A SAT+CAS Attack on the Kochen Specker Problem from Quantum Foundations

Brian Li, Curtis Bright, Vijay Ganesh

University of Waterloo, Waterloo, ON, Canada

brian.li@uwaterloo.ca

November 23, 2022

Spin of a Particle

One of the central ideas of quantum mechanics is the notion of spin. Certain subatomic particles have spin. Given a direction, a particle can spin up (positive), down (negative), or not at all.

UNIVERSITY OF WATERLOO

Image Source: S. Weinert

The Stern–Gerlach Experiment (1922)

We can observe the particle spinning by performing such experiment.

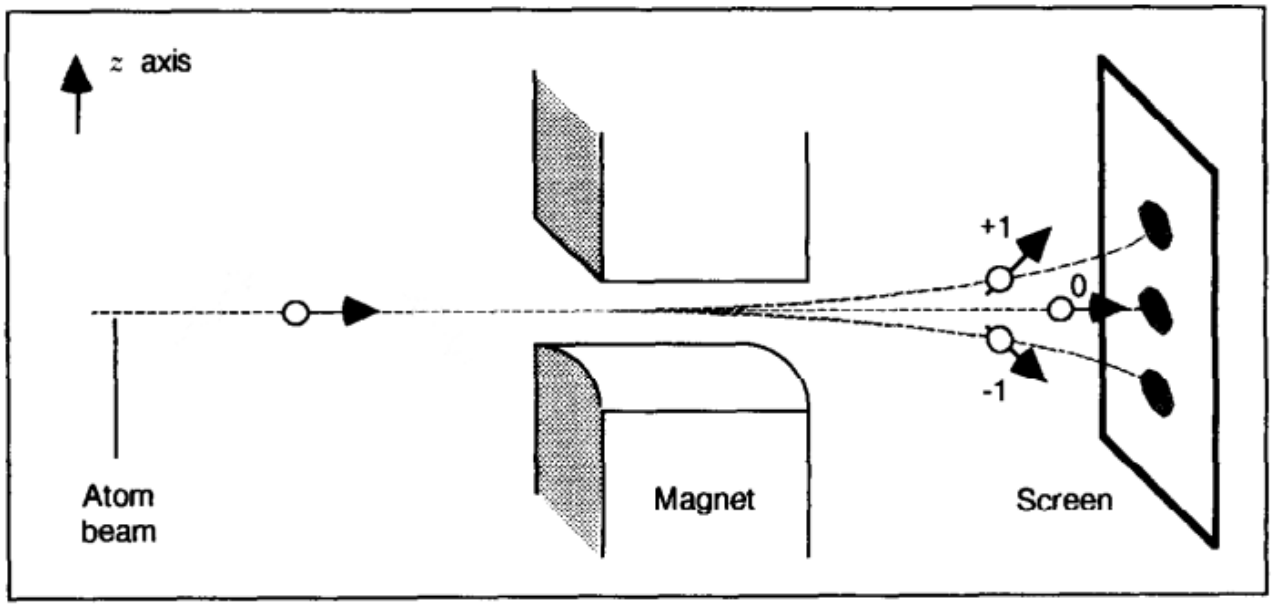

The *spin* of the atom (in the direction of the field) is $+1$, -1, or 0.

UNIVERSITY OF WATERLOO

Image Source: S. Weinert

The SPIN Axiom

Suppose the ± 1 beams are combined producing the *squared spin* which is either 1 or 0.

If you measure this in the x, y and z axes it will be zero **in exactly one of these directions**.

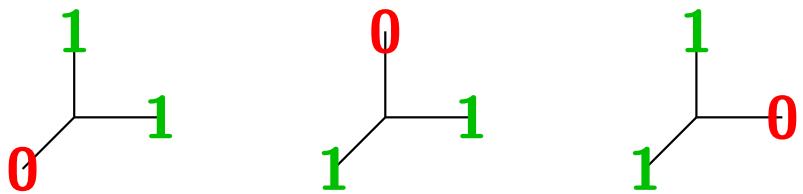

UNIVERSITY OF WATERLOO

The Kochen–Specker Theorem

It is impossible to assign $\{\mathbf{0}, \mathbf{1}\}$ values to the following 31 vectors in a way that does not violate the SPIN axiom.

