



# Constraint Programming

Table Constraint

Hélène Verhaeghe & Pierre Schaus

# Table Constraint

x	y	z
1	2	3
1	3	3
2	1	3
2	1	1
3	3	3
4	1	2
4	4	4

A table constraint has an enumeration of the possible assignments for its variables (here x, y, and z).

## Semantics

$(x=1 \wedge y=2 \wedge z=3) \vee (x=1 \wedge y=3 \wedge z=3) \vee (x=2 \wedge y=1 \wedge z=3) \vee \dots$

## Signature

```
/**  
 * Fixing  $x_0 = v_0, x_1 = v_1, \dots$  is only  
 * valid if there exists a row  $(v_0, v_1, \dots)$  in table.  
 */  
public Table(IntVar[] x, int[][] table)
```

# Intensional vs Extensional Formulation

- A constraint like `AllDifferent([x,y,z])` is said to be **intensional**. The solution set to the constraint is *implicit* with the semantics of the constraint.

- To make it *explicit*, via an **extensional** formulation, aka a Table constraint, we list *all* the solutions. For  $D(x) = D(y) = D(z) = \{0..2\}$  we have

- For  $n$  or  $n-1$  variables over a domain of size  $n$ , the extensional formulation of `AllDifferent` requires  $n!$  tuples.  
That is why intensional formulations are interesting.

- An extensional formulation can impose *any* relation on its variables.  
That is why Table constraints are interesting:  
this is the most flexible form of constraints.



If an efficient intensional constraint with a domain-consistent filtering exists in your CP solver, then you should probably prefer to use it rather than an extensional formulation of the problem constraint.

x	y	z
0	1	2
0	2	1
1	0	2
1	2	0
2	0	1
2	1	0

# Most flexible constraint of the universe

- ▶ Any predicate on  $k$  variables can be turned into a table constraint
- ▶ Just enumerate the solutions to the constraint into a table

**X+Y=Z**

X	Y	Z
0	0	0
0	1	1
0	2	2
1	0	1
1	1	2
1	2	3
2	0	2
2	1	3
2	2	4

**if X is even, then Y=2,  
else Z>0**

X	Y	Z
0	2	0
0	2	1
0	2	2
1	0	1
1	0	2
1	1	1
1	1	2
1	2	1
1	2	2
2	2	0
2	2	1
2	2	2

**AllDifferent(X,Y,Z)**

X	Y	Z
0	1	2
0	2	1
1	0	2
1	2	0
2	0	1
2	1	0

# Most flexible constraint of the universe

- ▶ A practical example is the solving of the Enigma machine.
- ▶ Given an output and tiny clues about the input, can we find the full input and the settings of the Enigma machine?
- ▶ Yes! Using the Modulo constraint  $X \bmod Y = Z$ .

Antuori V., Portoleau T., Rivière L. and Hébrard E.  
On How Turing and Singleton Arc Consistency Broke the Enigma Code.  
CP 2021.

# Application: Enigma machine

- Composed of rotors, input = [0,3,4], encoded as [2,4,0]

- After each input, the rotor rotates by one position

- Step 0

- Input 0

- Output 2

- Step 1

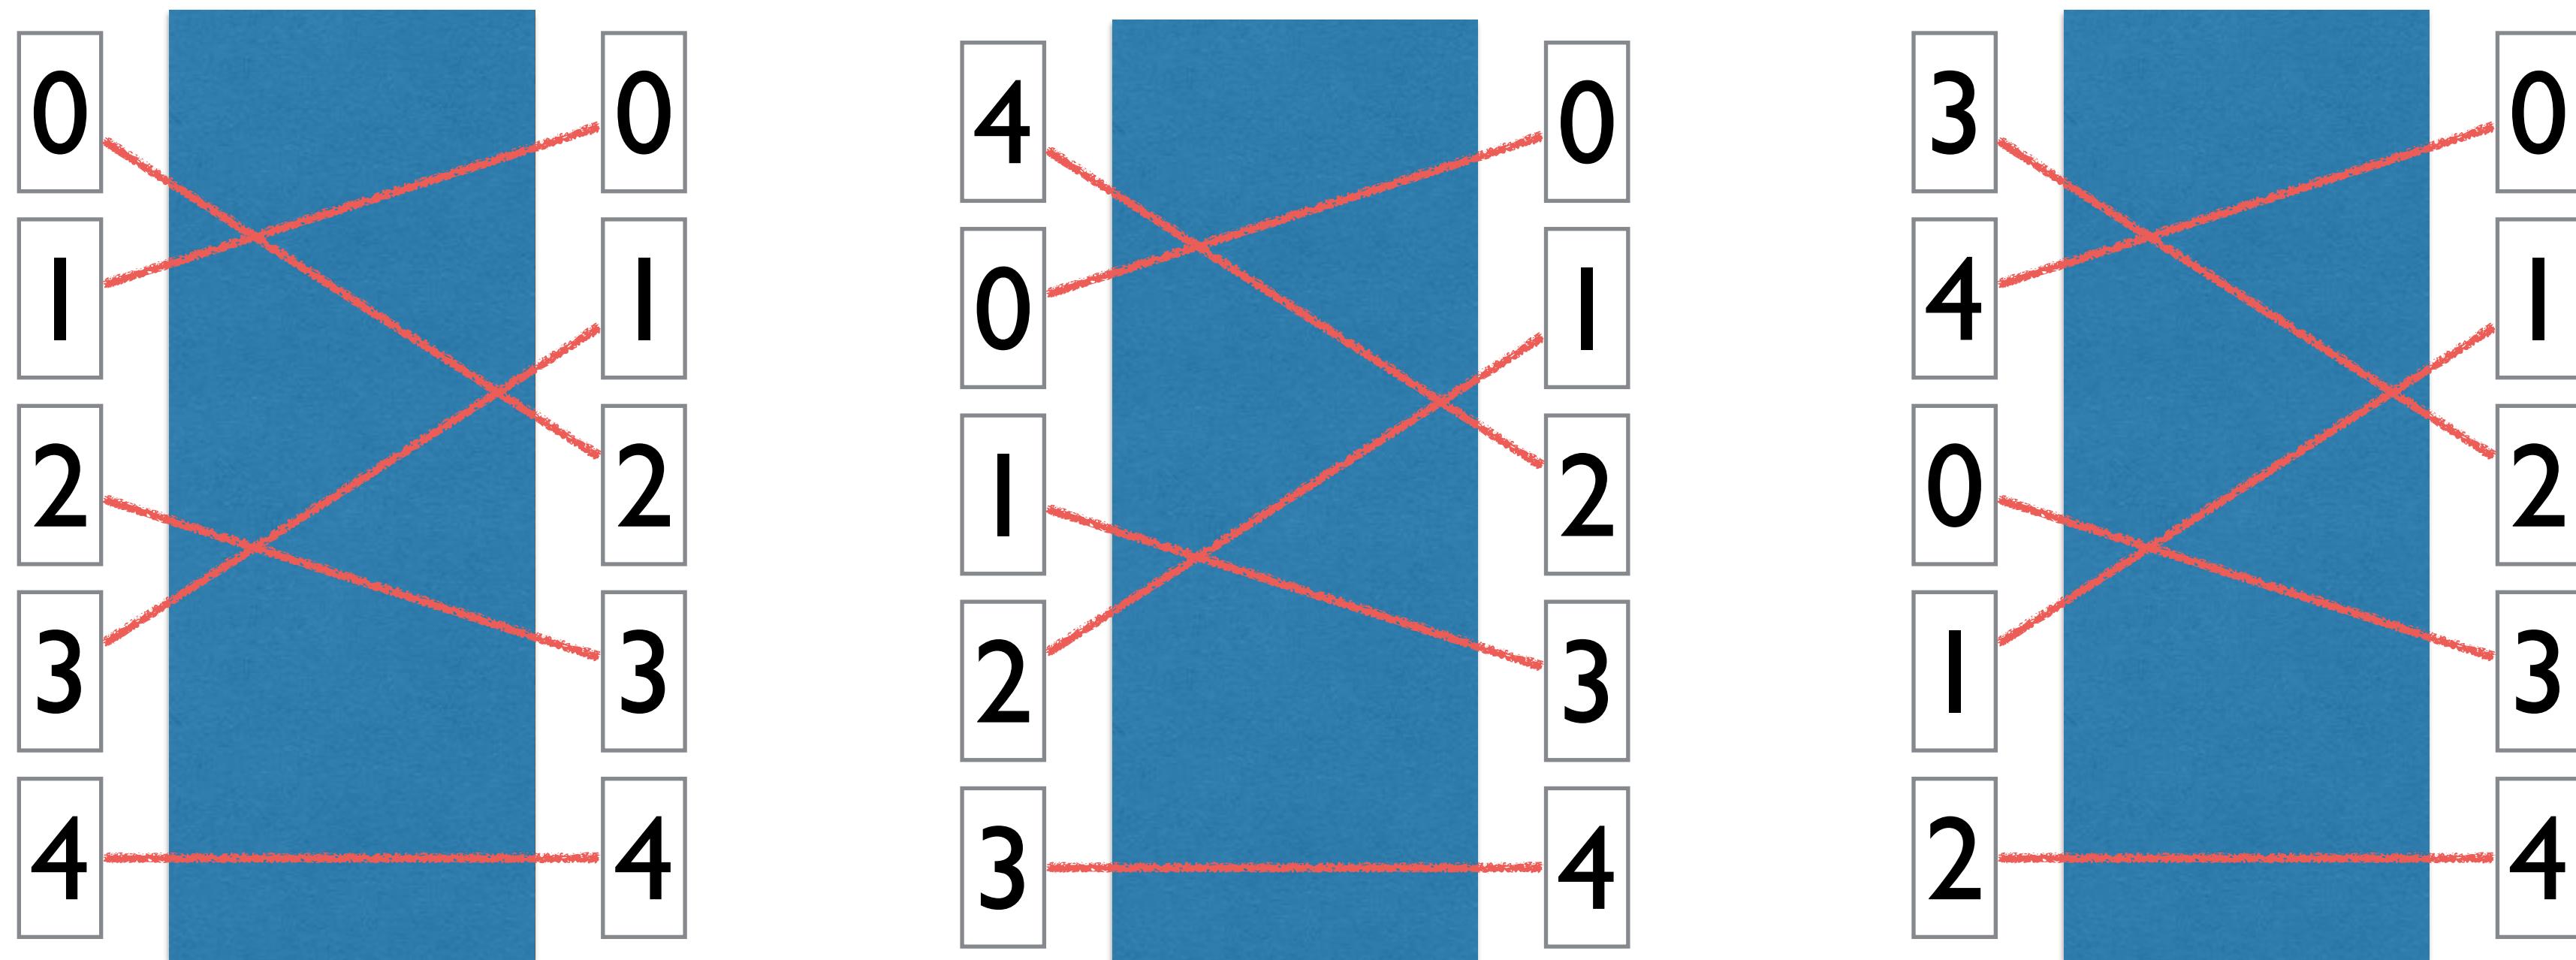
- Input 3

- Output 4

- Step 2

- Input 4

- Output 0



# Application: Enigma machine

- $T = [2,0,3,1,4]$  (mapping at initial position)
- What is the output at step  $i$  for input  $I$ ? Answer:  $\text{output} = T[(I+i)\%5]$

