

Constraint Satisfaction and Scheduling

Andrew W. Moore
Professor
School of Computer Science
Carnegie Mellon University

www.cs.cmu.edu/~awm

awm@cs.cmu.edu

412-268-7599

Note to other teachers and users of these slides. Andrew would be delighted if you found this source material useful in giving your own lectures. Feel free to use these slides verbatim, or to modify them to fit your own needs. PowerPoint originals are available. If you make use of a significant portion of these slides in your own lecture, please include this message, or the following link to the source repository of Andrew's tutorials: <http://www.cs.cmu.edu/~awm/tutorials>. Comments and corrections gratefully received.

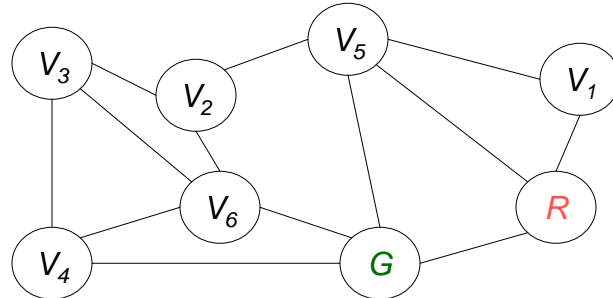
Slide 1

Overview

- CSPs defined
- Using standard search for CSPs
- Blindingly obvious improvements
 - Backtracking search
 - Forward Checking
 - Constraint Propagation
- Some example CSP applications
 - Overview
 - Waltz Algorithm
 - Job Shop Scheduling
- Variable ordering
- Value ordering
- Tedious Discussion

Slide 2

A Constraint Satisfaction Problem



Inside each circle marked $V_1 \dots V_6$ we must assign: R , G or B .
No two connected circles may be assigned the same symbol.
Notice that two circles have already been given an assignment.

Slide 3

Formal Constraint Satisfaction Problem

A CSP is a triplet $\{V, D, C\}$. A CSP has a finite set of variables $V = \{V_1, V_2 \dots V_N\}$.

Each variable may be assigned a value from a domain D of values.

Each member of C is a pair. The first member of each pair is a set of variables. The second element is a set of legal values which that set may take.

Example:

$V = \{V_1, V_2, V_3, V_4, V_5, V_6\}$

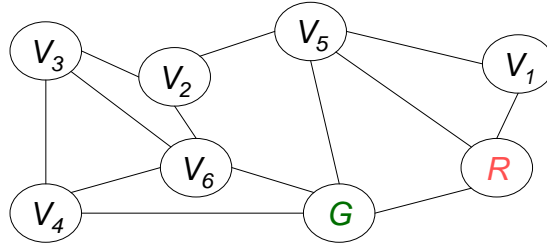
$D = \{R, G, B\}$

$C = \{(V_1, V_2) : \{(R,G), (R,B), (G,R), (G,B), (B,R), (B,G)\},$
 $\{(V_1, V_3) : \{(R,G), (R,B), (G,R), (G,B), (B,R), (B,G)\},$
 \vdots
 $\vdots\}$

Obvious point: Usually C isn't represented explicitly, but by a function.

Slide 4

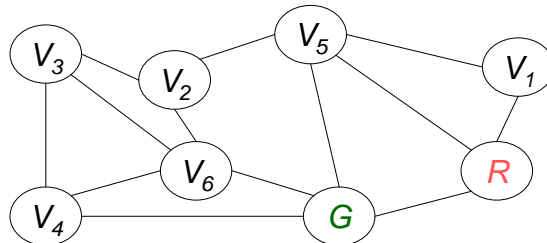
How to solve our CSP?



- How about using a search algorithm?
- Define: a search state has variables $1 \dots k$ assigned. Values $k+1 \dots n$, as yet unassigned.
- Start state: All unassigned.
- Goal state: All assigned, and all constraints satisfied.
- Successors of a state with $V_1 \dots V_k$ assigned and rest unassigned are all states (with $V_1 \dots V_k$ the same) with V_{k+1} assigned a value from D .
- Cost on transitions: 0 is fine. We don't care. We just want any solution.

Slide 5

How to solve our CSP?



START = $(V_1=? V_2=? V_3=? V_4=? V_5=? V_6=?)$

succs(START) =

$(V_1=R V_2=? V_3=? V_4=? V_5=? V_6=?)$

$(V_1=G V_2=? V_3=? V_4=? V_5=? V_6=?)$

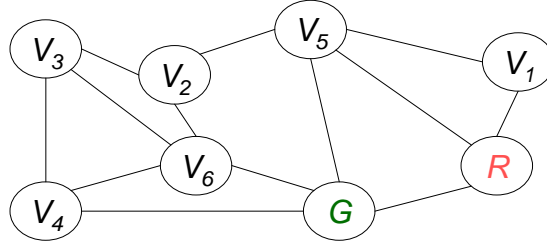
$(V_1=B V_2=? V_3=? V_4=? V_5=? V_6=?)$

What search algorithms could we use?

It turns out BFS is not a popular choice. Why not?