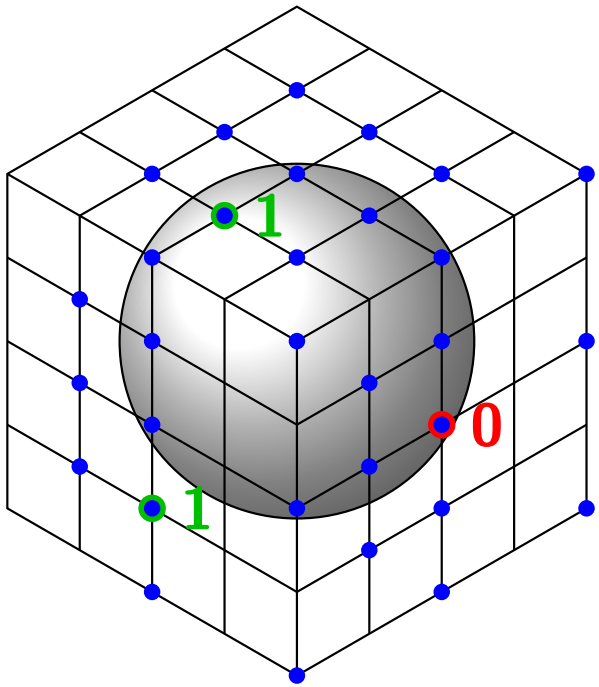

31 vector KS system of Conway and Kochen

The particle cannot have a predetermined spin in every direction.

Can We Do Better Than 31?

The best known result is that at least 22 vectors are required.[1]

This was shown by translating a hypothetical 21-vector KS system into a 21-vertex graph and performing an exhaustive search.

There are a huge number of such graphs and the computation took 75 CPU years using the best graph enumeration algorithms.

UNIVERSITY OF WATERLOO

[1] S. Uijlen, B. Westerbaan. A Kochen-Specker System Has at Least 22 Vectors. *New Generation Computing*, 2016.

Motivation

Theoretically, finding the minimum 3-D KS system is an interesting problem attempted by many renowned mathematicians and physicists using various different approaches.

Practically, finding the minimum KS systems have direct applications in quantum information processing, more specifically in

- security of quantum cryptographic protocols
- zero-error classical communication
- dimension witnessing

UNIVERSITY OF WATERLOO

Related Works on the KS Problem

Authors	Year	KS
Kochen, Specker	1967	≤ 117
Jost	1976	≤ 109
Conway, Kochen	1990	≤ 31
Arends, Ouaknine, Wampler	2009	≥ 18
Uijlen, Westerbaan	2016	≥ 22
Li, Bright, Ganesh	2022	≥ 23

Table: A history of the bounds on the size of the minimum KS system.

UNIVERSITY OF WATERLOO

Our Main Result

We improved the lower bound on the size of the 3-D KS system from 22 to 23 with a significant speed-up over previous computational approaches.

Our approach combines satisfiability (SAT) solvers, computer algebra systems (CASs), and SAT modulo theory (SMT) solvers.

UNIVERSITY OF WATERLOO

Introduction to Satisfiability (SAT) solver

A SAT solver is an algorithm for establishing satisfiability. It takes in CNF formulas as input, and returns

- SAT if it finds a combination of variables that can satisfy the formula
- UNSAT if it can demonstrate that no such combination exists

The Boolean logic formula are converted to a standard form that it is more amenable to algorithmic manipulation. Any formula can be re-written as a conjunction of disjunctions (i.e., the logical AND of statements containing OR relations). We call such form the conjunctive normal form (CNF).

$$\phi := (x_1 \vee x_2 \vee x_3) \wedge (\neg x_1 \vee x_2 \vee x_3) \wedge (x_1 \vee \neg x_2 \vee x_3)$$

UNIVERSITY OF WATERLOO

Introduction to Computer Algebra System (CAS)

Computer algebra systems can perform calculations and manipulate expressions from many branches of mathematics:

- Row reducing a matrix
- Evaluating sums, integrals, etc
- Computing symmetries of combinatorial objects

The two commercially successful CAS programs are Maple and Mathematica.

UNIVERSITY OF WATERLOO

Motivation for the SAT + CAS paradigm

We aim to provide the best aspects of both the SAT and CAS approaches while minimizing the weaknesses of each respective tool. For example, one of the primary drawbacks of SAT solvers is that they lack mathematical expressiveness—many mathematical concepts are difficult or even impossible to efficiently encode in Boolean logic. On the other hand, a huge variety of mathematical concepts can easily be expressed in a CAS. Thus, the SAT+CAS paradigm combines the search power of a SAT solver with the expressive power of a CAS.

UNIVERSITY OF WATERLOO

Converting SPIN axiom to 010-colorability

The squared spin components of a spin-1 particle are 1, 0, 1 in these three directions. Thus, the observable corresponding to the question "is the squared spin 0?" measured in three mutually orthogonal directions will always produce *yes* (or 1) in exactly one direction and *no* (or 0) in the other two orthogonal directions in 3-dimensional Euclidean space.

Satisfying the SPIN axiom is equivalent to being 010-colorable:

- Two orthogonal vectors are not both assigned to 1.
- Three mutually orthogonal vectors are not all assigned to 0.

