


# Parallelism in Constraint Programming

Carl Christian Rolf



# Welcome...



# 'Tiny Sudoku'

X, Y, Z, Q must have values between 1 and 2  
The values in each row and column must be unique

X	Y
Z	Q

# Tiny Sudoku in Constraint Programming

X, Y, Z, Q must have values between 1 and 2

The values in each row and column must be unique

X	Y
Z	Q



$\{X, Y, Z, Q\} \in \{1..2\},$

$X \neq Y,$

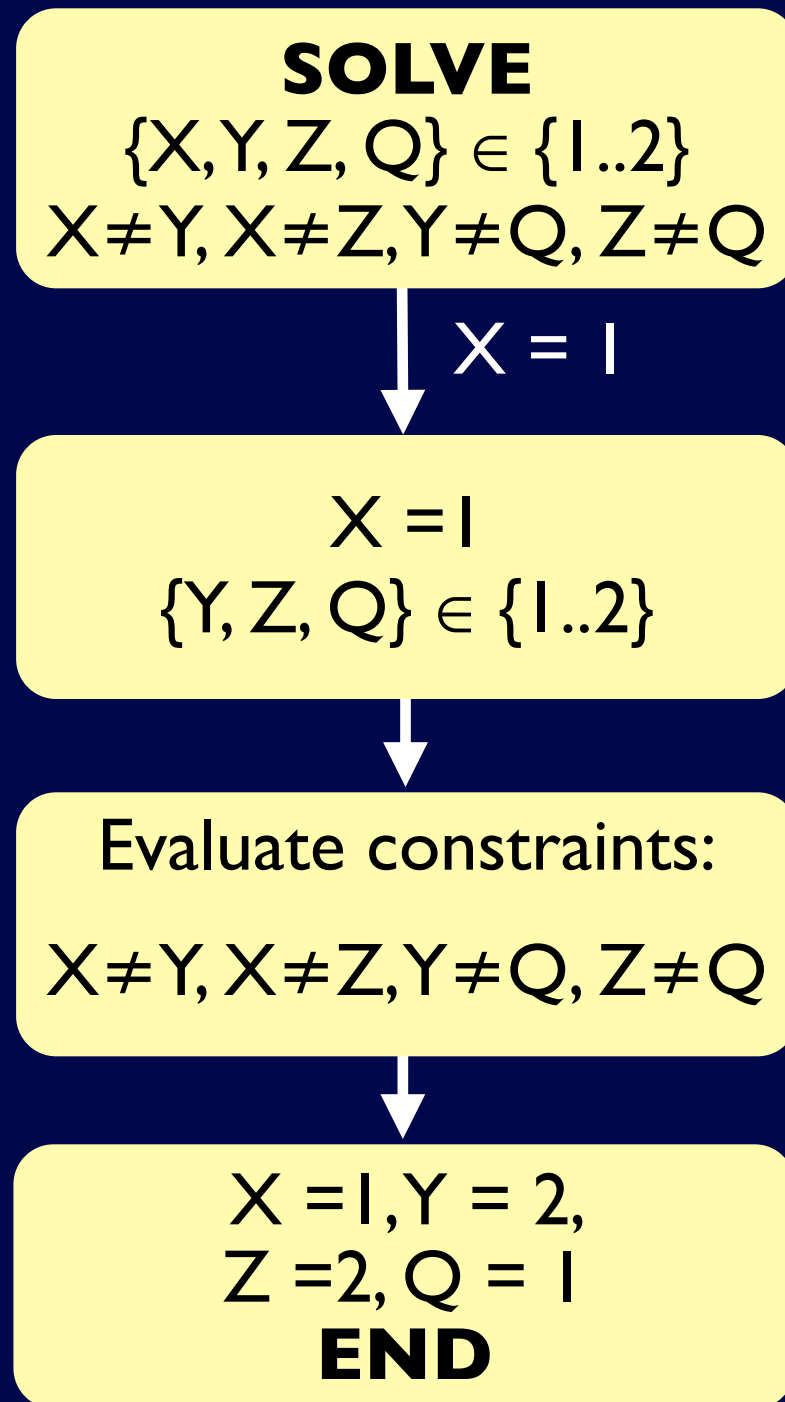
$X \neq Z,$

$Y \neq Q,$

$Z \neq Q,$

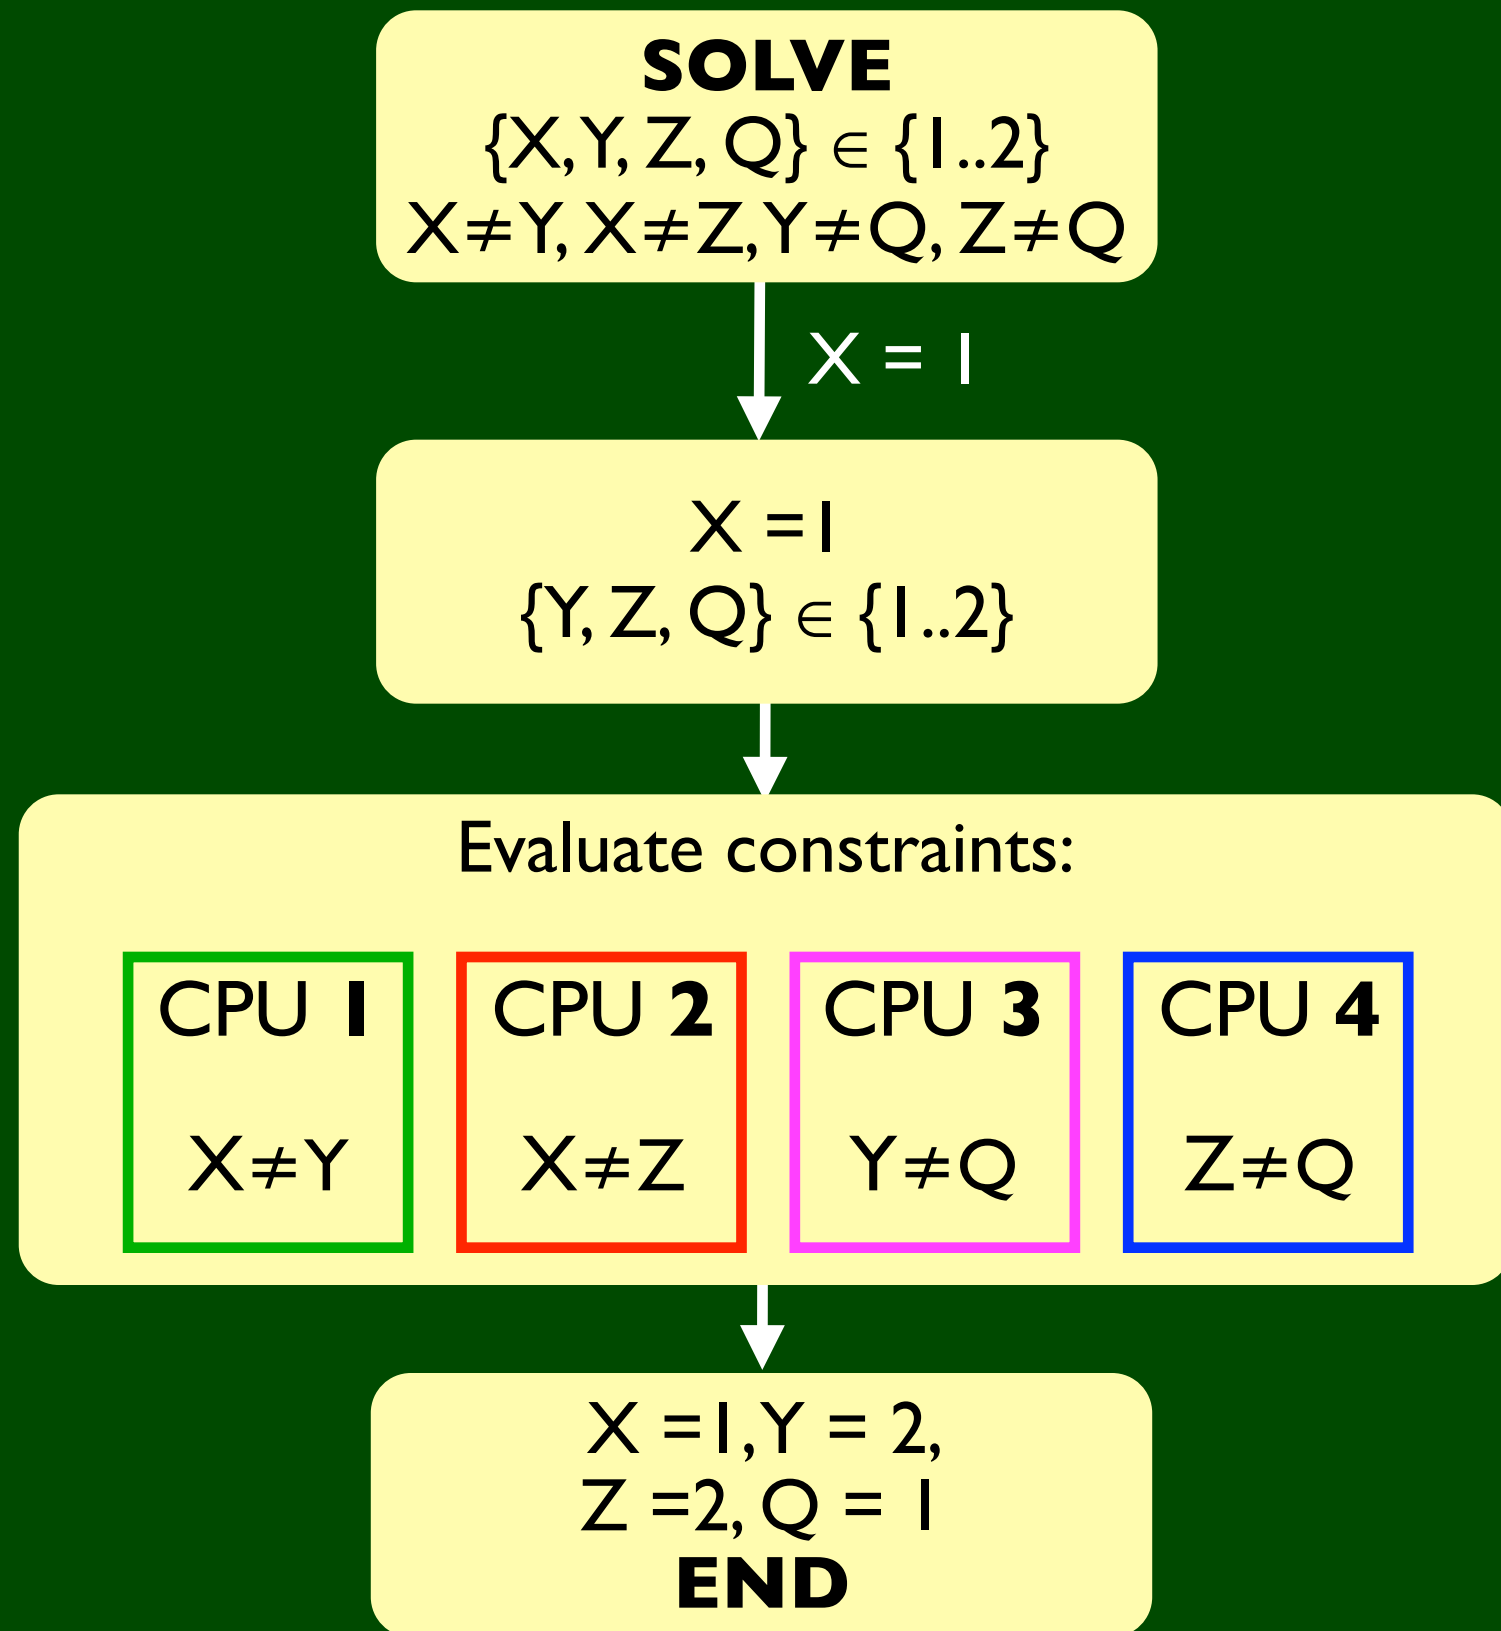
**solve.**

# Solving Tiny Sudoku in CP



# Parallel Consistency

# Parallel Consistency for Tiny Sudoku



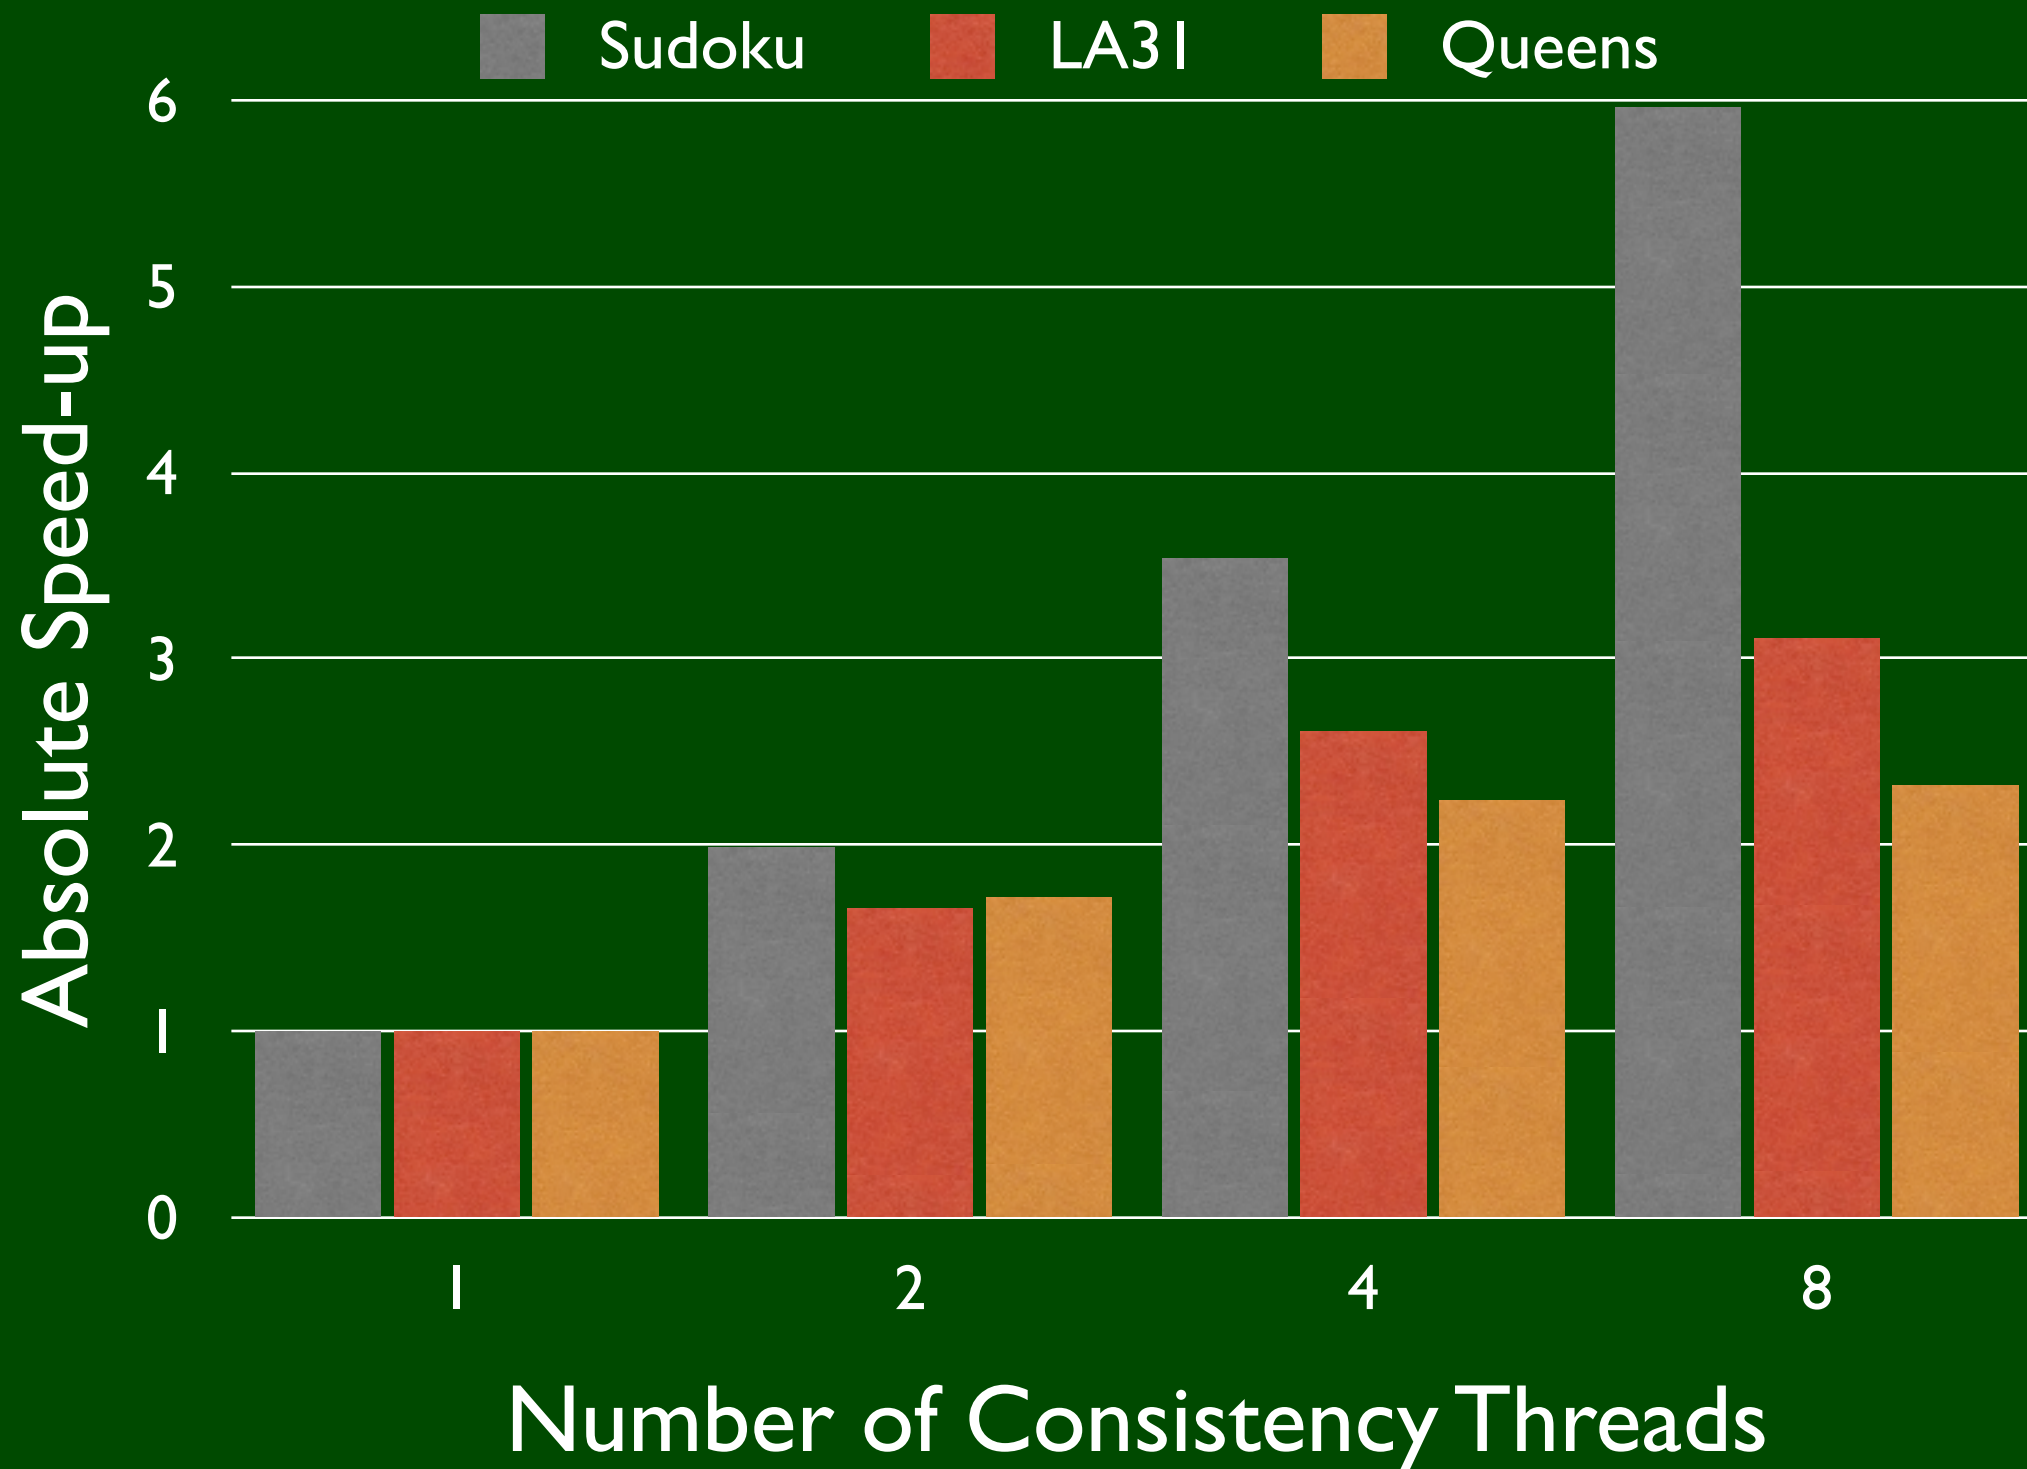