- Step 0

- Input 0

- Output 2

- Step 1

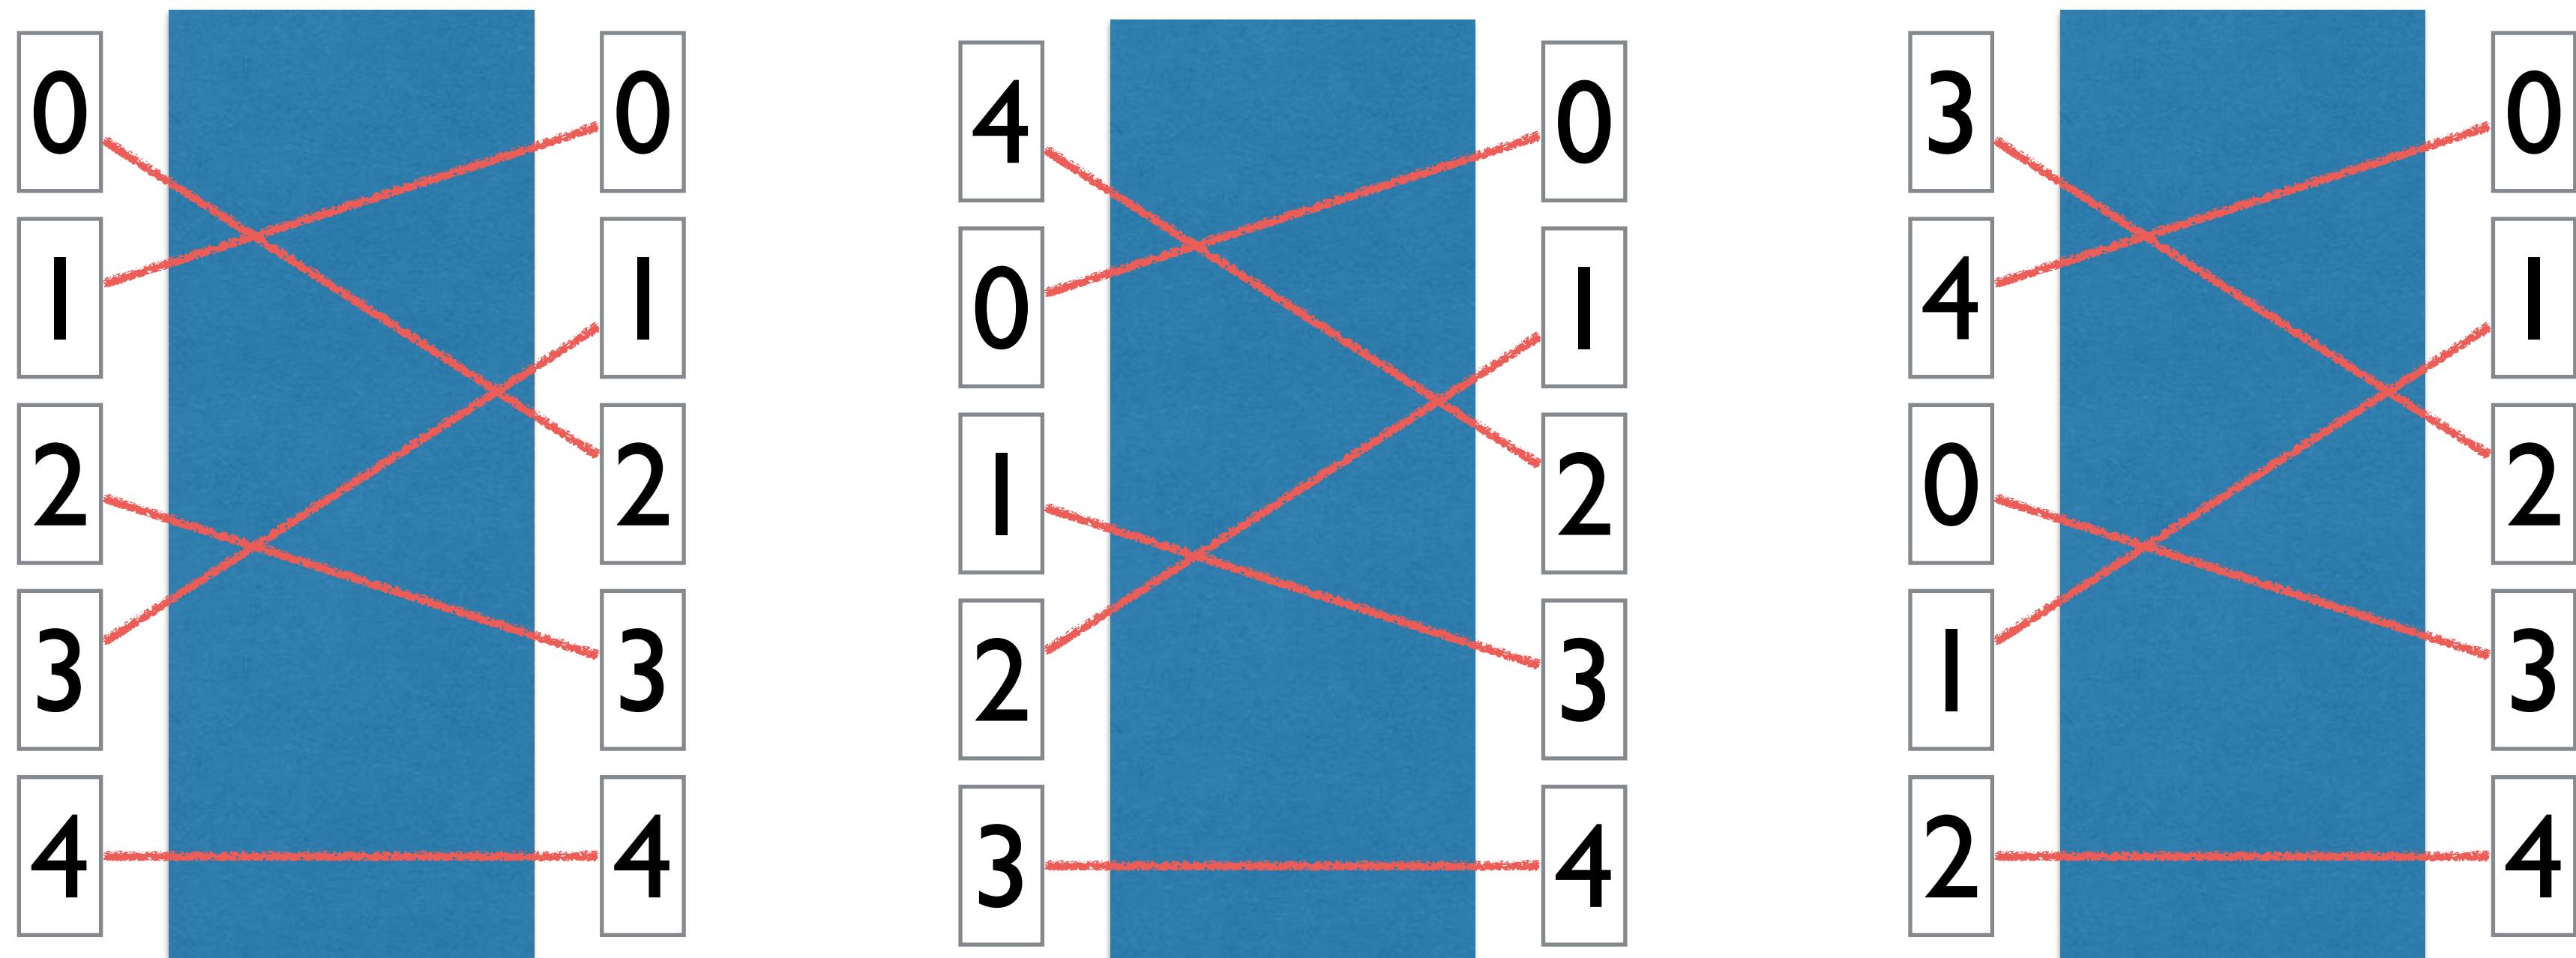
- Input 3

- Output 4

- Step 2

- Input 4

- Output 0



# Application: Enigma machine

- $T = [2,0,3,1,4]$  (mapping at initial position)
- What is the output at step  $i$  for input  $I$ ? Answer: output =  $T[(I+i)\%5] = A$

I	i	A
0	0	2
0	1	0
0	2	3
0	3	1
0	4	4
0	5	2
0	6	0
0	7	3
0	8	1
0	9	4
0	10	2
0	11	0
0	12	3
0	13	1
0	14	4
...	...	...

I	i	A
1	0	0
1	1	3
1	2	1
1	3	4
1	4	2
1	5	0
1	6	3
1	7	1
1	8	4
1	9	2
1	10	0
1	11	3
1	12	1
1	13	4
1	14	2
...	...	...

I	i	A
2	0	3
2	1	1
2	2	4
2	3	2
2	4	0
2	5	3
2	6	1
2	7	4
2	8	2
2	9	0
2	10	3
2	11	1
2	12	4
2	13	2
2	14	0
...	...	...

I	i	A
3	0	1
3	1	4
3	2	2
3	3	0
3	4	3
3	5	1
3	6	4
3	7	2
3	8	0
3	9	3
3	10	1
3	11	4
3	12	2
3	13	0
3	14	3
...	...	...

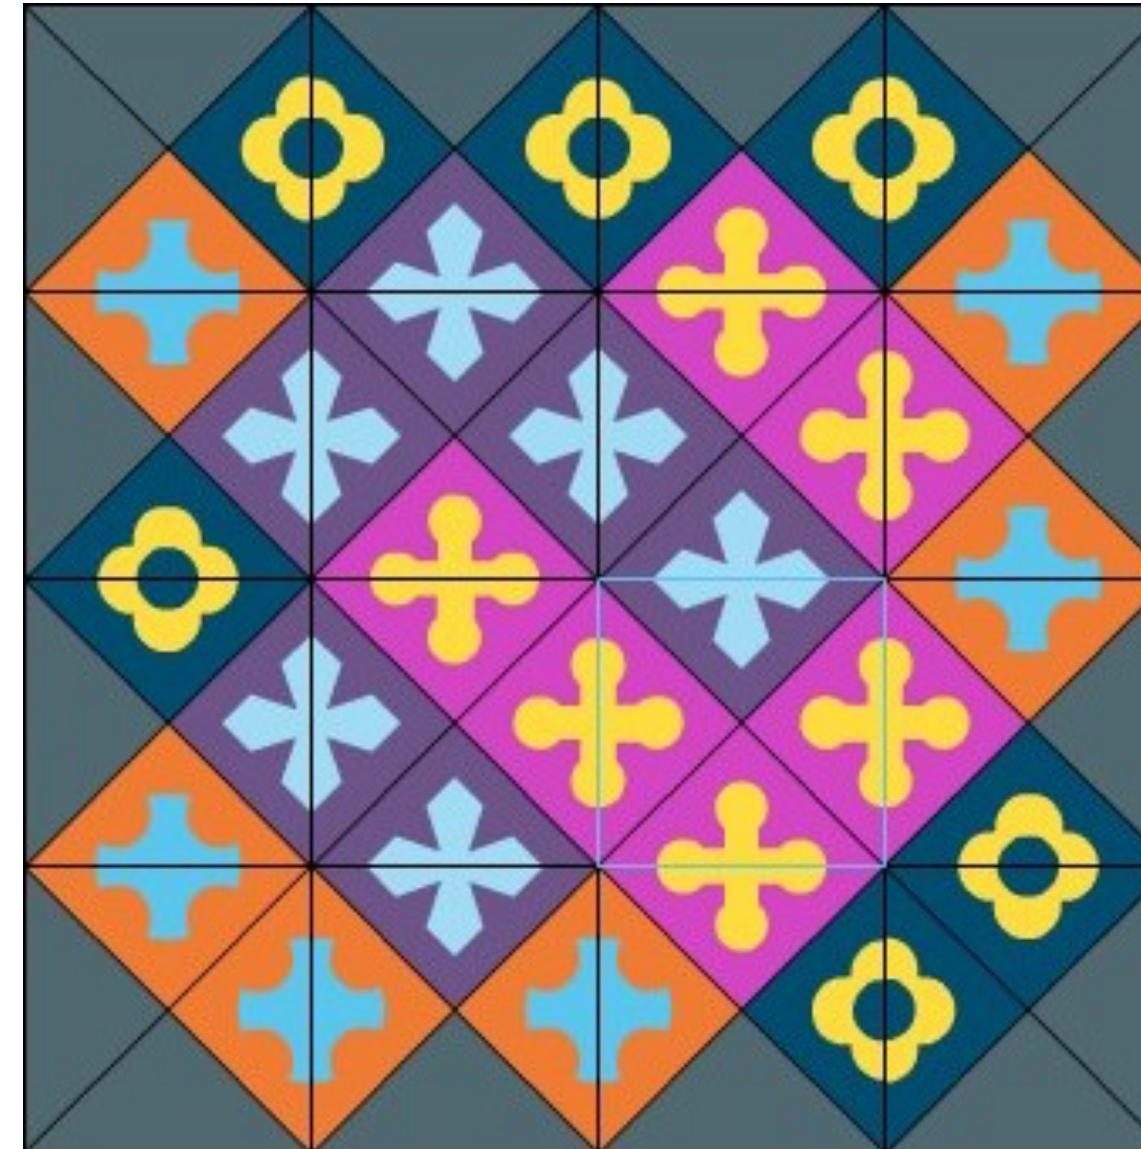
I	i	A
4	0	4
4	1	2
4	2	0
4	3	3
4	4	1
4	5	4
4	6	2
4	7	0
4	8	3
4	9	1
4	10	4
4	11	2
4	12	0
4	13	3
4	14	1
...	...	...



# Application of Table constraints: Eternity Puzzle

# Eternity II Puzzle

- Edge matching puzzle: place 256 square pieces into a 16x16 grid, constrained by the requirement to match adjacent edges.



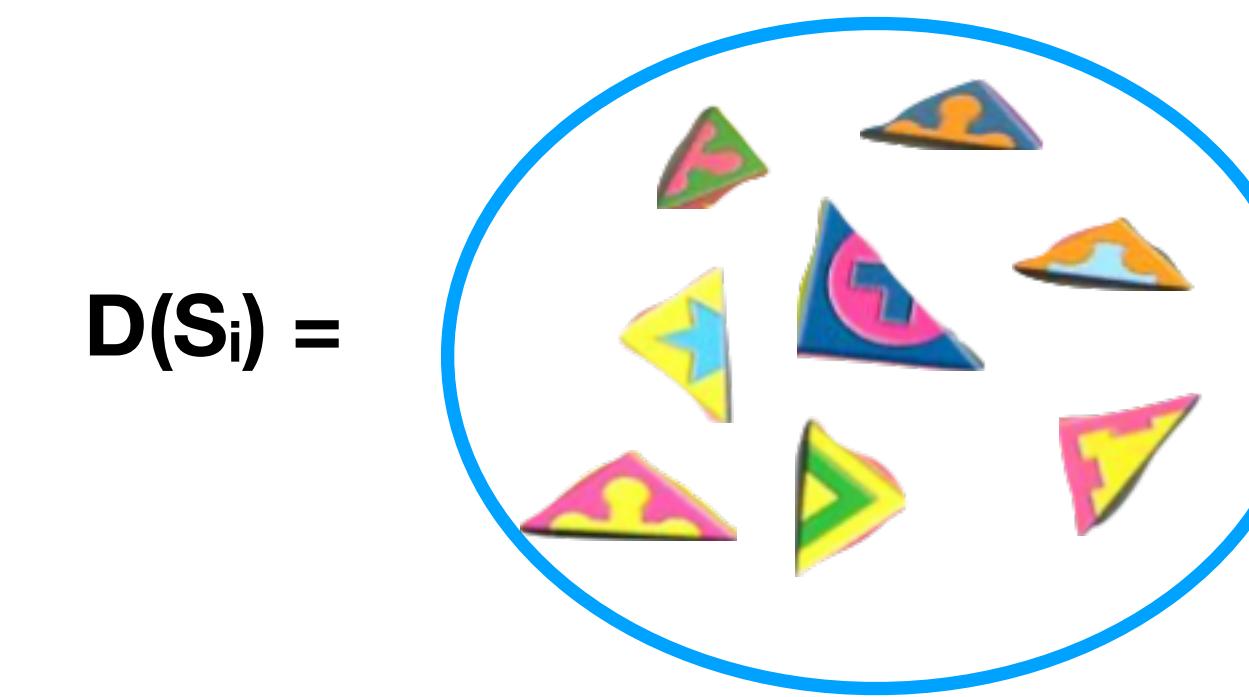
- How to model this puzzle?
- [https://en.wikipedia.org/wiki/Eternity\\_II\\_puzzle](https://en.wikipedia.org/wiki/Eternity_II_puzzle)

# Decision variables for Eternity II



For each position, we need 5 variables:

- 1 variable  $S_i$  (but  $Y_i$  on the picture) for each of the 4 sides
- 1 variable  $I$  for the identifier of the placed piece



How do we model that  
 $(I, S1, S2, S3, S4)$  corresponds to a  
valid piece of the game, such as  
**this one**?

