
On the Quest for an Acyclic Graph

Mikoláš Janota¹, Radu Grigore², and Vasco Manquinho¹

¹ IST/INESC-ID, Lisbon, Portugal

² School of Computing, University of Kent, UK

Abstract. The paper aims at finding acyclic graphs under a given set of constraints. More specifically, given a propositional formula ϕ over edges of a fixed-size graph, the objective is to find a model of ϕ that corresponds to a graph that is acyclic. The paper proposes several encodings of the problem and compares them in an experimental evaluation using state-of-the-art SAT solvers.

1 Introduction

SAT solvers have become popular means of solving computationally hard problems. However, most modern SAT solvers require the problem to be specified in conjunctive normal form (CNF). This may pose difficulty in problems comprising of constraints that do not have a straightforward translation to propositional logic. This paper targets one such constraint, *graph acyclicity*.

A graph is naturally modeled by a set of Boolean variables x_{jk} expressing the existence of an edge from j to k . So, for instance, a constraint $x_{12} \vee x_{13}$ expresses that we're looking for a graph where the node 1 is connected to at least one of the nodes 2 and 3. But how do we ensure that we are looking for a graph that is acyclic?

Acyclicity is a core concept from graph theory and naturally arises in a number of applications. In planning it is used to ensure causality [42]. Similarly, it is needed in networks used for computation, e.g. Bayesian networks [49]. More recently, software verification of concurrent programs relies on relaxed memory models, which are defined in terms of acyclic relations [34, 51].

There is a large body of work of translating constraints into CNF. Prime examples are cardinality constraints [5, 36, 43] pseudo-Boolean constraints [6, 17], XOR-constraints [24] and general CSP [46]. A number of graph-related problems approached by SAT can be found in the literature. For instance, encoding *graph connectivity* property [47], graph coloring [50], calculating *Steiner tree* [14, 30] or *maximum clique* [28]. It is also worth mentioning that graphs and SAT play an important role in the research on *Gene regulatory networks* [15].

The paper has the following contributions. It reviews CNF encodings that are found in the literature, and introduces a number of new ones. All these encodings were implemented and extensively evaluated. To generate a well-defined testbed we introduce The Supervisor Problem, which asks to assign supervisors to a group of persons, where there is a lower bound for each person's supervisors and an upper bound on how many persons he or she can supervise. This problem

is a generalization of the well-known ordering principle (also known as the GT family) studied in proof complexity [1, 10, 32, 44].

2 Preliminaries

Our basic setting is the following. We represent a directed graph on $n \geq 1$ vertices by Boolean variables x_{ij} , associated with edges. A family of graphs is defined by a formula $\phi_n(\vec{x}, \vec{y})$ over the edge variables \vec{x} and some auxiliary variables \vec{y} . Each graph corresponds to a satisfying assignment of $\exists \vec{y} \phi_n(\vec{x}, \vec{y})$. Given such a formula ϕ_n , we wish to construct another formula ψ_n , such that the satisfying assignments of $\exists \vec{y} \exists \vec{z} (\phi_n(\vec{x}, \vec{y}) \wedge \psi_n(\vec{x}, \vec{z}))$ correspond to the graphs from the given family that, in addition, are acyclic. We refer to the formula ψ_n as the *acyclicity checker* or, for short, the *checker*. In the rest of the section, we make this more precise, and we describe the particular formulas ϕ_n that we use in experiments. The acyclicity checkers ψ_n are discussed later (§3). Standard terminology from propositional logic and graph theory is assumed.

2.1 Propositional Logic and SAT Solving

Propositional formulas are constructed from propositional variables by Boolean connectives (\vee , \wedge , \Rightarrow) and negation (\neg). A propositional formula is in *conjunctive normal form (CNF)* if it is a conjunction of clauses, where a *clause* is a disjunction of literals, where a *literal* is a propositional variable or its negation.

Whenever convenient, we treat CNF as a set of clauses and a clause as a set of literals. Observe that the empty set of clauses corresponds semantically to true and the empty clause corresponds semantically to false. The empty clause is denoted as \perp .

Any propositional formula can be converted to an equisatisfiable CNF formula in linear time [41, 48].

The problem of satisfiability (SAT) is to decide for a given propositional formula whether there is a satisfying assignments for it or not.

The paper requires minimal understanding of SAT solving. We will assume that a SAT solver accepts a formula in CNF and gives the response “YES” if the formula is satisfiable and “NO” otherwise. Most modern SAT solvers also provide an actual satisfying assignment if the answer is “YES”.

When discussing the performance of SAT solvers, sometimes we refer to the underlying proof system, which is propositional resolution. Given two clauses $x \vee C$ and $\neg x \vee D$ their resolution is $C \vee D$. The two clauses are called antecedents and the resulting clause the resolvent. Resolution proof of a clause C from a formula ϕ is a sequence of clauses C_1, \dots, C_k such that $C_k = C$ and each clause C_i is either in ϕ or it is a resolvent of some C_j, C_k with $1 \leq j < k < i$. A resolution refutation is a resolution proof of the empty clause. A CNF formula is unsatisfiable if and only if there exists a resolution refutation for it. For most modern SAT solvers it holds that for any run of the SAT solver that responds “NO”, there exists a resolution refutation of the input formula linear to the run

of the solver. In fact, it even holds that for any resolution refutation there exists a polynomial run of the solver [39].

