

# SAT Encodings for Pseudo-Boolean Constraints Together With At-Most-One Constraints

Miquel Bofill<sup>a</sup>, Jordi Coll<sup>b,\*</sup>, Peter Nightingale<sup>c</sup>, Josep Suy<sup>a</sup>, Felix  
Ulrich-Oltean<sup>c</sup>, Mateu Villaret<sup>a</sup>

<sup>a</sup>*Universitat de Girona, Girona, Spain*

<sup>b</sup>*Aix Marseille Univ, Université de Toulon, CNRS, LIS, Marseille, France*

<sup>c</sup>*University of York, York, United Kingdom*

---

## Abstract

When solving a combinatorial problem using propositional satisfiability (SAT), the encoding of the problem is of vital importance. We study encodings of Pseudo-Boolean (PB) constraints, a common type of arithmetic constraint that appears in a wide variety of combinatorial problems such as timetabling, scheduling, and resource allocation. In some cases PB constraints occur together with at-most-one (AMO) constraints over subsets of their variables (forming PB(AMO) constraints). Recent work has shown that taking account of AMOs when encoding PB constraints using decision diagrams can produce a dramatic improvement in solver efficiency. In this paper we extend the approach to other state-of-the-art encodings of PB constraints, developing several new encodings for PB(AMO) constraints. Also, we present a more compact and efficient version of the popular Generalized Totalizer encoding, named Reduced Generalized Totalizer. This new encoding is also adapted for PB(AMO) constraints for a further gain. Our experiments show that the encodings of PB(AMO) constraints can be substantially smaller than those of PB constraints. PB(AMO) encodings allow many more instances to be solved within a time limit, and solving time is improved by more than one order of magnitude in some cases. We also observed that there is no single overall winner among the considered encodings, but efficiency of each encoding may depend on PB(AMO) characteristics such as the magnitude of coefficient values.

Published in Artificial Intelligence Journal: <https://doi.org/10.1016/j.artint.2021.103604>

*Keywords:* pseudo-Boolean constraints, encoding, at-most-one constraints, SAT

---

\*Corresponding author

*Email addresses:* [miquel.bofill@imae.udg.edu](mailto:miquel.bofill@imae.udg.edu) (Miquel Bofill), [jordi.coll@lis-lab.fr](mailto:jordi.coll@lis-lab.fr) (Jordi Coll), [peter.nightingale@york.ac.uk](mailto:peter.nightingale@york.ac.uk) (Peter Nightingale), [josep.suy@imae.udg.edu](mailto:josep.suy@imae.udg.edu) (Josep Suy), [fvuo500@york.ac.uk](mailto:fvuo500@york.ac.uk) (Felix Ulrich-Oltean), [mateu.villaret@imae.udg.edu](mailto:mateu.villaret@imae.udg.edu) (Mateu Villaret)

## 1. Introduction

Discrete decision-making problems crop up in many contexts in the modern world. Such problems can be expressed as constraint satisfaction (or optimisation) problems (CSPs or COPs), then solved using a variety of solver types. An increasingly popular and successful approach to solving CSPs and COPs is to encode them into Boolean formulas and then to apply an off-the-shelf SAT solver. This approach is attractive because of the power of modern conflict-directed clause learning (CDCL) SAT solvers, such as CaDiCaL [1] and Glucose [2], which incorporate conflict learning, powerful search heuristics, and fast propagation of the Boolean constraints.

Linear equations and inequalities are ubiquitous in constraint problems such as scheduling, routing, resource allocation, and many other hard combinatorial problems. Pseudo-Boolean (PB) constraints are a particular type of linear constraint – PB constraints are of the form  $\sum_{i=1}^n q_i x_i \# K$ , where  $\# \in \{<, \leq, =, \geq, >\}$ ,  $q_1, \dots, q_n$  and  $K$  are integer constants, and  $x_1, \dots, x_n$  are 0/1 variables. There has been a great deal of work on encoding PB constraints to SAT, some of which is reviewed by Philipp and Steinke [3]. State-of-the-art encodings are based on Binary Decision Diagrams [4, 5], Sequential Weight Counters [6], Generalized Totalizers [7, 8], and Polynomial Watchdog schemes [9, 10]. At-most-one (AMO) constraints (i.e. constraints of the form  $\sum_{i=1}^m x_i \leq 1$ ) are also very common, with the most basic being a mutual exclusion between two 0/1 variables. In this paper we extend several of the PB encodings mentioned above, and demonstrate substantially improved performance when PB constraints intersect with AMO constraints.

Bofill, Coll, Suy, and Villaret [11, 12] proposed a SAT encoding based on Multi-valued Decision Diagrams (MDDs) for a conjunction of a PB constraint with a set of AMO constraints over the variables of the PB constraint. Such conjunctions are referred to as PB(AMO) constraints. The AMO constraints, which can be encoded to SAT in any way, allow certain interpretations to be erased from decision diagrams, and to represent the PB constraint as an MDD instead of as a Binary Decision Diagram (BDD). The encoding of the MDD is notably smaller than the encoding of an equivalent BDD, and the solving time is substantially reduced. This technique has been used to provide efficient formulations of particular kinds of scheduling problems [11, 13]. Also, Ansótegui et al [14] integrated the MDD-based SAT encoding of PB(AMO) constraints into the automatic reformulation pipeline of Savile Row [15], showing important size and solving time improvements compared to a BDD-based encoding oblivious to the existence of AMO constraints.

Efficient encodings of the conjunction of PB and AMO constraints can have a significant impact on solving a wide range of CSPs. This combination of constraints appears in settings where one option has to be chosen among a set of incompatible options, and the decision has an associated cost. This pattern occurs in numerous applications, for example logistics [16], resource allocation [17], capital budgeting [18], telecommunications [19], combinatorial auctions [20], and routing [21], among many others. In short, any problem which is essentially a

multi-choice knapsack problem is likely to contain both PB and AMO constraints. Moreover, as stated in Bofill et al [12], any Linear Integer Arithmetic expression, which are also ubiquitous in CSP models, can be easily transformed to a PB(AMO) constraint. Therefore finding new and better SAT encodings of PB(AMO) constraints is of wide interest.

As a motivating example, consider the PB constraint  $2x_1 + 3x_2 + 4x_3 + 2x_4 + 3x_5 + 4x_6 \leq 7$ . Also, suppose there are two AMO constraints:  $x_1 + x_2 + x_3 \leq 1$  and  $x_4 + x_5 + x_6 \leq 1$ . Encoding the PB constraint alone would require several clauses and (depending on the chosen encoding) multiple additional variables. For example, the Generalized Totalizer encoding [7] has 23 additional variables and 56 clauses. However, the two AMO constraints rule out most of the values that the sum  $(2x_1 + \dots + 4x_6)$  could take, and almost all such values that break the PB constraint. Encoding the PB constraint together with the two AMO constraints requires just one clause to prevent  $x_3$  and  $x_6$  being assigned true together. This simple observation underpins all the PB(AMO) encodings presented in this paper.

### 1.1. Contributions

The main contribution of this paper is to generalize five state-of-the-art SAT encodings of PB constraints to encode PB(AMO) constraints. We generalize each of the following encodings: Sequential Weight Counter (SWC), Generalized Totalizer (GT), n-Level Modulo Totalizer (MTO), Global Polynomial Watchdog (GPW), and Local Polynomial Watchdog (LPW). In each case we demonstrate substantial reductions in size, and also improvements in solving time using two recent CDCL SAT solvers. Compared to their PB counterparts, PB(AMO) encodings allow many more instances to be solved within a time limit, and solving time is improved by more than one order of magnitude in some cases.

Since PB(AMO) constraints generalize PB constraints, we follow the convention of naming the new encodings after the original encoding, prefixing them with the word *Generalized*, e.g., from the Sequential Weight Counter (SWC) encoding we provide the Generalized Sequential Weight Counter (GSWC) encoding. We also show that the new encodings preserve the propagation properties of the original ones.