Slide 6

DFS for CSPs



What about DFS?

Much more popular. At least it has a chance of finding an easy answer quickly.

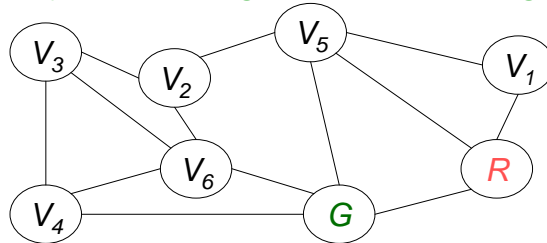
What happens if we do DFS with the order of assignments as *B* tried first, then *G* then *R*?

This makes DFS look very, very stupid!

Example: <http://www.cs.cmu.edu/~awm/animations/constraint/9d.html>

Slide 7

Blindly obvious improvement – Consistency Checking: “Backtracking Search”

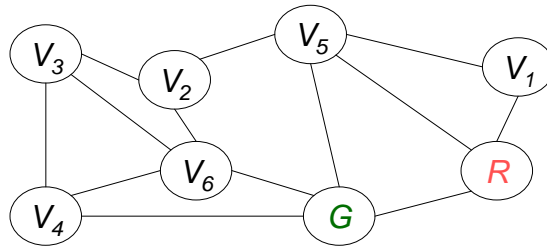


Don't ever try successor which causes inconsistency with its neighbors.

- Again, what happens if we do DFS with the order of assignments as *B* tried first, then *G* then *R*?
- What's the computational overhead for this?
- Backtracking still looks a little stupid!
- Examples: <http://www.cs.cmu.edu/~awm/animations/constraint/9b.html> and <http://www.cs.cmu.edu/~awm/animations/constraint/27b.html>

Slide 8

Obvious improvement – Forward Checking



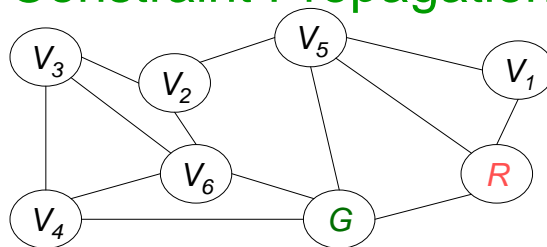
At start, for each variable, record the current set of possible legal values for it.

When you assign a value in the search, update set of legal values for all variables. Backtrack immediately if you empty a variable's constraint set.

- Again, what happens if we do DFS with the order of assignments as *B* tried first, then *G* then *R*?
- Example: <http://www.cs.cmu.edu/~awm/animations/constraint/27f.html>
- What's the computational overhead?

Slide 9

Constraint Propagation



Forward checking computes the domain of each variable independently at the start, and then only updates these domains when assignments are made in the DFS that are directly relevant to the current variable.

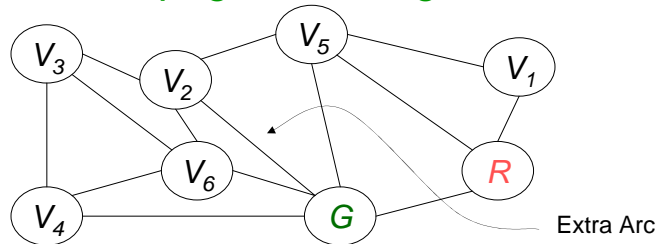
Constraint Propagation carries this further. When you delete a value from your domain, check all variables connected to you. If any of them change, delete all inconsistent values connected to them, etc...

In the above example it is useless

Web Example: <http://www.cs.cmu.edu/~awm/animations/constraint/27p.html>

Slide 10

Constraint Propagation being non-useless



- In this example, constraint propagation solves the problem without search ... **Not always that lucky!**
- Constraint propagation can be done as a preprocessing step. (Cheap).
- Or it can be maintained dynamically during the search. Expensive: when you backtrack, you must undo some of your additional constraints.

Slide 11

Graph-coloring-specific Constraint Propagation

In the case of Graph Coloring, CP looks simple: after we've made a search step (instantiated a node with a color), propagate the color at that node.

PropagateColorAtNode(node,color)

1. remove color from all of "available lists" of our uninstantiated neighbors.
2. If any of these neighbors gets the empty set, it's time to backtrack.
3. Foreach n in these neighbors: if n previously had two or more available colors but now has only one color c, run PropagateColorAtNode(n,c)

Slide 12

Graph-coloring-specific Constraint Propagation

In the case of Graph Coloring, CP looks simple: after we've made a search step (instantiated a node with a color), propagate the color at that node.

PropagateColorAtNode(node,color)

1. remove color from all of "available lists" of our uninstantiated neighbors.
2. If any of these neighbors has a unique value...
3. If any of these neighbors has a unique value...

But for General CSP problems, constraint propagation can do much more than only propagating when a node gets a unique value...