2.2 Graphs

Throughout the paper we will be discussing directed graphs. A *directed graph* $G = (V, E)$ consists of a set V of *vertices* and a set E of *edges*. In this paper, we work with finite graphs, and we fix $V = [n] = \{1, \dots, n\}$. We denote vertices by i, j, k, \dots and we abbreviate the (directed) edge (i, j) by ij . For an edge ij , we say that it has *source* i and *target* j ; conversely, we say that ij is an *outgoing* edge of i , and an *incoming* edge of j . The *in-degree* of a vertex is its number of incoming edges, and the *out-degree* of a vertex is its number of outgoing edges. A *sink* is a vertex with out-degree 0.

A *graph property* is a family \mathcal{G} of graphs. In this paper, we consider only properties \mathcal{G}_n that contain graphs with a fixed number n of vertices, and are defined by a formula ϕ_n . More precisely, ϕ_n has one variable x_{ij} for each pair $(i, j) \in V \times V$, and may also use some auxiliary variables \vec{y} . Then, a graph with edges E is in \mathcal{G}_n if and only if it corresponds to a satisfying assignment of $\exists \vec{y} \phi_n(\vec{x}, \vec{y})$. We say that a property is *monotone* if adding edges to a graph preserves the property. This corresponds to the Boolean function $\exists \vec{y} \phi_n(\vec{x}, \vec{y})$ being monotone. We remark that the property of having a cycle is monotone, and the intersection of two monotone properties is monotone.

2.3 The No-Sink Family

The graphs with no sinks are described by the following formula:

$$\phi_n(\vec{x}) := \bigwedge_{i \in [n]} \left(\bigvee_{j \in [n]} x_{ij} \right) \quad (\text{no-sink})$$

Observe that the no-sink property is monotone. Also, the *ordering principle* says that graphs with no sink have a cycle. It follows that, no matter what acyclicity check ψ_n we use, the formula $\exists \vec{y} \exists \vec{z} (\phi_n(\vec{x}, \vec{y}) \wedge \psi_n(\vec{x}, \vec{z}))$ will be unsatisfiable. Since we know the answer, the no-sink family is a good test case for the correctness of our implementation. In addition, as we shall see, these unsatisfiable formulas are not easy for state-of-the art SAT solvers.

2.4 The Supervisor Problem

We define a parameterized family of graphs, which will typically contain both cyclic and acyclic graphs. The idea is as follows. Consider a group of n people where each person must have a certain number of supervisors and at the same time there is a limit on how many people a person can supervise. The supervisor problem is formalized as follows.

Input: A positive integer n and a sequence of pairs of nonnegative integers $(u_1, l_1), \dots, (u_n, l_n)$.

Output: “YES” if and only if there exists an acyclic graph with the vertices $V = [n]$ such that vertex i has at most u_i incoming edges and at least l_i outgoing edges.

An instance n, \vec{u}, \vec{l} of the Supervisor problem reduces to the satisfiability of the following formula:

$$\begin{aligned} \phi_{n, \vec{u}, \vec{l}}(\vec{x}) := & \bigwedge_{i \in [n]} \text{atleast}(l_i; x_{i1}, \dots, x_{in}) \\ & \wedge \bigwedge_{j \in [n]} \text{atmost}(u_j; x_{1j}, \dots, x_{nj}) \end{aligned} \quad (\text{supervisor})$$

The cardinality constraint $\text{atleast}(l; \vec{x})$ evaluates to 1 when at least l of the Boolean variables \vec{x} evaluate to 1. There are several ways to encode such cardinality constraints in CNF [5, 17, 36, 43]. The cardinality constraint $\text{atmost}(u; \vec{x})$ is analogous.

The Supervisor problem has three interesting special cases: (a) the no-sink family; (b) the pigeonhole principle; and (c) the DAG realizability problem. We obtain a description of the no-sink family by requiring the bounds $(n, 1)$ for each of the n vertices. We obtain a description of the pigeonhole principle by requiring the bounds $(0, 1)$ for k_1 vertices and the bounds $(1, 0)$ for k_2 vertices, with $k_1 + k_2 = n$ and $k_1 > k_2$. Finally, any instance of the DAG realizability problem is also an instance of the Supervisor problem, which satisfies the additional constraint $\sum_{i=1}^n u_i = \sum_{i=1}^n l_i$.

Since the pigeonhole principle is a special case, it follows that (supervisor) sometimes requires large resolution refutations [25]. Since the DAG realizability problem is a special case, it follows that it is NP-complete to find a model of (supervisor) that corresponds to an acyclic graph [26]. Finally, we observe that (supervisor) is not monotone.

3 Encodings

We discuss several ways of constructing the acyclicity checker ψ_n , based on reachability (§3.1), based on labeling vertices with numbers (§3.2 and §3.3), and based on oblivious algorithms (§3.4 and §3.5).

3.1 Transitive Closure

A graph $G = (V, E)$ does not have cycles if and only if E^+ is irreflexive. Here, E^+ denotes the irreflexive transitive closure of E . Further, G does not have cycles if and only if there exists an irreflexive transitive relation R that contains E ; that is, if there exists $R \subseteq V \times V$ such that (a) $ii \notin R$ for all $i \in V$, (b) $ij \in R$ and $jk \in R$ only if $ik \in R$ for all $i, j, k \in V$, and (c) $ij \in R$ if $ij \in E$ for all $i, j \in V$.