# Parallel Consistency: Summary

- The ‘task-parallelism’ of constraint programming
- Previous work has not handled global constraints
- Might have to run more iterations of consistency
- Load-balancing depends on the problem model
- Hard to share data during consistency



# Parallel Consistency: Performance

Run on an eight core Mac Pro



# Parallel Consistency: Conclusions

- Excellent performance for regular problems
- Some problems do not scale well, they need constraint-specific parallel consistency algorithms

Combining  
Parallel Consistency  
with  
Parallel Search

# Parallel Search for Tiny Sudoku

CPU 1

**SOLVE**  
 $\{X, Y, Z, Q\} \in \{1..2\}$   
 $X \neq Y, X \neq Z, Y \neq Q, Z \neq Q$

$X = 1$

$X = 1$   
 $\{Y, Z, Q\} \in \{1..2\}$

Evaluate constraints:  
 $X \neq Y, X \neq Z, Y \neq Q, Z \neq Q$

$X = 1, Y = 2,$   
 $Z = 2, Q = 1$   
**END**

$X = 2$

CPU 2

$X = 2$   
 $\{Y, Z, Q\} \in \{1..2\}$

Evaluate constraints:  
 $X \neq Y, X \neq Z, Y \neq Q, Z \neq Q$

$X = 2, Y = 1,$   
 $Z = 1, Q = 2$   
**END**

# Combining Parallelisms: Tiny Sudoku

CPU 1

**SOLVE**  
 $\{X, Y, Z, Q\} \in \{1..2\}$   
 $X \neq Y, X \neq Z, Y \neq Q, Z \neq Q$

$X = 1$

$X = 1$   
 $\{Y, Z, Q\} \in \{1..2\}$

Evaluate constraints:

CPU 1

$X \neq Y$

CPU 2

$X \neq Z$

CPU 3

$Y \neq Q$

CPU 4

$Z \neq Q$

$X = 1, Y = 2,$   
 $Z = 2, Q = 1$   
**END**

$X = 2$

CPU 5

$X = 2$   
 $\{Y, Z, Q\} \in \{1..2\}$

Evaluate constraints:

CPU 5

$X \neq Y$

CPU 6

$X \neq Z$

CPU 7

$Y \neq Q$

CPU 8

$Z \neq Q$

$X = 2, Y = 1,$   
 $Z = 1, Q = 2$   
**END**

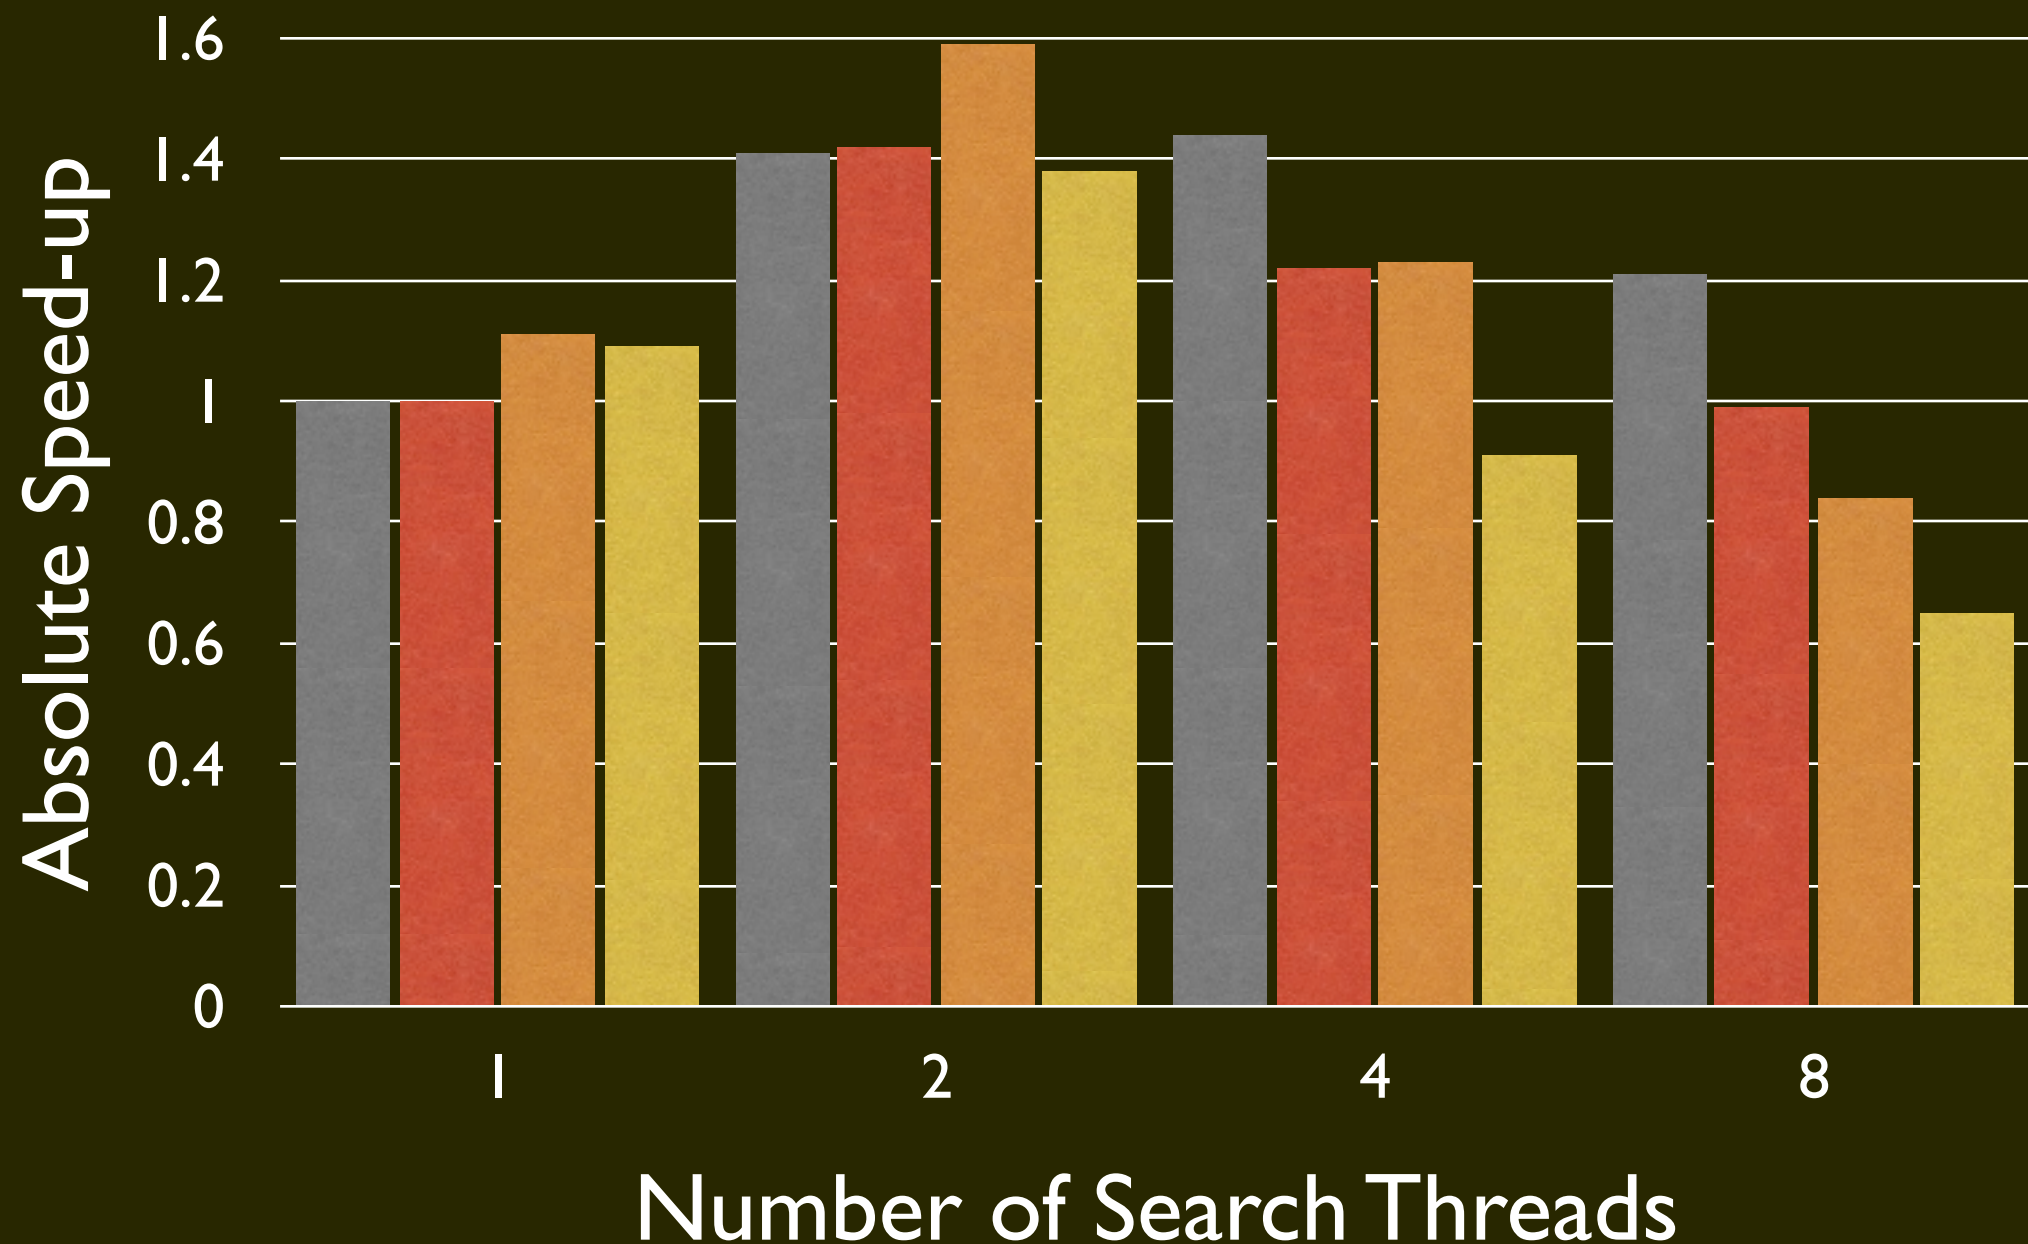
# Combining Parallelisms: Summary

- Never studied before in constraint programming
- Easier to achieve a speed-up than if the solver only offers one type of parallelism
- Does not fit all types of problems, often needs problem-specific optimization of, e.g., thread allocation