Slide 13

A New CSP (where fancier propagation is possible)

- The semi magic square
- Each variable can have value 1, 2 or 3

V_1	V_2	V_3	← This row must sum to 6
V_4	V_5	V_6	← This row must sum to 6
V_7	V_8	V_9	← This row must sum to 6
↑ This column must sum to 6	↑ This column must sum to 6	↑ This column must sum to 6	↙ This diagonal must sum to 6

Slide 14

General Constraint Propagation

Propagate(A_1, A_2, \dots, A_n)

finished = FALSE

while not finished

finished = TRUE

foreach constraint C

Assume C concerns variables V_1, V_2, \dots, V_k

Set $NewA_{V_1} = \{\}, NewA_{V_2} = \{\}, \dots, NewA_{V_k} = \{\}$

Foreach assignment $(V_1=x_1, V_2=x_2, \dots, V_k=x_k)$ in C

If x_1 in A_{V_1} and x_2 in A_{V_2} and ... x_k in A_{V_k}

Add x_1 to $NewA_{V_1}, x_2$ to $NewA_{V_2}, \dots, x_k$ to $NewA_{V_k}$

for $i = 1, 2 \dots k$

$A_{V_i} := A_{V_i}$ intersection $NewA_{V_i}$

If A_{V_i} was made smaller by that intersection

finished = FALSE

If A_{V_i} is empty, we're toast. Break out with "Backtrack" signal.

Specification: Takes a set of availability-lists for each and every node and uses all the constraints to filter out impossible values that are currently in availability lists

Details on next slide

Slide 15

General Constraint Propagation

Propagate(A_1, A_2, \dots, A_n)

finished = FALSE

while not finished

finished = TRUE

foreach constraint C

Assume C concerns variables V_1, V_2, \dots, V_k

Set $NewA_{V_1} = \{\}, NewA_{V_2} = \{\}, \dots, NewA_{V_k} = \{\}$

Foreach assignment $(V_1=x_1, V_2=x_2, \dots, V_k=x_k)$ in C

If x_1 in A_{V_1} and x_2 in A_{V_2} and ... x_k in A_{V_k}

Add x_1 to $NewA_{V_1}, x_2$ to $NewA_{V_2}, \dots, x_k$ to $NewA_{V_k}$

for $i = 1, 2 \dots k$

$A_{V_i} := A_{V_i}$ intersection $NewA_{V_i}$

If A_{V_i} was made smaller by that intersection

finished = FALSE

If A_{V_i} is empty, we're toast. Break out with "Backtrack" signal.

A_i denotes the current set of possible values for variable i . This is call-by-reference. Some of the A_i sets may be changed by this call (they'll have one or more elements removed)

We'll keep iterating until we do a full iteration in which none of the availability lists change. The "finished" flag is just to record whether a change took place.

Slide 16

General Constraint Propagation

Propagate(A_1, A_2, \dots, A_n)

finished = FALSE

while not finished

finished = TRUE

foreach constraint C

Assume C concerns variables V_1, V_2, \dots, V_k

Set $NewA_{V_1} = \{\}, NewA_{V_2} = \{\}, \dots, NewA_{V_k} = \{\}$

Foreach assignment $(V_1=x_1, V_2=x_2, \dots, V_k=x_k)$ in C

If x_1 in A_{V_1} and x_2 in A_{V_2} and ... x_k in A_{V_k}

Add x_1 to $NewA_{V_1}$, x_2 to $NewA_{V_2}$, ... x_k to $NewA_{V_k}$

for $i = 1, 2, \dots, k$

$A_{V_i} := A_{V_i}$ intersection $NewA_{V_i}$

If A_{V_i} was made smaller by that intersection

finished = FALSE

If A_{V_i} is empty, we're toast. Break

$NewA_i$ is going to be filled up with the possible values for variable V_i taking into account the effects of constraint C

After we've finished all the iterations of the foreach loop, $NewA_i$ contains the full set of possible values of variable V_i taking into account the effects of constraint C .

General Constraint Propagation

Propagate(A_1, A_2, \dots, A_n)

finished = FALSE

while not finished

finished = TRUE

foreach constraint C

Assume C concerns variables V_1, V_2, \dots, V_k

Set $NewA_{V_1} = \{\}, NewA_{V_2} = \{\}, \dots, NewA_{V_k} = \{\}$

Foreach assignment $(V_1=x_1, V_2=x_2, \dots, V_k=x_k)$ in C

If x_1 in A_{V_1} and x_2 in A_{V_2} and ... x_k in A_{V_k}

Add x_1 to $NewA_{V_1}$, x_2 to $NewA_{V_2}$, ... x_k to $NewA_{V_k}$

for $i = 1, 2, \dots, k$

$A_{V_i} := A_{V_i}$ intersection $NewA_{V_i}$

If A_{V_i} was made smaller by that intersection

finished = FALSE

If A_{V_i} is empty, we're toast. Break out with "Backtrack" signal.

If A_{V_i} is empty we've proved that there are no solutions for the availability-lists that we originally entered the function with

If this test is satisfied that means that there's at least one value q such that we originally thought q was an available value for V_i but we now know q is impossible.

