction Problems: Overview

Chapter overview: constraint satisfaction problems

- D1–D2. Introduction
- D3–D5. Basic Algorithms
  - D3. Backtracking
  - D4. Arc Consistency
  - D5. Path Consistency
- D6–D7. Problem Structure

D4.1 Inference

Inference

Inference

Derive additional constraints (here: unary or binary)
that are implied by the given constraints,
i.e., that are satisfied in all solutions.

Inference: Example

Running Example

binary constraints:

- $R_{wx} = \{\langle 2, 1\rangle, \langle 4, 2\rangle\}$
- $R_{wz} = \{\langle 1, 2\rangle, \langle 1, 3\rangle, \langle 2, 3\rangle\}$
- $R_{yz} = \{\langle 2, 1\rangle, \langle 3, 1\rangle, \langle 3, 2\rangle, \langle 4, 1\rangle, \langle 4, 2\rangle, \langle 4, 3\rangle\}$

domains (unary constraints):

- $\mathrm{dom}(w) = \{1, 2, 3, 4\}$
- $\mathrm{dom}(x) = \{1, 2, 3\}$
- $\mathrm{dom}(y) = \{1, 2, 3, 4\}$
- $\mathrm{dom}(z) = \{1, 2, 3\}$

Can we use the constraint R_{wz} $(w < z)$ to come up with a unary constraint R_w?

Inference: Example

Running Example

binary constraints:

- $R_{wx} = \{\langle 2, 1\rangle, \langle 4, 2\rangle\}$
- $R_{wz} = \{\langle 1, 2\rangle, \langle 1, 3\rangle, \langle 2, 3\rangle\}$
- $R_{yz} = \{\langle 2, 1\rangle, \langle 3, 1\rangle, \langle 3, 2\rangle, \langle 4, 1\rangle, \langle 4, 2\rangle, \langle 4, 3\rangle\}$

domains (unary constraints):

- $\mathrm{dom}(w) = \{1, 2\}$
- $\mathrm{dom}(x) = \{1, 2, 3\}$
- $\mathrm{dom}(y) = \{1, 2, 3, 4\}$
- $\mathrm{dom}(z) = \{1, 2, 3\}$

Can we use the constraint R_{wz} ($w < z$) to come up with a unary constraint R_w?
$\rightsquigarrow$ tighten domain with unary constraint
(sometimes called node consistency)

Inference: Example

Running Example

binary constraints:

- $R_{wx} = \{\langle 2, 1\rangle\}$
- $R_{wz} = \{\langle 1, 2\rangle, \langle 1, 3\rangle, \langle 2, 3\rangle\}$
- $R_{yz} = \{\langle 2, 1\rangle, \langle 3, 1\rangle, \langle 3, 2\rangle, \langle 4, 1\rangle, \langle 4, 2\rangle, \langle 4, 3\rangle\}$

domains (unary constraints):

- $\mathrm{dom}(w) = \{1, 2\}$
- $\mathrm{dom}(x) = \{1, 2, 3\}$
- $\mathrm{dom}(y) = \{1, 2, 3, 4\}$
- $\mathrm{dom}(z) = \{1, 2, 3\}$

How does this affect the binary constraint R_{wx}?

Inference: Example

Running Example

binary constraints:

- $R_{wx} = \{\langle 2, 1 \rangle\}$
- $R_{wz} = \{\langle 1, 2 \rangle, \langle 1, 3 \rangle, \langle 2, 3 \rangle\}$
- $R_{yz} = \{\langle 2, 1 \rangle, \langle 3, 1 \rangle, \langle 3, 2 \rangle, \langle 4, 1 \rangle, \langle 4, 2 \rangle, \langle 4, 3 \rangle\}$
- $R_{wy} = \{\langle 1, 3 \rangle, \langle 1, 4 \rangle, \langle 2, 4 \rangle\}$

domains (unary constraints):

- $\mathrm{dom}(w) = \{1, 2\}$
- $\mathrm{dom}(x) = \{1, 2, 3\}$
- $\mathrm{dom}(y) = \{1, 2, 3, 4\}$
- $\mathrm{dom}(z) = \{1, 2, 3\}$

Can we generate a "new" binary constraint from $w < z$ and $z < y$?
(i.e., tighten a trivial constraint)

Trade-Off Search vs. Inference

Inference formally

For a given constraint network $\mathcal{C}$, replace $\mathcal{C}$ with an equivalent, but tighter constraint network.