Another contribution of this paper is a new encoding of PB(AMO) constraints called Reduced Generalized Generalized Totalizer (RGGT). This encoding does not directly generalize any existing PB encoding. RGGT first constructs a Generalized Generalized Totalizer (GGT) tree, then applies a reduction algorithm in order to obtain a more compact representation that replaces individual numeric values with intervals. In some cases RGGT will detect that terms in the PB(AMO) constraint are redundant. When this occurs the redundant terms are removed and the entire encoding process is repeated (until a fixpoint is reached). RGGT is frequently substantially smaller than GGT and this translates to improved solver efficiency. RGGT can also be used as an encoding of PB constraints (without collateral AMO constraints). In this case, we refer to it as the Reduced Generalized Totalizer (RGT) encoding. We also present a new heuristic called *minRatio* to build the binary trees required by

GT, GGT, RGT, and RGGT encodings. The experimental results show that the minRatio heuristic has a substantial positive effect on the size and performance of GT and GGT encodings when compared to a simple balanced tree.

Our experimental results show that the size of the SAT encodings of PB constraints can be dramatically reduced thanks to taking AMO constraints into account, and that there can be a huge improvement in solving time when using the new generalized encodings. We provide new benchmarks (and use others from the literature) which contain AMO constraints and PB constraints in different configurations, and we show empirically that some encodings are better than others for particular kinds of PB(AMO).

The rest of the paper is structured as follows:

- Section 2 presents preliminary concepts and notations used in this paper.
- Section 3 describes the normalisation processes that we perform before encoding a PB(AMO) constraint.
- For the sake of completeness, Section 4 summarises the decision diagram encoding of PB constraints [5] and its generalization to PB(AMO) constraints [12].
- Sections 5, 6, 7, 8, and 9 present the new PB(AMO) encodings. Each section describes the existing PB encoding followed by its generalization for PB(AMO) constraints.
- Section 10 describes the propagation properties of the presented encodings.
- Section 11 presents our experimental results comparing the PB(AMO) encodings to each other and also to their corresponding PB encoding.
- Section 12 surveys other work related to PB(AMO) encoding.
- Section 13 is devoted to conclusions and future work.

This work is an extension of the previous work [22]; here we provide more detailed explanations and examples, and the following additional content: the set of preprocesses used to simplify PB(AMO) constraints before they are encoded, in Section 2; a summary of the PB(AMO) based on MDDs presented in [12] (in order to collect all existing PB(AMO) encodings in this paper), in Section 4; the Reduced Generalized Totalizer (RGT) encoding for PB constraints and its generalized version for PB(AMO) constraints in Section 6.3; the min-Ratio heuristic to build generalized totalizers in Section 6.4; the Generalized n-Level Modulo Totalizer (GMTO) encoding in Section 7; and the Generalized Local Polynomial Watchdog (GLPW) encoding in Section 9. We give further proofs of propagation properties of GGPW and the new encodings in Section 10. The experimental section has also been extended: we study in detail the size reduction achieved by the new encodings of PB and PB(AMO) constraints, namely RGT and RGGT; we study the results of the other new PB(AMO) encodings, namely GMTO and GLPW; we consider additional benchmark sets

from the Combinatorial Auctions problem, the resource-constrained scheduling problems MRCPSP and RCPSP/t, and the Nurse Scheduling Problem; and in addition to the original experiments run with Glucose, we run all experiments with the SAT solver CaDiCaL.

## 2. Preliminaries

A *Boolean variable* is a variable that can take truth values 0 (false) and 1 (true). A *literal* is a Boolean variable  $x$  or its negation  $\bar{x}$ . A *clause* is a disjunction of literals. A *propositional formula in conjunctive normal form* (CNF) is a conjunction of clauses. We will assume that all formulas are in CNF. Clauses are usually seen as sets of literals, and formulas as sets of clauses. A *Boolean function* is a function of the form  $f : \{0, 1\}^n \rightarrow \{0, 1\}$ .

We will consider constraints that are defined over a finite set of Boolean variables, i.e., Boolean functions. An *assignment* is a mapping of Boolean variables to truth values; it can also be seen as a set of literals, e.g.  $\{x = 1, y = 0, z = 0\}$  is usually denoted  $\{x, \bar{y}, \bar{z}\}$ . By  $scope(C)$  we denote the set of variables occurring in a constraint  $C$ .

A *satisfying assignment* of a Boolean function  $f$  is an assignment that makes  $f$  evaluate to 1. In particular, an assignment  $A$  satisfies a formula  $F$  in CNF if at least one literal  $l$  of each clause in  $F$  belongs to  $A$ . Such an assignment is called a *model* of the formula. Given two Boolean functions  $F$  and  $G$ , we say that  $G$  is logically implied by  $F$ , iff every model of  $F$  is also a model of  $G$ .

*Definition 1.* An *at-most-one* (AMO) constraint is a Boolean function of the form  $\sum_{i=1}^n x_i \leq 1$ , where all  $x_i$  are Boolean variables.

*Definition 2.* A *pseudo-Boolean* (PB) constraint is a Boolean function of the form  $\sum_{i=1}^n q_i x_i \# K$  where  $K$  and all  $q_i$  are integer constants, all  $x_i$  are Boolean variables, and  $\# \in \{<, \leq, =, \geq, >\}$ .

*Definition 3.* By *PB(AMO)* constraint we refer to a constraint of the form  $P \wedge M_1 \wedge \dots \wedge M_N$ , where  $P$  is a PB constraint, and  $M_1, \dots, M_N$  are AMO constraints such that  $\{scope(M_1), \dots, scope(M_N)\}$  is a partition of  $scope(P)$ .

Since a single variable constitutes a trivial at-most-one constraint, PB constraints are a particular case of PB(AMO) constraints, i.e., a PB constraint of the form  $\sum_{i=1}^n q_i x_i \leq K$  is a PB(AMO) constraint of the form  $\sum_{i=1}^n q_i x_i \leq K \wedge x_1 \leq 1 \wedge \dots \wedge x_n \leq 1$ .

*Example 1.*  $2x_1 + 3x_2 + 3x_3 + 7x_4 \leq 8 \wedge x_1 + x_2 \leq 1 \wedge x_3 + x_4 \leq 1$  is a PB(AMO) constraint. Notice that, for instance, the assignment  $x_1, x_2, x_3, \bar{x}_4$  satisfies the PB constraint  $2x_1 + 3x_2 + 3x_3 + 7x_4 \leq 8$  but does not satisfy the PB(AMO) constraint because of  $x_1 + x_2 \leq 1$ .

*Definition 4* (Encoding). We say that a formula  $G$  is an *encoding* of a Boolean function  $F$  if the following holds: given an assignment  $A$  over the variables of  $F$ ,  $A$  satisfies  $F$  iff  $A$  can be extended to a satisfying assignment of  $G$ .




Figure 1: Tree representation of a totalizer for variables  $x_1, x_2, x_3, x_4, x_5, x_6, x_7$ .

A large number of encodings have been proposed for AMO constraints and PB constraints in the literature, as well as encodings for *cardinality constraints*, i.e., PB constraints with all coefficients  $q_i = 1$ . Here we revisit some basics of the Totalizer encoding for cardinality constraints [23], since it is closely related to many of the encodings presented in the paper. Given a set of variables  $x_1, \dots, x_n$ , a *totalizer* is a binary tree that contains a different variable  $x_i$  at each leaf. An example is given in Figure 1. For any subtree, the root of the subtree contains a list with as many Boolean variables as leaves in the subtree. These variables are constrained to represent in unary notation the sum of the variables of the leaves of the subtree or, equivalently, the list will be a decreasing ordering of the values of the leaves. In particular, the list of variables  $a_1, \dots, a_n$  contained in the root of the totalizer represents the decreasing ordering of the values of  $x_1, \dots, x_n$ . Then, an encoding of the cardinality constraint  $\sum_{i=1}^n x_i \leq K$  is completed by forbidding that variable  $a_{K+1}$  is true.