Slide 18

Propagate on Semi-magic Square

- The semi magic square
- Each variable can have value 1, 2 or 3

1	123	123	← This row must sum to 6
123	123	123	← This row must sum to 6
123	123	123	← This row must sum to 6
↑ This column must sum to 6	↑ This column must sum to 6	↑ This column must sum to 6	↙ This diagonal must sum to 6

Slide 19

Propagate on Semi-magic Square

- The semi magic square
- Each variable can have value 1, 2 or 3

(V_1, V_2, V_3) must be one of
 (1,2,3)
 (1,3,2)
 (2,1,3)
 (2,2,2)
 (2,3,1)
 (3,1,2)
 (3,2,1)

1	123	123	← This row must sum to 6
123	123	123	← This row must sum to 6
123	123	123	← This row must sum to 6
↑ This column must sum to 6	↑ This column must sum to 6	↑ This column must sum to 6	↙ This diagonal must sum to 6

Slide 20

Propagate on Set

- $\text{NewAL}_{V_1} = \{ 1 \}$
 - $\text{NewAL}_{V_2} = \{ 2, 3 \}$
 - $\text{NewAL}_{V_3} = \{ 2, 3 \}$
- are
- Each variable can have value

(V_1, V_2, V_3) must be one of

- (1,2,3)
- (1,3,2)
- (2,1,3)
- (2,2,2)
- (2,3,1)
- (3,1,2)
- (3,2,1)

1	123	123	← This row must sum to 6
123	123	123	← This row must sum to 6
123	123	123	← This row must sum to 6
↑ This column must sum to 6	↑ This column must sum to 6	↑ This column must sum to 6	← This diagonal must sum to 6

Slide 21

After doing first row constraint...

1	23	23	← This row must sum to 6
123	123	123	← This row must sum to 6
123	123	123	← This row must sum to 6
↑ This column must sum to 6	↑ This column must sum to 6	↑ This column must sum to 6	← This diagonal must sum to 6

Slide 22

After doing all row constraints and column constraints...

1	23	23	← This row must sum to 6
23	123	123	← This row must sum to 6
23	123	123	← This row must sum to 6
↑ This column must sum to 6	↑ This column must sum to 6	↑ This column must sum to 6	↙ This diagonal must sum to 6

Slide 23

And after doing diagonal constraint...

1	23	23	← This row must sum to 6
23	23	123	← This row must sum to 6
23	123	23	← This row must sum to 6
↑ This column must sum to 6	↑ This column must sum to 6	↑ This column must sum to 6	↙ This diagonal must sum to 6

CP has now iterated through all constraints once. But does it make further progress when it tries iterating through them again?

Slide 24

And after doing another round of constraints...

1	23	23	← This row must sum to 6
23	23	12	← This row must sum to 6
23	12	23	← This row must sum to 6
↑ This column must sum to 6	↑ This column must sum to 6	↑ This column must sum to 6	← This diagonal must sum to 6

YES! And this showed a case of a constraint applying even when none of the variables involved was down to a unique value.

So.. any more changes on the next iteration?

Slide 25

CSP Search with Constraint Propagation

CPSearch(A_1, A_2, \dots, A_n)

Let i = lowest index such that A_i has more than one value

foreach available value x in A_i

 foreach k in $1, 2.. n$

 Define $A'_k := A_k$

$A'_i := \{x\}$

 Call Propagate(A'_1, A'_2, \dots, A'_n)

 If no "Backtrack" signal

 If A'_1, A'_2, \dots, A'_n are all unique we're done!

 Recursively Call CPSearch(A'_1, A'_2, \dots, A'_n)

Details on next slide

Slide 26

CSP Search with Constraint Propagation

Specification: Find out if there's any combination of values in the combination of the given availability lists that satisfies all constraints.

CPSearch(A_1, A_2, \dots, A_n)

Let i = lowest index such that A_i has more than one value

foreach available value x in A_i

foreach k in $1, 2.. n$

Define $A'_k := A_k$

$A'_i := \{x\}$

Call Propagate(A'_1, A'_2, \dots, A'_n)

If no "Backtrack" signal

If A'_1, A'_2, \dots, A'_n are all unique we're done!

Recursively Call CPMsearch(A'_1, A'_2, \dots, A'_n)

At this point the A-primes are a copy of the original availability lists except A'_i has committed to value x .

This call may prune away some values in some of the copied availability lists

Assuming that we terminate deep in the recursion if we find a solution, the CPMsearch function only terminates normally if no solution is found.

Slide 27

CSP Search with Constraint Propagation

CPSearch(A_1, A_2, \dots, A_n)

Let i = lowest index such that A_i has more than one value

foreach available value x in A_i

foreach k in $1, 2.. n$

Define $A'_k := A_k$

$A'_i := \{x\}$

Call Propagate(A'_1, A'_2, \dots, A'_n)

If no "Backtrack" signal

If A'_1, A'_2, \dots, A'_n are all unique we're done!

Recursively Call CPMsearch(A'_1, A'_2, \dots, A'_n)

What's the top-level call?