# Decision variables for Eternity II



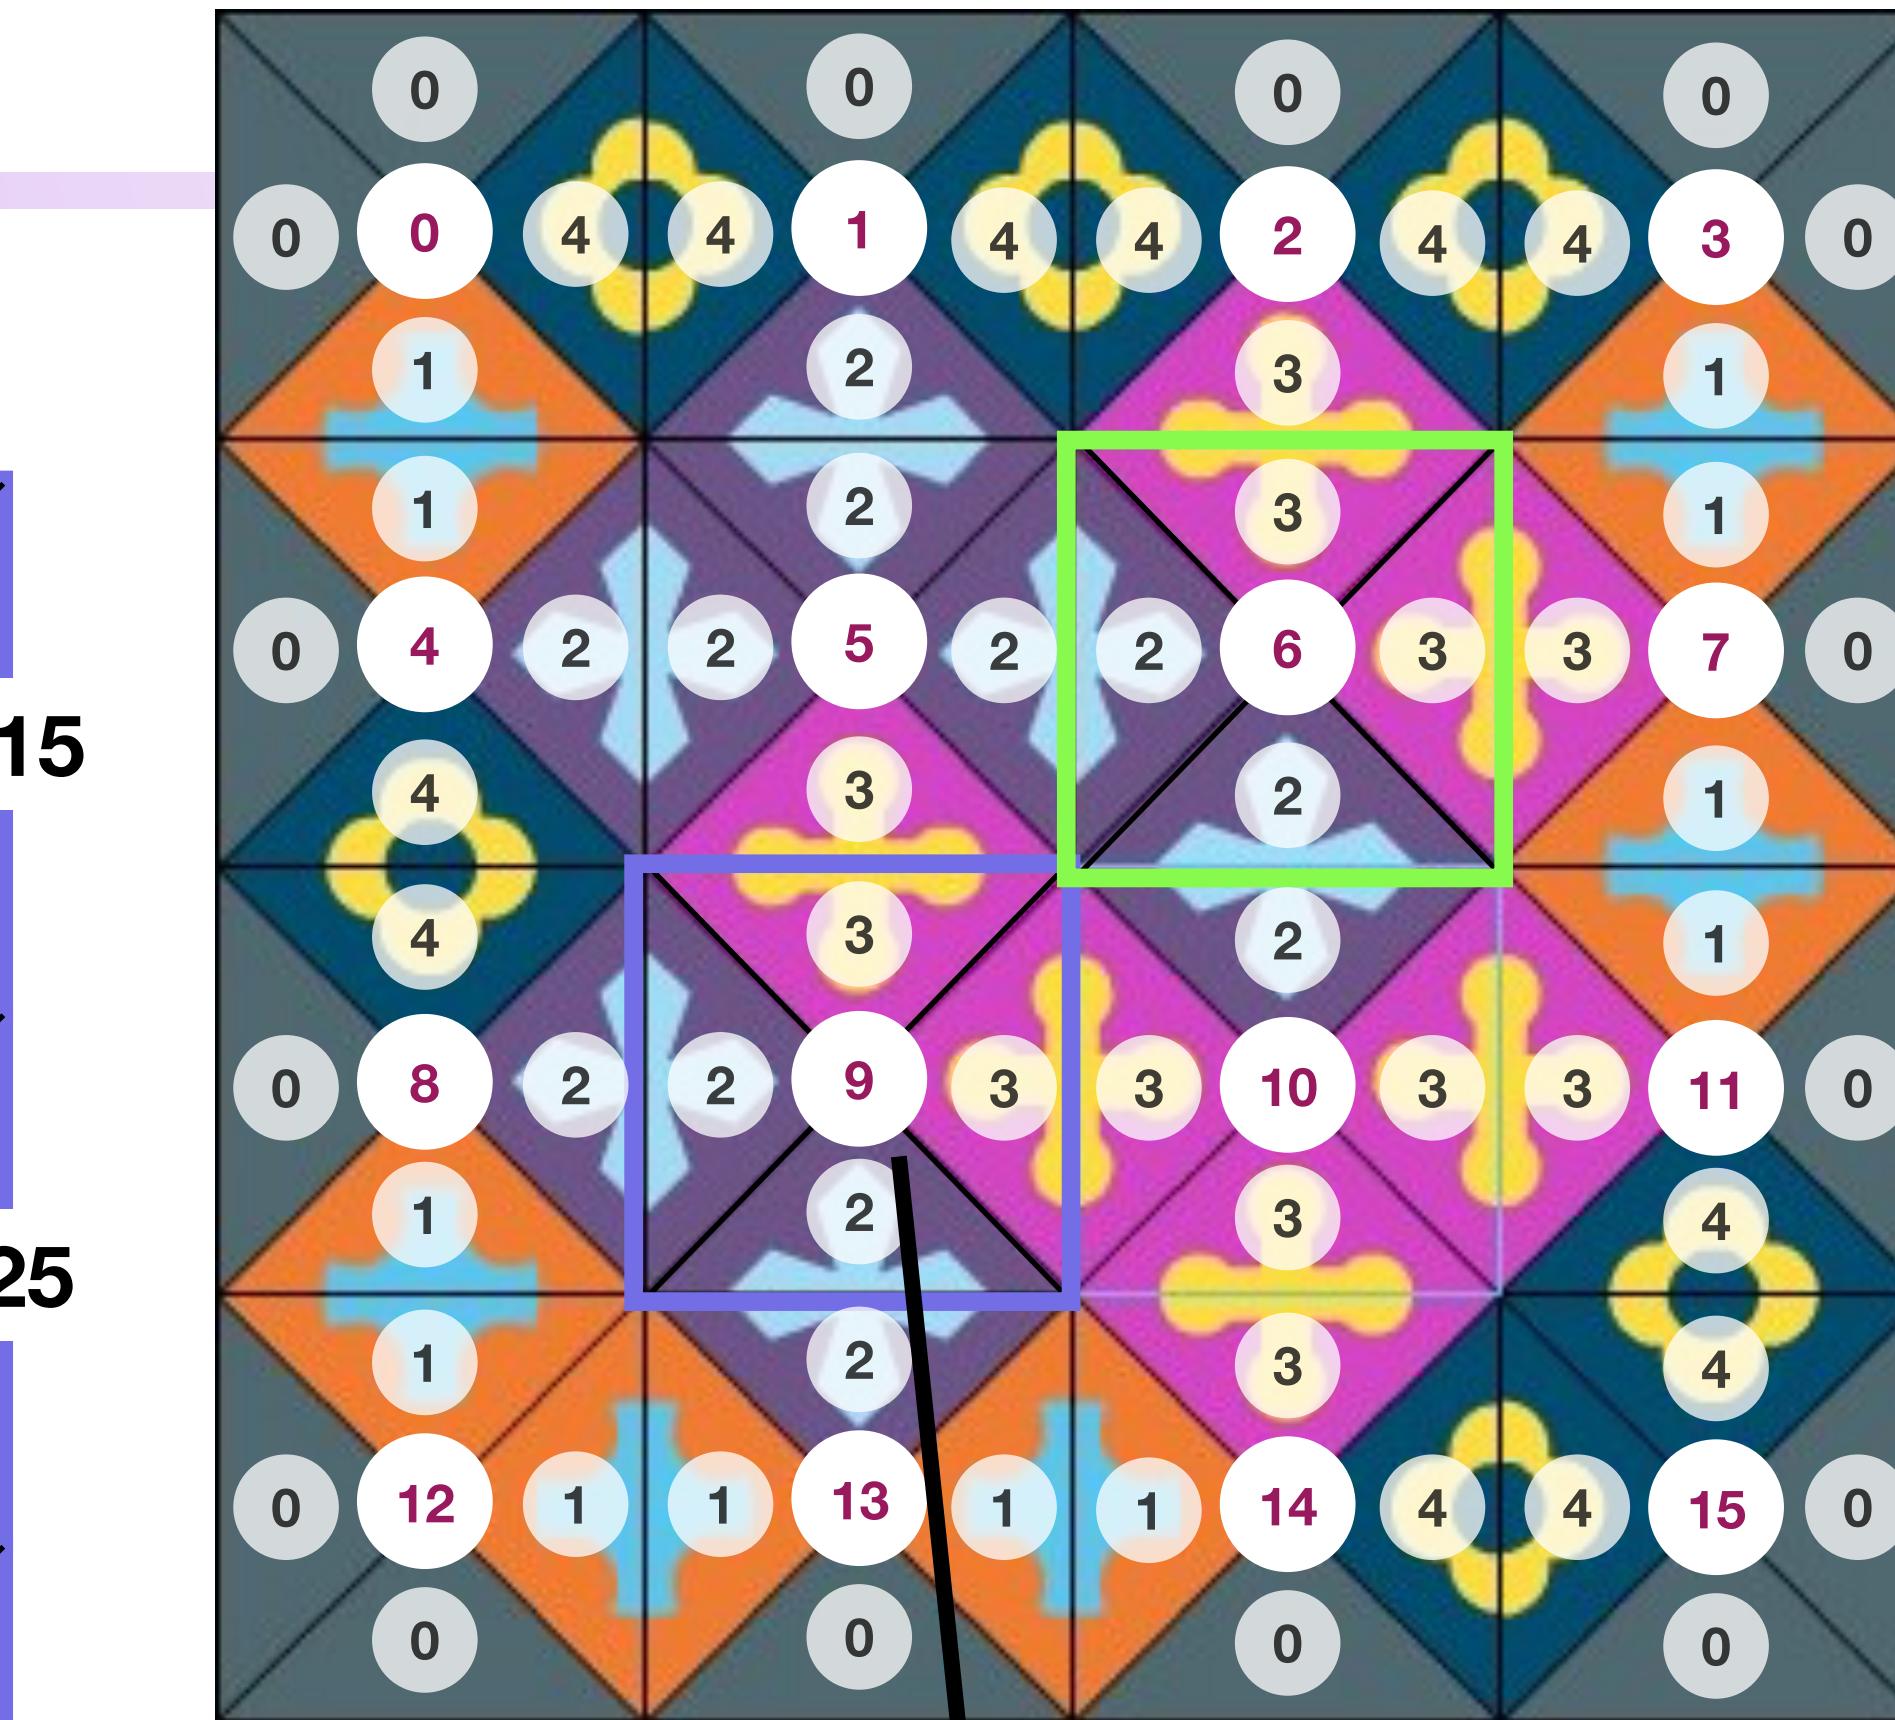
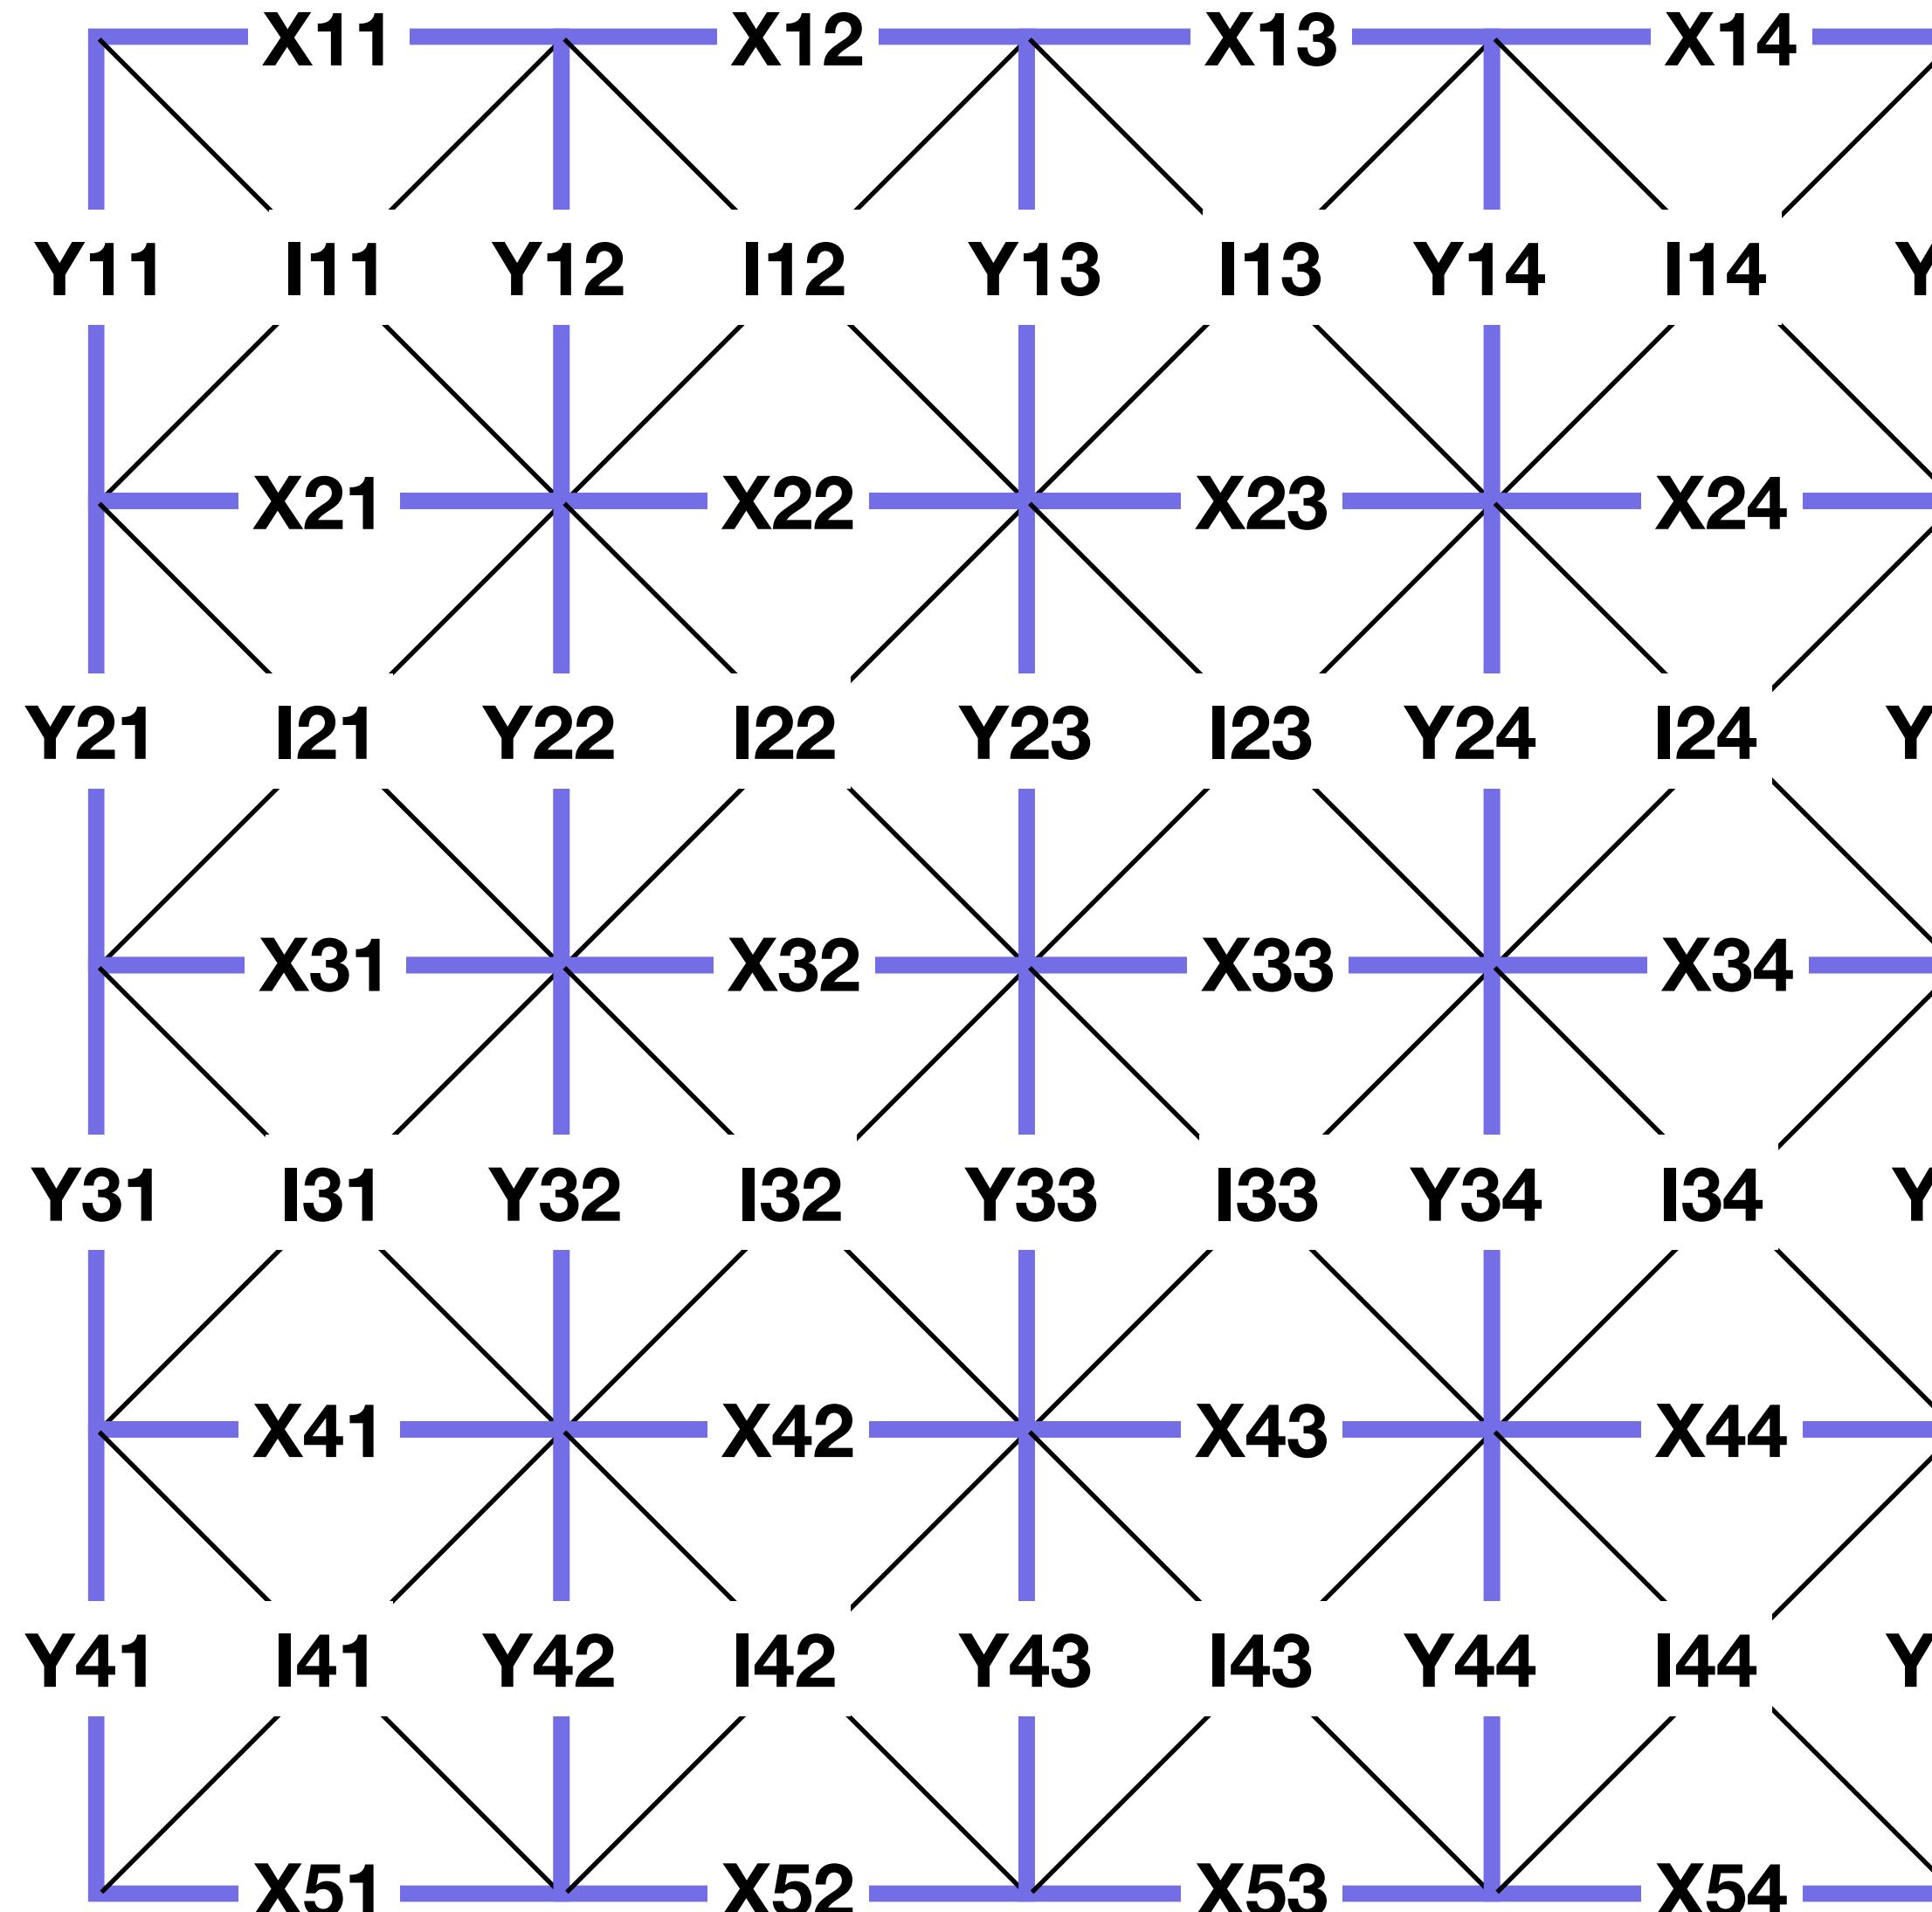
**Answer: With a Table constraint!**  
Let us build it together