### 3. Encoding Technique and Normalisation of PB(AMO) Constraints

Given a PB(AMO) constraint of the form  $P \wedge M_1 \wedge \dots \wedge M_N$ , a straightforward approach to encode it is to generate a formula  $F$  of the form  $G \wedge H_1 \wedge \dots \wedge H_N$ , where  $G$  is an encoding of  $P$ , and  $H_i$  is an encoding of  $M_i$  for all  $i \in 1..N$ . Instead, similarly to the MDD-based approach of [12], we propose to encode PB(AMO) constraints in a combined way. We encode the conjunction of AMO constraints in the usual way, i.e., each AMO is encoded separately and we use the conjunction of all the resulting clauses. However, we encode the PB constraint assuming that the accompanying AMO constraints are already enforced. This is what will let us significantly reduce the size of the PB constraint encoding. Lemma 1 states the correctness of our encoding technique.

*Lemma 1.* Let  $\mathcal{P}$  be a PB(AMO) constraint of the form  $P \wedge M_1 \wedge \dots \wedge M_N$ , where  $\mathcal{X} = \{X_1, \dots, X_N\}$  is a partition of the variables in  $P$  such that  $X_i = \text{scope}(M_i)$ . Let  $E$  be a Boolean formula such that, given any assignment  $A$  of the variables of  $\mathcal{P}$  that satisfies  $M_1 \wedge \dots \wedge M_N$ ,  $A$  can be extended to a model of  $E$  iff  $A$  satisfies  $P$ . Then the conjunction of  $E$  with an encoding of  $M_1 \wedge \dots \wedge M_N$  is an encoding of  $\mathcal{P}$ .

The encoding procedures described in the following sections generate the  $E$

formula of Lemma 1. It is assumed that the AMO constraints  $M_1$  to  $M_N$  are encoded elsewhere.

We will not restrict ourselves to any particular encoding for the AMO constraints when encoding a PB(AMO) constraint. Furthermore, in the context of a larger formula, if the AMO constraints are not explicit but logically implied by the formula, then the encoding of the PB constraint will suffice to obtain a correct encoding of the PB(AMO) constraint.

Given a PB(AMO) constraint  $\mathcal{P}$  of the form  $P \wedge M_1 \wedge \dots \wedge M_N$ , the encoding procedures described in the following sections take as input the pair  $(P, \mathcal{X})$ , where  $\mathcal{X} = \{X_1, \dots, X_N\} = \{scope(M_1), \dots, scope(M_N)\}$  is a partition of the variables of the PB constraint  $P$ . Moreover, we soundly preprocess the PB constraint  $P$  and the partition  $\mathcal{X}$  so that they satisfy a set of properties. Some of the properties are required for some of the encodings presented in this paper, while others are used to reduce the size of the constraints or to deal efficiently with trivial constraints or monomials in the PB constraints. The properties are enforced in the order they are given, and clauses and variables may be introduced at each step.

*Property 1.*  $P$  is of the form  $\sum_{i=1}^n q_i x_i \leq K$ , with  $q_i \geq 0$ .

A PB(AMO) satisfying this property is a monotonic decreasing Boolean function, i.e. any of its models remains a model after flipping assignments from 1 to 0. The PB constraint  $P$  of a PB(AMO) constraint can always be transformed to have this property [12]. If  $P$  is of the form  $\sum_{i=1}^n q_i x_i = K$ , it is replaced by the conjunction of two constraints  $\sum_{i=1}^n q_i x_i \leq K$  and  $\sum_{i=1}^n q_i x_i \geq K$ , both with partition  $\mathcal{X}$ . The operators  $\geq$  and  $>$  can be transformed to  $\leq$  or  $<$ , respectively, by multiplying both sides of the inequality by  $-1$ . Furthermore,  $<$  can be replaced by  $\leq$  by subtracting 1 from the right-hand side. Finally, in order to get only positive coefficients, we proceed as follows. For each AMO group  $X_i$  with some negative coefficient in the PB constraint, consider such minimum negative coefficient  $q$  and perform the following transformation: increase by  $-q$  all coefficients of the variables of the AMO-group, as well as  $K$ . Moreover, add a new variable  $y$  defined by  $y \leftrightarrow \bigwedge_{x_i \in X_i} \overline{x_i}$  and add it to the PB constraint with coefficient  $-q$ . This way the inequality is preserved, since the left-hand side will be incremented exactly by  $-q$  whatever variables are set to true, thanks to the AMO constraint. As an example, consider the PB constraint  $4x_1 - 7x_2 + \dots \leq 15$ , and assume  $x_1$  and  $x_2$  constitute an AMO group, i.e., we have  $x_1 + x_2 \leq 1$ . According to our transformation we would get  $11x_1 + 0x_2 + 7y + \dots \leq 22$ , with  $y \leftrightarrow \overline{x_1} \wedge \overline{x_2}$ . Observe that, on the one hand, the contribution of each variable of the AMO group is incremented by 7, so the right-hand side must be incremented by 7 as well and, on the other hand, if neither of the variables of the AMO group is set to true, then  $y$  contributes 7 to the left-hand side, so the original relation is preserved. Note moreover that the variables originally with the minimum coefficient ( $x_2$  in the example) can be removed from the PB constraint, as their coefficient becomes 0 after the transformation.

*Property 2.*  $K > 0$ .

If  $K < 0$  (the constraint is unsatisfiable) the encoding of  $(P, \mathcal{X})$  is just the

empty clause. If  $K = 0$ , the encoding is just the union of the unit clauses  $\bar{x}_l$  for all terms  $q_l x_l$  where  $q_l > 0$ .

*Property 3.* All coefficients satisfy  $q_l > 0$ .

Otherwise, for any  $q_l = 0$  where  $x_l \in X_i$ , remove the term  $q_l x_l$  from  $P$ , remove  $x_l$  from  $X_i$ , and remove any set  $X_i \in \mathcal{X}$  that becomes empty.

*Property 4.* All coefficients satisfy  $q_l \leq K$ .

This can be achieved by adding the unit clause  $\bar{x}_l$  for all  $q_l > K$ . For such  $q_l$ , where  $x_l \in X_i$ , remove the corresponding monomial  $q_l x_l$  from  $P$ , remove  $x_l$  from  $X_i$ , and remove any set  $X_i \in \mathcal{X}$  that becomes empty.

*Property 5.*  $N > 1$ .

Otherwise, the encoding of  $(P, \mathcal{X})$  is just the empty CNF (note that this is sound given Property 4 and the AMOs assumption).

*Property 6.*  $\sum_{i=1}^N \max_{x_l \in X_i} (q_l) > K$ .

This property means that the constraint is not trivially true when all AMO constraints hold. Otherwise the encoding of  $(P, \mathcal{X})$  is just the empty CNF.

*Property 7.* There are no two variables  $x_l, x_{l'}$  belonging to the same  $X_i$  with the same coefficient  $q_l = q_{l'}$ .

Otherwise, for all maximal subsets of variables  $X'_i \subseteq X_i$  with the same coefficient  $q$  (i.e.  $\forall x_l \in X'_i, q_l = q$ ): add an auxiliary variable  $y$ ; update  $X_i$  as  $X_i = (X_i \setminus X'_i) \cup \{y\}$ ; remove from  $P$  all monomials  $q_l x_l$ , for all  $x_l \in X'_i$ ; add to  $P$  the monomial  $qy$ ; and add to the encoding of  $(P, \mathcal{X})$  the clauses  $\bar{x}_l \vee y$ , for all  $x_l \in X'_i$ .

**Remark:** Henceforth, we assume that PB(AMO) constraints satisfy all these properties.

## 4. Decision Diagram encoding

In order to cover all existing encodings of PB(AMO) constraints, here we provide an overview of using Binary Decision Diagrams (BDDs) to encode PB constraints following [5], and in Subsection 4.2 we provide an overview of using Multivalued Decision Diagrams (MDDs) to encode PB(AMO) constraints as is done in [12].

