

Towards Distributed Constraint Solving with CRDT

Hakan Hasan Pierre Talbot

Interdisciplinary Centre for Security, Reliability and Trust (SnT)
University of Luxembourg
hakan.hasan@uni.lu, pierre.talbot@uni.lu



Workshop on Principles and Practice of Consistency for Distributed Data
27 April 2026



Constraint Programming

Constraint programming (CP) is about modeling and solving problems under constraints.

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



australia.mzn × +



```
1 include "nvalue_fn.mzn";
2
3 enum REGION = { WA, NT, SA, Q, NSW, V, T };
4 enum COLOR = Color(1..card(REGION));
5
6 % Neighboring regions
7 array [_] of tuple(REGION, REGION): neighbors = [
8   (WA, NT),
9   (WA, SA),
10  (NT, SA),
11  (NT, Q),
12  (SA, Q),
13  (SA, NSW),
14  (SA, V),
15  (Q, NSW),
16  (NSW, V)
17 ];
18
19 % Color of each Region
20 array [REGION] of var COLOR: color;
21 % Number of colors used
22 var 1..card(COLOR): n_colors :: output = nvalue(color);
23
24 % Neighboring regions have different colours
25 constraint forall (pair in neighbors)
26   (color[pair.1] != color[pair.2]);
27
28 % Use as few colors as possible
29 solve minimize n_colors;
```



```
australia.mzn x +
1 include "nvalue_fn.mzn";
2
3 enum REGION = { WA, NT, SA, Q, NSW, V, T };
4 enum COLOR = Color(1..card(REGION));
5
6 % Neighboring regions
7 array [_] of tuple(REGION, REGION): neighbors = [
8   (WA, NT),
9   (WA, SA),
10  (NT, SA),
11  (NT, Q),
12  (SA, Q),
13  (SA, NSW),
14  (SA, V),
15  (Q, NSW),
16  (NSW, V)
17 ];
18
19 % Color of each Region
20 array [REGION] of var COLOR: color;
21 % Number of colors used
22 var 1..card(COLOR): n_colors :: output = nvalue(color);
23
24 % Neighboring regions have different colours
25 constraint forall (pair in neighbors)
26   (color[pair.1] != color[pair.2]);
27
28 % Use as few colors as possible
29 solve minimize n_colors;
```

variables



australia.mzn x


+



```
1 include "nvalue_fn.mzn";
2
3 enum REGION = { WA, NT, SA, Q, NSW, V, T };
4 enum COLOR = Color(1..card(REGION));
5
6 % Neighboring regions
7 array [_] of tuple(REGION, REGION): neighbors = [
8   (WA, NT),
9   (WA, SA),
10  (NT, SA),
11  (NT, Q),
12  (SA, Q),
13  (SA, NSW),
14  (SA, V),
15  (Q, NSW),
16  (NSW, V)
17 ];
18
19 % Color of each Region
20 array [REGION] of var COLOR: color;
21 % Number of colors used
22 var 1..card(COLOR): n_colors :: output = nvalue(color);
23
24 % Neighboring regions have different colours
25 constraint forall (pair in neighbors)
26   (color[pair.1] != color[pair.2]);
27
28 % Use as few colors as possible
29 solve minimize n_colors;
```

domains



```
australia.mzn x + 
```

```
1 include "nvalue_fn.mzn";
2
3 enum REGION = { WA, NT, SA, Q, NSW, V, T };
4 enum COLOR = Color(1..card(REGION));
5
6 % Neighboring regions
7 array [_] of tuple(REGION, REGION): neighbors = [
8   (WA, NT),
9   (WA, SA),
10  (NT, SA),
11  (NT, Q),
12  (SA, Q),
13  (SA, NSW),
14  (SA, V),
15  (Q, NSW),
16  (NSW, V)
17 ];
18
19 % Color of each Region
20 array [REGION] of var COLOR: color;
21 % Number of colors used
22 var 1..card(COLOR): n_colors :: output = nvalue(color);
23
24 % Neighboring regions have different colours
25 constraint forall (pair in neighbors)
26   (color[pair.1] != color[pair.2]);
27
28 % Use as few colors as possible
29 solve minimize n_colors;
```

constraint



australia.mzn x +

```
1 include "nvalue_fn.mzn";
2
3 enum REGION = { WA, NT, SA, Q, NSW, V, T };
4 enum COLOR = Color(1..card(REGION));
5
6 % Neighboring regions
7 array [_] of tuple(REGION, REGION): neighbors = [
8   (WA, NT),
9   (WA, SA),
10  (NT, SA),
11  (NT, Q),
12  (SA, Q),
13  (SA, NSW),
14  (SA, V),
15  (Q, NSW),
16  (NSW, V)
17 ];
18
19 % Color of each Region
20 array [REGION] of var COLOR: color;
21 % Number of colors used
22 var 1..card(COLOR): n_colors :: output = nvalue(color);
23
24 % Neighboring regions have different colours
25 constraint forall (pair in neighbors)
26   (color[pair.1] != color[pair.2]);
27
28 % Use as few colors as possible objective
29 solve minimize n_colors;
```