I	S1	S2	S3	S4
14	1	2	3	4
14	2	3	4	1
14	3	4	1	2
14	4	1	2	3

$(I, S1, S2, S3, S4) \in \text{table}$  ensures  
that it corresponds to one of the  
four rotations for this piece

# Model for 4x4 Eternity II

$$D(X21)=\{0..4\} \quad D(I11)=\{0..15\}$$

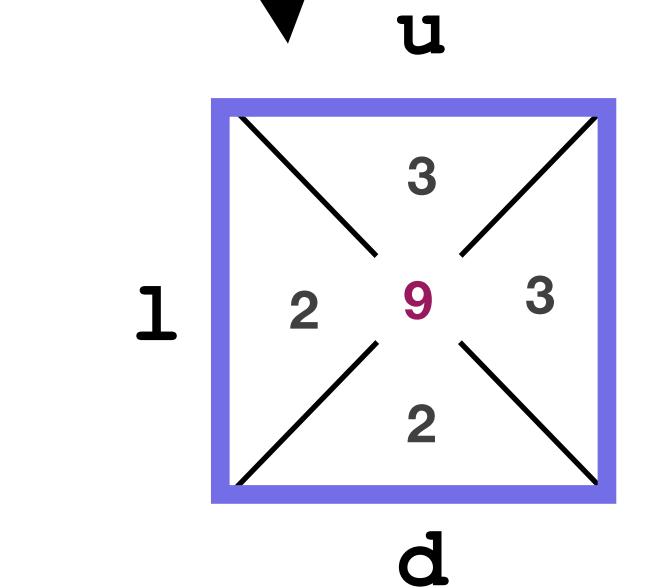


tableOfPieces

i u r d l

9	3	3	2	2
9	3	2	2	3
9	2	2	3	3
9	2	3	3	2

6	3	3	2	2
6	3	2	2	3
6	2	2	3	3
6	2	3	3	2

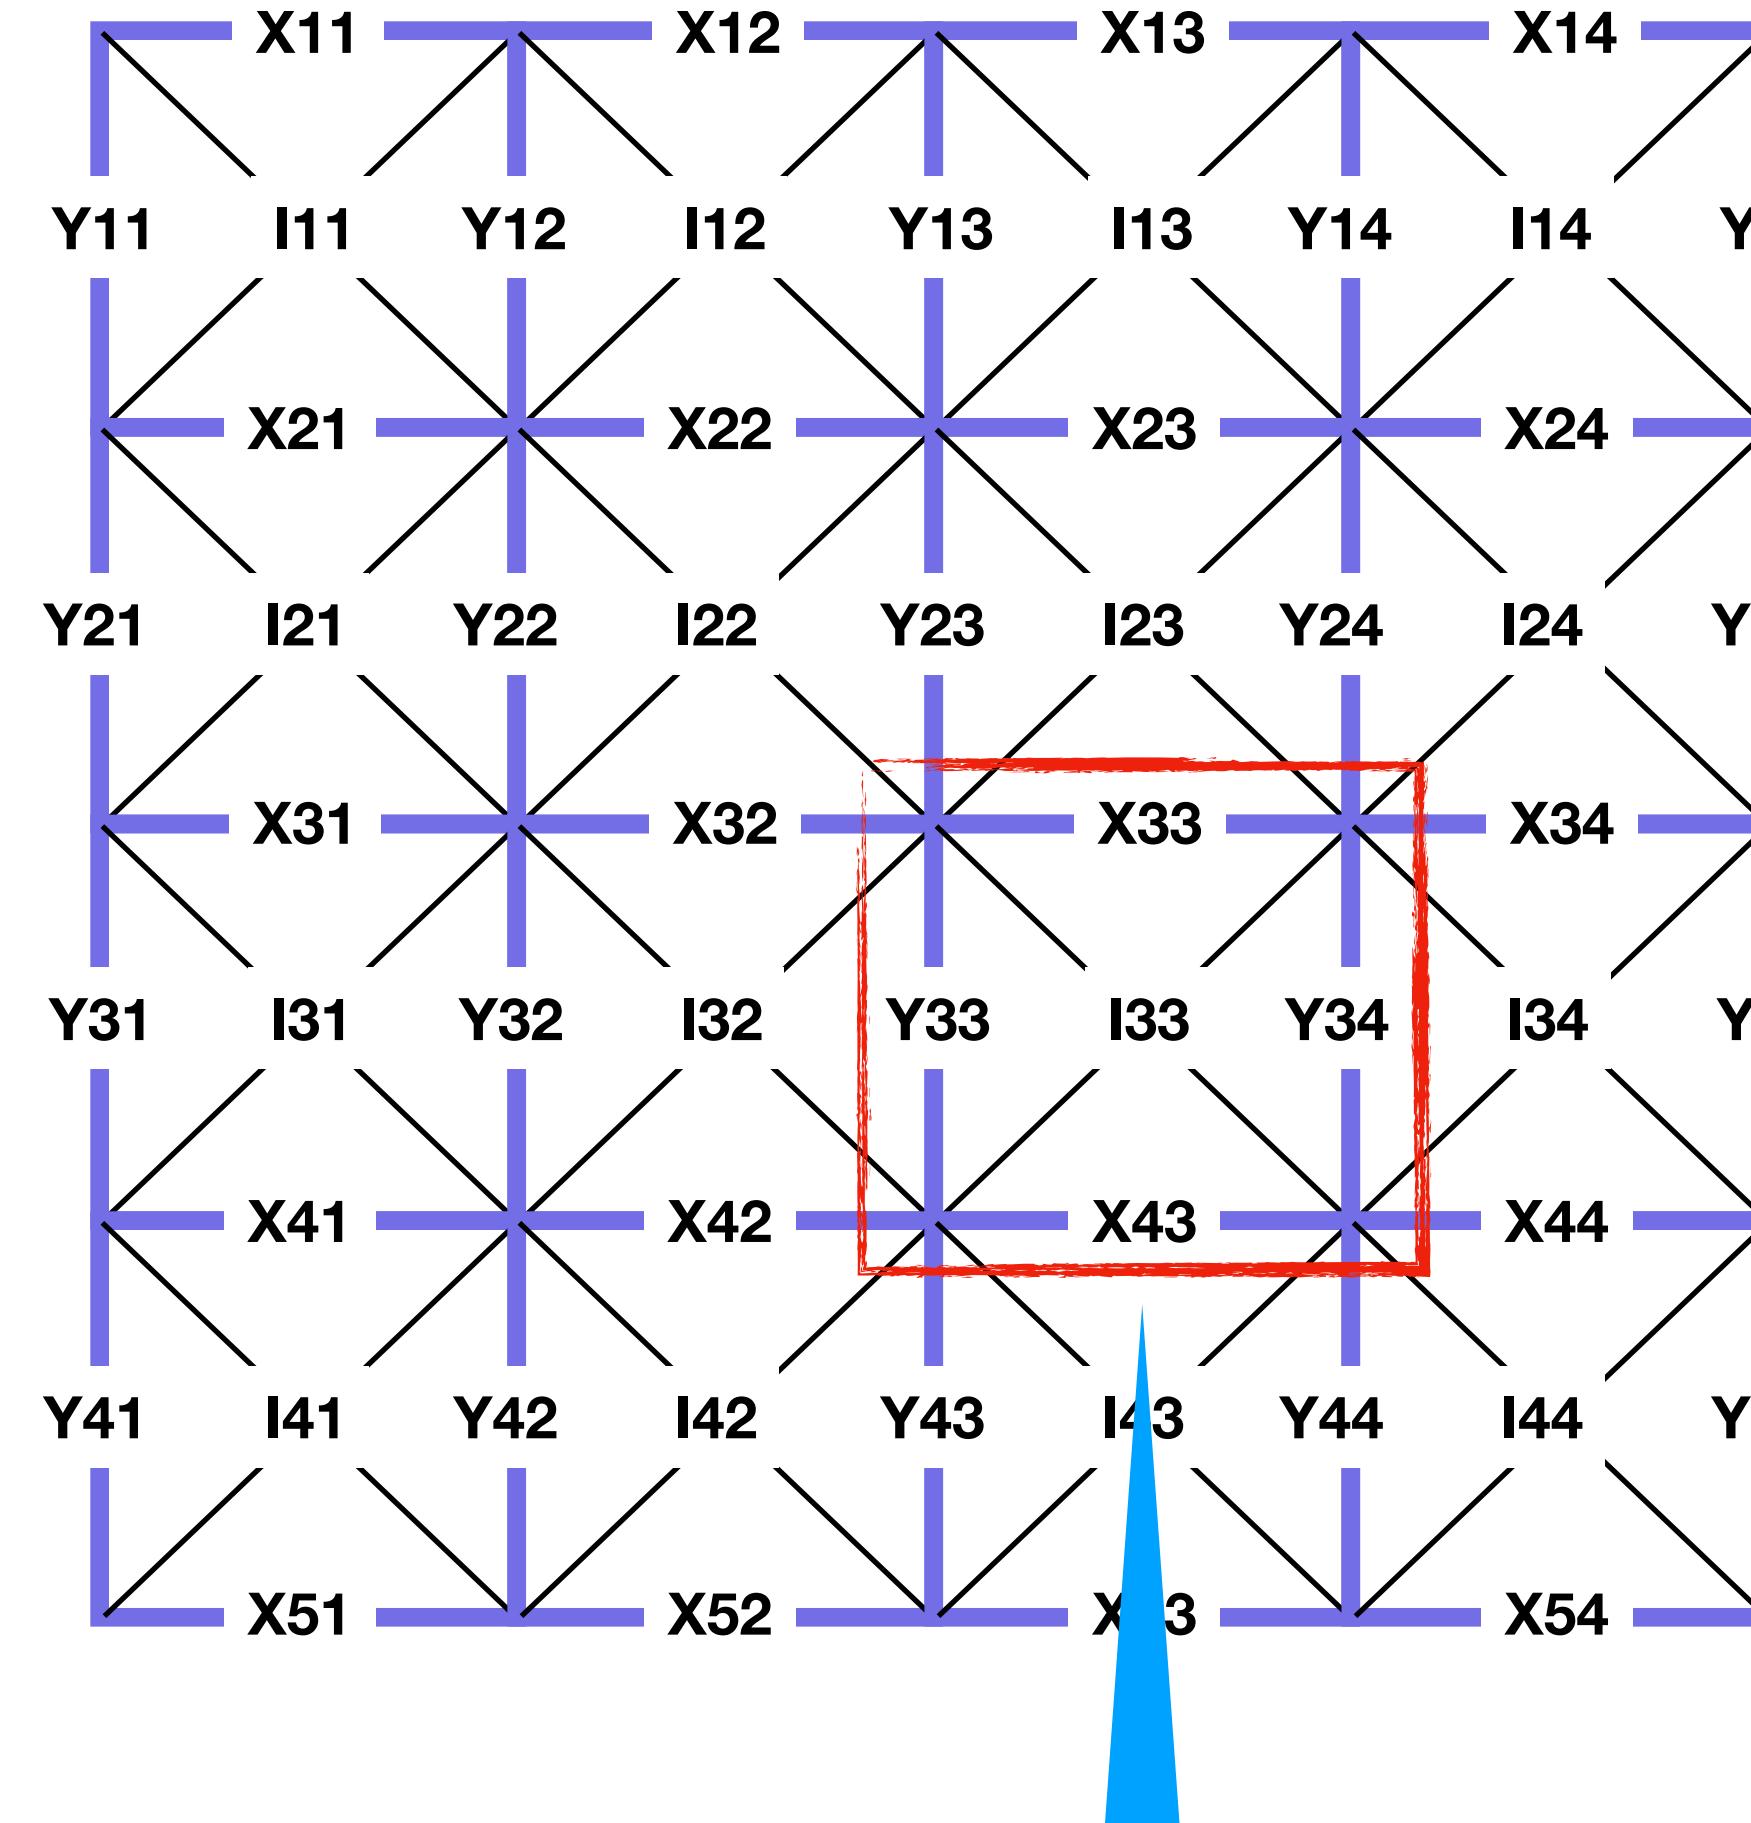


Every piece has 4 possible rotations,  
hence 4 entries per piece are created

# Model for 4x4 Eternity

$D(X21)=\{0..4\}$

$D(I11)=\{0..15\}$



Each square contains a valid piece:  
 $[I33, X33, Y34, X43, Y33] \in \text{tableOfPieces}$

All the pieces are placed and each can be placed only once:  
 $\text{AllDifferent}(I11, I12, \dots, I44)$

All the positions are occupied with valid pieces  
 $(I_{ij}, X_{i,j}, Y_{i,j+1}, X_{i+1,j}, Y_{ij}) \in \text{tableOfPieces} \ \forall i, j \in [1..4] \times [1..4]$