UNIVERSITY OF WATERLOO

Reduction to Satisfiability (SAT)

An **orthogonality graph** $G_{\mathcal{K}}$ of a vector system has vertices corresponding to the vectors, and two vertices are connected if and only if their corresponding vectors are orthogonal.

To find a KS system, we want to find graphs G such that

- G is **Non-010-colorable**: G has no possible 010-coloring
- G is **Embeddable**: G is an orthogonality graph for a 3-d vector system

In addition, previous research has proven mathematically that G satisfies

- **Squarefree Constraint**: G must not contain a C_4 subgraph
- **Minimum Degree Constraint**: every vertex of G must have minimum degree 3
- **Triangle Constraint**: every vertex of G is contained in at least one C_3 subgraph.

Reduction to Satisfiability (SAT)

We assign a Boolean variable e_{ij} for edge between vertex i and j. There are in total $\binom{n}{2}$ such variables. If e_{ij} is True, then vertex i and j are connected.

We assign a Boolean variable t_{ijk} for the C_3 subgraph containing vertex i, j, k. There are in total $\binom{n}{3}$ such triangle variables. If t_{ijk} is True, then vertex i, j, k are mutually connected.

This is expressed as $t_{ijk} \leftrightarrow (e_{ij} \wedge e_{ik} \wedge e_{jk})$.

UNIVERSITY OF WATERLOO

Encoding the Squarefree Constraint in CNF

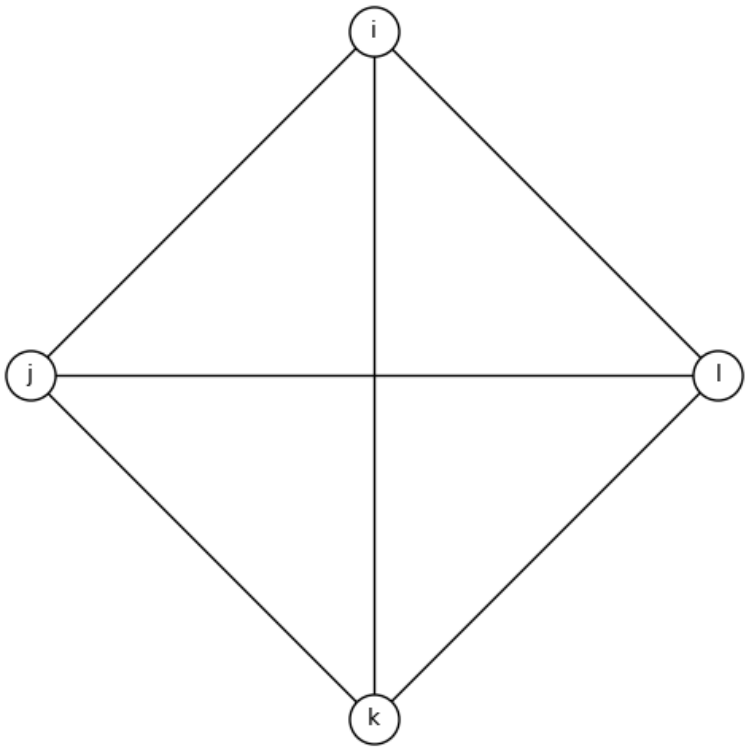

For every choice of four vertices i, j, k, l, we block all possible squares using the clause

$$\neg e_{ij} \vee \neg e_{jk} \vee \neg e_{kl} \vee \neg e_{li}$$
$$\neg e_{ij} \vee \neg e_{jl} \vee \neg e_{lk} \vee \neg e_{ki}$$
$$\neg e_{il} \vee \neg e_{lj} \vee \neg e_{jk} \vee \neg e_{ki}$$

connected by conjunction.

UNIVERSITY OF WATERLOO

Encoding the Triangle Constraint

For each vertex i, we require 2 other distinct vertices to form a triangle, and there are $\binom{n-1}{2}$ possible triangles containing i. At least one of those triangles must be present in the graph. We use the clause

$$\bigvee_{j,k \in V-i, j<k} t_{ijk}$$

Encoding the Minimum Degree Constraint

For each vertex i, define the vertex set $V - i$.
For each subset V' of $V - i$ with $|V'| = n - 3$, we construct the clause

$$\bigvee_{j \in V'} e_{ij}$$

This enforces at least three edge variables containing i are true.

UNIVERSITY OF WATERLOO

Encoding the Colourability Constraint

A graph is non-010-colourable if and only if for all $\{0, 1\}$-colourings of the graph a pair of colour-1 vertices is connected or a set of three colour-0 vertices are mutually connected.