### 4.1. Binary Decision Diagram

Many BDD-based encodings of PB constraints have been proposed. In this work we consider the encoding for monotonic decreasing PB constraints presented in [5], which in this paper will be referred to as the *BDD* encoding. Given a PB constraint of the form  $\sum_{i=1}^n q_i x_i \leq K$  with  $q_i \geq 0$  for all  $i$ , the BDD encoding consists of, first of all, representing the constraint as a Reduced Ordered BDD (ROBDD), and then encoding that BDD as a set of clauses. We use the following definition of a BDD:



*Definition 5* (Definition 5 in [12]). A *Binary Decision Diagram* (BDD) is a rooted, directed, acyclic graph which represents a Boolean function. BDDs have two terminal nodes, namely  $\perp$ -terminal and  $\top$ -terminal. Each nonterminal node has an associated Boolean variable (*selector*), and two outgoing edges, representing the *true* and the *false* assignment of the selector. Every truth assignment of the variables follows a path from the root to the  $\top$ -terminal when it satisfies the formula, or to the  $\perp$ -terminal otherwise.

A BDD is called *ordered* if different variables appear in the same order on all paths from the root. A BDD is said to be *reduced* if it satisfies the following two conditions: it contains no isomorphic sub-BDDs, and there is no node whose true and false child are the same.

Figure 2a contains an example BDD representing a PB constraint. There exist algorithms to construct a ROBDD representing a PB constraint that run in polynomial time w.r.t. the size of the resulting ROBDD.

Once the ROBDD has been constructed, an auxiliary variable  $v$  is introduced for each node. The encoding enforces  $v$  to be false whenever the sub-ROBDD rooted at that node follows a path to the  $\perp$ -terminal with a given assignment. That is, it enforces  $\bar{v}_0 \wedge \bar{x}_i \rightarrow \bar{v}$  and  $\bar{v}_1 \wedge x_i \rightarrow \bar{v}$ , where  $x_i$  is the variable in the node of the ROBDD, and  $v_0$  and  $v_1$  are the auxiliary variables of the false child and the true child of that node, respectively. However, since the considered constraints are monotonic, we can simplify the constraint  $\bar{v}_0 \wedge \bar{x}_i \rightarrow \bar{v}$  to  $\bar{v}_0 \rightarrow \bar{v}$ , because if  $x_i = 0$  falsifies the PB constraint, so will  $x_i = 1$ . Therefore, for each nonterminal node, the encoding introduces the following clauses:

$$v_0 \vee \bar{v} \tag{1}$$

$$v_1 \vee \bar{x}_i \vee \bar{v} \tag{2}$$

The encoding is completed by adding three unary clauses:

$$v_r \wedge \bar{v}_\perp \wedge v_\top \tag{3}$$

where  $v_r$ ,  $v_\perp$  and  $v_\top$  are the auxiliary variables of the root node, the  $\perp$ -terminal and the  $\top$ -terminal respectively.

#### 4.2. Multivalued Decision Diagram

Bofill, Coll, Suy, and Villaret [11, 12] presented a generalization of the BDD encoding for PB(AMO) constraints. Like the new encodings presented in this paper, the MDD encoding receives as input a PB constraint  $P$  and a partition  $\mathcal{X}$  of its variables. This encoding uses at-most-one Multivalued Decision Diagrams (AMO-MDD), which is a more generic structure than a BDD. In an AMO-MDD, as defined in [12], each node contains a set of selector variables. Given the partition  $\mathcal{X} = \{X_1, \dots, X_N\}$ , the nodes of the  $i$ -th layer of the AMO-MDD contain the variables of  $X_i$  as selectors. An example is given in Figure 2b. For each node there are multiple possible choices (outgoing edges): assigning one of the selectors to true, or assigning all of them to false (the *else* edge). Therefore, the AMO-MDDs correctly represent all assignments where at most one variable

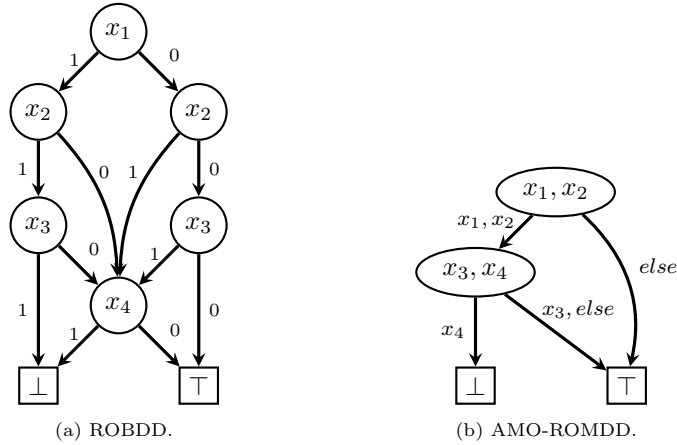


Figure 2: (a): ROBDD for the PB constraint  $2x_1 + 3x_2 + 4x_3 + 7x_4 \leq 8$  with variable ordering  $x_1 < x_2 < x_3 < x_4$ . (b): AMO-ROMDD for  $2x_1 + 3x_2 + 4x_3 + 7x_4 \leq 8$ , with ordered partition  $\{x_1, x_2\} < \{x_3, x_4\}$ . Multiple edges between two nodes are represented as a single edge with multiple labels.

in each set  $X_i$  is true. Similarly to BDD, first of all an AMO-ROMDD is constructed, given an order of the elements of  $\mathcal{X}$ . Then, one auxiliary variable is introduced for each node and Clauses (3) are introduced. Also, for each nonterminal node, the following clauses are introduced:

$$v_0 \vee \bar{v} \tag{4}$$

$$v_j \vee \bar{x}_j \vee \bar{v} \quad \forall x_j \in X_i \text{ s.t. } v_j \neq v_0 \tag{5}$$

where  $v$  is the auxiliary variable of the encoded node,  $v_j$  is the auxiliary variable of the child node selected by the selector variable  $x_j$ , and  $v_0$  is the auxiliary variable of the *else* child.

The BDD encoding requires  $O(nK)$  auxiliary variables and  $O(nK)$  clauses, while the AMO-MDD encoding requires  $O(NK)$  auxiliary variables and  $O(nK)$  clauses.

## 5. Sequential Weight Counter Encoding

In this section we first recall the Sequential Weight Counter encoding for PB constraints from [6] and in Subsection 5.2 we provide its generalization to encode PB(AMO) constraints.

### 5.1. Sequential Weight Counter

The idea of the *Sequential Weight Counter* (SWC) encoding for PB constraints is to build a circuit that sequentially sums from left to right the coefficients (a.k.a. weights)  $q_i$  whose variable  $x_i$  is set to true. Specifically, given a PB constraint  $\sum_{i=1}^n q_i x_i \leq K$ , there is a sequence of  $n$  counters of  $K$  inputs and

$K$  outputs, where the  $i$ -th counter is associated to the variable  $x_i$ . Each counter receives as input a vector of Boolean variables, which is the unary representation of an integer value, and adds the weight  $q_i$  to the output if the associated variable  $x_i$  is set to true. Therefore, the  $i$ -th counter receives as input  $\sum_{j=1}^{i-1} q_j x_j$  and outputs  $\sum_{j=1}^i q_j x_j$ . Note that the output of the counter number  $i - 1$  is the input of the  $i$ -th counter.

An example of a sequence of counters is shown in Figure 3a. The encoding introduces  $n \cdot K$  variables, denoted  $s_{i,j}$ , with  $1 \leq i \leq n$ ,  $1 \leq j \leq K$ , where  $s_{i,j}$  is the  $j$ -th output of the  $i$ -th counter and also the  $j$ -th input of the  $(i + 1)$ -th counter. The encoding introduces the following clauses:

$$\overline{s_{i-1,j}} \vee s_{i,j} \quad 2 \leq i < n, 1 \leq j \leq K \quad (6)$$

$$\overline{x_i} \vee s_{i,j} \quad 1 \leq i < n, 1 \leq j \leq q_i \quad (7)$$

$$\overline{s_{i-1,j}} \vee \overline{x_i} \vee s_{i,j+q_i} \quad 2 \leq i < n, 1 \leq j \leq K - q_i \quad (8)$$

$$\overline{s_{i-1,K+1-q_i}} \vee \overline{x_i} \quad 2 \leq i \leq n \quad (9)$$

where  $s_{0,j}$  is the constant 0 for all  $j$ , to represent the input of the first counter which is the empty sum. Clauses (6) state that  $\sum_{j=1}^i q_j x_j \geq \sum_{j=1}^{i-1} q_j x_j$ . Clauses (7) and (8) enforce that if a variable  $x_i$  is true then its coefficient is added to the input of the next counter. Finally, Clauses (9) enforce that the sum never exceeds  $K$ .