We represent the relation R by variables y_{ij} , and we define ψ_n by clauses that directly correspond to the conditions (a), (b), and (c):

$$\psi_n(\vec{x}, \vec{y}) := \bigwedge_{i \in [n]} \neg y_{ii} \wedge \bigwedge_{i, j, k \in [n]} (y_{ij} \wedge y_{jk} \rightarrow y_{ik}) \wedge \bigwedge_{i, j \in [n]} (x_{ij} \rightarrow y_{ij}) \quad (\text{tc1})$$

A possible alternative is the following:

$$\psi_n(\vec{x}, \vec{y}) := \bigwedge_{i \in [n]} \neg y_{ii} \wedge \bigwedge_{i, j, k \in [n]} (y_{ij} \wedge x_{jk} \rightarrow y_{ik}) \wedge \bigwedge_{i, j \in [n]} (x_{ij} \rightarrow y_{ij}) \quad (\text{tc2})$$

In relational language, (tc1) corresponds to the conditions (a) $R \cap \text{id} = \emptyset$, (b) $RR \subseteq R$, and (c) $E \subseteq R$. Here, id is defined to be $\{ii \mid i \in V\}$. In (tc2), we replace condition (b) by $RE \subseteq R$. Juxtaposing relations corresponds to normal relational composition:

$$RE := \{(i, j) \mid (i, k) \in R \text{ and } (k, j) \in E, \text{ for some } k\}$$

If the family of graphs satisfying ϕ_n is monotone, in the sense that adding edges preserves ϕ_n , then we do not need to use auxiliary variables y_{ij} , as we can simply reuse the variables x_{ij} :

$$\psi_n(\vec{x}) := \bigwedge_{i \in [n]} \neg x_{ii} \wedge \bigwedge_{i, j, k \in [n]} (x_{ij} \wedge x_{jk} \rightarrow x_{ik}) \quad (\text{tc3})$$

The conjunction of (tc3) and (no-sink) has $\Omega(2^n)$ regular tree resolution proofs but $O(n^3)$ dag resolution proofs; see [10, 32, 44] and [31, §7.2.2.2, Theorem R]. Such conjunctions express the *ordering principle*.

Observe that the formulas in each of (tc1), (tc2), and (tc3) have size $\Theta(n^3)$.

3.2 Binary Labeling

A graph is acyclic when it can be embedded in a linear order; that is, when there exists a labeling $l : V \rightarrow \mathbb{Z}$ of the vertices such that $ij \in E$ implies $l(i) < l(j)$ for all $i, j \in V$. Recall that $V = [n]$. Clearly, it is sufficient to use $b := \lceil \log_2 n \rceil$ bits for labels, so $l : V \rightarrow \{0, 1\}^b$. Accordingly, we use n groups $\vec{y}_1, \dots, \vec{y}_n$ of auxiliary variables, each containing b Boolean variables. The acyclicity checker is then [42]:

$$\psi_n(\vec{x}, \vec{y}_1, \dots, \vec{y}_n) := \bigwedge_{i, j \in [n]} (x_{ij} \rightarrow \text{less}(\vec{y}_i, \vec{y}_j)) \quad (\text{bin})$$

The implementation of `less` depends on how we represent numbers. Because we use binary numbers, we set `less` := `lexb`, and define the latter to test lexicographic ordering of the bits:

$$\text{lex}_0() := 0 \quad (1)$$

$$\text{lex}_b(\vec{y}y, \vec{z}z) := (\neg y \wedge z) \vee ((\neg y \vee z) \wedge \text{lex}_{b-1}(\vec{y}, \vec{z})) \quad (2)$$

Converting this circuit to CNF will require some more auxiliary variables, which we do not denote explicitly. The conversion is standard, and the size of the resulting CNF is $\Theta(b)$.

Observe that the formula in (bin) has size $\Theta(n^2 \log n)$.

3.3 Unary Labeling

We can also use the labeling idea with unary numbers. To represent a number, instead of $\lceil \log_2 n \rceil$ bits we use $n - 1$ bits, and force them to follow the pattern 0^*1^* . We could still compare these numbers by their lexicographic order using `lex`, but we can also write directly a CNF formula, thus avoiding the (small) overhead associated with converting a circuit into CNF.

The acyclicity checker is defined by

$$\psi_n(\vec{x}, \vec{y}_1, \dots, \vec{y}_n) := \bigwedge_{i,j \in [n]} (x_{ij} \rightarrow \text{less}(\vec{y}_i, \vec{y}_j)) \wedge \bigwedge_{i \in [n]} \text{unary}(\vec{y}_i) \quad (\text{unr})$$

where each \vec{y}_i uses $n - 1$ bits.

We could use `less` := `lexn-1`, but we prefer `less`(\vec{y}, \vec{z}) := $\exists \vec{u}$ `lessunr`($\vec{y}, \vec{z}, \vec{u}$), where

$$\text{lessunr}(\vec{y}, \vec{z}, \vec{u}) := \bigwedge_{i=1}^{n-1} ((\neg y_i \vee \neg u_i) \wedge (z_i \vee \neg u_i)) \wedge \bigvee_{i=1}^{n-1} u_i \quad (3)$$

The auxiliary \vec{u} variables will identify a position i where $y_i = 0$ and $z_i = 1$. This is sufficient because we enforce the pattern 0^*1^* .

$$\text{unary}(\vec{y}) := \bigwedge_{i=2}^{n-1} (y_{i-1} \rightarrow y_i) \quad (4)$$

Observe that the formula in (unr) has size $\Theta(n^3)$.