# Combining Parallelisms: Results

Finding 200 solutions to 100x100 Sudoku. Run on 8-core Mac Pro

Consistency threads per search thread: ■ 1 ■ 2 ■ 4 ■ 8



# Combining Parallelisms: Conclusions

- Needs better control of mutual exclusion than currently offered by Java
- The problem must suit both types of parallelism to get a large performance benefit
- Problem-specific optimizations are necessary for good performance



# Relative-Measured Load-Balancing

# Load-Balancing for Tiny Sudoku

CPU 1

**SOLVE**  
 $\{X, Y, Z, Q\} \in \{1..2\}$   
 $X \neq Y, X \neq Z, Y \neq Q, Z \neq Q$

$X = 1$

$X = 1$   
 $\{Y, Z, Q\} \in \{1..2\}$

Evaluate constraints:  
 $X \neq Y, X \neq Z, Y \neq Q, Z \neq Q$

$X = 1, Y = 2,$   
 $Z = 2, Q = 1$   
**END**

Should CPU 1  
or CPU 2  
be allowed to  
send work?

CPU 2

Waiting for work

CPU 3

**SOLVE**  
 $\{X, Y, Z, Q\} \in \{1..2\}$   
 $X \neq Y, X \neq Z, Y \neq Q, Z \neq Q$

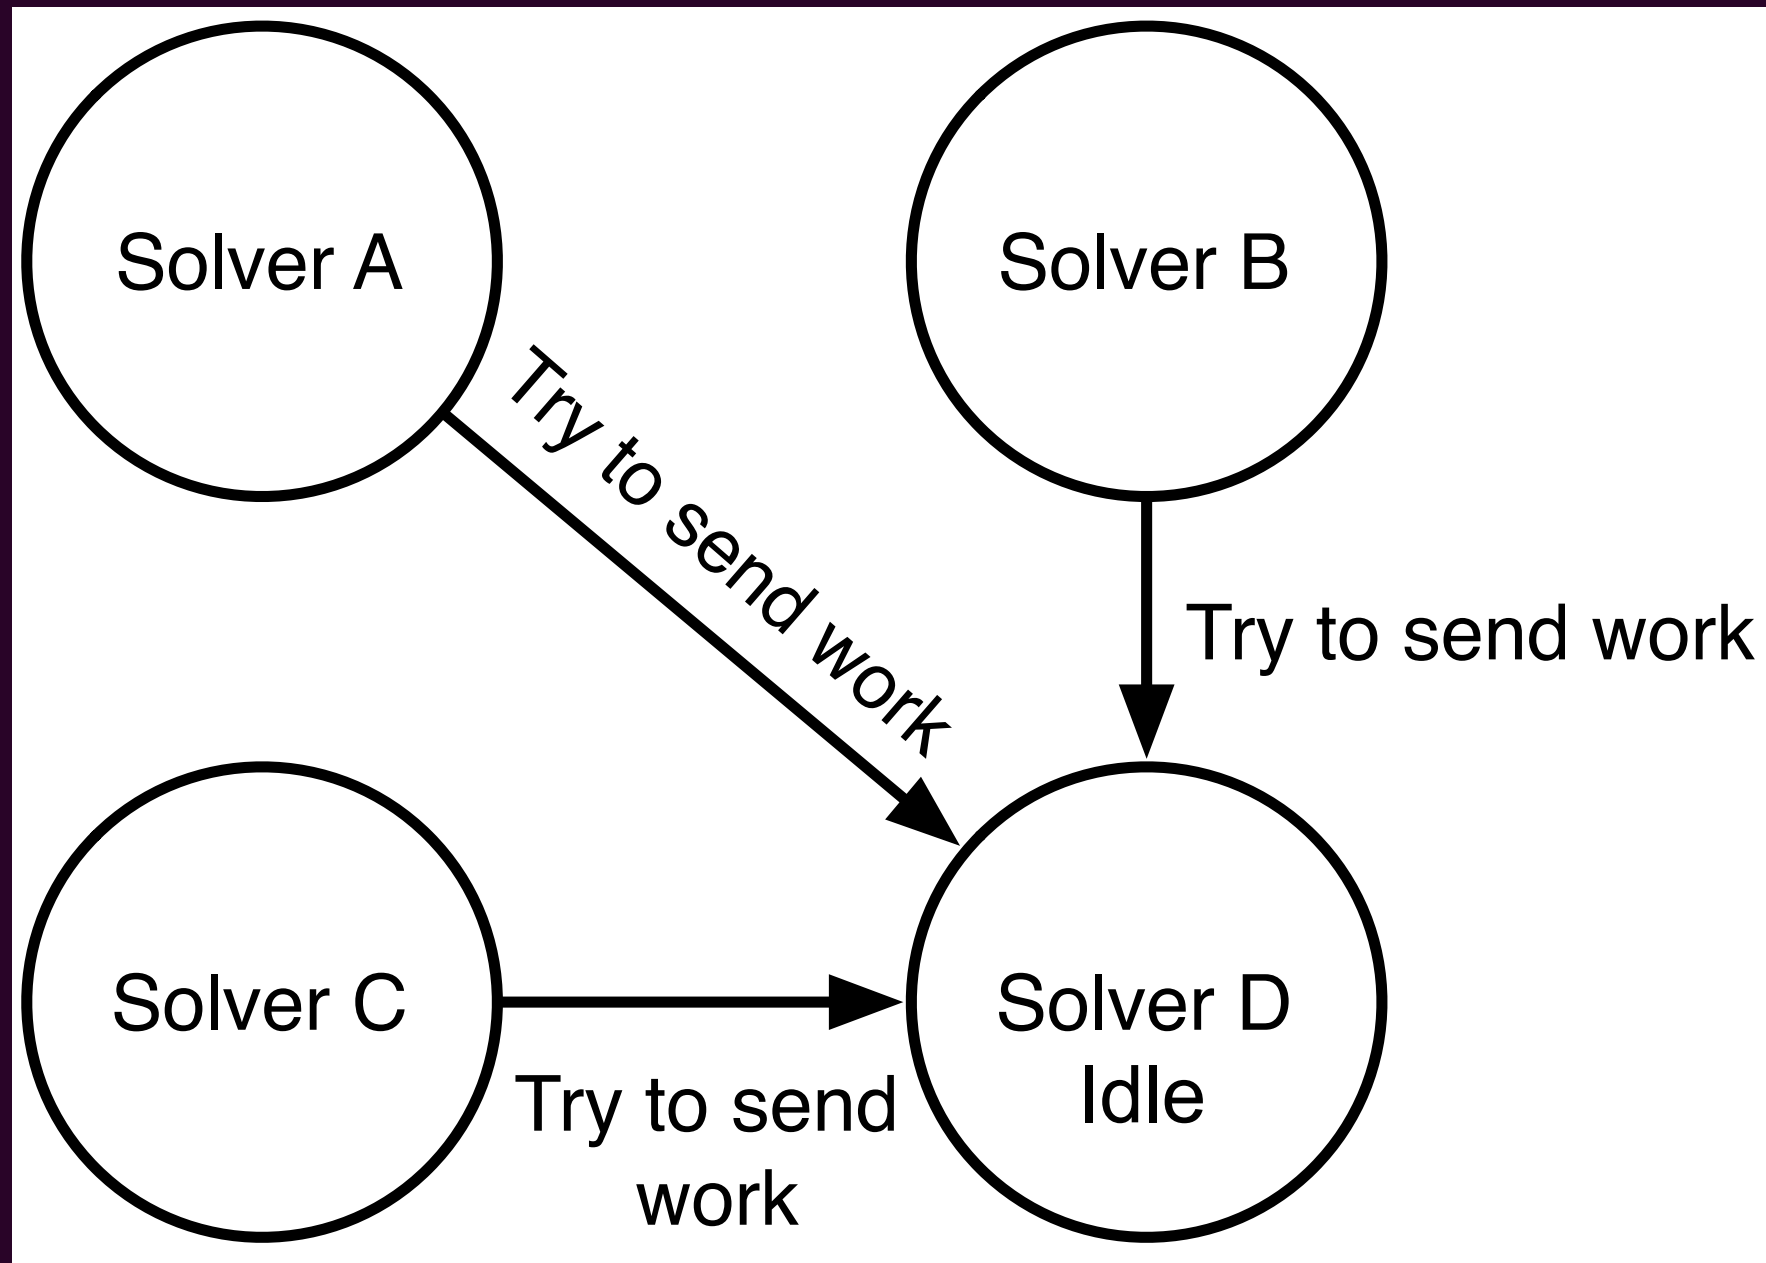
$X = 1$

$X = 1$   
 $\{Y, Z, Q\} \in \{1..2\}$

Evaluate constraints:  
 $X \neq Y, X \neq Z, Y \neq Q, Z \neq Q$

$X = 1, Y = 2,$   
 $Z = 2, Q = 1$   
**END**

# Relative-Measured Load-Balancing



Busy solvers always compete for sending work to idle solvers


# Relative-Measured Load-Balancing: Summary

- Infeasible to get an exact measure of the work-size due to the way CP solves problems
- We let solvers compete based on their work-size estimates
- We can use any measure that can be partially ordered
- Using measures from several solvers increases accuracy by eliminating systematic errors

# Relative-Measured Load-Balancing: Performance

Golomb-12, proving optimality. Slowest and fastest measures

■ Random Polling   ■ Least Labeled First   ■ Most Labeled First



# Relative-Measured Load-Balancing: Conclusions

- Relative measures lets even simple estimates outperform random polling by over 20%
- Advanced measures can easily be used
- Performance benefit increases with the number of solvers