Trade-off:

- the more complex the inference, and
- the more often inference is applied,
- the smaller the resulting state space, but
- the higher the complexity per search node.

When to Apply Inference?

different possibilities to apply inference:

- once as preprocessing before search
- combined with search: before recursive calls during backtracking procedure
  - already assigned variable $v \mapsto d$ corresponds to $\mathrm{dom}(v) = \{d\}$
    $\rightsquigarrow$ more inferences possible
  - during backtracking, derived constraints have to be retracted because they were based on the given assignment

  $\rightsquigarrow$ powerful, but possibly expensive

Backtracking with Inference

function BacktrackingWithInference($\mathcal{C}, \alpha$):

if α is inconsistent with $\mathcal{C}$:
  return inconsistent

if α is a total assignment:
  return α

$\mathcal{C}' := \langle V, \mathrm{dom}', (R'_{uv}) \rangle :=$ copy of $\mathcal{C}$
apply inference to $\mathcal{C}'$

if $\mathrm{dom}'(v) \neq \emptyset$ for all variables v:

  select some variable v for which α is not defined

  for each $d \in$ copy of $\mathrm{dom}'(v)$ in some order:
    $\alpha' := \alpha \cup \{v \mapsto d\}$
    $\mathrm{dom}'(v) := \{d\}$
    $\alpha'' :=$ BacktrackingWithInference($\mathcal{C}', \alpha'$)
    if $\alpha'' \neq$ **inconsistent**:
      return α''

return inconsistent

Backtracking with Inference: Discussion

- Inference is a placeholder:
  different inference methods can be applied.
- Inference methods can recognize unsolvability (given α)
  and indicate this by clearing the domain of a variable.
- Efficient implementations of inference are often incremental:
  the last assigned variable/value pair $v \mapsto d$ is taken
  into account to speed up the inference computation.

D4.2 Forward Checking

Forward Checking

We start with a simple inference method:

Forward Checking

Let α be a partial assignment.

Inference: For all unassigned variables v in α, remove all values from the domain of v that are in conflict with already assigned variable/value pairs in α.

$\rightsquigarrow$ definition of conflict as in the previous chapter

Incremental computation:

- When adding $v \mapsto d$ to the assignment, delete all pairs that conflict with $v \mapsto d$.

Forward Checking: Example

Running Example

Removing values in conflict with $\alpha = \{w \mapsto 2\}$:

binary constraints:

- $R_{wx} = \{\langle 2, 1\rangle, \langle 4, 2\rangle\}$
- $R_{wz} = \{\langle 1, 2\rangle, \langle 1, 3\rangle, \langle 2, 3\rangle\}$
- $R_{yz} = \{\langle 2, 1\rangle, \langle 3, 1\rangle, \langle 3, 2\rangle, \langle 4, 1\rangle, \langle 4, 2\rangle, \langle 4, 3\rangle\}$

domains:

- w is already assigned
- $\mathrm{dom}(x) = \{1, 2, 3\}$
- $\mathrm{dom}(y) = \{1, 2, 3, 4\}$
- $\mathrm{dom}(z) = \{1, 2, 3\}$

Forward Checking: Example

Running Example

Removing values in conflict with $\alpha = \{w \mapsto 2\}$:

binary constraints:

- $R_{wx} = \{\langle 2, 1\rangle, \langle 4, 2\rangle\}$
- $R_{wz} = \{\langle 1, 2\rangle, \langle 1, 3\rangle, \langle 2, 3\rangle\}$
- $R_{yz} = \{\langle 2, 1\rangle, \langle 3, 1\rangle, \langle 3, 2\rangle, \langle 4, 1\rangle, \langle 4, 2\rangle, \langle 4, 3\rangle\}$

domains:

- w is already assigned
- $\mathrm{dom}(x) = \{1\}$
- $\mathrm{dom}(y) = \{1, 2, 3, 4\}$
- $\mathrm{dom}(z) = \{3\}$

Forward Checking: Discussion

properties of forward checking:

- correct inference method (retains equivalence)
- affects domains (= unary constraints), but not binary constraints
- consistency check at the beginning of the backtracking procedure no longer needed (Why?)
- cheap, but often still useful inference method

$\rightsquigarrow$ apply at least forward checking in the backtracking procedure

In the following, we will consider more powerful inference methods.

D4.3 Arc Consistency

Arc Consistency: Definition

Definition (Arc Consistent)

Let $\mathcal{C} = \langle V, \mathrm{dom}, (R_{uv}) \rangle$ be a constraint network.

1. The variable $v \in V$ is arc consistent with respect to another variable $v' \in V$, if for every value $d \in \mathrm{dom}(v)$ there exists a value $d' \in \mathrm{dom}(v')$ with $\langle d, d' \rangle \in R_{vv'}$.
2. The constraint network $\mathcal{C}$ is arc consistent, if every variable $v \in V$ is arc consistent with respect to every other variable $v' \in V$.

German: kantenkonsistent

remarks:

- definition for variable pair is not symmetrical
- v always arc consistent with respect to v' if the constraint between v and v' is trivial

Arc Consistency: Example

Running Example

Consider variables w and z from our running example:

- $\mathrm{dom}(w) = \{1, 2, 3, 4\}$
- $\mathrm{dom}(z) = \{1, 2, 3\}$
- $R_{wz} = \{\langle 1, 2\rangle, \langle 1, 3\rangle, \langle 2, 3\rangle\}$

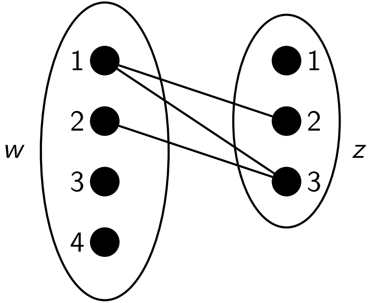

Arc consistency
of w with respect to z and
of z with respect to w
is violated.

Enforcing Arc Consistency

- Enforcing arc consistency, i.e., removing values from dom(v) that violate the arc consistency of v with respect to v', is a correct inference method. (Why?)
- more powerful than forward checking (Why?)
  $\rightsquigarrow$ Forward checking is a special case: enforcing arc consistency of all variables with respect to the just assigned variable corresponds to forward checking.

We will next consider algorithms that enforce arc consistency.

Processing Variable Pairs: revise

function revise($\mathcal{C}, v, v'$):
$\langle V, \text{dom}, (R_{uv}) \rangle := \mathcal{C}$
for each $d \in \text{dom}(v)$:
    if there is no $d' \in \text{dom}(v')$ with $\langle d, d' \rangle \in R_{vv'}$:
        remove d from $\text{dom}(v)$

input: constraint network $\mathcal{C}$ and two variables v, v' of $\mathcal{C}$

effect: v arc consistent with respect to v'.
All violating values in $\text{dom}(v)$ are removed.