**tableOfPieces**

i u r d l

9 3 3 2 2

9 3 2 2 3

9 2 2 3 3

9 2 3 3 2

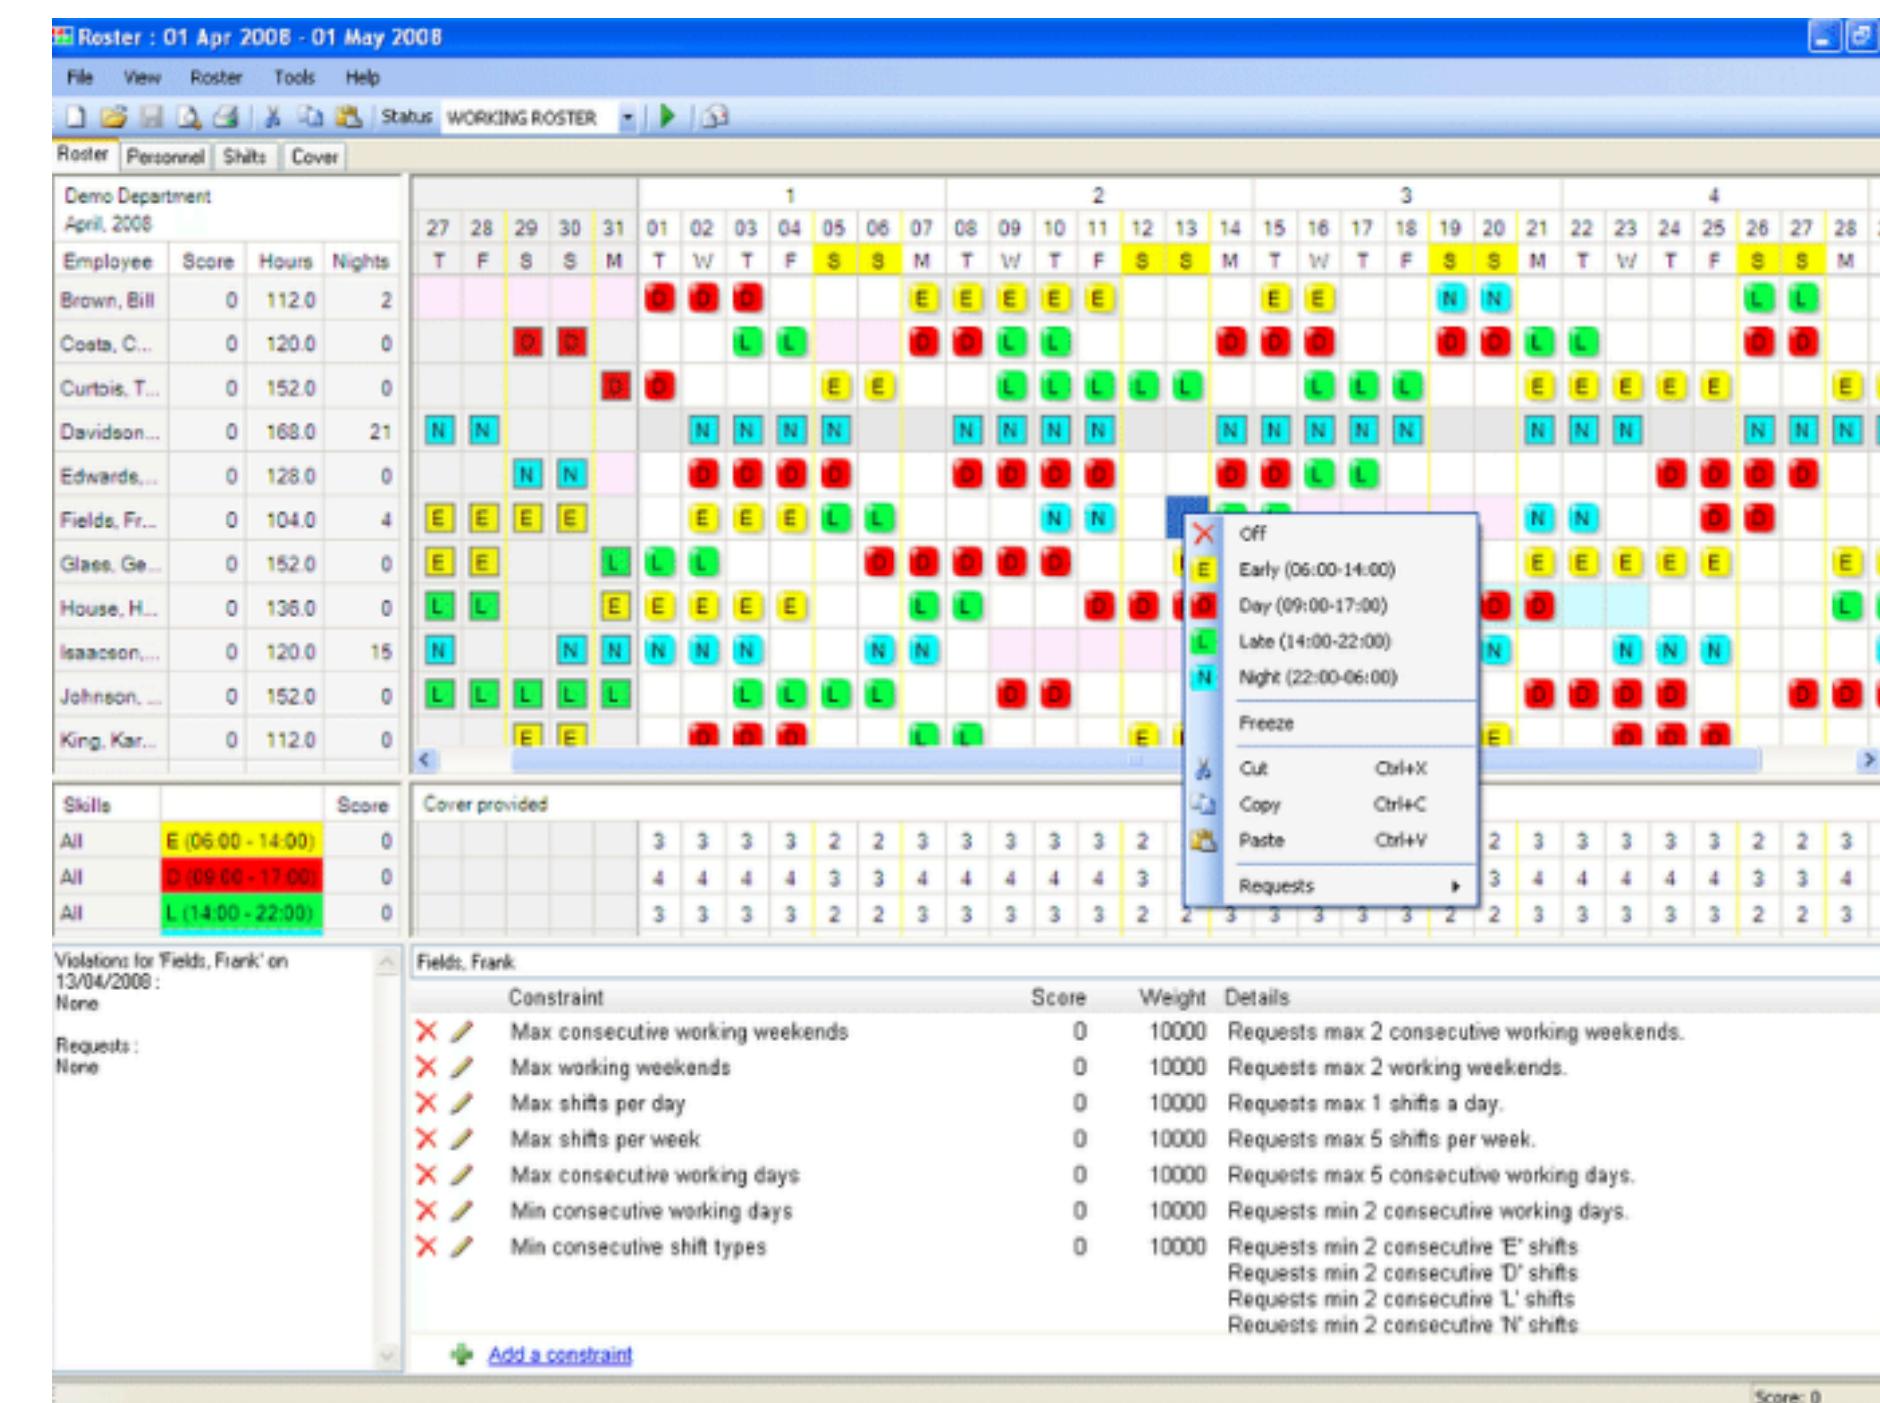
6 3 3 2 2

6 3 2 2 3

6 2 2 3 3

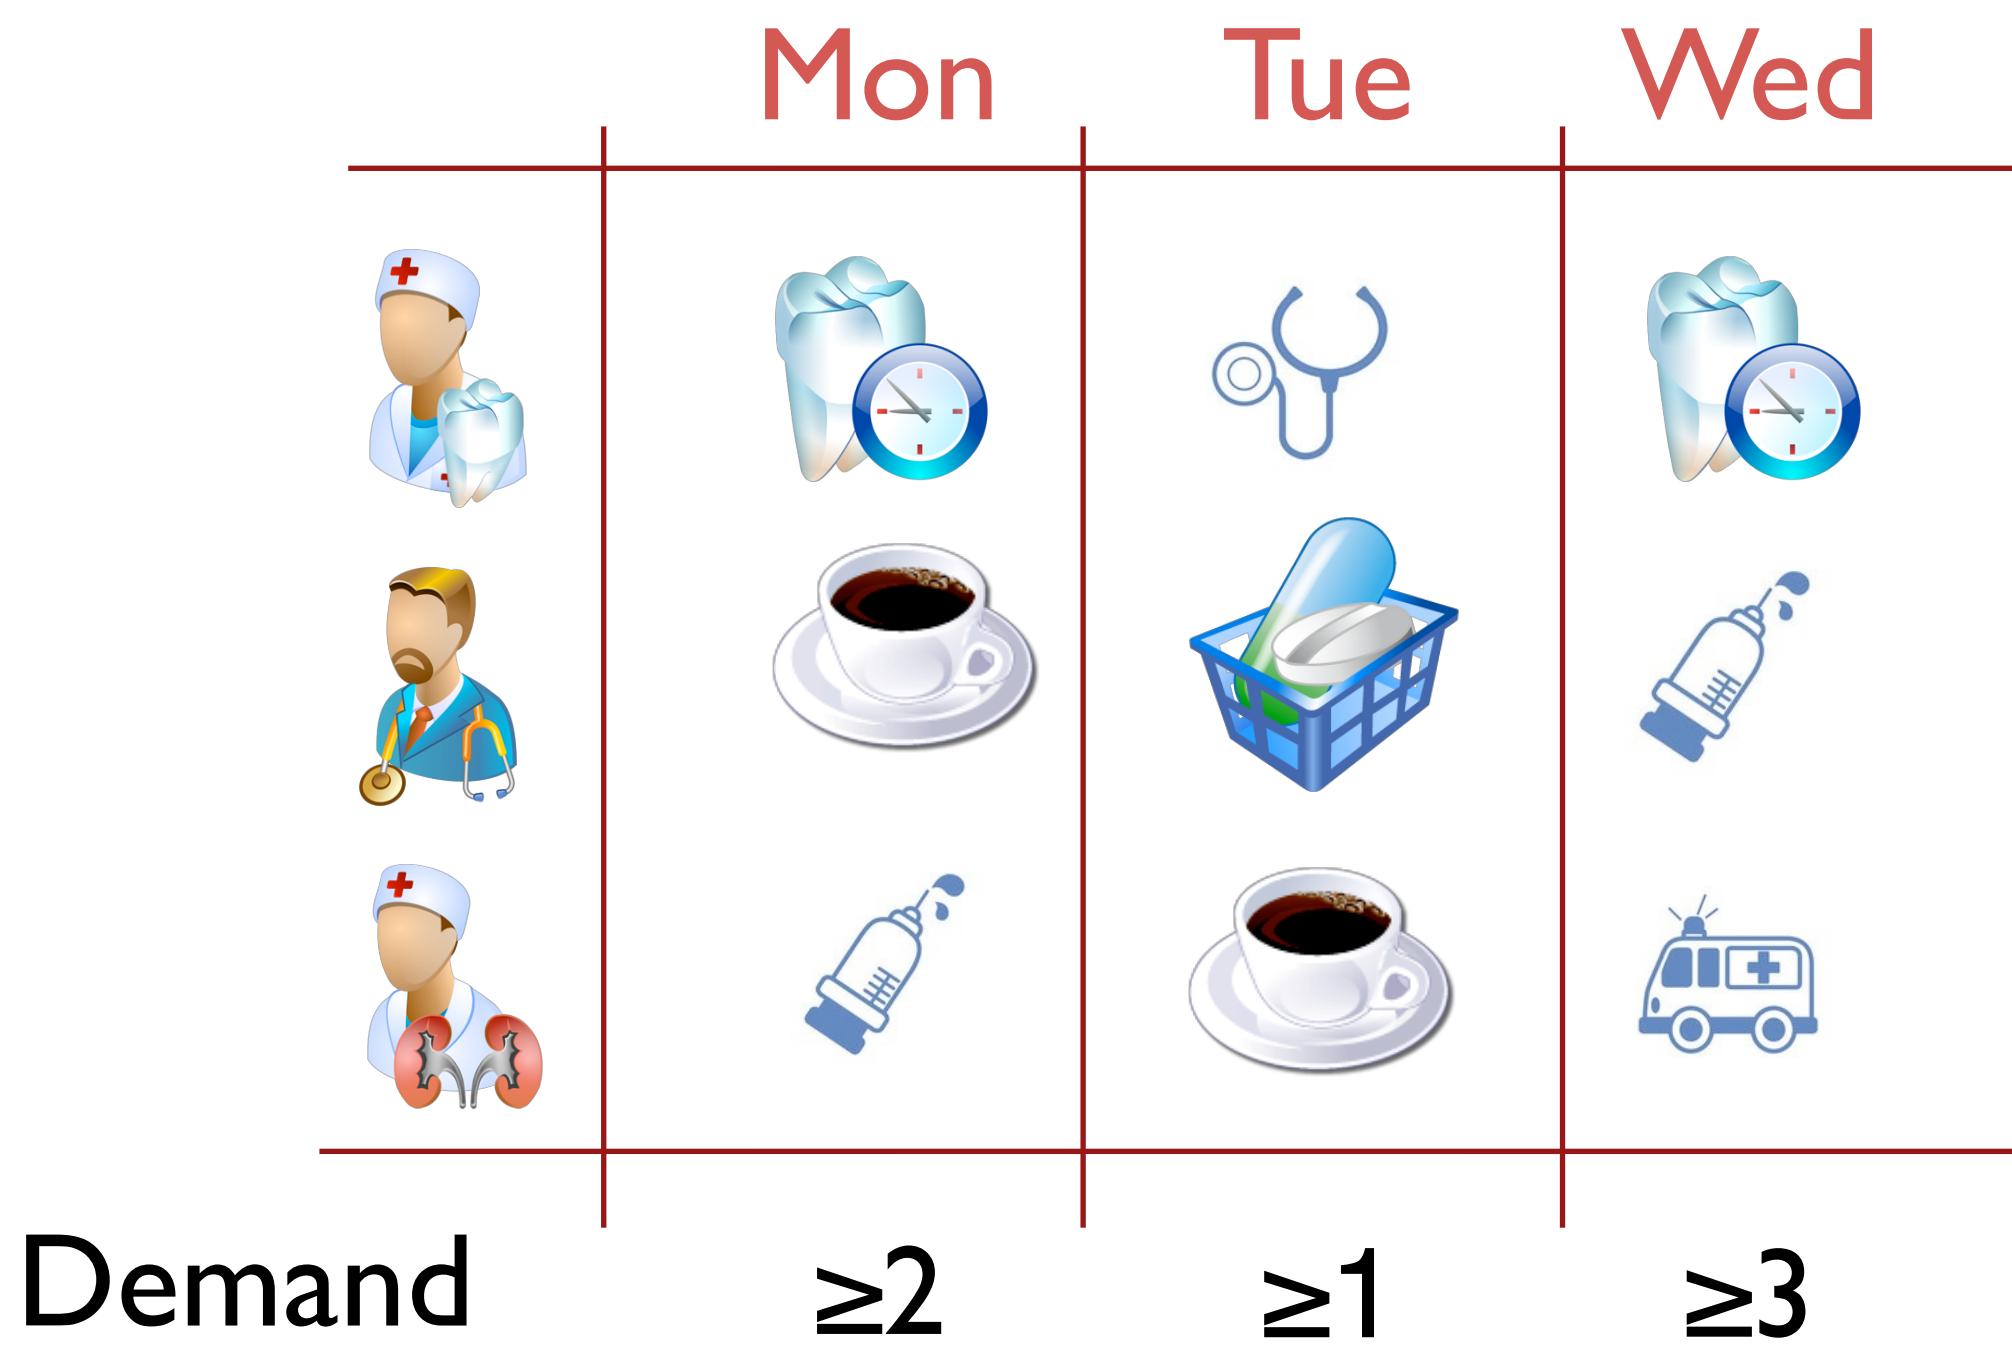
6 2 3 3 2

# Application of Table constraints: Regular (Automaton) constraint for rostering problems



# Rostering problems

## Nurse Rostering Problem (NRP)



NRP is the problem of finding an optimal way to assign nurses to shifts, typically with a set of hard constraints which all valid solutions must follow, and a set of soft constraints which define the relative quality of valid solutions.  
[https://en.wikipedia.org/wiki/Nurse\\_scheduling\\_problem](https://en.wikipedia.org/wiki/Nurse_scheduling_problem)

### Examples of (horizontal) constraints for NRP:

- ✓ A nurse cannot work the day shift, night shift, and late-night shift on the same day (i.e., no 24-hour duties).
- ✓ A nurse may go on a holiday and will not work shifts then.
- ✓ A nurse cannot do a late-night shift followed by a day shift the next day.
- ✓ ...

Typically, each such constraint gives rise to a regular expression.