Slide 28

CSP Search with Constraint Propagation

CPSearch(A_1, A_2, \dots, A_n)

Let i = lowest index such that A_i has more than one value

foreach available value x in A_i

 foreach k in $1, 2.. n$

 Define $A'_k := A_k$

$A'_i := \{x\}$

 Call Propagate(A'_1, A'_2, \dots, A'_n)

 If no "Backtrack" signal

 If A'_1, A'_2, \dots, A'_n are all unique we're done!

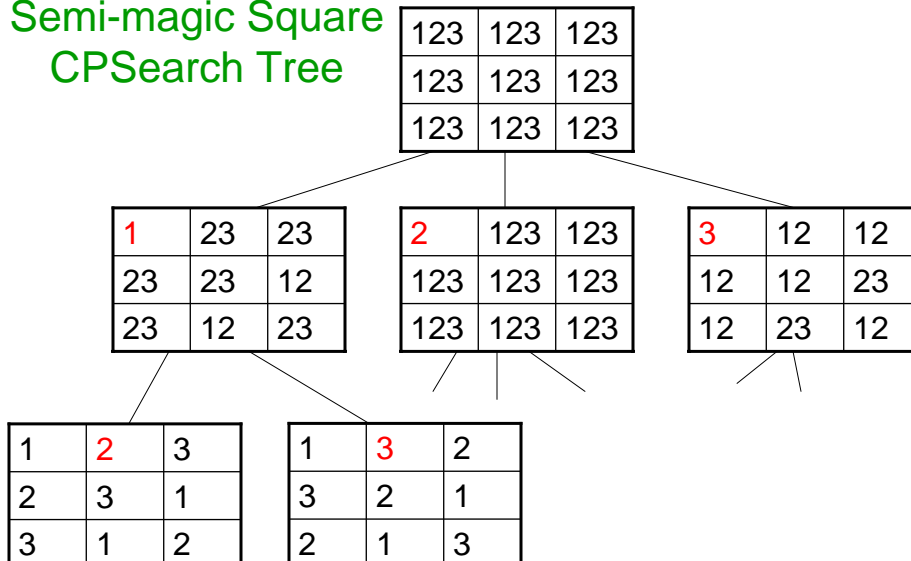
 Recursively Call CPMsearch(A'_1, A'_2, \dots, A'_n)

What's the top-level call?

Call with that A_i = complete set of possible values for V_i .

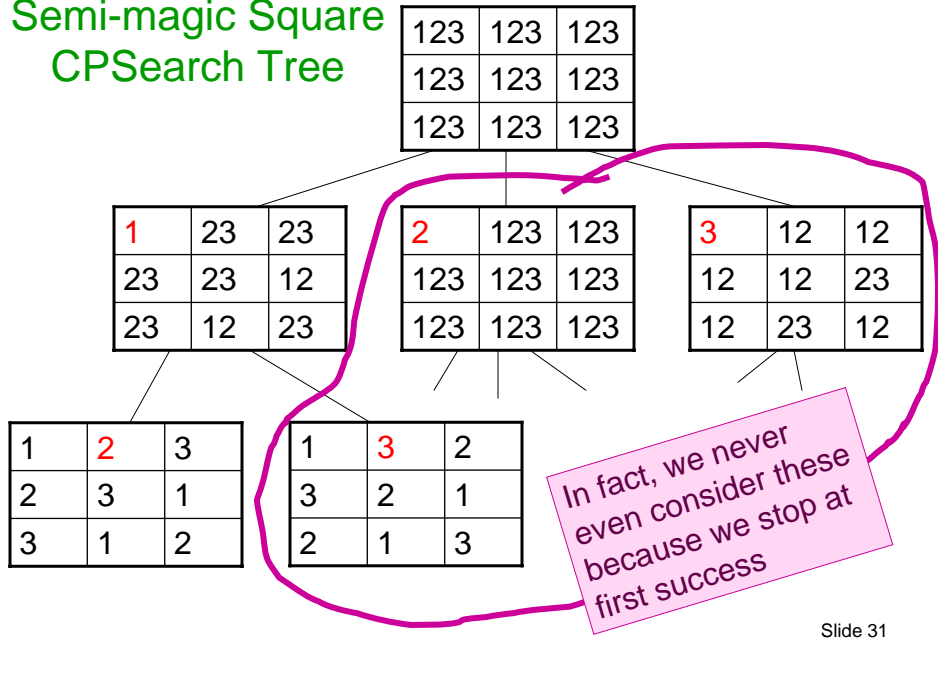
Slide 29

Semi-magic Square CPSearch Tree



Slide 30

Semi-magic Square CPSearch Tree



Some real CSPs

- Graph coloring is a real, and useful, CSP. Applied to problems with many hundreds of thousands of nodes. Not very AI-esque.
- VLSI or PCB board layout.
- Selecting a move in the game of “minesweeper”.