3.4 Warshall

The Warshall algorithm computes transitive closure, in cubic time.

WARSHALL

```

1  for  $k \in [n]$ 
2      for  $i \in [n]$ 
3          for  $j \in [n]$ 
4               $a_{ij} := \text{Or}(a_{ij}, \text{And}(a_{ik}, a_{kj}))$ 
```

Initially, a_{ij} says if the graph contains an edge ij ; and, at the end a_{ij} says if the graphs contains a path from i to j . In particular, at the end, a_{ii} says if there is a cycle containing vertex i .

In a usual implementation, the procedures `Or`, `And` are simply performing the corresponding Boolean operations. Instead, we could replace each of them

with a function that constructs the corresponding circuit. Because the Warshall algorithm is simple enough, we can also build the acyclicity checker directly. We use $n^2(n+1)$ auxiliary variables y_{ijk} , for $i, j \in [n]$ and $k \in \{0, \dots, n\}$. The variable y_{ijk} will be 1 only if there exists a path from i to j that uses intermediate vertices only from the set $[k]$.

$$\begin{aligned} \psi_n(\vec{x}, \vec{y}) := & \bigwedge_{i \in [n]} \neg y_{iin} \wedge \bigwedge_{i, j \in [n]} (x_{ij} \rightarrow y_{ij0}) \wedge \bigwedge_{i, j, k \in [n]} (y_{ij(k-1)} \rightarrow y_{ijk}) \\ & \wedge \bigwedge_{i, j, k \in [n]} (y_{ik(k-1)} \wedge y_{kj(k-1)} \rightarrow y_{ijk}) \end{aligned} \tag{fw}$$

We remark that the formula above is close but *not* the same as what we would obtain by replacing the basic operations **Or**, **And** with circuit builders. One difference is that we use implications rather than equivalences: It is possible that $y_{ijk} = 1$ even if there is no corresponding path. Another difference is that the last conjunct built by Warshall would be $(y_{ik(k-[k \geq j])} \wedge y_{kj(k-[k \geq i])}) \leftrightarrow y_{ijk}$, where $[p]$ evaluates to 1 when p is true, and 0 when p is false. Either form is correct.

The formula in (fw) is similar to the one in (tc1). It has the same size $\Theta(n^3)$, but uses more auxiliary variables.

3.5 Matrix Multiplication

Let us take stock of the encodings we have so far. Four encodings have size $\Theta(n^3)$, namely (tc1), (tc2), (unr), and (fw); another encoding has size $\Theta(n^2 \log n)$, namely (bin). But, size is not the only criterion on which to judge a satisfiability formula. We note that (tc1), (tc2), and (fw) have the following desirable property: *if the edge variables \vec{x} are fixed, then satisfiability is decided by unit propagation*. The smaller encoding (bin) does not have this property. Is it possible to have a formula of sub-cubic size that is solved by unit propagation once the edge variables \vec{x} are fixed? We answer this question affirmatively by showing an encoding based on matrix multiplication. We first describe the algorithm. The formula ψ_n will be constructed by replacing the basic operations with circuit builders, and then converting the circuit to CNF.

We will consider $n \times n$ matrices of two kinds, over integers and over Booleans. Integers form a commutative ring, so the corresponding matrix operations are unrestricted. Booleans do not form a ring: We take addition to be logical-or and we take multiplication to be logical-and, but there is no subtraction. Since we use both types of operations and sometimes reinterpret the same matrix as being over integers or over Booleans, we distinguish Boolean matrices with a hat. When we reinterpret an integer matrix X as a Boolean matrix \hat{X} , all non-zero integers collapse to 1.

Let X be the adjacency matrix of a graph, with entries in $\{0, 1\}$. Entry (i, j) of X^l is the number of paths of length l from i to j , and entry (i, j) of \hat{X}^l tells us if there exists a path of length l from i to j . Thus, to decide if the graph

has a cycle it suffices to check the diagonal of the Boolean matrix $\sum_{l=1}^n \hat{X}^l$. We compute this matrix by repeated squaring:

$$\begin{aligned} \hat{A}_0 &:= \hat{X} & \hat{A}_{k+1} &:= \hat{A}_k^2 & (5) \\ \hat{B}_0 &:= \hat{X} & \hat{B}_{k+1} &:= \hat{B}_k + \hat{A}\hat{B}_k & (6) \end{aligned}$$

It is easy to check that $\hat{A}_k = \hat{X}^{2^k}$ and $\hat{B}_k = \sum_{l=1}^{2^k} \hat{X}^l$, and also that \hat{B}_k equals $\sum_{l=1}^n \hat{X}^l$ whenever $k \geq \log_2 n$. Thus we can decide if a directed graph is acyclic in the time needed to perform $\Theta(\log n)$ multiplications of $n \times n$ Boolean matrices. One could perform a naive matrix multiplication in $\Theta(n^3)$ time, giving an acyclicity checker ψ_n of size $\Theta(n^3 \log n)$.