# Dynamic Balancing of Communication and Computation

# Dynamic Balancing for Tiny Sudoku

CPU 1

**SOLVE**  
 $\{X, Y, Z, Q\} \in \{1..2\}$   
 $X \neq Y, X \neq Z, Y \neq Q, Z \neq Q$

$X = 1$

$X = 1$   
 $\{Y, Z, Q\} \in \{1..2\}$

Evaluate constraints:  
 $X \neq Y, X \neq Z, Y \neq Q, Z \neq Q$

$X = 1, Y = 2,$   
 $Z = 2, Q = 1$   
**END**

$X = 2$

CPU 2

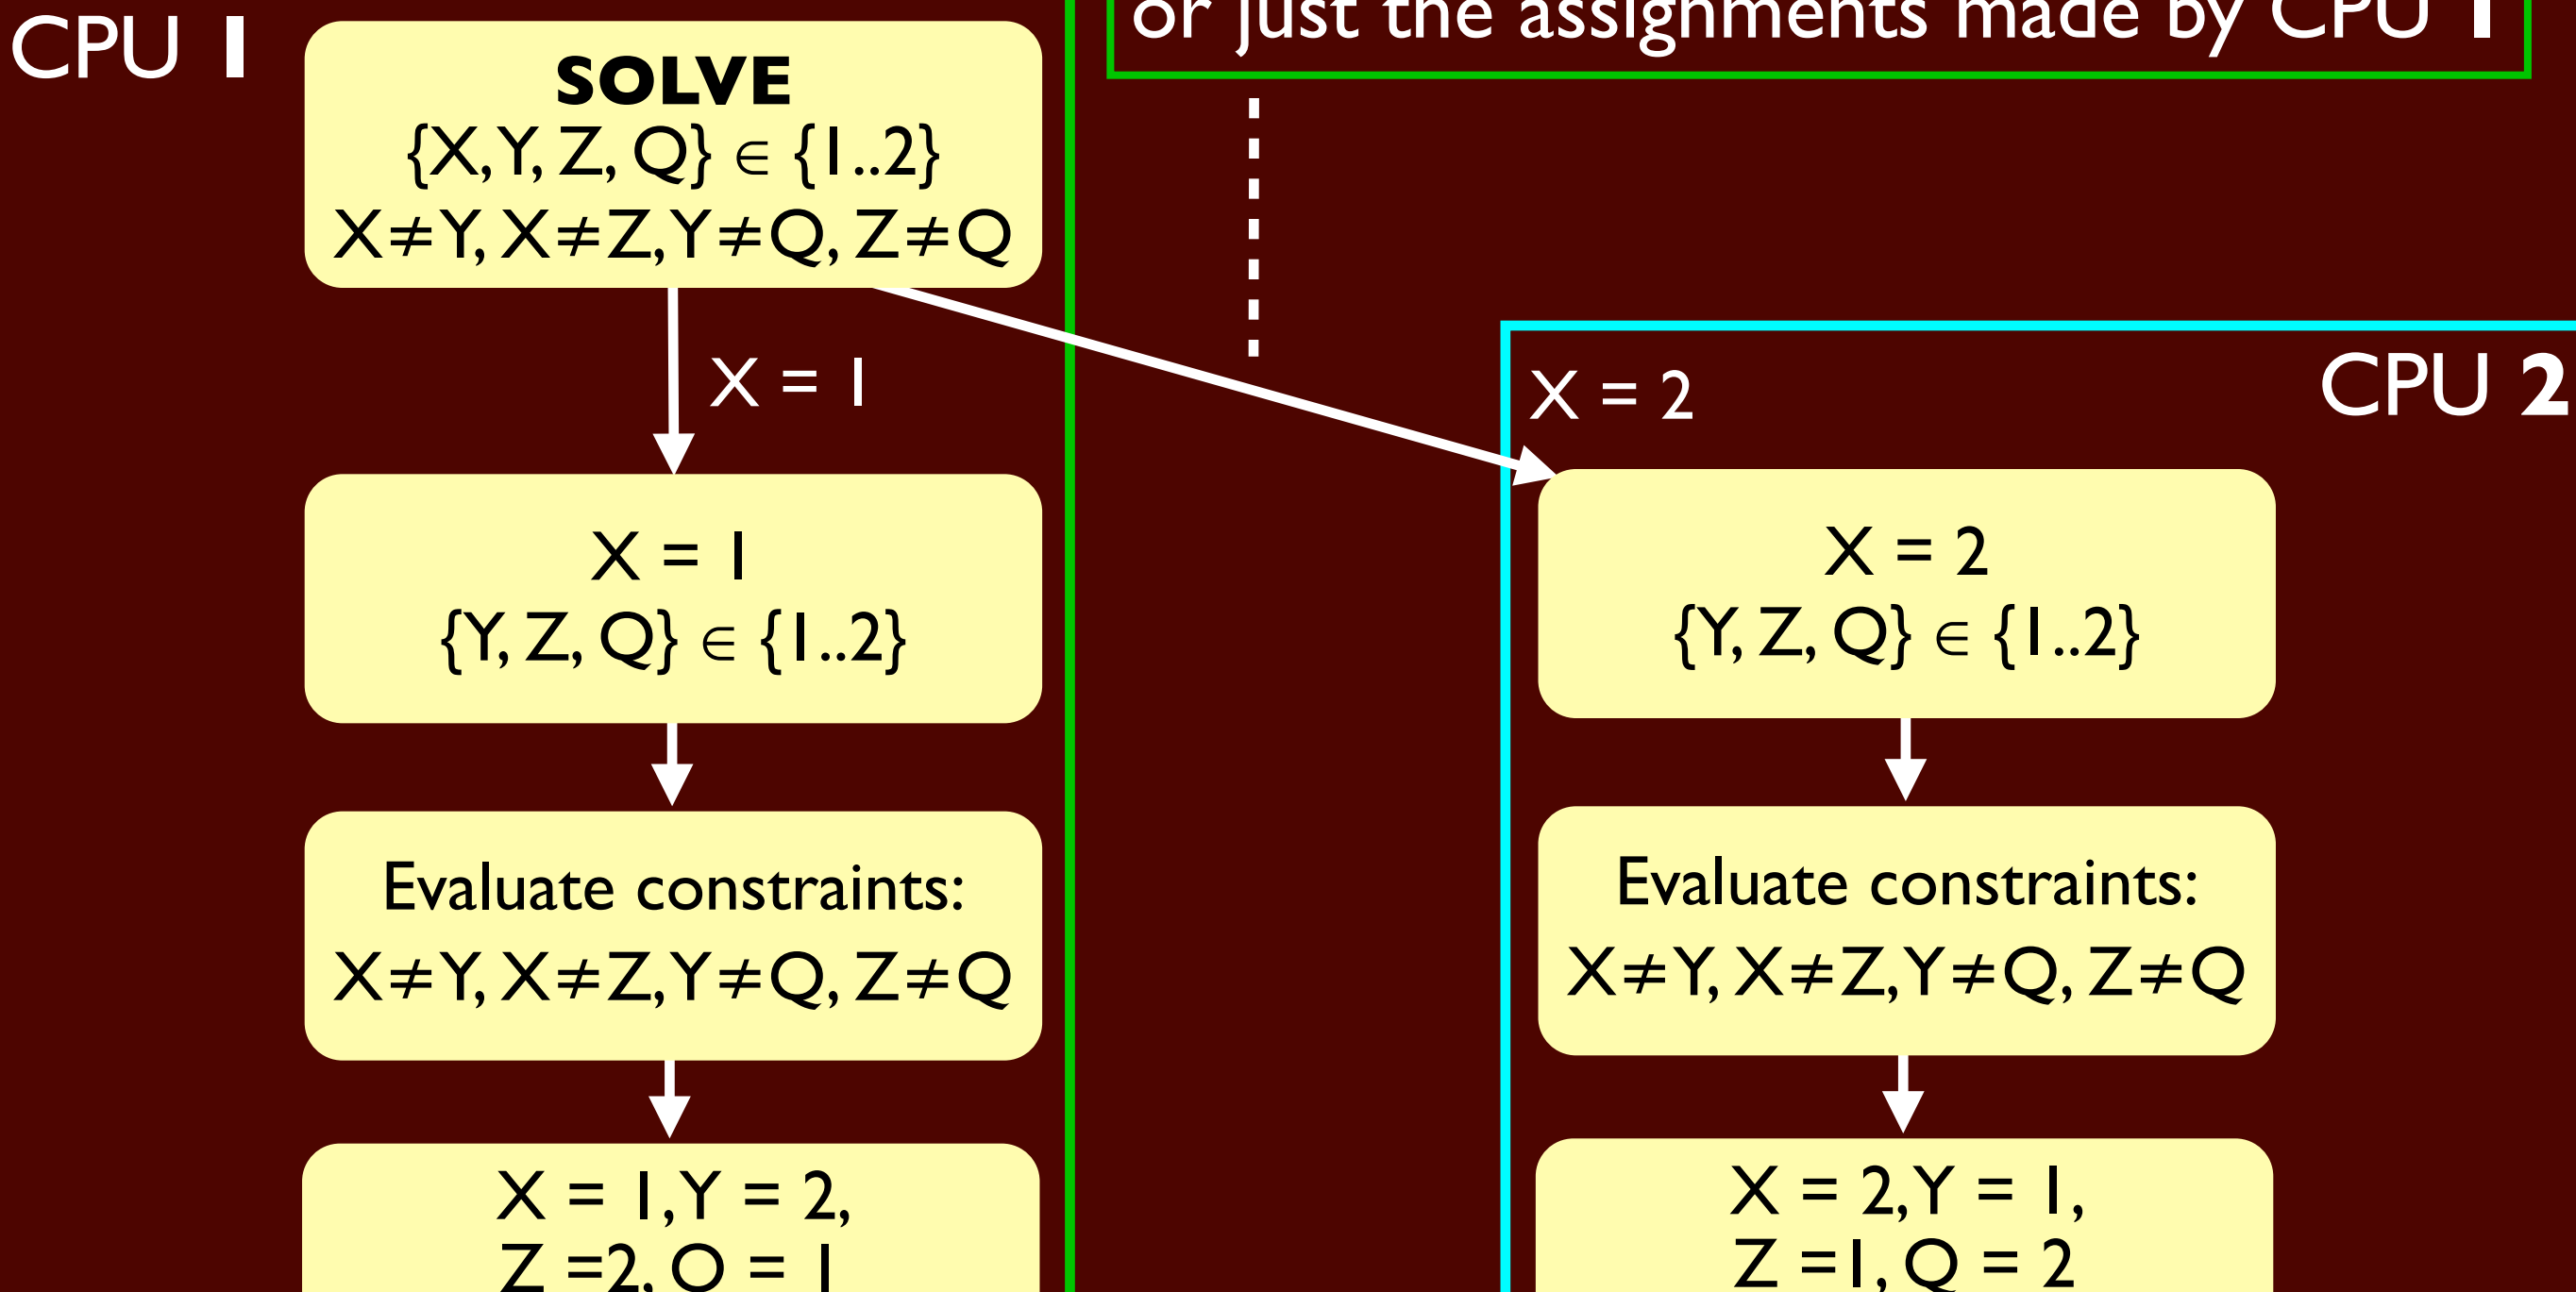
$X = 2$   
 $\{Y, Z, Q\} \in \{1..2\}$

Evaluate constraints:  
 $X \neq Y, X \neq Z, Y \neq Q, Z \neq Q$

$X = 2, Y = 1,$   
 $Z = 1, Q = 2$   
**END**



# Tiny Sudoku: Zoomed in




# Balancing Communication and Computation: Summary

- Copying sends a lot of data, but needs very little processing
- Recomputation often needs a lot of processing, but sends little information
- We estimate the network load to avoid getting stuck in performance bottlenecks