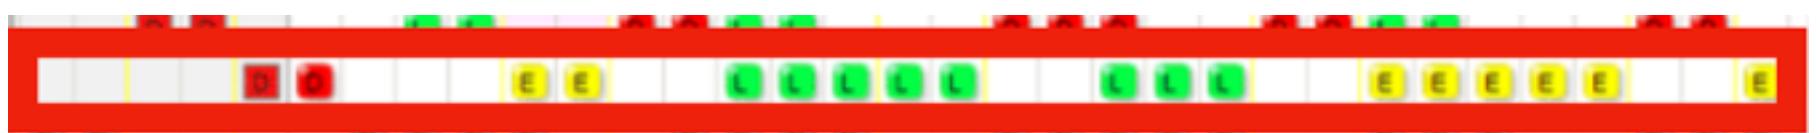
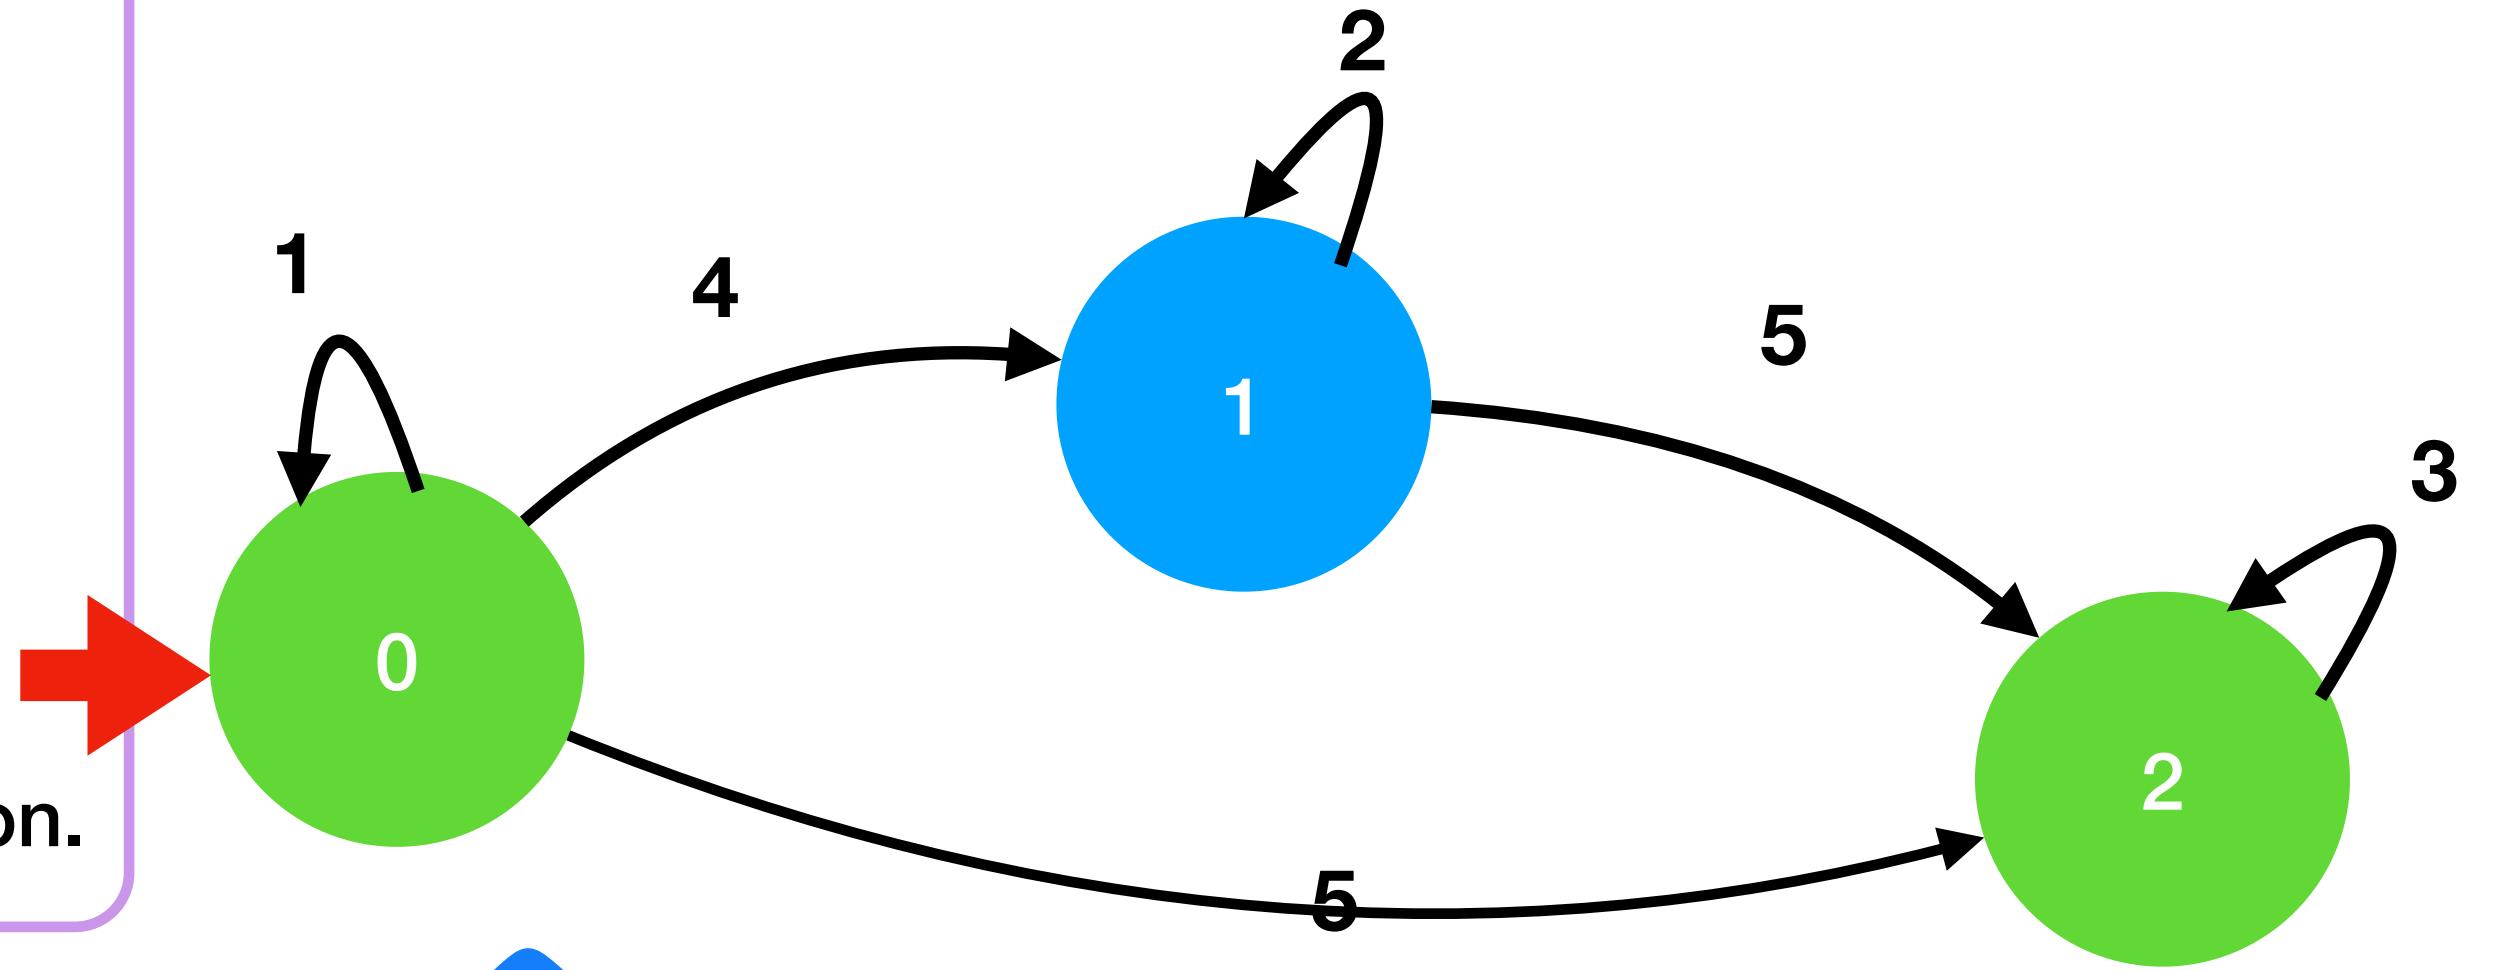
# How to enforce rostering constraints?

Some examples of constraints are:

- ✓ A nurse cannot work the day shift, night shift, and late-night shift on the same day (i.e., no 24-hour duties).
- ✓ A nurse may go on a holiday and will not work shifts then.
- ✓ A nurse cannot do a late-night shift followed by a day shift the next day.
- ✓ ...

Typically, each such constraint gives rise to a regular expression.

💡 By aggregating them into *one* automaton (implementation of regular expression), with transitions and accepting states

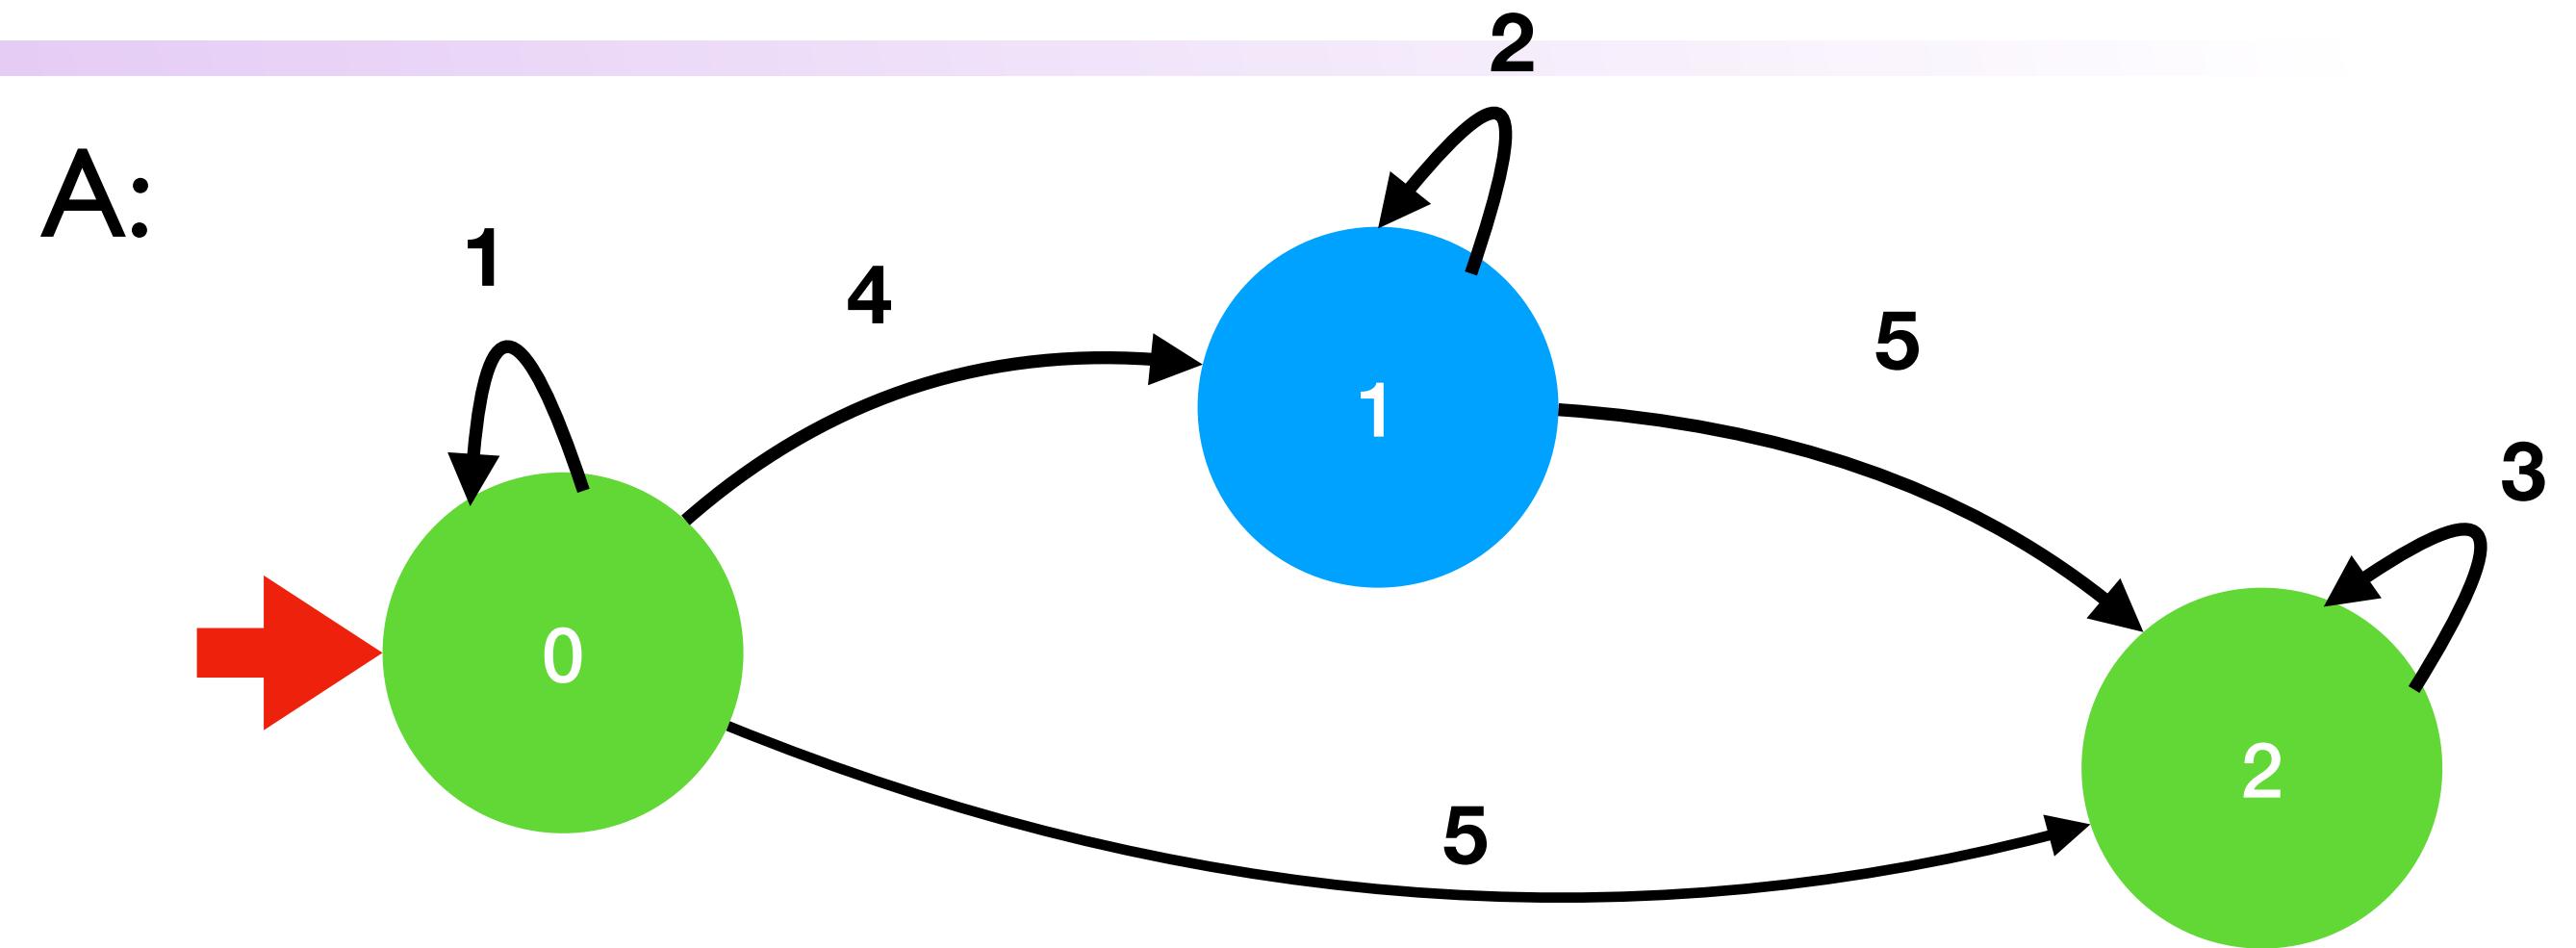


ok iff accepted by automaton

The question is: How do we model in CP an automaton and the acceptance of a string of variables by an automaton?

# Regular constraint

$(x,y,z) \in \text{Language}(A)$



Constraint:  $(x,y,z)$  is a word accepted (at a green state) by the deterministic automaton A, with start state 0.

symbols

T	0	1	2	3	4	5
0		0			1	2
1			1			2
2				2		

states

The model:

$S_x$  = state after reading variable  $x$

$S_y$  = state after reading variables  $x$  and  $y$

$S_z$  = state after reading variables  $x$ ,  $y$ , and  $z$

$T[s][v]$  = state after reading at state  $s$  the symbol  $v$

$S_x = T[0][x]$

$S_y = T[S_x][y]$ , encode Element2D via Table

$S_z = T[S_y][z]$ , encode Element2D via Table

$S_z \in \{0, 2\}$

# Element2D: $T[x][y] = z$

- Can be modelled with  $(x,y,z) \in \text{table}$

		y				
		0	1	2	3	
		0	1	8	9	6
		1	1	9	2	4
		2	9	8	9	8
		3	1	9	2	5

table

x	y	$T[x][y]$
0	0	1
0	1	8
0	2	9
0	3	6
1	0	1
1	1	9
1	2	2
1	3	4
2	0	9
2	1	8
2	2	9
2	3	8
3	0	1
3	1	9
3	2	2
3	3	5

- If we have a domain-consistent filtering for Table, then we also have one for Element2D.
- Element2D can be encoded with a Table constraint.