Alternatively, to multiply Boolean matrices \hat{A} and \hat{B} , we multiply them as integer matrices obtaining $C = AB$; the product of \hat{A} and \hat{B} is then \hat{C} . (The idea of using intermediate integer matrices is old [18, 19, 37].) To perform the multiplication of integer matrices, we use Strassen's algorithm [45], which works in $O(n^{2.81})$ time. To achieve this improved runtime, Strassen's algorithm makes essential use of subtraction, which is not available over Booleans. However, the runtime $O(n^{2.81})$ counts operations over integers, while we must count operations over bits since we ultimately want to build a circuit. Since the matrices A and B , which we multiply, have entries from $\{0, 1\}$, the product matrix C has entries from $\{0, \dots, n\}$. Some intermediate values in Strassen's algorithm are outside this range. Nevertheless, one can check that $\Theta(\log n)$ bits are sufficient to represent all intermediate values. So, the size of the circuit is $O(n^{2.81} \log n \cdot f(\log n))$, where $f(w)$ is the size of a circuit necessary to perform an arithmetic operation on integers with w bits. In our implementation, we use a simple quadratic multiplication algorithm, leading to a circuit of size $O(n^{2.81} \log^3 n)$.

In theory, one can use the same approach to achieve $O(n^\omega \log^2 n \log \log n \cdot 8^{\log^* \log n})$ by using better algorithms for the multiplication of matrices [20] and for the multiplication of integers [27]. But such algorithms are known to not be practical. Our current implementation based on Strassen is not competitive either (§4). Improving it seems to require new ideas.

Discussion. The approach we took for both Warshall and for matrix multiplication was to replace the basic operations in the algorithm with circuit builders. This approach is fairly general, but it does not apply to all algorithms, just to the oblivious ones. An algorithm is said to be *oblivious* when the memory locations it accesses do not depend on the input data. If an algorithm runs in $t(n)$ time in the Turing model of computation, then one can construct an oblivious algorithm that runs in $O(t(n) \log t(n))$ time, and hence a Boolean circuit of that size [40]. In contrast, a similar construction does *not* exist in the deterministic RAM model of computation [23].

The standard (and optimal) algorithm for finding cycles in directed graphs is DFS (depth-first search). However, we are not aware of an implementation of DFS that works in sub-cubic time in the (oblivious, deterministic, multi-tape) Turing machine model of computation.

| solver/checker | tc1 | unr | bin | fw | tc2 |
|-----------------------|------------|------------|------------|-----------|------------|
| lingeling | 49 | 48 | 10 | 38 | 11 |
| glucose | 32 | 37 | 14 | 15 | 11 |
| minisat | 25 | 49 | 12 | 12 | 11 |
| minisat-prepro | 26 | 49 | 11 | 19 | 11 |

(a) GT Family with 49 instances

| solver/checker | tc1 | unr | bin | fw | tc2 |
|-----------------------|------------|------------|------------|------------|------------|
| lingeling | 436 | 429 | 426 | 435 | 435 |
| glucose | 437 | 434 | 427 | 437 | 437 |
| minisat | 435 | 425 | 424 | 435 | 435 |
| minisat-prepro | 435 | 426 | 422 | 435 | 435 |

(b) Supervisor Family with 441 instances

Fig. 1. Number of solved instances for sizes 2-50.

4 Experimental Evaluation

Overall setup. Three popular state-of-the-art SAT solvers were used for the evaluation: Lingeling Version 276 [8], glucose 4.0 [3], and minisat 2.2 (from github) [16]. Further, minisat was a run in two modes: the default mode, and with additional preprocessing techniques enabled (asymmetric breaking and redundancy checking). The encodings we consider are Transitive closure (tc1), Unary labeling (unr), Binary labeling (bin), Warshall (fw), and an alternate version of Transitive closure (tc2). The techniques based on matrix multiplication are not included, because they quickly lead to formulas in the realm of hundred of megabytes, and more (Figure 3).

The experimental results were obtained on an Intel Xeon 5160 3GHz, with 4GB of memory, time limit 500seconds, and memory limit 2GB.

Generation of Graph Families. The generation of the no-sink graph family is trivial: we simply instantiate (no-sink) for different values of n . We also generate random instances of the Supervisor problem. Each upper bound u_i on the in-degree of vertex i is chosen uniformly at random from the set $\{0, \dots, n-1\}$. Each lower bound l_i is 0 with probability p , which is a parameter, and otherwise is chosen uniformly at random from the set $\{1, \dots, n-1\}$. If the generated sequence $(l_1, u_1), \dots, (l_n, u_n)$ is not satisfied by any directed graph, we throw it away and try again. This filtering is easy, because it is possible to check in linear time whether a sequence is satisfiable by some directed graph, using the Fulkerson–Chen–Anstee theorem [2]; but, as we mentioned, checking whether a sequence is satisfiable by an *acyclic* directed graph is NP-complete [26].

For each $n \in \{2, 3, \dots, 50\}$ and for each $p \in \{10\%, 20\%, \dots, 90\%\}$, we generate a sequence $(l_1, u_1), \dots, (l_n, u_n)$ using the method just described. For each such sequence, we generate the formula (supervisor). The *atmost* and *atleast* constraints were encoded into CNF using the well-known *Sequential Counter* encoding [43].