0	0	1			
0	0	1			
0	0	1			
1	1	2			

Which squares have a bomb? Squares with numbers don't. Other squares might. Numbers tell how many of the eight adjacent squares have bombs. We want to find out if a given square can possibly have a bomb....

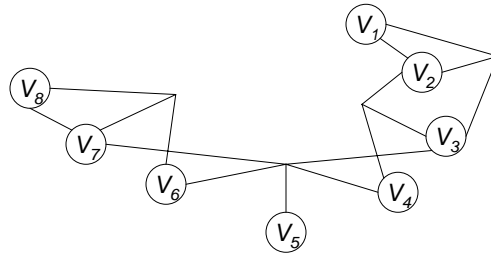
Slide 32

“Minesweeper” CSP

0	0	1	V_1		
0	0	1	V_2		
0	0	1	V_3		
1	1	2	V_4		
V_8	V_7	V_6	V_5		

$V = \{ V_1, V_2, V_3, V_4, V_5, V_6, V_7, V_8 \}, D = \{ B \text{ (bomb)}, S \text{ (space)} \}$

$C = \{ (V_1, V_2) : \{ (B,S), (S,B) \}, (V_1, V_2, V_3) : \{ (B,S,S), (S,B,S), (S,S,B) \}, \dots \}$

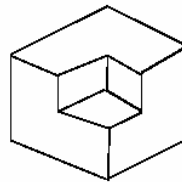
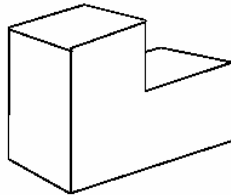


Slide 33

The Waltz algorithm

One of the earliest examples of a computation posed as a CSP.

The Waltz algorithm is for interpreting line drawings of solid polyhedra.



Look at all intersections.



What kind of intersection could this be? A concave intersection of three faces? Or an external convex intersection?

Adjacent intersections impose constraints on each other. Use CSP to find a unique set of labelings. Important step to “understanding” the image.

Slide 34

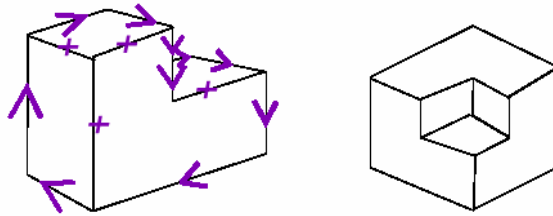
Waltz Alg. on simple scenes

Assume all objects:

- Have no shadows or cracks
- Three-faced vertices
- "General position": no junctions change with small movements of the eye.


Then each line on image is one of the following:

- Boundary line (edge of an object) (<) with right hand of arrow denoting "solid" and left hand denoting "space"
- Interior convex edge (+)
- Interior concave edge (-)



Slide 35

18 legal kinds of junctions

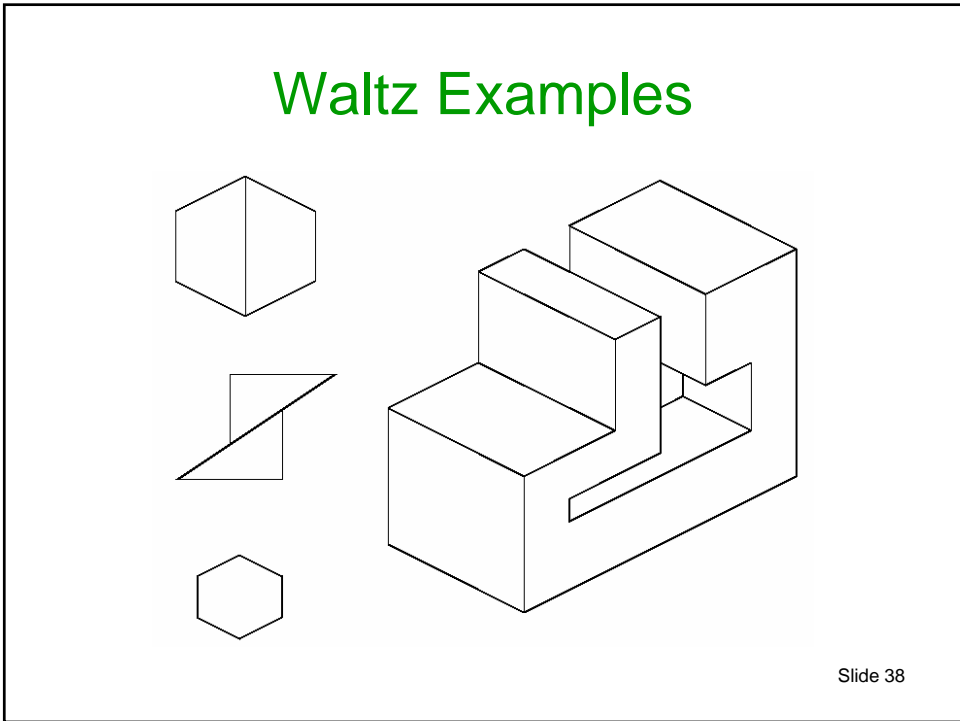



Given a representation of the diagram, label each junction in one of the above manners.