For each $\{0, 1\}$-colouring, we have a set of colour-0 vertices V_0 and a set of colour-1 vertices V_1. Given a specific such colouring, the clause

$$\bigvee_{\substack{i,j \in V_1 \\ i<j}} e_{ij} \vee \bigvee_{\substack{i,j,k \in V_0 \\ i<j<k}} t_{ijk}$$

enforces that the colouring is not a 010-colouring of the graph since either a pair of colour-1 vertices is connected or a set of three colour-0 vertices is mutually connected.

UNIVERSITY OF WATERLOO

Embeddability Checking

If the SAT solver outputs solutions, these solutions are not guaranteed to be KS graphs until proven to be embeddable.

To decide embedability of graphs, the following methods were attempted by previous papers:

- Homotopy Continuation: performance quickly degraded with larger order
- Interval Arithmetic: effective but performance varies greatly depending on the graph
- Cubic Grids: highly efficient but the absence of a grid embedding does not allow one to draw any conclusion regarding embeddability.

UNIVERSITY OF WATERLOO

Introduction to SMT Solver

SMT solver generalizes the Boolean satisfiability problem (SAT) to more complex formulas involving real numbers, integers, and/or various data structures such as lists, arrays, bit vectors, and strings.

We will be using the SMT solver 'Z3' to determine the satisfiability of a system of nonlinear equations.

Embeddability Checking

The embeddability checking pipeline contains two parts.

- Finds all possible interpretations of the orthogonal relations between vectors. Each edge of the graph are being encoded in either a cross product or dot product. Choose the interpretation with the least number of unassigned vectors.
- SMT solver Z3 determines the satisfiability of an intepretation by converting it into a system of (nonlinear) cross product and dot product equations.

UNIVERSITY OF WATERLOO

SAT Symmetry Breaking

A SAT approach outperformed the previously used graph enumeration approach. However, a SAT solver generates many isomorphic copies of the same graph.

Thus, we combine SAT with isomorph-free exhaustive generation (also previously used to solve Lam's problem).[2]

UNIVERSITY OF WATERLOO

[2] C. Bright, K. Cheung, B. Stevens, I. Kotsireas, V. Ganesh. A SAT-based Resolution of Lam's Problem. *AAAI 2021*.

Isomorphisms

When generating combinatorial objects we only care about generating them *up to isomorphism*. Unfortunately, objects usually have many isomorphic representations.

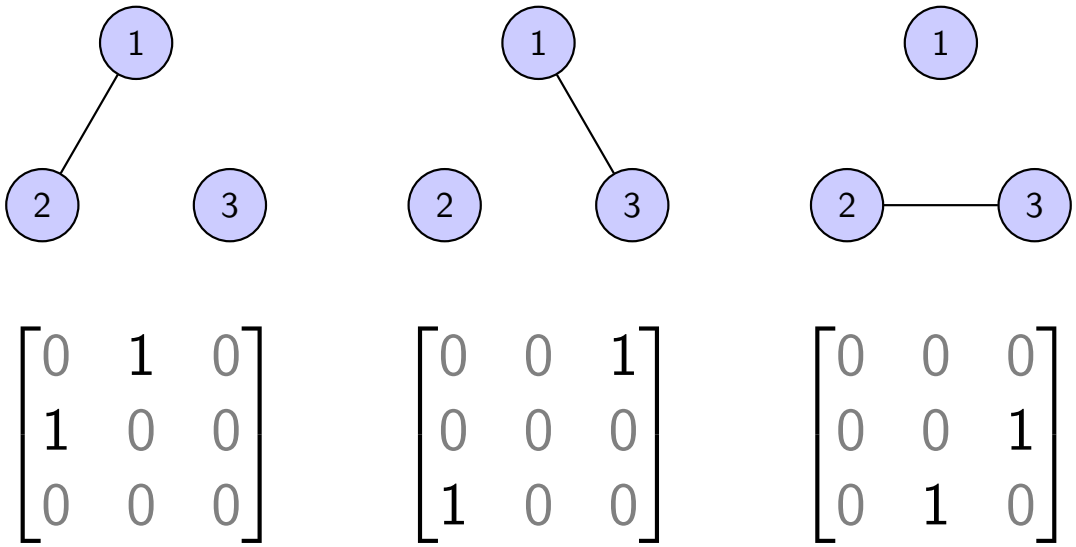

$$\begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix} \quad \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}$$

A graph with n vertices has up to $n!$ distinct isomorphic adjacency matrices. For efficiency, these should be detected and removed.

UNIVERSITY OF WATERLOO

SAT Symmetry Breaking

A typical SAT approach is to add "symmetry breaking" constraints that remove as many isomorphic solutions as possible.