## 5.2. Generalized Sequential Weight Counter (GSWC)

We define the GSWC encoding by, instead of associating a single monomial  $q_i x_i$  from the PB constraint to each counter, associating a set of monomials to each of them. In our generalization, given a partition  $\mathcal{X} = \{X_1, \dots, X_N\}$  of the variables of the PB constraint, the resulting formulation will have just  $N$  counters, where the  $i$ -th counter will handle all the monomials  $q_l x_l$  for the variables  $x_l$  in  $X_i$ . If the variables in each set  $X_i$  are subject to an AMO constraint then, given an assignment satisfying those constraints, at most one coefficient  $q_l$  will be added by each counter, and the output of the whole circuit will correspond to the value of  $\sum_{i=1}^n q_i x_i$ . As in the original encoding, we will enforce that a sum exceeding  $K$  is not reached. The GSWC encoding introduces the following clauses:

$$\overline{s_{i-1,j}} \vee s_{i,j} \quad 2 \leq i < N, 1 \leq j \leq K \quad (10)$$

$$\overline{x_l} \vee s_{i,j} \quad 1 \leq i < N, x_l \in X_i, 1 \leq j \leq q_l \quad (11)$$

$$\overline{s_{i-1,j}} \vee \overline{x_l} \vee s_{i,j+q_l} \quad 2 \leq i < N, x_l \in X_i, 1 \leq j \leq K - q_l \quad (12)$$

$$\overline{s_{i-1,K+1-q_l}} \vee \overline{x_l} \quad 2 \leq i \leq N, x_l \in X_i \quad (13)$$

Clauses (10) propagate the accumulated sum in the same way as Clauses (6). Clauses (11) and (12) enforce  $S_i \geq S_{i-1} + q_l x_l$ , for all  $x_l \in X_i$ , where  $S_{i-1}$  and  $S_i$  are respectively the input and output value of the  $i$ -th counter. Clauses (13)

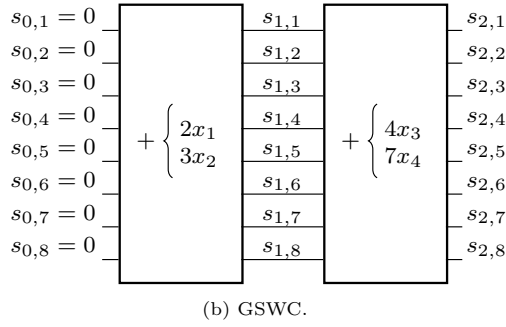



Figure 3: (a): high level circuit representation of  $SWC(2x_1 + 3x_2 + 4x_3 + 7x_4 \leq 8)$ . (b): high level circuit representation of  $GSWC(2x_1 + 3x_2 + 4x_3 + 7x_4 \leq 8, \{\{x_1, x_2\}, \{x_3, x_4\}\})$ .

enforce that the sum never exceeds  $K$ . A high level circuit representation of a GSWC encoding is shown in Figure 3b.

The main difference between the SWC and GSWC encodings is that the latter has only  $N$  counters, instead of  $n$ , and therefore introduces fewer auxiliary variables (assuming  $N < n$ ). Also, the number of Clauses (10) in the GSWC encoding is smaller than the number of Clauses (6) in the SWC encoding. The SWC encoding requires  $O(nK)$  auxiliary variables and  $O(nK)$  clauses, while the GSWC encoding requires  $O(NK)$  auxiliary variables and  $O(nK)$  clauses.

### 5.2.1. Comparison with decision diagrams

There is a close relationship between sequential weight counter circuits and decision diagrams, i.e. between BDD and SWC encodings, and between MDD and GSWC encodings. In particular, a reduced ordered decision diagram can be seen as a sequential weight counter where the  $i$ -th counter only has output pins representing sum values that can be obtained from a subset of the first  $i$  variable coefficients. Moreover, two sum values may share the same decision diagram node if this does not affect satisfiability, and this reduction is not performed in sequential counters. Also, reduced ordered decision diagrams shortcut intermediate nodes when all their output edges point to the same child. This is usually referred to as a long edge, and is also not used by sequential weight

counters. Therefore, there are possible sum values that are not represented by any variable in decision diagram encodings. This means that decision diagram encodings are smaller, but in some models it might be useful having a variable representing each sum value. Both the results of the original paper of SWC [6] and our results show that sometimes encodings based on sequential counters perform better than those based on decision diagrams.

## 6. Generalized Totalizer Encoding

In this section we first revisit the Generalized Totalizer encoding for PB constraints introduced in [7], then in Subsection 6.2 we provide its generalization to encode PB(AMO) constraints. In Subsection 6.3 we introduce a new encoding for PB(AMO) constraints, the Reduced Generalized Generalized Totalizer, which is able to merge sets of equivalent values into intervals, potentially reducing the size of the encoding. Subsection 6.4 provides a heuristic to build the binary trees used in the totalizer encodings.

### 6.1. Generalized Totalizer


The *Generalized Totalizer* (GT) encoding was presented in [7] as a generalization of the Totalizer encoding for cardinality constraints [23]. The overall idea of GT is to represent a PB constraint  $\sum_{i=1}^n q_i x_i \leq K$  as a binary tree where each leaf represents a term of the sum, and non-leaf nodes represents the sum of the terms beneath. In each non-leaf node, every possible value of the sum is represented with one Boolean variable.

In the original presentation, every node of the tree has a distinct label and an attribute *vars* which consists of a set of Boolean variables. We add another attribute *vals*, a set of values. In GT, *vals* corresponds to *vars* (each non-zero value in *vals* has a corresponding variable in *vars*), but this is not the case in one of our generalizations of GT. Each variable  $x_i$  of the PB constraint is placed into the attribute *vars* of a different leaf node, and is renamed after the label of the node and its associated coefficient  $q_i$  (e.g., given the monomial  $3x_1$ , if the variable  $x_1$  is inserted into a leaf node labelled by letter  $O$ , then the variable is named  $o_3$ ). The attribute *vals* of the leaf node is simply  $\{0, q_i\}$ . The attribute *vars* of any non-leaf node labelled  $O$  contains a variable  $o_w$  for every variable  $l_w$  of its left child and for every variable  $r_w$  of its right child. Moreover,  $O.vars$  also contains a variable  $o_w$  for every value  $w$  in the range  $[1, K]$  resulting from summing any pair of values  $w_1$  and  $w_2$ , where  $l_{w_1}$  and  $r_{w_2}$  are variables of the left and right children respectively. Also, *vars* contains a variable  $o_{K+1}$  iff any of the sums is larger than  $K$ . For each variable  $o_c \in O.vars$ , value  $c$  is inserted into  $O.vals$ , along with value 0. Figure 4a illustrates an example binary tree.


Once the tree is properly constructed, the GT encoding introduces the following clauses for each non-leaf node  $O$  with children  $L$  and  $R$ :

$$\overline{t_w} \vee o_w \quad t_w \in L.vars \cup R.vars \quad (14)$$

$$\overline{l_{w_1}} \vee \overline{r_{w_2}} \vee o_w \quad l_{w_1} \in L.vars, r_{w_2} \in R.vars, w = \min(w_1 + w_2, K + 1) \quad (15)$$



(a) GT.



(b) GGT.

Figure 4: (a): binary tree of  $GT(2x_1 + 3x_2 + 4x_3 + 5x_4 + 3x_5 + 4x_6 + 6x_7 + 8x_8 \leq 10)$ . (b): binary tree of  $GGT(2x_1 + 3x_2 + 4x_3 + 5x_4 + 3x_5 + 4x_6 + 6x_7 + 8x_8 \leq 10, \{\{x_1, x_2, x_3, x_4\}, \{x_5, x_6, x_7, x_8\}\})$ .