Presentation of the results. Figure 1 presents the number of instances solved. Figure 1(a) shows the results for graphs with no sinks, which lead to the GT family of formulas when coupled with the acyclicity checker (tc1). Figure 1(b) shows the results for the general supervisor problem. Figure 2 shows a cactus for the (no-sink) family for all the considered combinations of solvers and encodings.




Fig. 2. Cactus plot for the (no-sink) family.

Recall that if a cactus plot contains a point (i, s) it means there are i instances such that each was solved within s seconds. (Therefore, the graph is necessarily increasing.) A more detailed presentation of the results can be found on the authors' homepage.³

The first thing to observe is that the supervisor family turns out to be easy for all the approaches. Indeed, all approaches solved over 90% benchmarks. Even though the differences are rather small, the tables suggests that the Transitive closure encoding ((tc1)) is good for all solvers. In contrast, Binary labeling (bin) appears to be the worst.

In the case of (no-sink) graphs, the results are widely different for each encoding but also for each solver. Most notably, lingeling performs extremely well on (tc1). Both versions of minisat perform well on Unary labeling (unr). Glucose also performs the best with (unr), but once the graph size gets over 30 nodes, the performance quickly deteriorates. The worst encoding appears to be Binary labeling (bin) and the alternative Transitive closure (tc2), with which all solvers could only handle less than half of all instances. Warshall has a peculiar behavior, as it performs poorly for all the solvers except lingeling.

³ <http://sat.inesc-id.pt/~mikolas/sat-cycles/>

Discussion. The results are surprisingly uneven across different solvers. In a particular, lingeling performs extremely well with (tc1), while the other solvers do not perform very well with the same encoding. It is possible that lingeling benefits from some simplification technique (pre- or in-processing) that is missing from the other solvers, but this remains to be understood. Binary labeling (bin) is overall the most compact but with the poorest performance. This is not entirely surprising. First, binary labeling does not offer many opportunities for applying unit propagation. And second, binary labeling makes it virtually impossible to learn anything about any particular variable.

In contrast, Unary labeling (unr) fares quite well over the considered solvers. Further, we observed that minisat with preprocessing solves the (no-sink) family with no conflicts, i.e., merely with preprocessing. This also holds for lingeling but only up to a certain size of the graph. It is not difficult to see how these formulas can be completely solved by preprocessing, in particular by failed literal detection and asymmetric breaking.

In unary labeling (§3.3), the ordering of labels is ensured by the following constraint. For every vertex i labeled by a number $l(i) \in \{0, \dots, n-1\}$ there must be a neighbor j labeled by a number $l(j) \in \{0, \dots, n-1\}$ such that $l(i) < k \leq l(j)$ for some $k \in \{0, \dots, n-1\}$. To show that such a labeling does not exist for no-sink graphs, we reason as follows. If $k = n-1$, it immediately follows that vertex j has no neighbor labeled with a greater number since it has reached the maximum. Performing this reasoning for every vertex yields that the labels must be in the interval $\{0, \dots, n-1\}$. Repeating this reasoning n times leads to a contradiction.

We argue that to perform this reasoning, it is sufficient to use unit propagation on unary labeling. In definition (3), setting $u_k := 1$ imposes exactly the condition as listed above, i.e., the label of the source node is at most $k-1$ and the target node at least k . It is easy to verify that setting $u_{n-1} := 1$ yields a conflict by unit propagation.

This idea is realized in the preprocessing techniques failed literal detection [12, 35] and asymmetric breaking [38]. More generally, if setting a literal l to true yields a conflict by unit propagation it is called a failed literal and the literal is set to false. In this manner, these two techniques acting in tandem can gradually remove the u_i literal corresponding to the highest number, and eventually derive an empty clause.

Conclusion. The experimental results support the following conclusions. The combination of choice for checking acyclicity is lingeling with Transitive closure (tc1). (Do *not* use the seemingly very similar alternative Transitive closure (tc2), which leads to poor performance.) Minisat is the second solver of choice but *not* with the (tc1) encoding—instead, use (unr). SAT solvers appear to be good at finding an acyclic graph if it exists; proving that an acyclic graph does not exist can easily become very difficult.

Formula Size. Figure 3 presents the size of the acyclicity checkers (§3). The cubic methods (tc1, tc2, unr, fw) are all within a factor of 3 of each-other. The super-




Fig. 3. The size of the acyclicity checker formula ψ_n as a function of the number of vertices n . Five lines correspond to equations (tc1), (tc2), (bin), (unr), and (fw). The remaining two correspond to Boolean matrix multiplication methods, using the definition (mm), and using Strassen’s algorithm (ss). The size of a CNF formula consisting of clauses C_1, \dots, C_m is defined to be $m + \sum_{i=1}^m |C_i|$.

cubic one (mm, which is based on the definition of Boolean matrix multiplication) is already worse than the cubic ones at $n = 3$, and it only gets worse for larger n . There are two sub-cubic methods: ss with size $O(n^{2.81} \log^3 n)$, and bin with size $O(n^2 \log n)$. The method based on Strassen (ss) generates large formulas even for $n = 10$ and, looking at Figure 3, there is no overtaking in sight. Indeed, even if the constants hidden in the O -notation were just as good for ss as they are for tc1, the crossing point would be at $n \sim 10^{32}$. The $\log^3 n$ factor is important, and it seems that without new ideas this method will never be practical. On the other hand, the method based on binary labeling (bin), generates the smallest acyclicity checkers for $n \geq 100$.