# Balancing Communication and Computation: Results

Proving optimal Golomb ruler of size 12

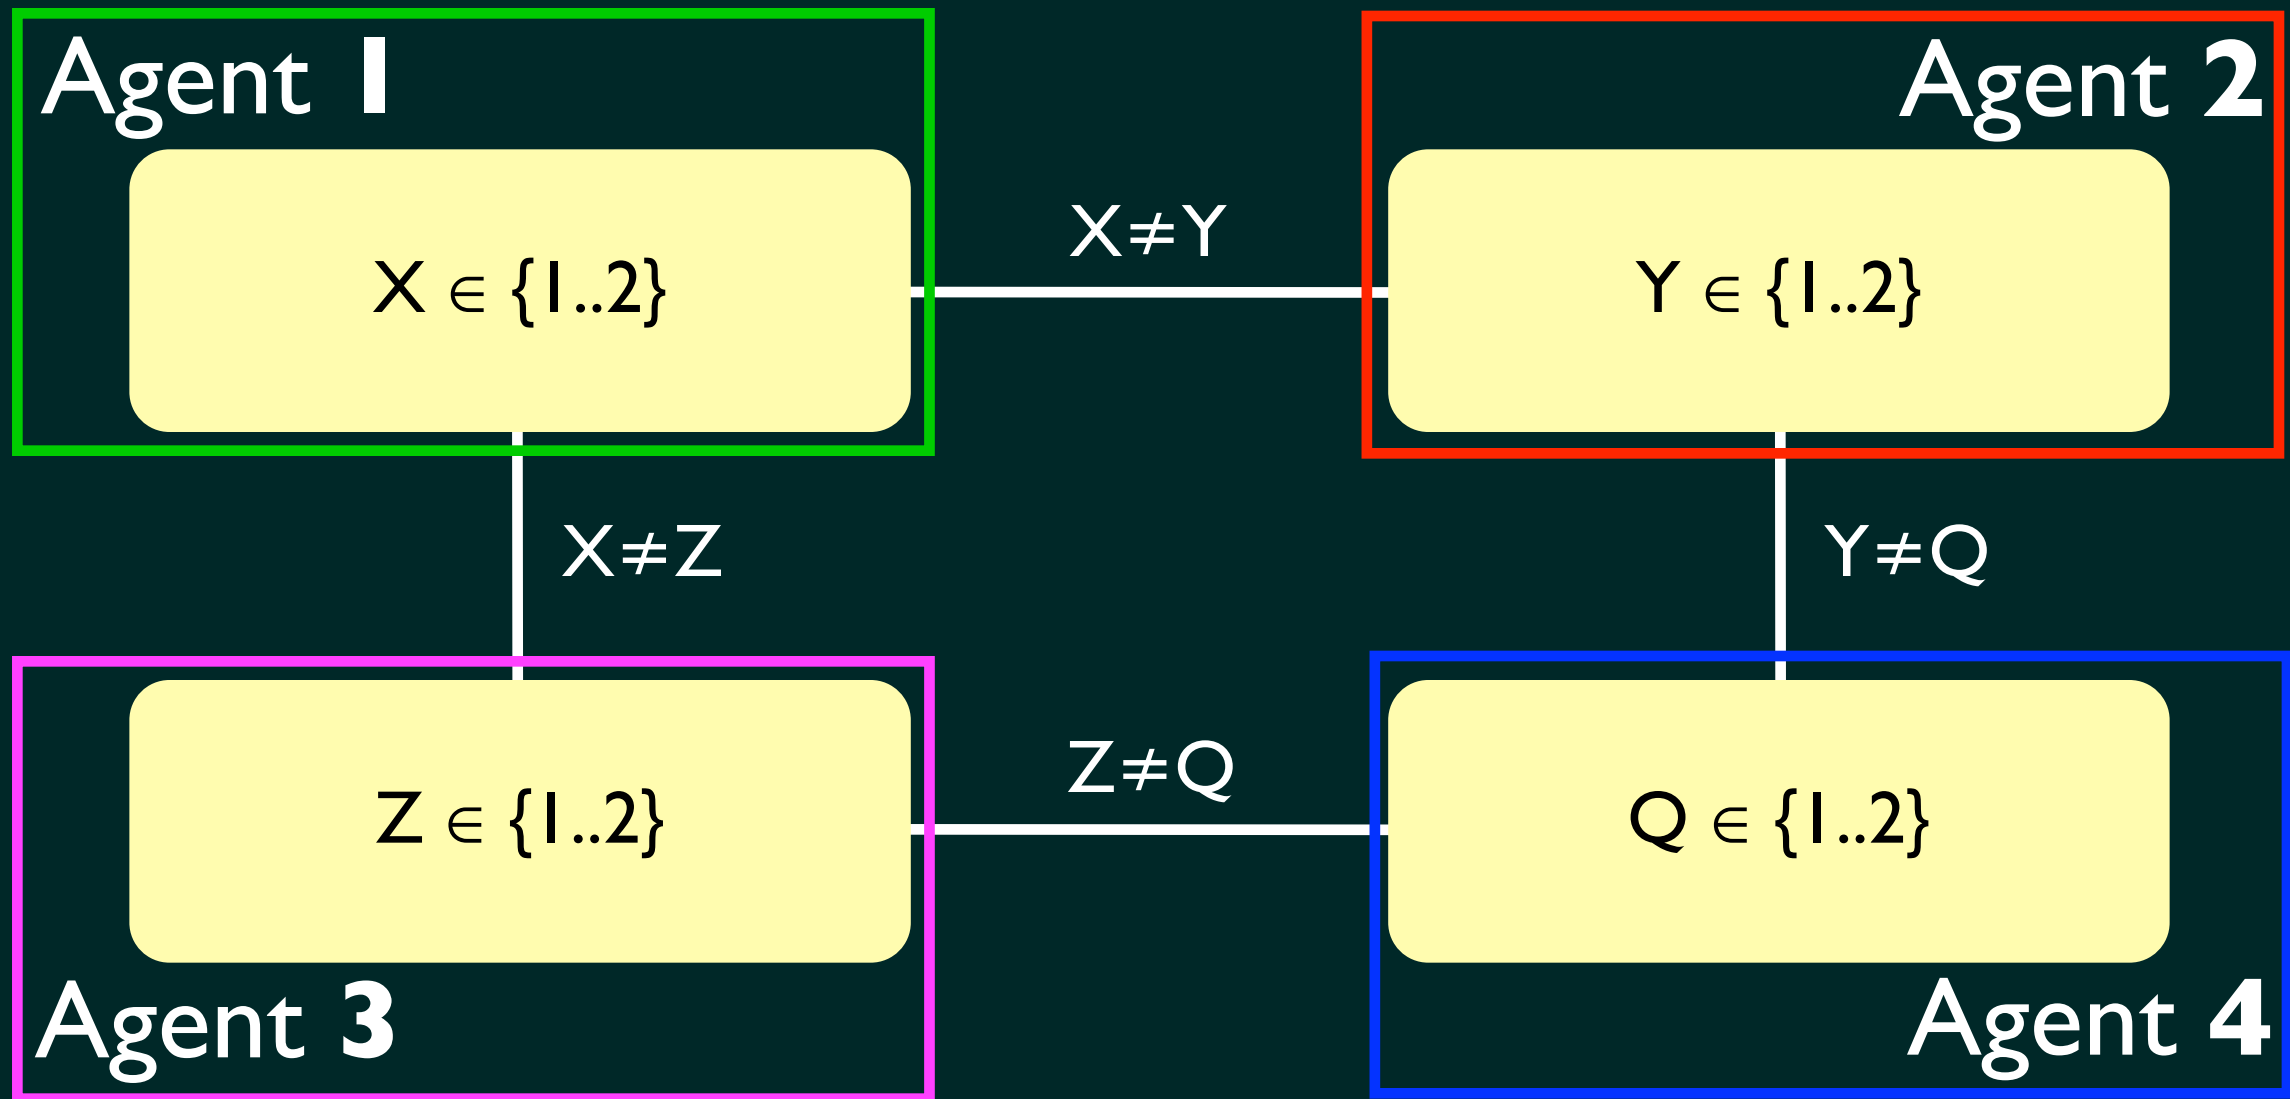


# Balancing Communication and Computation: Conclusions

- Switching dynamically between copying and recomputation often increases performance
- Simple measure to estimate where the performance bottlenecks are

# Distributed Constraint Programming with Agents (DCP)

# Tiny Sudoku in DCP



Solving starts in one agent, constraints  
communicate prunings




# Our use of DCP

- To be used in UAVs in catastrophe areas
- We want independent agents to cooperate, to for instance share a heat camera
- We want to find good, preferably optimal, schedules

# Example of Job Shop Scheduling

Two jobs, consisting of three tasks. Tasks have to execute on specific machines in a certain order





# DCP: In Contrast to the Traditional Approach

- A full constraint solver in each agent
- A set of variables in each agent
- A set of  $n$ -ary (global) constraints in each agent
- No use of memory-demanding table constraints
- Advanced search methods

# DCP: Experimental Results

## Traditional Model

Problem	Time	Solution
LA01	>30min	936
LA04	>30min	976
LA05	>30min	720
MT06	87.7s	55

## Our Model

Problem	Time	Solution
LA01	3.8s	666
LA04	10.8s	590
LA05	0.7s	593
MT06	3.0s	55

We proved the optimum of all problems, the traditional model of none

# DCP: Conclusions

- Our model outperforms traditional models by orders of magnitude
- The best traditional approaches will remain slower, as speed-up is limited

# Overall Conclusions

- We've developed and evaluated several new ways to parallelize constraint solving with global constraints
- Our model of distributed constraint programming vastly outperforms traditional approaches
- Parallelism in constraint programming can, with no understanding of parallel programming, give well-scaling performance for many kinds of problems

# Lastly, a nice Quote

“We finish this review with a single paper, probably one that best represents the state of the art [31].”

[31] C.C. Rolf and K. Kuchcinski. *Combining parallel search and parallel consistency in constraint programming*. In TRICS workshop at CP2010, pages 38–52, 2010.

From:

I. P. Gent, C. Jefferson, I. Miguel, N. C.A. Moore, P. Nightingale, P. Prosser, C. Unsworth. *A Preliminary Review of Literature on Parallel Constraint Solving*. In Parallel Methods for Constraint Solving workshop at CP2011, pages 7–19, 2011.

Thank You!