It also introduces the unary clause:

$$\overline{a_{K+1}} \tag{16}$$

where  $A$  is the root node of the tree and  $a_{K+1} \in A.vars$  (otherwise the constraint would be trivially satisfied).

Clauses (14) enforce that the variable  $o_w$  will be set to true by UP if some child has a variable  $t_w$  set to true. Clauses (15) enforce that the variable  $o_w$  will be set to true by UP if there exists a pair of variables  $l_{w_1}, r_{w_2}$  from the children nodes that are set to true and such that  $w = \min(w_1 + w_2, K + 1)$ . Finally, Clause (16) states that the sum of the tree (i.e., the value of the left hand side expression of the PB constraint) cannot be greater than  $K$ .

We apply a minor optimisation to the GT encoding as well as the GGT encoding presented below. Variables  $a_0, \dots, a_K$  of the root node only appear in Clauses (14) and (15) and they are never negated. Therefore, we do not introduce variables  $a_0, \dots, a_K$  of the root node nor their associated clauses.

### 6.1.1. Comparison with Decomposable Negation Normal Form

The encodings described in Section 4 use decision diagrams (ROBDD in particular). Ordered decision diagrams are one target language for knowledge compilation [24]. It is natural to consider whether GT and its generalizations in this section (GGT, RGT, and RGGT) also correspond to a knowledge compilation language. The GT encoding and its generalizations (GGT, RGT, and RGGT) can be translated straightforwardly to the knowledge compilation language Decomposable Negation Normal Form (DNNF) [24]. For each non-leaf,

non-root node  $O$ , for each value  $w$  in  $O.vals - \{0\}$ ,  $w$  can be represented in DNNF as a disjunction with one disjunct for each pair of values  $w_1$  and  $w_2$  (from the left and right children respectively) where  $\min(w_1 + w_2, K + 1) = w$ . Each disjunct would be a conjunction of the two DNNF terms representing  $w_1$  and  $w_2$  respectively (reusing terms where possible). The root node is treated similarly but all values in the range  $[1, K]$  are represented by a single disjunction (since the exact value is irrelevant). For a leaf  $q_i x_i$ , value  $q_i$  is represented by  $x_i$  and value 0 by  $true$  (subsequently removed by partial evaluation). This construction is in DNNF but not d-DNNF (it does not have the *determinism* property [24]). DNNF can be encoded into SAT using the Tseitin encoding, and in this case it would be completed by asserting  $\overline{a_{K+1}}$  where  $a_{K+1}$  is the Tseitin variable for the root node value  $K + 1$ . GT has a specialised SAT encoding for the tree which is more compact than the Tseitin encoding of this DNNF formula. To adapt this DNNF formula for the RGT and RGGT encodings (described below), values are replaced with intervals.

### 6.2. Generalized Generalized Totalizer (GGT)

In our generalization of the GT encoding, we will use the same definition of the binary tree, but the leaves will be instantiated differently. Instead of introducing a leaf node for each variable of the PB constraint, we introduce a leaf node for each of the sets in the partition  $\mathcal{X}$ . The leaf node  $O$  associated with set  $X_i$  will contain a variable  $o_{q_i}$  in its *vars* attribute for each distinct coefficient  $q_i$  such that  $x_l \in X_i$ . Note that, due to Property 7,  $q_l \neq q_{l'}$  for every two distinct variables  $x_l, x_{l'} \in X_i$ . As in the GT encoding, every variable  $x_l$  is renamed as  $o_{q_l}$  and placed in  $O.vars$ . Each distinct coefficient  $q_l$  is placed in  $O.vals$ , along with value 0. Then, the GGT encoding introduces Clauses (14), (15) and (16) as in the GT encoding. Figure 4b depicts the binary tree of a GGT encoding.

Note that assuming that an AMO constraint over each set  $X_i$  is satisfied, at most one of the variables in each leaf node will be true, and therefore the encoding correctly evaluates  $\sum_{i=1}^n q_i x_i \leq K$ .

The GT encoding requires  $O(nK)$  auxiliary variables and  $O(nK^2)$  clauses, while the GGT encoding requires  $O(NK)$  auxiliary variables and  $O(NK^2)$  clauses. However, as stated in [7], this size depends on the number of unique sums of coefficients, and hence the previously given size bounds are only accurate when the number of unique sums is close to  $K$ .

### 6.3. Reduced Generalized Generalized Totalizer (RGGT)

In the GT and GGT encodings it is possible to have two values  $i, j$  of a tree node  $O$  such that the choice of  $i$  or  $j$  does not affect the semantics of the constraint: informally, it makes no difference whether the node takes value  $i$  or  $j$ . We present a new encoding that makes use of this idea, named Reduced Generalized Generalized Totalizer (RGGT). Since PB constraints are a particular case of PB(AMO) constraints, this encoding is also useful to encode PB constraints. We refer to the PB encoding as Reduced Generalized Totalizer

(RGT), and given an input PB constraint  $\sum_{i=1}^n q_i x_i \leq K$ , RGT consists of the RGGT encoding with partition  $\mathcal{X} = \{\{x_1\}, \dots, \{x_n\}\}$ .

The key idea of the RGGT is to merge sets of equivalent values into intervals and use only one variable per interval in the encoding. The RGGT was inspired by the ROBDD encoding of pseudo-Boolean constraints [5], which never contains two equivalent states by construction. In RGGT, tree nodes have an attribute *intervals* which is a sequence of integer intervals. Within each interval, the choice of value does not change the semantics of the constraint.

Take for example the constraint  $20x_1 + 30x_2 + 20x_3 + 40x_4 + 10x_5 + 20x_6 + x_7 \leq 55$  with AMO partition  $\mathcal{X} = \{\{x_1, x_2\}, \{x_3, x_4\}, \{x_5, x_6\}, \{x_7\}\}$ . Figure 5a shows a GGT for this constraint (in this case the *vals* attribute is shown in each node). First, all values of the root node (A) that are  $\leq 55$  are equivalent: they all unconditionally satisfy the constraint. Therefore the root node has two intervals:  $[0, 51], [56, \infty]$ . Next, in node B, values 0, 20, and 30 are equivalent: for all values  $c \in C.vals$ , the sums  $0 + c$ ,  $20 + c$ , and  $30 + c$  are in the same interval of A therefore the choice of 0, 20, or 30 is not significant regardless of the value of C. Collecting values into intervals is called *reduction* and parents are always reduced before children. An interval is non-trivial if it contains more than one value.

Figure 5b shows the RGGT tree created by one pass of reduction. Nodes A, B, C, and G have non-trivial intervals. Node G is a leaf node, representing term  $x_7$ . Non-trivial intervals in leaf nodes trigger changes to the PB(AMO) constraint, the details of these changes will be described later. In this case, term  $x_7$  and its AMO group are deleted because the term does not affect the satisfaction of the constraint (values 0 and 1 are in the same interval). Whenever the PB constraint is modified the entire RGGT process is repeated from scratch. Figure 5c shows the GGT tree constructed in the second pass, with 2 fewer nodes than Figure 5a. Finally, Figure 5d shows the RGGT tree produced by the second pass. The final RGGT tree would require 6 SAT variables to encode, whereas the original GGT in Figure 5a would require 11.

The reduction algorithm is presented in Algorithm 1. The first section shows the outer loop, the body of which builds a GGT tree with only *vals* attributes. No additional Boolean variables are created at this point. Then `makeChildIntervals` is called to create the *intervals* attribute of each node. After all intervals have been constructed, for each non-trivial interval of a leaf node, corresponding terms of the PB constraint are adjusted. First, coefficients are reduced to the smallest value in the interval (lines 5–8). Any value in the interval could be used; the smallest is a heuristic choice. Terms with coefficient 0 are removed, then empty AMO sets are removed from  $\mathcal{X}$  (on lines 9–10), thus restoring Property 3.