For example, you can order the rows of an adjacency matrix of a graph lexicographically.[3] However, many distinct isomorphic representations still exist, like

$$\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix} \text{ and } \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}.$$

UNIVERSITY OF WATERLOO

[3] M. Codish, A. Miller, P. Prosser, P. Stuckey. Constraints for symmetry breaking in graph representation. *Constraints*, 2019.

Isomorph-free Orderly Generation

Only "canonical" intermediate objects are recorded. The notion of canonicity is defined so that:

1. Every isomorphism class has exactly one canonical representative.
2. If an object is canonical then its parent in the search tree is also canonical.

Developed independently by Faradžev and Read in 1978.

UNIVERSITY OF WATERLOO

Canonicity Example

An adjacency matrix of a graph is *canonical* if the vector of its entries below the diagonal is lexicographically smallest (among all matrices in the same isomorphism class).

For example,

$$\begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}, \text{ and } \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}$$

are isomorphic adjacency matrices but only the last is canonical.

UNIVERSITY OF WATERLOO

Orderly Generation of Graphs

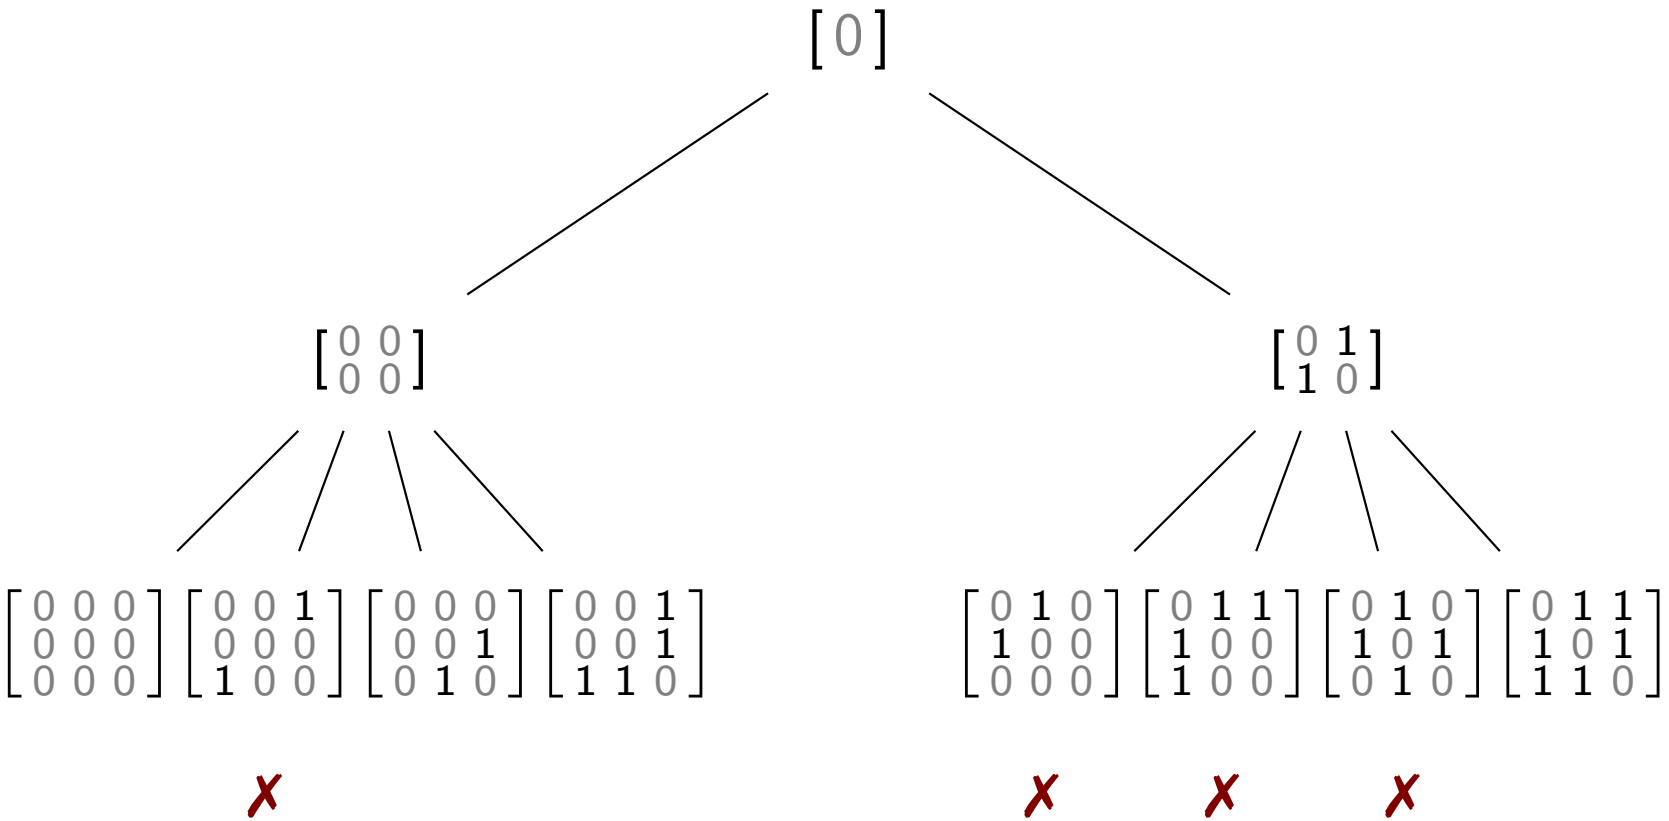