time complexity: $O(k^2)$, where k is maximal domain size

revise($\mathcal{C}$, w, z) in Running Example

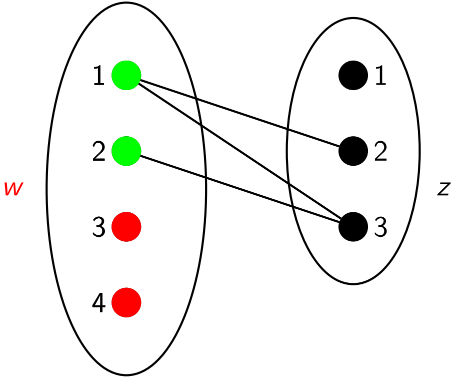

Enforcing Arc Consistency: AC-1

function AC-1($\mathcal{C}$):
$\langle V, \mathrm{dom}, (R_{uv}) \rangle := \mathcal{C}$
repeat
  for each nontrivial constraint R_{uv}:
    revise($\mathcal{C}$, u, v)
    revise($\mathcal{C}$, v, u)
until no domain has changed in this iteration

input: constraint network $\mathcal{C}$

effect: transforms $\mathcal{C}$ into equivalent arc consistent network

time complexity: $O(n \cdot e \cdot k^3)$, with n variables, e nontrivial constraints and maximal domain size k

AC-1: Discussion

- AC-1 does the job, but is rather inefficient.
- Drawback: Variable pairs are often checked again and again although their domains have remained unchanged.
- These (redundant) checks can be saved.

$\rightsquigarrow$ more efficient algorithm: AC-3

Enforcing Arc Consistency: AC-3

idea: store potentially inconsistent variable pairs in a queue

function AC-3($\mathcal{C}$):

$\langle V, \mathrm{dom}, (R_{uv}) \rangle := \mathcal{C}$

$queue := \emptyset$

for each nontrivial constraint R_{uv}:

  insert $\langle u, v \rangle$ into *queue*

  insert $\langle v, u \rangle$ into *queue*

while $queue \neq \emptyset$:

  remove an arbitrary element $\langle u, v \rangle$ from *queue*

  revise($\mathcal{C}, u, v$)

  if dom(u) changed in the call to revise:

    for each $w \in V \setminus \{u, v\}$ where R_{wu} is nontrivial:

      insert $\langle w, u \rangle$ into *queue*

AC-3: Discussion

- *queue* can be an arbitrary data structure that supports insert and remove operations (the order of removal does not affect the result)

$\rightsquigarrow$ use data structure with fast insertion and removal, e.g., stack

- AC-3 has the same effect as AC-1: it enforces arc consistency
- proof idea: invariant of the **while** loop: If $\langle u, v \rangle \notin \mathit{queue}$, then u is arc consistent with respect to v

AC-3: Time Complexity

Proposition (time complexity of AC-3)

Let $\mathcal{C}$ be a constraint network with e nontrivial constraints and maximal domain size k.

The time complexity of AC-3 *is $O(e \cdot k^3)$.*

AC-3: Time Complexity (Proof)

Proof.

Consider a pair $\langle u, v \rangle$ such that there exists a nontrivial constraint R_{uv} or R_{vu}. (There are at most $2e$ of such pairs.)

Every time this pair is inserted to the queue (except for the first time) the domain of the second variable has just been reduced.

This can happen at most k times.

Hence every pair $\langle u, v \rangle$ is inserted into the queue at most $k + 1$ times $\rightsquigarrow$ at most $O(ek)$ insert operations in total.

This bounds the number of **while** iterations by $O(ek)$, giving an overall time complexity of $O(ek) \cdot O(k^2) = O(ek^3)$. $\square$

D4.4 Summary

Summary: Inference

- inference: derivation of additional constraints that are implied by the known constraints

$\rightsquigarrow$ tighter equivalent constraint network

- trade-off search vs. inference
- inference as preprocessing or integrated into backtracking

Summary: Forward Checking, Arc Consistency

- cheap and easy inference: forward checking
  - remove values that conflict with already assigned values
- more expensive and more powerful: arc consistency
  - iteratively remove values without a suitable "partner value" for another variable until fixed-point reached
  - efficient implementation of AC-3: $O(ek^3)$ with e: #nontrivial constraints, k: size of domain