# Filtering a Table constraint: slow algorithm

# Table constraint

```
// x.length = n, dim(table) = m x n
public Table(IntVar[] x, int[][] table)
```

- ▶ A tuple (table row) is **valid** iff all its values are in the domains of the corresponding variables:
  - $\text{valid}(\text{table}[r]) \equiv \forall i : \text{table}[r][i] \in D(x[i])$
- ▶ Literal  $(x[i], v)$  is **supported** iff there is a valid tuple with value  $v$  in column  $i$ :
  - $\exists r : \text{valid}(\text{table}[r]) \wedge \text{table}[r][i] = v$
- ▶ Example:  $D(x) = \{1,2\}$ ,  $D(y) = \{1,2,3\}$ ,  $D(z) = \{1,2,3\}$ 
  - $(z,3)$  is supported,
  - but  $(z,2)$  is *not* supported and hence 2 must be removed from  $D(z)$ .

x	y	z
1	2	3
1	3	3
2	2	3
3	3	3
2	1	1
4	1	2
4	4	4

invalid

# A first, slow (but domain-consistent) filtering

```
// x.length=n, dim(table)= m x n
public Table(IntVar[] x, int[][] table)
```

```
SlowTableFiltering(x,table) {
    for (xi <- x){
        for (v <- D(xi)){
            if ( $\exists r: \forall j \neq i: \text{table}(r, j) \in D(x_j) \wedge \text{table}(r, i) = v$ ){
                D(xi) <- D(xi) \ {v}
            }
        }
    }
}
```

# Slow Table filtering: implementation in MiniCP

```
public void propagate() throws InconsistencyException {
    for (int i = 0; i < x.length; i++) {
        for (int v = x[i].getMin(); v <= x[i].getMax(); v++) {
            if (x[i].contains(v)) {You
                boolean supported = false;
                for (int tupleIdx = 0; tupleIdx < table.length &&
                     !supported; tupleIdx++) {
                    if (table[tupleIdx][i] == v) {
                        boolean allSupported = true;
                        for (int j = 0; j < x.length && allSupported; j++) {
                            if (!x[j].contains(table[tupleIdx][j])) {
                                allSupported = false;
                            }
                        }
                        supported = allSupported;
                    }
                }
                if (!supported)
                    x[i].remove(v);
            }
        }
    }
}
```

should use  
your fillArray  
here



# Filtering a Table constraint: the STR algorithm family

# Simple Tabular Reduction (STR) algorithms

- Lecoutre, Christophe. STR2: Optimized simple tabular reduction for table constraints. *Constraints*, 2011.

Simple Tabular Reduction (STR) algorithms:

1. For each tuple in the table:

- The tuple is **valid**  $\Rightarrow$  all its values are in supported literals, so:  
collect the supported literals in a set.
  - Example: (1,2,3) is valid  $\Rightarrow$  (x,1), (y,2), and (z,3) are supported.
- The tuple is **invalid**:  
remove it (in a stateful way) from the table,  
giving a smaller table, hence incrementality.

2. For each literal  $(x_i, v)$ :

if it is not supported (check in the collected set of literals), then remove  $v$  from  $D(x_i)$ .

x	y	z
1	2	3
1	3	3
2	2	3
3	3	3
2	1	1
4	1	2
4	4	4

# STR2 algorithm

```

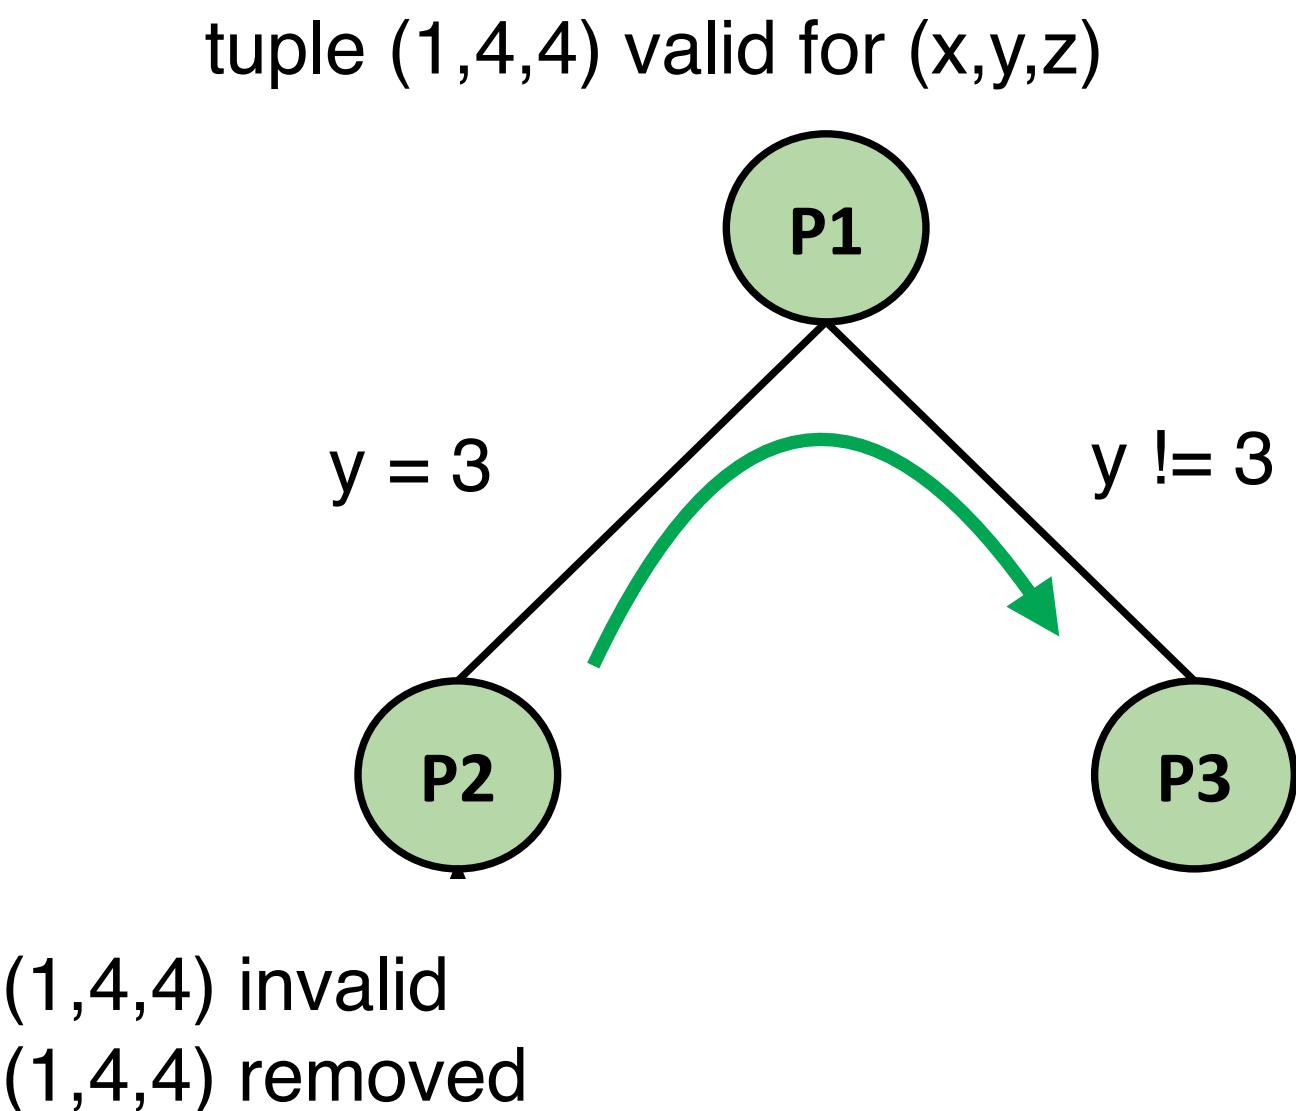
STR2Filtering(x,table) {
    supported =  $\emptyset$ 
    for (t <- table) {
        if ( $\forall i: x_i \text{.contains}(t(i))$ ) {
            for (xi <- x) {
                supported += (xi,t(i))
            }
        } else {
            table.remove(t)
        }
    }
    for (xi <- x) {
        for (v <- D(xi)) {
            if ( $(x_i, v) \notin \text{supported}$ ) {
                D(xi) <- D(xi) \ {v}
            }
        }
    }
}

```

- if tuple is **valid**
- literals **collected**
- else: tuple **removed**
- remove **unsupported**

# Incrementality of STR2

- ▶ Incrementality of STR2 comes from the table.
  - Invalid tuples are removed from the table.
  - If a tuple is removed, then it is not inspected in future executions.
- ▶ The table has to be stateful (aka reversible), using the state manager, which restores the state on backtrack.



# STR2 needs a stateful table

- ▶ Use a stateful table (details omitted here) to represent the table.
- ▶ All that needs to be backtracked on is just **one integer**, a **StateInt**, denoting the current number of valid tuples in the table, which are stored before the row having that StateInt as index

# Stateful table

Assume  $D(x)=\{1,2\}$ ,  $D(y)=\{1,2,3\}$ ,  $D(z)=\{1,3\}$  now:

x	y	z
1	2	3
1	3	3
2	2	3
2	1	1
<hr/>		
3	3	3
4	1	2
4	4	4

The table is partitioned into two sets:

the **remaining** tuples are  
before **size**

the **removed** tuples are  
from **size** on

# Stateful table

Assume 3 is removed from  $D(y)$ , giving  $D(x) = \{1,2\}$ ,  $D(y) = \{1,2\}$ ,  $D(z) = \{1,3\}$ :

x	y	z
1	2	3
1	3	3
2	2	3
2	1	1
3	3	3
4	1	2
4	4	4

When a tuple is removed:

- it is swapped with the one in position **size-1**
- **size** is decremented

# Stateful table

Assume 3 is removed from  $D(y)$ , giving  $D(x) = \{1,2\}$ ,  $D(y) = \{1,2\}$ ,  $D(z) = \{1,3\}$ :

x	y	z
1	2	3
2	1	1
2	2	3
<b>size</b>		
1	3	3
3	3	3
4	1	2
4	4	4

When a tuple is removed:

- it is swapped with the one in position **size-1**
- **size** is decremented

# Stateful table

Assume restoration to previous state, with  $D(x) = \{1,2\}$ ,  $D(y) = \{1,2,3\}$ ,  $D(z) = \{1,3\}$ :

x	y	z
1	2	3
2	1	1
2	2	3
1	3	3
<b>size</b>		
3	3	3
4	1	2
4	4	4

On backtracking (`sm.restoreState()`): restoring **size**

- restores the removed tuples
- at possibly different positions in the table



# Filtering a Table constraint: the Compact Table algorithm

# Compact Table: filtering to domain consistency

Demeulenaere, J., Hartert, R., Lecoutre, Ch., Perez, G., Perron, L., Régis, J.-C., & Schaus, P. Compact-table: Efficiently filtering table constraints with reversible sparse bit-sets. *CP 2016*.

- ▶ It is the most efficient known algorithm for filtering a Table constraint to domain consistency.
- ▶ It relies on bitwise operations using a data structure called ***reversible/stateful sparse bit set***.
- ▶ It is easy to implement (quite similarly to STR2).

# Compact Table

index	x	y	z
0	7	5	8
1	2	1	4
2	1	3	2
3	2	4	2
4	6	5	9
5	7	7	8
6	4	2	1
7	1	1	1
8	7	8	9
9	8	9	6
10	2	2	3
11	0	0	0
12	3	3	1
13	5	8	5
14	9	7	7
15	2	3	1

Assume  $D(x) = D(y) = D(z) = \{1,2,3,4,5\}$  in the meantime:  
what filtering happens now?

Initial list of  
allowed tuples

# Precomputation of support bit sets

index	x	y	z	supports				
				x=1	x=2	x=3	...	z=5
0	7	5	8	0	0	0		0
1	2	1	4	0	1	0		0
2	1	3	2	1	0	0		0
3	2	4	2	0	1	0		0
4	6	5	9	0	0	0		0
5	7	7	8	0	0	0		0
6	4	2	1	0	0	0		0
7	1	1	1	1	0	0		0
8	7	8	9	0	0	0		0
9	8	9	6	0	0	0		0
10	2	2	3	0	1	0		0
11	0	0	0	0	0	0		0
12	3	3	1	0	0	1		0
13	5	8	5	0	0	0		1
14	9	7	7	0	0	0		0
15	2	3	1	0	1	0		0

Bit Sets

$$D(x) = D(y) = D(z) = \{1, 2, 3, 4, 5\}$$

Every bit supports( $x_i, v$ )( $r$ ) is computed when posting the constraint:

- 1 if  $\text{table}[r][i] = v$
- 0 otherwise

Can we identify the valid tuples (green ones) from the supports( $x_i, v$ ) bit sets?

# Computation of each validTuples(r)

index	validTuples	x	y	z	supports				
					x=1	x=2	x=3	...	z=5
0	0	7	5	8	0	0	0		0
1	1	2	1	4	0	1	0		0
2	1	1	3	2	1	0	0		0
3	1	2	4	2	0	1	0		0
4	0	6	5	9	0	0	0		0
5	0	7	7	8	0	0	0		0
6	1	4	2	1	0	0	0		0
7	1	1	1	1	1	0	0		0
8	0	7	8	9	0	0	0		0
9	0	8	9	6	0	0	0		0
10	1	2	2	3	0	1	0		0
11	0	0	0	0	0	0	0		0
12	1	3	3	1	0	0	1		0
13	0	5	8	5	0	0	0		1
14	0	9	7	7	0	0	0		0
15	1	2	3	1	0	1	0		0

If row r is supported, then 1, else 0

validTuples =

$(\text{supports}(x,1) \mid \text{supports}(x,2) \mid \text{supports}(x,3) \mid \text{supports}(x,4) \mid \text{supports}(x,5)) \ \&$   
 $(\text{supports}(y,1) \mid \text{supports}(y,2) \mid \text{supports}(y,3) \mid \text{supports}(y,4) \mid \text{supports}(y,5)) \ \&$   
 $(\text{supports}(z,1) \mid \text{supports}(z,2) \mid \text{supports}(z,3) \mid \text{supports}(z,4) \mid \text{supports}(z,5))$

# Compact Table: domain-consistency filtering

Goal: remove values not supported anymore

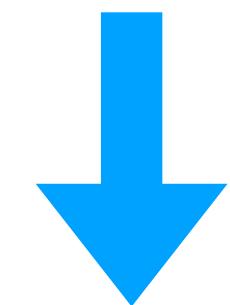
```
CompactTableFiltering(x,table) {  
    for ( $x_i \leftarrow x$ ) {  
        for ( $v \leftarrow D(x_i)$ ) {  
            if ( $validTuples \& supports(x_i, v) = \emptyset$ ) {  
                 $D(x_i) \leftarrow D(x_i) \setminus \{v\}$   
            }  
        }  
    }  
}
```

$validTuples \& supports(x_i, v) = \emptyset$   
(where  $\emptyset$  is the all-zero bit set) is implemented  
in the Java class `BitSet` using method `intersects`

# Compact Table: filtering example

index	validTuples	x	y	z	supports		
					x=5	y=5	z=5
0	0	7	5	8	0	1	0
1	1	2	1	4	0	0	0
2	1	1	3	2	0	0	0
3	1	2	4	2	0	0	0
4	0	6	5	9	0	1	0
5	0	7	7	8	0	0	0
6	1	4	2	1	0	0	0
7	1	1	1	1	0	0	0
8	0	7	8	9	0	0	0
9	0	8	9	6	0	0	0
10	1	2	2	3	0	0	0
11	0	0	0	0	0	0	0
12	1	3	3	1	0	0	0
13	0	5	8	5	1	0	1
14	0	9	7	7	0	0	0
15	1	2	3	1	0	0	0

$\text{validTuples} \& \text{supports}(x, 5) = 0$   
 $\text{validTuples} \& \text{supports}(y, 5) = 0$   
 $\text{validTuples} \& \text{supports}(z, 5) = 0$



$D(x) = \{1, 2, 3, 4, \cancel{5}\}$   
 $D(y) = \{1, 2, 3, 4, \cancel{5}\}$   
 $D(z) = \{1, 2, 3, 4, \cancel{5}\}$

# Update of validTuples when a domain change occurs

- ▶ Assume 1 and 2 are now removed from  $D(x) = \{1,2,3,4\}$ .
- ▶ We first need to update **validTuples**. There are two possible strategies:

1. **From scratch**, based on the remaining values of *all* variable domains:  
 $D(x) = \{3,4\}$  and  $D(y) = D(z) = \{1,2,3,4\}$ . Same as the initial computation:

**validTuples** = (supports(x,3)  $\sqcap$  supports(x,4))  $\&$  (supports(y,1)  $\sqcap$  supports(y,2)  $\sqcap$  supports(y,3)  $\sqcap$  supports(y,4))  $\&$  (supports(z,1)  $\sqcap$  supports(z,2)  $\sqcap$  supports(z,3)  $\sqcap$  supports(z,4))

2. **Incrementally**, based on *the* modified variable domain:  $D(x) = \{3,4\}$ .