The junctions must be labeled so that lines are labeled consistently at both ends.

Can you formulate that as a CSP? *FUN FACT: Constraint Propagation always works perfectly.*

Slide 36



Scheduling

A very big, important use of CSP methods.

- Used in many industries. Makes many multi-million dollar decisions.
- Used extensively for space mission planning.
- Military uses.

People *really care* about improving scheduling algorithms!

Problems with phenomenally huge state spaces. But for which solutions are needed very quickly.

Many kinds of scheduling problems e.g.:

- ❖ *Job shop*: Discrete time; weird ordering of operations possible; set of separate jobs.
- ❖ *Batch shop*: Discrete or continuous time; restricted operation of ordering; grouping is important.
- ❖ *Manufacturing cell*: Discrete, automated version of open job shop.

Slide 39

Job Shop scheduling

At a job-shop you make various products. Each product is a "job" to be done. E.G.

Job₁ = Make a polished-thing-with-a-hole

Job₂ = Paint and drill a hole in a widget

Each job requires several operations. E.G.

Operations for Job₁: Polish, Drill

Operations for Job₂: Paint, Drill

Each operation needs several resources. E.G.

Polishing needs the Polishing machine

Polishing needs Pat (a Polishing expert)

Drilling needs the Drill

Drilling needs Pat (also a Drilling expert)

Or Drilling can be done by Chris

Some operations need to be done in a particular order (e.g. Paint after you've Drilled)

Slide 40

Job Shop Formalized

A Job Shop problem is a pair (J, RES)

J is a set of jobs $J = \{j_1, j_2, \dots, j_n\}$

RES is a set of resources $RES = \{R_1 .. R_m\}$

Each job j_i is specified by:

- a set of operations $O^i = \{O^i_1, O^i_2 \dots O^i_{n(i)}\}$
- and must be carried out between release-date rd_i and due-date dd_i .
- and a partial order of operations: $(O^i_r \text{ before } O^i_j)$, $(O^i_r \text{ before } O^i_j)$, etc...

Each operation O^i_j has a variable start time st^i_j and a fixed duration du^i_j and requires a set of resources. e.g.: O^i_j requires $\{R^i_{j1}, R^i_{j2} \dots\}$.

Each resource can be accomplished by one of several possible physical resources, e.g. R^i_{j1} might be accomplished by any one of $\{r^i_{ij1}, r^i_{ij2}, \dots\}$. Each of the r^i_{ijk} s are a member of RES .

Slide 41

Job Shop Example

$j_1 = \text{polished-hole-thing} = \{O^1_1, O^1_2\}$

$j_2 = \text{painted-hole-widget} = \{O^2_1, O^2_2\}$

$RES = \{Pat, Chris, Drill, Paint, Drill, Polisher\}$

$O^1_1 = \text{polish-thing: need resources...}$

$\{R^1_{11} = Pat, R^1_{12} = Polisher\}$

$O^1_2 = \text{drill-thing: need resources...}$

$\{R^1_{21} = (r^1_{211}=Pat \text{ or } r^1_{212}=Chris), R^1_{22} = Drill\}$

$O^2_1 = \text{paint-widget: need resources...}$

$\{R^2_{11} = Paint\}$

$O^2_2 = \text{drill-widget: need resources...}$

$\{R^2_{21} = (r^2_{211}=Pat \text{ or } r^2_{212}=Chris), R^2_{22} = Drill\}$

Precedence constraints : O^2_2 before O^2_1 . All operations take one time unit $du^i_j = 1$ for all i, j . Both jobs have release-date $rd^i = 0$ and due-date $dd^i = 1$.

Slide 42

Job-shop: the Variables and Constraints

Variables

- The operation state times st_i^t
- The resources R_{ij}^l (usually these are obvious from the definition of O_i^l . Only need to be assigned values when there are alternative physical resources available, e.g. *Pat* or *Chris* for operating the *drill*).

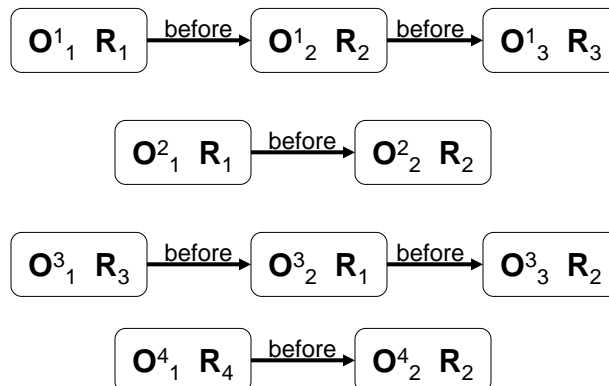
Constraints:

- Precedence constraints. (Some O_i 's must be before some other O_j 's).
- Capacity constraints. There must never be a pair of operations with overlapping periods of operation that use the same resources.

Non-challenging question. Can you schedule our Job-shop?