australia.mzn x +


```
1 include "nvalue_fn.mzn";
2
3 enum REGION = { WA, NT, SA, Q, NSW, V, T };
4 enum COLOR = Color(1..card(REGION));
5
6 % Neighboring regions
7 array [_] of tuple(REGION, REGION): neighbors = [
8   (WA, NT),
9   (WA, SA),
10  (NT, SA),
11  (NT, Q),
12  (SA, Q),
13  (SA, NSW),
14  (SA, V),
15  (Q, NSW),
16  (NSW, V)
17 ];
18
19 % Color of each Region
20 array [REGION] of var COLOR: color;
21 % Number of colors used
22 var 1..card(COLOR): n_colors :: output = nvalue(color);
23
24 % Neighboring regions have different colours
25 constraint forall (pair in neighbors)
26   (color[pair.1] != color[pair.2]);
27
28 % Use as few colors as possible
29 solve minimize n_colors;
```

▼ Running australia.mzn

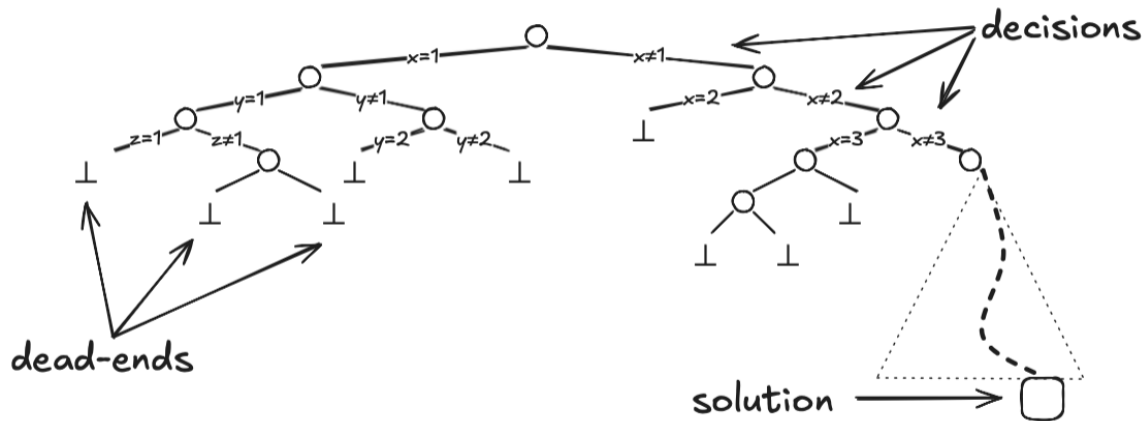
```
color = [WA: Color(3), NT: Color(2), SA:
Color(1), Q: Color(3), NSW: Color(2), V:
Color(3), T: Color(1)];
n_colors = 3;
```

=====

Finished in 177msec.



Backtracking Algorithm



Constraint Solvers



chuffed/**chuffed**

The Chuffed CP solver



16 Contributors 20 Issues 123 Stars 44 Forks

ptal/**turbo**

A constraint solver purely on GPUs (CUDA)



4 Contributors 5 Issues 40 Stars 13 Forks

Constraint Solvers



chuffed/**chuffed**

The Chuffed CP solver



16 Contributors 20 Issues 123 Stars 44 Forks

ptal/**turbo**

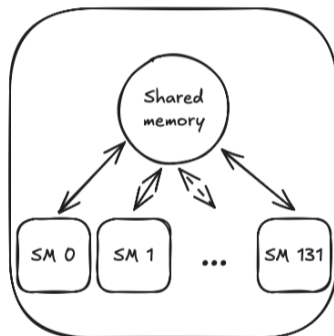
A constraint solver purely on GPUs (CUDA)



4 Contributors 5 Issues 40 Stars 13 Forks

A general purpose GPU CP solver.¹²

The underlying theoretical model is based on lock-free, asynchronous fixpoint iteration.




¹P. Talbot et al., 'A Variant of Concurrent Constraint Programming on GPU', AAI, 2022

²P. Talbot, 'A GPU-based Constraint Programming Solver', AAI, 2026


- Implementation of the first multi-GPU constraint solver.
- Investigation of state-based CRDT³ for distributed constraint solving.
- High-performance implementation using MPI showing near-linear scaling.
- Establishing a connection between CRDT semantics and prior work on asynchronous iteration.