The function `makeChildIntervals` traverses the GGT tree top-down, creating the *intervals* attribute of each node  $V$  using *intervals* of the parent  $U$  and *vals* of the sibling  $W$ . Two adjacent values  $a, b$  of  $V$  are placed in the same interval iff (for each value  $w$  of the sibling node)  $a + w$  and  $b + w$  are in the same interval of the parent – i.e. the choice of  $a$  or  $b$  is not significant.



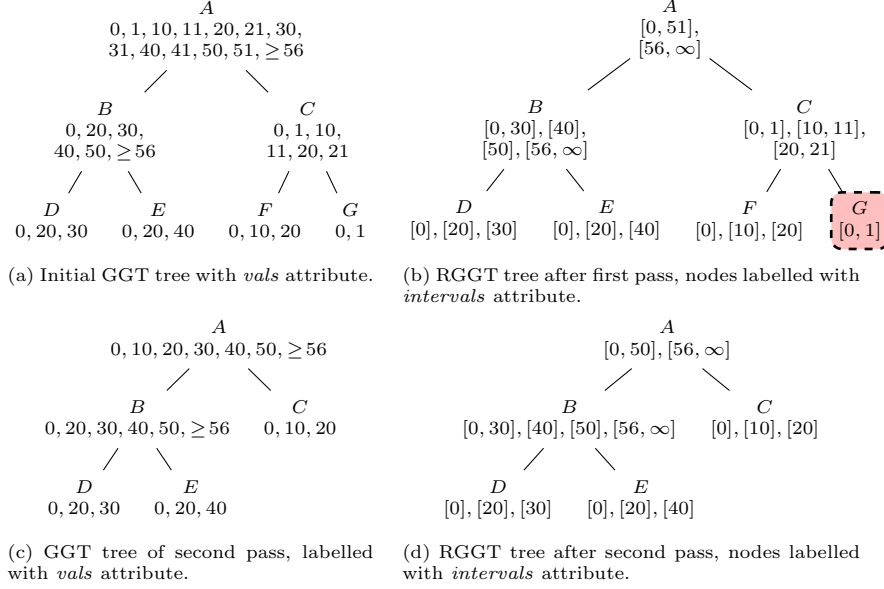


Figure 5: Two passes of the RGGT algorithm reach a stable tree for the PB(AMO)  $20x_1 + 30x_2 + 20x_3 + 40x_4 + 10x_5 + 20x_6 + x_7 \leq 55$ ,  $\mathcal{X} = \{\{x_1, x_2\}, \{x_3, x_4\}, \{x_5, x_6\}, \{x_7\}\}$ .

*Property 8.* For each non-leaf node  $O$  with children  $L$  and  $R$ , for each pair of intervals  $[a, b] \in L.intervals$  and  $[c, d] \in R.intervals$  there must exist an interval  $[e, f] \in O.intervals$  that contains  $[a + c, b + d]$ .

Note that Property 8 will hold by construction, otherwise either  $[a, b]$  or  $[c, d]$  could not have been constructed by lines 16–17 of Algorithm 1.

Finally the RGGT tree is encoded into SAT. For each non-leaf node  $U$ , for each interval  $[a, b] \in U.intervals$  where  $a > 0$ , one Boolean variable  $u_{a,b}$  is created and placed in  $U.vars$ . For each leaf node  $L$ , note that all intervals must be trivial (since Algorithm 1 has completed) but some intervals may be linked to more than one term. For each interval  $[a, a] \in L.intervals$  where  $a > 0$ , if exactly one term  $q_i x_i$  is linked to  $L$  and  $q_i = a$ , then  $x_i$  is named  $l_{a,a}$  and added to  $L.vars$ . Otherwise a new Boolean variable  $l_{a,a}$  is created, and for all terms  $q_i x_i$  linked to  $L$  where  $q_i = a$  the clause  $\neg x_i \vee l_{a,a}$  is added.

The RGGT encoding introduces the following clauses for each non-leaf node  $O$  with children  $L$  and  $R$ . The following two sets of clauses are comparable to Clauses 14 and 15 of (G)GT.

$$\overline{t_{a,b}} \vee o_{e,f} \quad t_{a,b} \in L.vars \cup R.vars, o_{e,f} \in O.vars, e \leq a, b \leq f \quad (17)$$

$$\overline{l_{a,b}} \vee \overline{r_{c,d}} \vee o_{e,f} \quad l_{a,b} \in L.vars, r_{c,d} \in R.vars, o_{e,f} \in O.vars, e \leq a + c, b + d \leq f \quad (18)$$

For the root node A the unary clause  $\overline{a_{K+1, \infty}}$  is added.

---

**Algorithm 1:** RGGT reduction algorithm.

---

**Data:** A PB(AMO) constraint  $\sum_{i=1}^n q_i x_i \leq K$ , with variable partition  $\mathcal{X}$

**Result:** A Reduced Generalized Generalized Totalizer tree

```

1 repeat
2   Build a GGT tree with root node  $A$  without vars attributes
3    $A.intervals \leftarrow [[0, \max(A.vals \setminus \{K + 1\})], [K + 1, \infty]]$ 
4   makeChildIntervals( $A$ )
5   foreach  $L$ : leaf node with non-trivial interval do
6     foreach  $[l, u]$ : non-trivial interval ( $u > l$ ) in  $L.intervals$  do
7       foreach  $q_i x_i$  linked to  $L$  where  $l < q_i \leq u$  do
8          $q_i \leftarrow l$ 
9       Delete terms  $q_i x_i$  where  $q_i = 0$ 
10      Delete empty cells of partition  $\mathcal{X}$ 
11 until No changes to PB(AMO) constraint

12 Procedure makeChildIntervals( $U$ ) is
13   foreach child  $V$  (with sibling  $W$ ) of  $U$  do
14      $V.intervals \leftarrow [[a, a] \mid a \in V.vals, a < K + 1] + [[a, \infty] \mid a \in$ 
15        $V.vals, a = K + 1]$ 
16     foreach adjacent pair of intervals  $[a, b], [c, d]$  in  $V.intervals$  in
17       ascending order do
18       if  $\forall w \in W.vals. \exists I \in U.intervals. (w + b) \in I \wedge (w + c) \in I$ 
19         then
20         Replace  $[a, b], [c, d]$  with  $[a, d]$  in  $V.intervals$ 
21     if  $V$  has children then
22       makeChildIntervals( $V$ )

```

---

#### 6.4. A Heuristic To Build GT, GGT, RGT, and RGGT Trees

Each of the totalizer encodings described above (GT, GGT, RGT, and RGGT) requires a method to construct a binary tree given a set of leaf nodes. In this section we propose a new heuristic named minRatio to build the binary tree required by these encodings. The minRatio heuristic compares the number of values of a (proposed) internal node to the product of the numbers of values of its two children, greedily minimising the ratio of these two quantities. Suppose we have three tree nodes  $C, D, E$  with value sets  $C.vals = \{0, 3, 6\}$ ,  $D.vals = \{0, 3, 6\}$ , and  $E.vals = \{0, 4, 8\}$ , and  $K = 15$ . If  $C$  and  $E$  shared a parent  $A$ , then  $A.vals = \{0, 3, 4, 6, 7, 8, 10, 11, 14\}$ . The size of  $A.vals$  is the product of  $|C.vals|$  and  $|E.vals|$ . However, if  $C$  and  $D$  shared a parent  $B$ , then  $B.vals = \{0, 3, 6, 9, 12\}$  – smaller than the product of its children. The minRatio heuristic would generate  $B$  in this case.