Canonical testing introduces overhead, but every negative test prunes a large part of the search space.

UNIVERSITY OF WATERLOO

Orderly Generation in Practice

Each canonical test is independent, making the method easy to parallelize.

Verifying a matrix is *noncanonical* is often fast—it requires finding a single permutation of the vertices giving a lex-smaller matrix.

SAT and Isomorph-free Generation

There have been surprisingly few attempts at combining isomorph-free generation and SAT solving.[4,5]

This is perhaps a result of the historical separation between the SAT and symbolic computation communities. We will now discuss applying orderly generation and SAT to the minimum Kochen–Specker problem.

[4] T. Junttila, M. Karppa, P. Kaski, J. Kohonen. An adaptive prefix-assignment technique for symmetry reduction. *Journal of Symbolic Computation*, 2020.

[5] J. Savela, E. Oikarinen, M. Järvisalo. Finding periodic apartments via Boolean satisfiability and orderly generation. *LPAR 2020*.

Orderly Generation in SAT

During the search the SAT solver will find partial solutions (complete definitions for the edges in some subgraphs)...

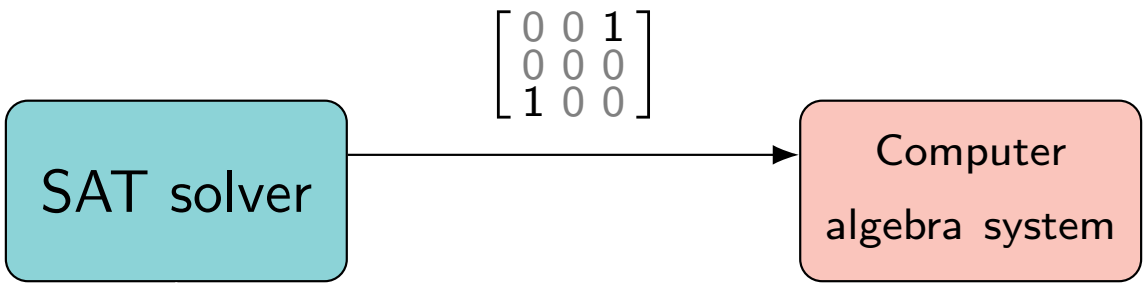

Orderly Generation in SAT

During the search the SAT solver will find partial solutions (complete definitions for the edges in some subgraphs)...