**validTuples** = **validTuples**  $\&$  (supports(x,3)  $\sqcap$  supports(x,4))

Indeed, dropping the influence of bit set  $a$  on the bit set  $(a \sqcap b) \& c$ ,  
so as to get  $b \& c$ , can also be done by computing  $((a \sqcap b) \& c) \& b$ .



# Underlying data structure: the StateSparseBitSet API

# Compact Table and state restoration

index	validTuples
0	0
1	0
2	0
3	0
4	0
5	0
6	1
7	0
8	0
9	0
10	0
11	0
12	1
13	0
14	0
15	0

words[0] words[1]

words[2] words[3]

- The update requires having a stateful validTuples bit set: it must recover on backtrack (sm.restoreState, ...).
- We introduce a data structure called StateBitSet that encapsulates an array of StateLong (of 64 bits each).
- This data structure represents validTuples:

validTuples: StateBitSet

words = StateLong[ ]

In practice, assume  
Long of 64 bits ;-)

# Can we further improve the efficiency?

Yes, as bitwise operations do not need to be computed on words that are zero!

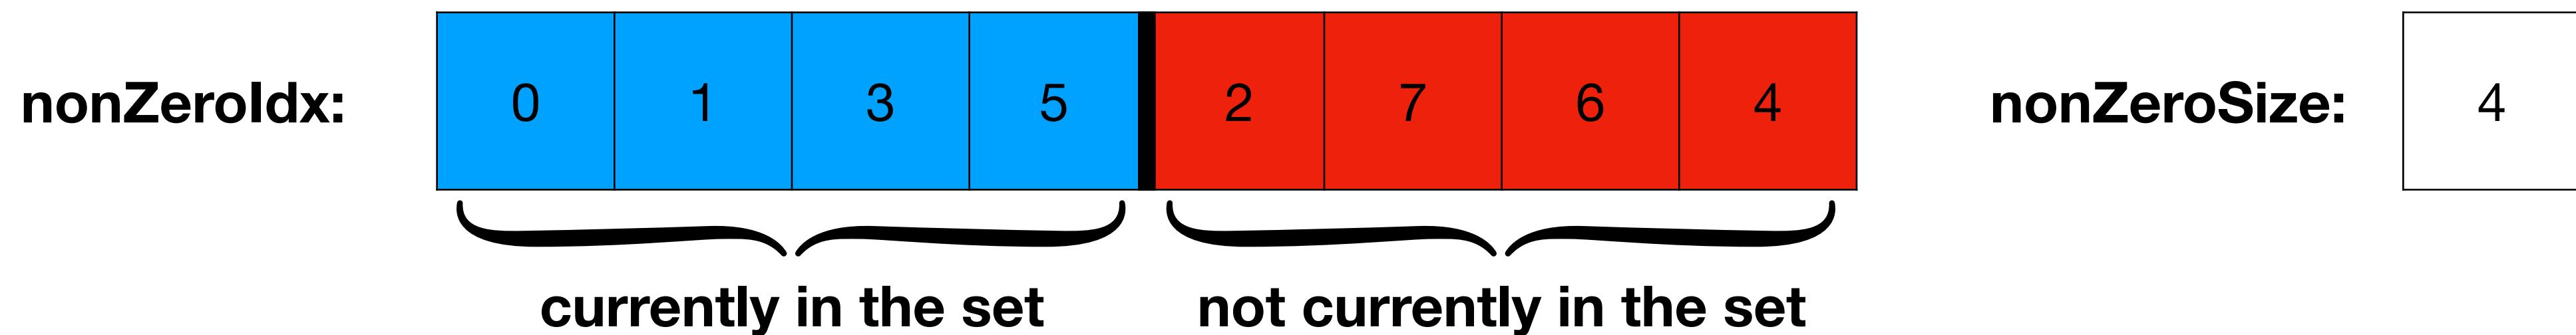
Index	validTuples	x	y	z	supports(y,1)
words[0]	0	7	5	8	0
	0	2	1	4	1
	0	1	3	2	0
	0	2	4	2	0
words[1]	0	6	5	9	0
	0	7	7	8	0
	1	4	2	1	0
	0	1	1	1	1
words[2]	0	7	8	9	0
	0	8	9	6	0
	0	2	2	3	0
	0	0	0	0	0
words[3]	1	3	3	1	0
	0	5	8	5	0
	0	9	7	7	0
	0	2	3	1	0

$$D(x) = \{3,4\}, D(y) = \{1,2,3,4\} = D(z)$$

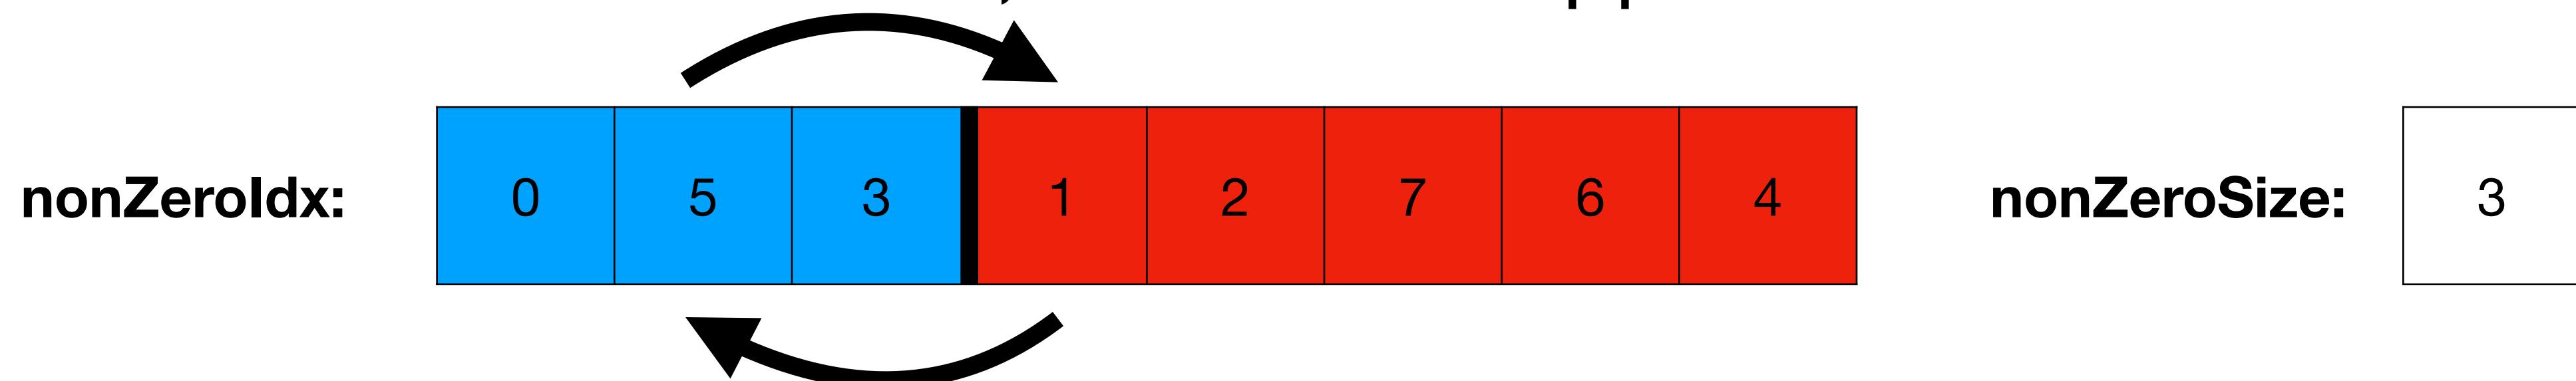
```
CompactTableFiltering(x,table) {
    for (xi <- x) {
        for (v <- D(xi)) {
            if (validTuples & supports(xi,v) = 0) {
                D(xi) <- D(xi) \ {v}
            }
        }
    }
}
```

# How can we do that efficiently?

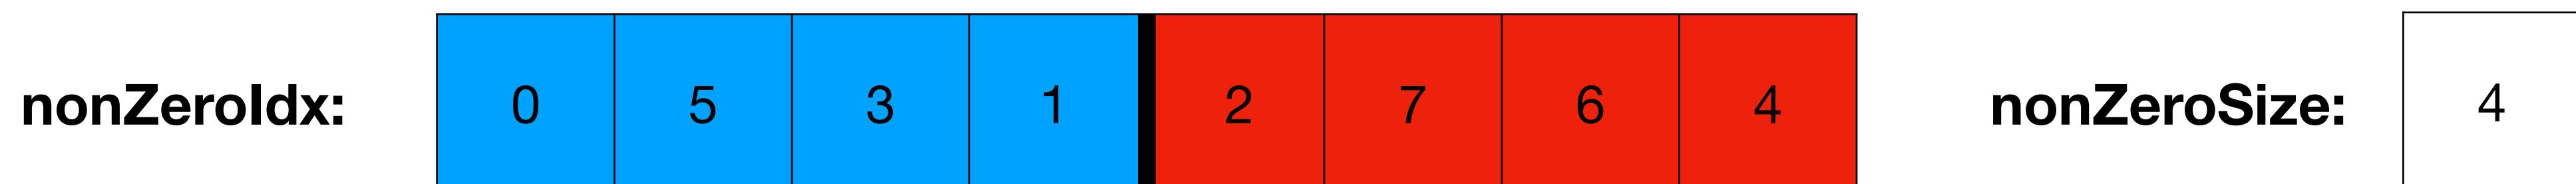
- ▶ Use an internal sparse set for the state
- ▶ Goal: maintain the set of indices of non-zero words



- ▶ If a word becomes zero, then it is swapped with the last non-zero word



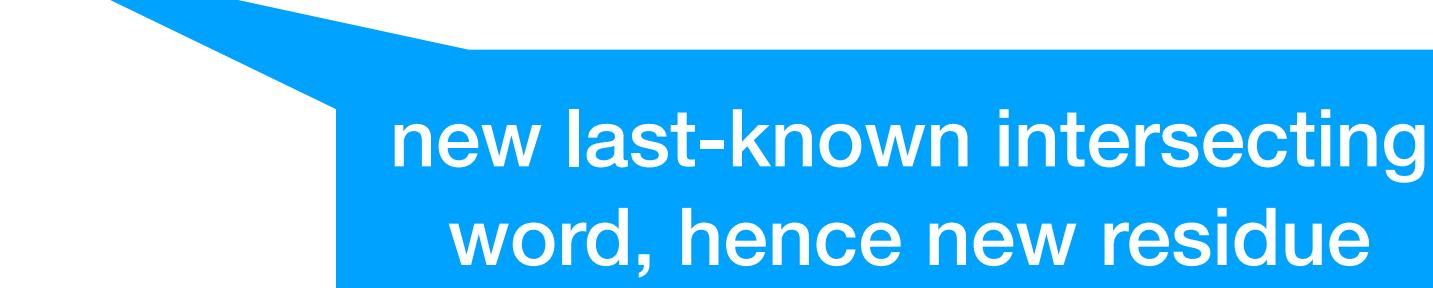
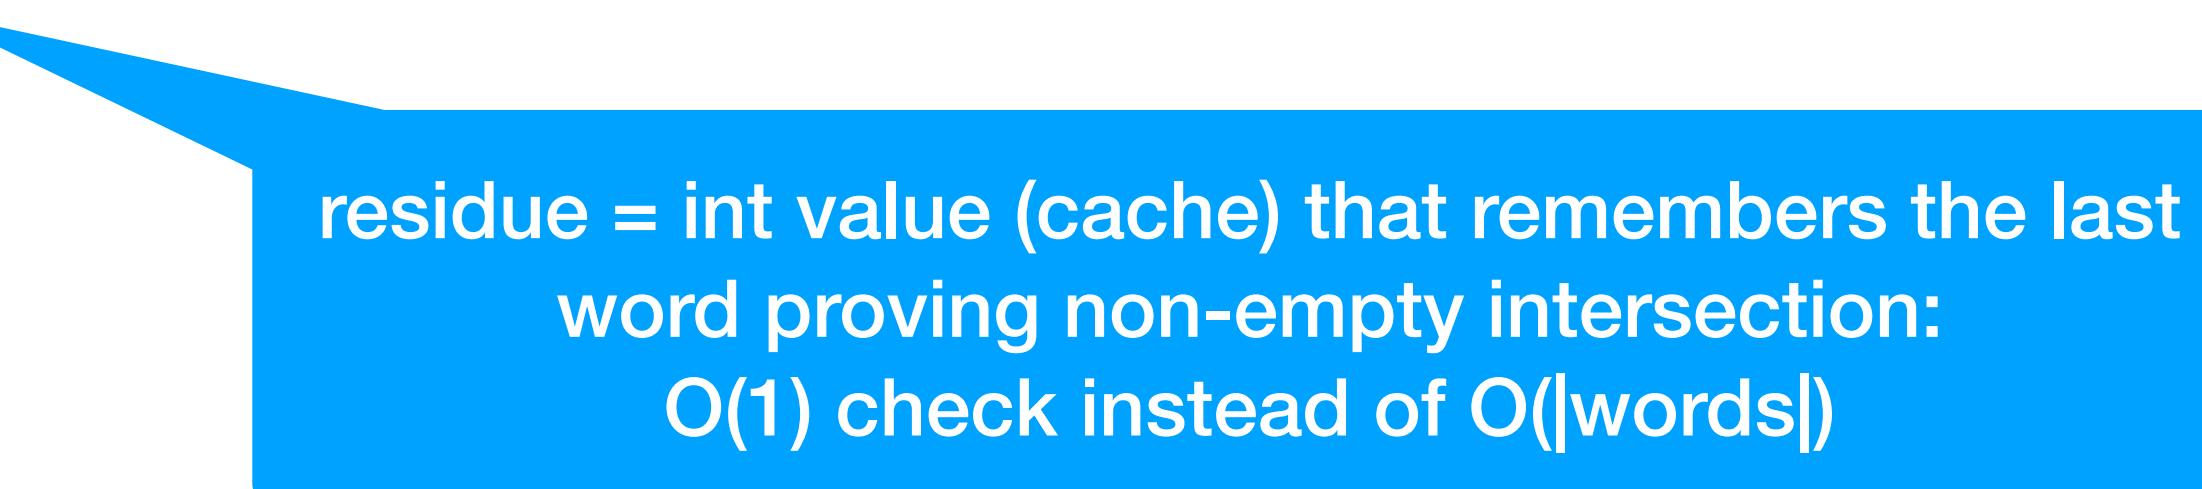
- ▶ Restoration requires only the size variable to be restored



# Can we even further improve the efficiency?

- ▶ Yes, as the last intersecting word is more likely to intersect again
- ▶ Remember the index, called ***residue***, of the last word that led to a non-empty intersection and try it first for the next intersection test with some supports( $x_i, v$ )

```
boolean intersects(validTuples, supports[x, v]) {  
    residue = support[x, v].residue  
    if (validTuples.words[residue] & supports[x, v].words[residue] != 0L)  
        return true  
    i = nonZeroSize  
    while (i > 0){  
        i = i - 1  
        wordIdx = nonZeroIdx[i]  
        if (validTuples.words[wordIdx] & supports[x, v].words[wordIdx] != 0L){  
            supports[x, v].residue = wordIdx  
            return true  
        }  
    }  
    return false  
}
```



# How can we do that efficiently?

All this can be implemented in a data structure called StateSparseBitSet, with the following API:

```
mask.clear() // empty the bit set
mask.or(supports[x,a]) // mask = mask | supports[x,a]
validTuples.and(mask) // words = words & mask
validTuples.intersects(supports[x,c]) // words & supports[x,c] != 0L
```

# Update validTuples with StateSparseBitSet API



```
validTuples = validTuples & (supports(x,3) | supports(x,4))
```

1. mask.clear()
2. mask.or(supports[x,3])
3. mask.or(supports[x,4])
4. validTuples.and(mask)

# Filtering validTuples with StateSparseBitSet API



Is  $x=3$  still possible?

1. `answer = validTuples.intersects(supports[x,3])`
2. `if (!answer) { x.remove(3)}`

- ▶ **Set**: collection of objects
- ▶ **BitSet**: uses bits (one dedicated bit per object) to represent the presence (bit set to 1) or absence (bit set to 0) of an object in the set
- ▶ **Sparse**: optimized to avoid computations on empty parts of the data structure
- ▶ **State**: allow automatic restoration to a previous saved state