³M. Shapiro et al., 'Conflict-free replicated data types', SSS, 2011


Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing



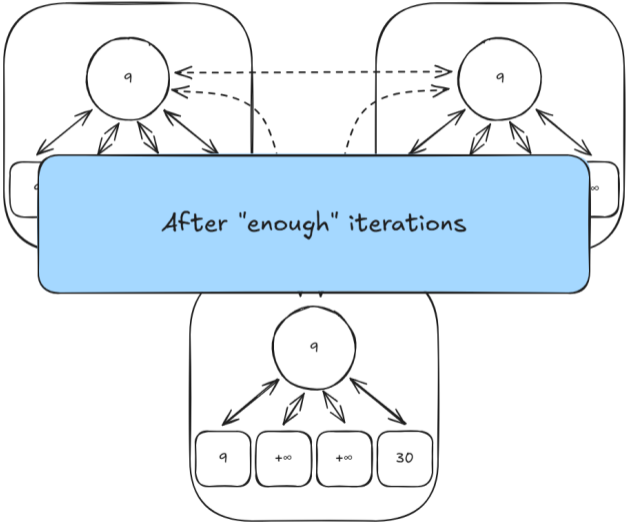
Counter CRDT for Bound Sharing




Counter CRDT for Bound Sharing



Counter CRDT for Bound Sharing



Counter CRDT for Bound Sharing

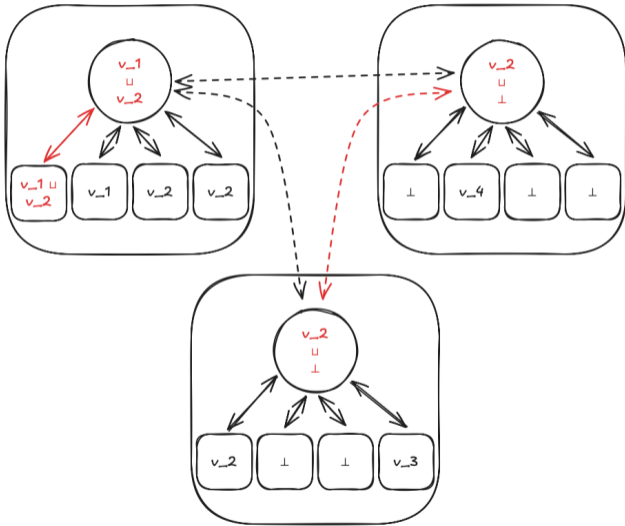


Beyond Bound Sharing: General Model

The solving process produces bounds that update the state monotonically. The accumulation of global information in distributed constraint optimization can be expressed as a fixpoint computation over monotone shared state.

The value of the best bound is the least fixpoint of the combined effect of all of join operators. Operationally, asynchronous execution corresponds to chaotic iteration of these per-update operators.

Fixpoint-Based Coordination Model




Implementation: CRDT-based Distributed Constraint Solving


- Local bound sharing and solver coordination are handled via one-sided MPI operations, ensuring asynchronous execution (in regard to the solving process).
- Hierarchical gossip-style protocol for periodic information exchange between nodes.

Single-GPU and Multi-GPU: One Coordination Mechanism

Inter-GPU



Intra-GPU



- **Single GPU (Intra-GPU):**
 - Threads share a common objective bound stored in an atomic variable.
- **Multi-GPU (Inter-GPU/Inter-Node):**
 - Independent solver instances exchange bounds asynchronously through one-sided MPI communication.
- **Same Logic, Different Scale:**
 - Both cases rely on the same principles: monotonicity, idempotence, and order-independent updates.
 - Fixpoint computation of the best bound.

Experimental Results

Benchmarks: 16 instances from a MiniZinc Challenge.

Time limit: 300s per run

Solver	nodes/sec	S_p	E_p	MiniZinc score
Turbo (16 GPUs)	6174470	14.67	92%	30.7
Turbo (8 GPUs)	2625910	6.24	78%	22.3
Turbo (4 GPUs)	1384340	3.29	82%	21.9
Turbo (2 GPUs)	711093	1.69	84%	19.2
Turbo (1 GPU)	420828	1.00	100%	8.1

Sharing nogoods

Definition

A nogood is an instantiation of a subset of variables that cannot be extended to a solution.

Example

$$\neg(x = a \wedge y = b \wedge z = c)$$

$$\neg(w = c \wedge x = b \wedge y = b \wedge z = a)$$

Recording nogoods permits to improve the process of exploring the search space.

Challenge: need to find a CRDT representation of nogoods.

- The first multi-GPU distributed CP solver.
- Reframing the sharing of global knowledge in constraint solving explicitly in CRDT semantics.
- Implementation of the CRDT bound-sharing concept using MPI over GPU constraint solver.
- Establishing a connection between CRDT semantics and prior work on asynchronous iteration.

Thank you for your attention!
Questions?