5 Summary and Future Work

The paper presents a number of encodings for the graph acyclicity constraint into conjunctive normal form. This is useful when calculating acyclic graphs using SAT solvers but also other logic-based solvers, such as quantified Boolean formula (QBF) and Satisfiability modulo theories (SMT) solvers. The experimental evaluation suggests that the performance is highly sensitive both to the encoding but also to the SAT solver used. This suggests that for practical purposes it is useful to maintain a portfolio of encodings and solvers.

There are a number of directions for further research. All the encodings that are asymptotically better than cubic perform poorly in practice. Hence, it is a challenge for future research to identify encodings with asymptotic complexity better than cubic but that also perform well.

Ultimately, all the considered encodings are not feasible for large graphs. Indeed, any cubic encoding generates formulas that are too large once the number of nodes gets over 1000. A possible solution is to consider a lazy approach where only some portion of the encoding is considered. However, this will probably require an iterative sequence of calls to the SAT solver. Another approach, tried recently [9,13], is to integrate specialized algorithms for acyclicity inside the SAT solver. This is also closely related to *loop formulas*, which are used in Answer Set Programming (ASP) to avoid cyclic reasoning cf. [21,29,33].

Some questions about properties of the encodings with respect to *unit propagation* remain open. In particular, is it the case that any partial assignment that contains a cycle produces a conflict by unit propagation in a given encoding? In theory this is possible due to monotonicity [7] (adding edges preserves cycles). We believe this holds for the encodings Transitive closure, Warshall, and matrix multiplication, but does not hold for Unary and Binary labeling. Hence, this also begs the question whether the labeling methods can be made more deterministic. Other properties related to unit propagation are also worth exploring, e.g., *generalized arc consistency* and *propagation completeness*, cf. [4,7,11,22,24].

The experimental results lead to several interesting questions. Why is *lingeling* so efficient on the transitive closure encoding (tc1)? Why does the alternative encoding for transitive closure (tc2) perform so poorly across all solvers? Do formulas (tc2) lack short resolution refutations?

Last but not least, this research opens avenues for exploring other graph-related properties, e.g., connectiveness and reachability in general.

Acknowledgments

This work was supported by national funds through Fundação para a Ciência e a Tecnologia (FCT) with reference UID/CEC/50021/2013, and by the *PrideMM Web Interface* grant from VeTSS.

References

1. Alekhovich M, Johannsen J, et al. *An exponential separation between regular and general resolution*. Theory of Computing, 3(5):81–102, 2007.
2. Anstee RP. *Triangular (0, 1)-matrices with prescribed row and column sums*. Discrete Mathematics, 40(1):1–10, 1982.
3. Audemard G, Simon L. *Predicting learnt clauses quality in modern SAT solvers*. In *International Joint Conference on Artificial Intelligence (IJCAI)*. 2009.
4. Babka M, Balyo T, et al. *Complexity issues related to propagation completeness*. Artif Intell, 203:19–34, 2013.
5. Bailleux O, Boufkhad Y. *Efficient CNF encoding of boolean cardinality constraints*. In *Principles and Practice of Constraint Programming (CP)*. Springer, 2003.
6. Bailleux O, Boufkhad Y, et al. *A translation of pseudo-boolean constraints to SAT*. Journal on Satisfiability, Boolean Modeling and Computation, 2006.
7. Bessiere C, Katsirelos G, et al. *Circuit complexity and decompositions of global constraints*. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 412–418. 2009.

8. Biere A. *Lingeling, Plingeling and Treengeling entering the SAT competition*. Proceedings of SAT Competition, 2013.
9. Bomanson J, Gebser M, et al. *Answer set programming modulo acyclicity*. Fundam Inform, 2016.
10. Bonet ML, Galesi N. *Optimality of size-width tradeoffs for resolution*. Computational Complexity, 10(4):261–276, 2001.
11. Bordeaux L, Marques-Silva J. *Knowledge compilation with empowerment*. In *International Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM)*, pages 612–624. Springer, 2012.
12. Crawford JM, Auton LD. *Experimental results on the crossover point in satisfiability problems*. In *AAAI*, volume 93, pages 21–27. 1993.
13. Cuteri B, Dodaro C, et al. *Constraints, lazy constraints, or propagators in ASP solving: An empirical analysis*. In *International Conference on Logic Programming (ICLP)*. 2017. <http://arxiv.org/abs/1707.04027>.
14. de Oliveira RT, Silva F. *SAT and MaxSAT encodings for trees applied to the Steiner tree problem*. In *Brazilian Conference on Intelligent Systems (BRACIS)*, pages 192–197. IEEE, 2014.
15. Dubrova E, Teslenko M. *A SAT-based algorithm for finding attractors in synchronous boolean networks*. IEEE/ACM transactions on computational biology and bioinformatics, 2011.
16. Eén N, Sörensson N. *An extensible SAT-solver*. In *International Conference on Theory and Applications of Satisfiability Testing (SAT)*, pages 502–518. 2003.
17. Eén N, Sörensson N. *Translating pseudo-boolean constraints into SAT*. Journal on Satisfiability, Boolean Modeling and Computation, 2:1–26, 2006.
18. Fisher MJ, Meyer AR. *Boolean matrix multiplication and transitive closure*. Switching and Automata, 1971.
19. Furman ME. *Application of a method of fast multiplication of matrices in the problem of finding the transitive closure of a graph*. Soviet Math Dokl, 1970.
20. Gall FL. *Powers of tensors and fast matrix multiplication*. In *ISSAC*. 2014.
21. Gebser M, Janhunen T, et al. *SAT modulo graphs: Acyclicity*. In *Logics in Artificial Intelligence (JELIA)*. 2014.
22. Gent IP. *Arc consistency in SAT*. In *European Conference on Artificial Intelligence (ECAI)*, pages 121–125. IOS Press, 2002.
23. Goldreich O, Ostrovsky R. *Software protection and simulation on oblivious RAMs*. J ACM, 43(3):431–473, 1996.
24. Gwynne M, Kullmann O. *On SAT representations of XOR constraints*. In *Language and Automata Theory and Applications Conference (LATA)*, pages 409–420. 2014.
25. Haken A. *The intractability of resolution*. Theoretical Computer Science, 39:297–308, 1985.
26. Hartung S, Nichterlein A. *NP-hardness and fixed-parameter tractability of realizing degree sequences with directed acyclic graphs*. SIAM Journal on Discrete Mathematics, 29(4):1931–1960, 2015.
27. Harvey D, van der Hoeven J, et al. *Even faster integer multiplication*. Journal of Complexity, 2016.
28. Ignatiev A, Morgado A, et al. *Cardinality encodings for graph optimization problems*. In *International Joint Conference on Artificial Intelligence (IJCAI)*. 2017.
29. Janhunen T, Niemelä I. *Compact translations of non-disjunctive answer set programs to propositional clauses*. In *Logic Programming, Knowledge Representation, and Nonmonotonic Reasoning: Essays Dedicated to Michael Gelfond on the Occasion of His 65th Birthday*. 2011.