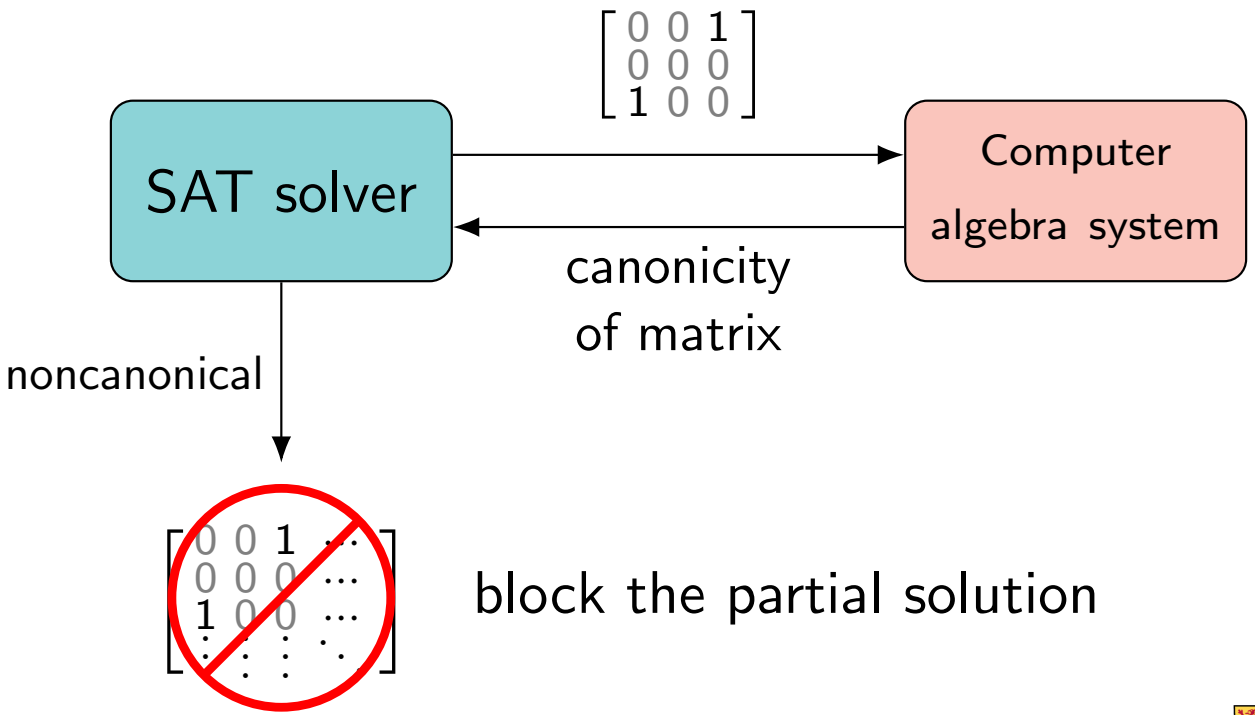

KS Search Results

The speedup factor that we found when using SAT-based orderly generation in the search for KS systems of a given order:

order	speedup factor
16	6.5
17	13.6
18	37.8
19	104.5

Implementation – Cube-and-Conquer

The cube-and-conquer satisfiability solving paradigm was developed to solve hard combinatorial problems.

1. A "cubing solver" splits a SAT instance into a large number of distinct sub-problems specified by cubes-formulas.
2. For each cube a "conquering solver" solves the original instance under the assumption that the cube is true.

For large orders, parallelization is applied by dividing the instance into smaller subproblems using the cube-and-conquer approach. During the splitting, cube-and-conquer finds the next variable that splits the search space the most evenly.