Slide 43

A slightly bigger example



Example from [Sadeh and Fox, 96]: Norman M. Sadeh and Mark S. Fox, Variable and Value Ordering Heuristics for the Job Shop Scheduling Constraint Satisfaction Problem, Artificial Intelligence Journal, Number Vol 86, No1, pages 1-41, 1996. Available from citeseer.nj.nec.com/sadeh96variable.html

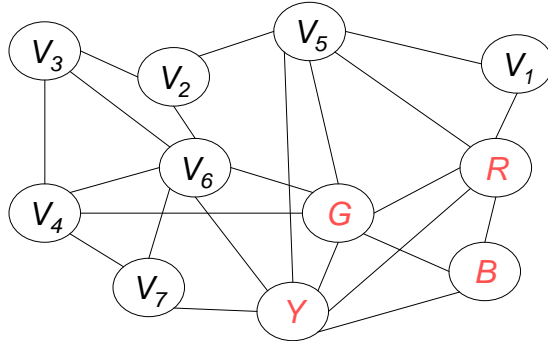
4 jobs. Each 3 units long. All jobs have release date 0 and due date 15. All operations use only one resource each.

Slide 44

Further CSP techniques

Let's look at some other important CSP methods. Keep the job-shop example in mind.

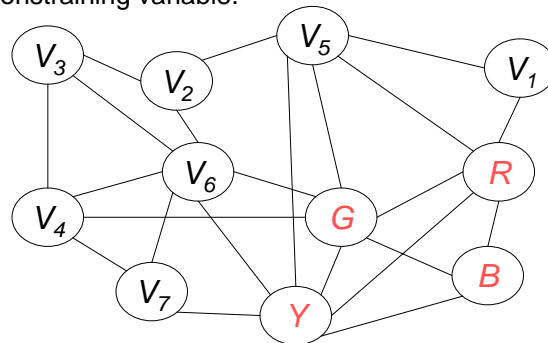
Here's another graph-coloring example (you're now allowed R , G , B and Y)



Slide 45

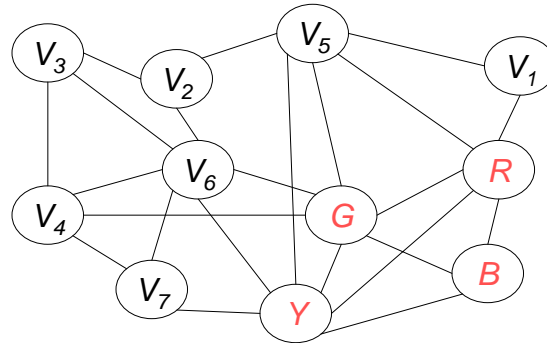
General purpose Variable Ordering Heuristics

1. Most constrained variable.
2. Most constraining variable.



Slide 46

General purpose Value Ordering Heuristics



A good general purpose one is “least-constrained-value”. Choose the value which causes the smallest reduction in number of available values for neighboring variables

Slide 47

General purpose CSP algorithm

(From Sadeh+Fox)

1. If all values have been successfully assigned then stop, else go on to 2.
2. Apply the consistency enforcing procedure (e.g. forward-checking if feeling computationally mean, or constraint propagation if extravagant. There are other possibilities, too.)
3. If a deadend is detected then backtrack (simplest case: DFS-type backtrack. Other options can be tried, too). Else go on to step 4.
4. Select the next variable to be assigned (using your variable ordering heuristic).
5. Select a promising value (using your value ordering heuristic).
6. Create a new search state. Cache the info you need for backtracking. And go back to 1.

Slide 48

Job-shop example. Consistency enforcement

Sadeh claims that generally forward-checking is better, computationally, than full constraint propagation. But it can be supplemented with a Job-shop specific **TRICK**.

The precedence constraints (i.e. the available times for the operations to start due to the ordering of operations) can be computed exactly, given a partial schedule, very efficiently.

Slide 49

Reactive CSP solutions

- Say you have built a large schedule.
- Disaster! Halfway through execution, one of the resources breaks down. We have to reschedule!
- Bad to have to wait 15 minutes for the scheduler to make a new suggestion.

Important area of research: efficient schedule repair algorithms.

- Question: If you expect that resources may sometimes break, what could a scheduling program do to take that into account?
- Unrelated Question: Why has none of this lecture used A*?

Slide 50

Other approaches. And What You Should Know

Other Approaches:

- Hill-climbing, Tabu-search, Simulated annealing, Genetic Algorithms. (to be discussed later)

What you should know:

- ✓ How to formalize problems as CSPs
- ✓ Backtracking Search, Forward Checking, Constraint Propagation
- ✓ The Waltz algorithm
- ✓ You should understand and appreciate the way job-shop scheduling is formalized. It is an excellent representative example of how important well-studied constraint satisfaction problems are represented.
- ✓ Understand examples of Variable ordering and Value ordering heuristics

In those cases where your lecturer or these handouts are too incomprehensible, consult Chap 5 of the Russell handout. Winston's "Artificial Intelligence" book has good discussion of constraint satisfaction and Waltz algorithm.

Slide 51