30. Jiang Y, Kautz H, et al. *Solving problems with hard and soft constraints using a stochastic algorithm for MAX-SAT*. In *1st International Joint Workshop on Artificial Intelligence and Operations Research*. 1995.
31. Knuth DE. *The Art of Computer Programming, Volume 4, Fascicle 6: Satisfiability*. Addison-Wesley Professional, 1st edition, 2015.
32. Krishnamurthy B. *Short proofs for tricky formulas*. *Acta Informatica*, 22(3):253–275, 1985.
33. Lifschitz V, Razborov A. *Why are there so many loop formulas?* *ACM Transactions on Computational Logic (TOCL)*, 7(2):261–268, 2006.
34. Lustig D, Sethi G, et al. *COATCheck: Verifying memory ordering at the hardware-OS interface*. *ACM SIGOPS Operating Systems Review*, 2016.
35. Lynce I, Marques-Silva J. *Probing-based preprocessing techniques for propositional satisfiability*. In *ICTAI*, pages 105–110. 2003.
36. Marques-Silva J, Lynce I. *Towards robust CNF encodings of cardinality constraints*. In *Principles and Practice of Constraint Programming (CP)*, pages 483–497. Springer, 2007.
37. Munro I. *Efficient determination of the transitive closure of a directed graph*. *Inf Process Lett*, 1(2):56–58, 1971.
38. Piette C, Hamadi Y, et al. *Vivifying propositional clausal formulae*. In *European Conference on Artificial Intelligence (ECAI)*, pages 525–529. 2008.
39. Pipatsrisawat K, Darwiche A. *On the power of clause-learning SAT solvers as resolution engines*. *Artificial Intelligence*, 175(2):512–525, 2011.
40. Pippenger N, Fischer MJ. *Relations among complexity measures*. *J ACM*, 26(2):361–381, 1979.
41. Plaisted DA, Greenbaum S. *A structure-preserving clause form translation*. *J Symb Comput*, 2(3):293–304, 1986.
42. Rintanen J, Heljanko K, et al. *Planning as satisfiability: parallel plans and algorithms for plan search*. *Artificial Intelligence*, 170(12):1031 – 1080, 2006.
43. Sinz C. *Towards an optimal cnf encoding of boolean cardinality constraints*. In *Principles and Practice of Constraint Programming (CP)*, pages 827–831. Springer, 2005.
44. Stålmarck G. *Short resolution proofs for a sequence of tricky formulas*. *Acta Informatica*, 33(3):277–280, 1996.
45. Strassen V. *Gaussian elimination is not optimal*. *Numer Math*, 1969.
46. Tamura N, Taga A, et al. *Compiling finite linear CSP into SAT*. *Constraints*, 2009.
47. Tavares de Oliveira R, Silva F, et al. *On modeling connectedness in reductions from graph problems to extended satisfiability*. In *Advances in Artificial Intelligence (IBERAMIA)*, pages 381–391. Springer, 2012.
48. Tseitin GS. *On the complexity of derivations in the propositional calculus*. *Studies in Constructive Mathematics and Mathematical Logic, Part II*, 1968.
49. van Beek P, Hoffmann HF. *Machine learning of Bayesian networks using constraint programming*. In *International Conference on Principles and Practice of Constraint Programming (CP)*, pages 429–445. 2015.
50. Van Gelder A. *Another look at graph coloring via propositional satisfiability*. *Discrete Appl Math*, 2008.
51. Wickerson J, Batty M, et al. *Automatically comparing memory consistency models*. In *Symposium on Principles of Programming Languages (POPL)*. ACM, 2017.