UNIVERSITY OF WATERLOO

Pipeline Overview

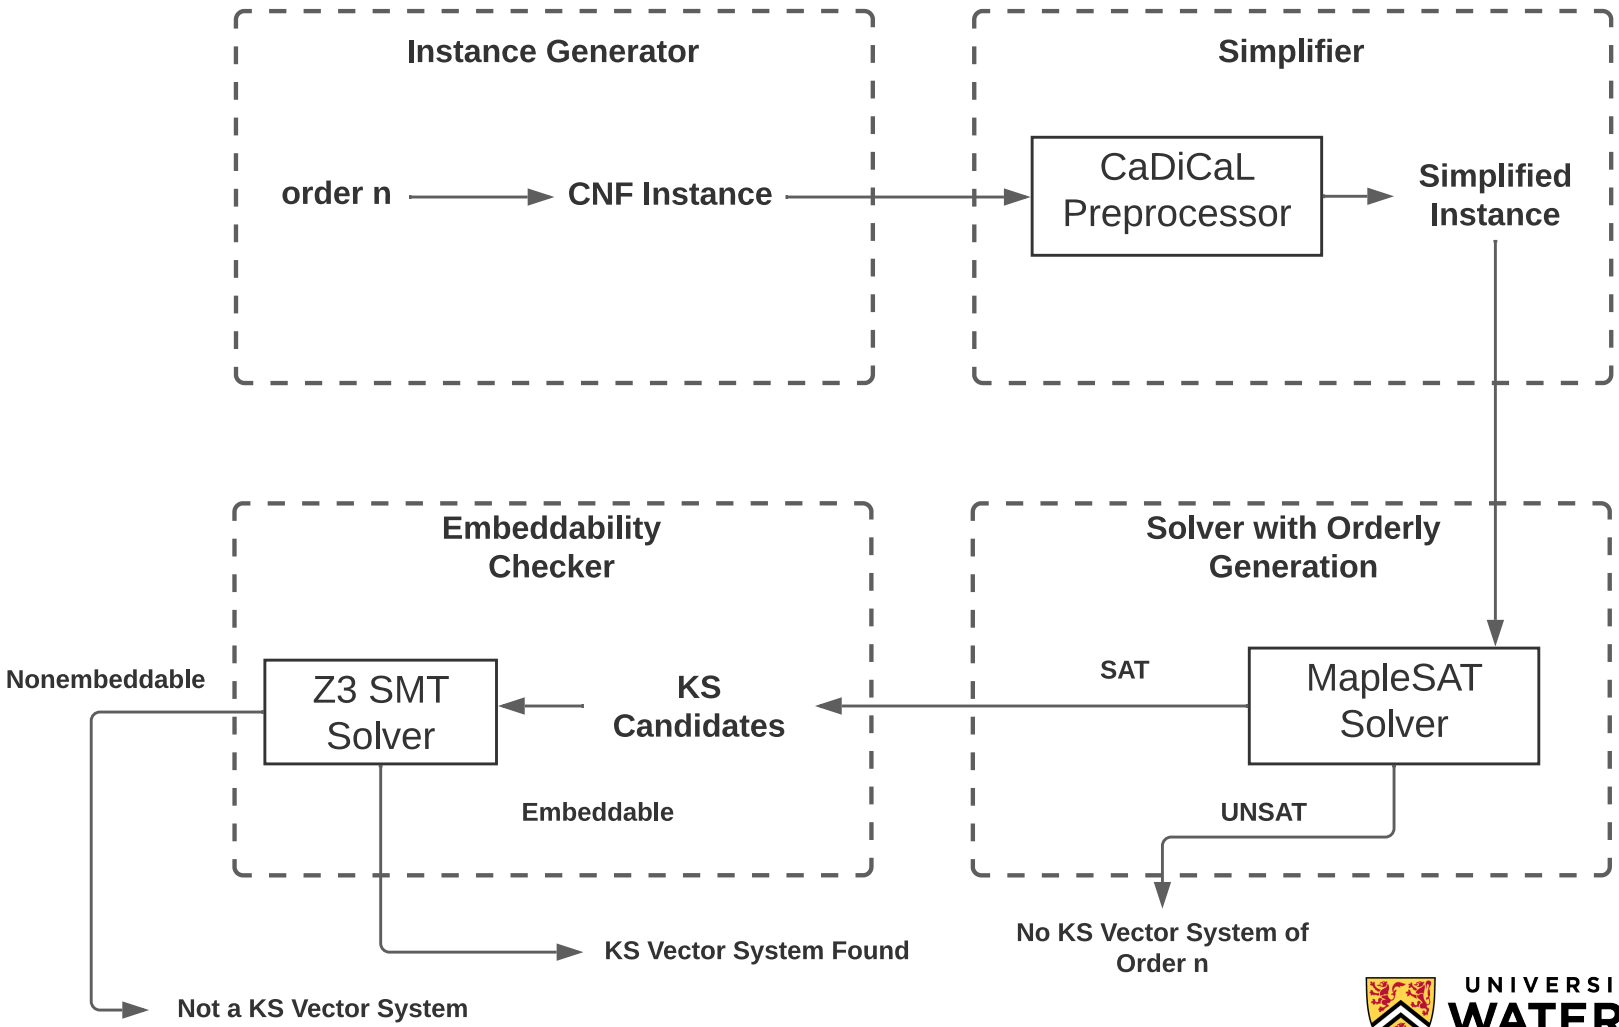

UNIVERSITY OF WATERLOO

KS Candidates and Results

Given the CNF file with the encoded constraints, we used the aforementioned techniques combined with the SAT + CAS approach to verify all previous results on KS systems up to order 21 and improve the best known lower bound with a significant speed-up factor.

All computations are done on Intel E5-2683 CPUs @ 2.1GHz administrated by Compute Canada, and measured in total CPU time.

Our computation on order 21 is over 1000 times faster than the previous computational by Uijlen and Westerbaan.

UNIVERSITY OF WATERLOO

KS Candidates and Results

Order	KS Candidates	Simplification	Cubing	Solving
17	1	0.02 hrs	N/A	0.02 hrs
18	0	0.02 hrs	N/A	0.13 hrs
19	8	0.31 hrs	N/A	2.46 hrs
20	147	0.54 hrs	N/A	39.71 hrs
21	2,497	1.50 hrs	38 hrs	1,019 hrs
22	88,282	2.54 hrs	953.7 hrs	46,079 hrs

Table: A summary of our results in the Kochen–Specker problem on orders $17 \leq n \leq 22$.

All 90,935 KS candidates of order less than 23 are unembeddable, so a KS system must contain at least 23 vectors.

Conclusion

We improve the lower bound of the KS vector system and the search efficiency by orders of magnitude. We found four additional order-20 KS candidates that were not present in Uijlen and Westerbaan's search, implying a bug in their pipeline.

We demonstrate the benefits of the SAT + CAS paradigm, showing that it is less error-prone as we uncover inconsistencies with the most recent results.

UNIVERSITY OF WATERLOO

A Promising Future

The SAT + CAS paradigm is able to produce a hybrid solver capable of exponential speedups over a pure SAT or computer algebra approach.

The approach is very general and can be applied to many combinatorial generation problems. I believe it has yet to be used to its full potential.

Thank You!

404briannotfound.tech

UNIVERSITY OF WATERLOO