MinRatio works on a set  $S$  of tree nodes. Initially  $S$  contains all leaf nodes. MinRatio has completed when  $S$  contains a single node (which will be the root). At each step, two nodes  $B$  and  $C$  are removed from  $S$  and their parent

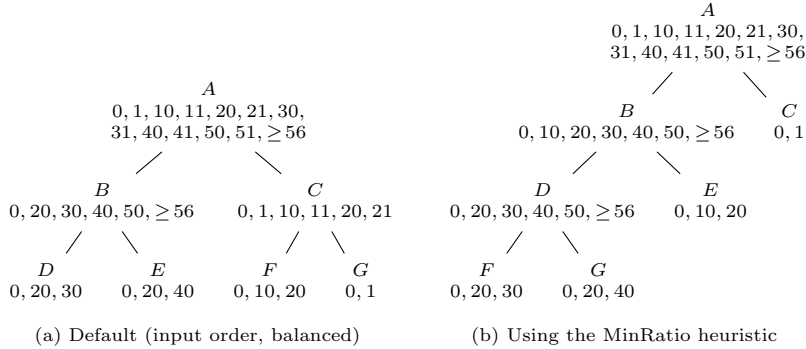


Figure 6: Comparing tree-building heuristics for  $GGT(20x_1 + 30x_2 + 20x_3 + 40x_4 + 10x_5 + 20x_6 + x_7 \leq 55, \{x_1, x_2\}, \{x_3, x_4\}, \{x_5, x_6\}, \{x_7\})$ .

$A$  is created and added to  $S$ .  $B$  and  $C$  are selected to minimise the quotient  $|A.vals|/(|B.vals| \times |C.vals|)$ . Note that  $K + 1$  is counted as a single value. In the experiments described below, we refer to the encodings using the minRatio heuristic simply as GT, GGT, RGT, and RGGT.

In order to evaluate the impact of the minRatio heuristic, in Section 11 we will compare its performance to that of a natural default heuristic with the GT and GGT encodings. The default heuristic constructs a balanced binary tree where the elements of  $\mathcal{X}$  are placed in the leaf nodes from left to right in input order. When the number of leaf nodes is not a power of two, the leftmost leaves are the ones of higher depth. In the experimental section we refer to the encodings using these heuristics as GTd and GGTD (where d indicates *default*). Our results show that the minRatio heuristic has a substantial positive effect on the size and performance of GT and GGT encodings when compared to GTd and GGTD.

The default heuristic produces balanced trees of height  $\lceil \log_2(|\mathcal{X}|) \rceil$ , whereas minRatio can produce unbalanced trees of much greater height. In fact we observed that minRatio tends to produce highly unbalanced trees. In our experiments with the Multi-Choice Multidimensional Knapsack Problem (MMKP) (described in Section 11.1), when using GGT the number of leaf nodes is  $|\mathcal{X}| = 15$  and the tree height is 14 in all cases except for one constraint of one instance in MMKP3 which has height 13. With just one exception, the trees are as high as possible (i.e. maximally unbalanced). When using GT, considering MMKP1, the trees have height at least 112 (with 150 leaf nodes). In MMKP2, trees have height at least 97 (with 150 leaves), and in MMKP3, trees have height 57 or more (with 75 leaf nodes). Minimal heights are 8 for MMKP1 and MMKP2, and 7 for MMKP3, so it is clear that minRatio is building extremely unbalanced trees. For the Combinatorial Auctions problem (see Section 11.1) PB(AMO) constraints are of different sizes so we report percentages. All GGT trees are a minimum of 74% of the maximum possible height, and GT trees are at least 38% of the maximum possible height — in both

cases substantially higher than a balanced tree. Figure 6 compares the default heuristic to minRatio on a small example with the GGT encoding. In this case, minRatio produces an unbalanced tree of maximal height by first combining the three leaf nodes with larger values (i.e. values that are multiples of 10) in one subtree. The combination of node C with any other node has a poor ratio, so node C is the last leaf node to be incorporated into the tree.

#### 6.4.1. Comparison with decision diagrams

Given that minRatio produces extremely unbalanced trees, it is natural to compare the RGT and RGGT encodings with minRatio to the BDD and MDD encodings in Section 4. Suppose that the tree for RGT/RGGT is of maximal height, and that the order of terms or AMO groups in the tree (from the deepest leaves to the root) is consistent with the order of the decision diagram. The nodes  $\top$  and  $\perp$  of the decision diagram correspond to the two intervals of the root node of the RGT/RGGT. The two encodings collapse equivalent states in very similar ways. The most significant difference is that the decision diagram may have long edges (i.e. edges that bypass one or more terms or AMO groups) whereas this is not possible in the RGT/RGGT encoding.

## 7. n-Level Modulo Totalizer Encoding

In this section we first revisit the n-Level Modulo Totalizer encoding for PB constraints from [8] and in Subsection 7.2 we provide its generalization to encode PB(AMO) constraints. Subsection 7.3 gives some construction details.

### 7.1. n-Level Modulo Totalizer

The *n-Level Modulo Totalizer* (MTO) encoding generalises the Weighted Totalizer encoding which, as stated in [8], is essentially an equivalent definition to the GT encoding that was presented in a parallel work. For the sake of readability we unify the nomenclature of the MTO encoding with the one of the GT encoding, but the overall idea of the encoding summarised here and the resulting clauses are the same as in the original work.

The MTO encoding, equally to the GT encoding, consists in building a (generalization of a) totalizer in which the root of every subtree represents the sum of the monomials associated to its leaf nodes and, in particular, the root of the whole totalizer represents the sum of the monomials in the PB constraint. The main difference with respect to GT is that the values of the nodes are represented in a mixed radix base.

A *mixed radix base* is a vector  $\Lambda = \langle \lambda_0, \dots, \lambda_{\beta-1} \rangle$  where  $\beta \in \mathbb{N}_0$  and  $\forall i \in 0..\beta - 1 : \lambda_i \in \mathbb{N}, \lambda_i > 1$ . A number  $I \in \mathbb{N}_0$  is represented in base  $\Lambda$  as  $d_\beta d_{\beta-1} \dots d_0$  where  $\forall i \in 0..\beta - 1 : 0 \leq d_i < \lambda_i$ ,  $d_\beta \in \mathbb{N}_0$ , and  $I = d_\beta \times (\lambda_0 \times \lambda_1 \times \dots \times \lambda_{\beta-1}) + \dots + d_2 \times (\lambda_0 \times \lambda_1) + d_1 \times (\lambda_0) + d_0$ .

For instance, numbers  $0, 1, \dots, 10$  in base  $\langle 3, 2 \rangle$  would be represented as  $0\ 0\ 0, 0\ 0\ 1, 0\ 0\ 2, 0\ 1\ 0, 0\ 1\ 1, 0\ 1\ 2, 1\ 0\ 0, 1\ 0\ 1, 1\ 0\ 2, 1\ 1\ 0, 1\ 1\ 1$ . Number 100 would be represented in base  $\langle 3, 2 \rangle$  as  $16\ 1\ 1$ , that is  $16 \times (2 \times 3) + 1 \times 3 + 1$ , and it would

be represented in base  $\langle 3, 2, 2 \rangle$  as 8 0 1 1, that is  $8 \times (2 \times 2 \times 3) + 0 \times (2 \times 3) + 1 \times 3 + 1$ . As a particular case, a binary base can be defined as a vector containing an infinite number of 2s, that we denote by  $\Lambda = \langle 2^* \rangle$ , and similarly a decimal base can be defined as  $\Lambda = \langle 10^* \rangle$ .

The MTO encoding builds a totalizer where the value of a node is represented in a mixed radix base, and each digit is represented with a distinct symbol as in the GT encoding. An example is given in Figure 7a. For each node  $O$  and digit  $d_h$ , with  $h \in 0..\beta - 1$ , we define a list of variables  $O_h$  of maximum length  $\lambda_h$ , which will contain a subset of the variables  $o_0^h, \dots, o_{\lambda_h-1}^h$ . We also define a list  $O_\beta$  (i.e., the list for the digit of most weight), that can contain as many variables as required to represent the value of the node. If a variable  $o_i^h$  is true, it means that the  $d_h$  digit of the value of  $O$  is at least  $i$ . The element  $o_0^h$  is always present in a list  $O_h$ , and is defined to be a 1 constant (it is omitted in Figure 7). The other variables only appear if required to represent the node values, similarly to GT, as follows:

- At leaf nodes we place the variables of the PB constraint  $P$  transformed into variables of  $O_h$  lists. Note that each variable in  $P$  can correspond to more than one variable in the  $O_h$  lists. For instance, in Figure 7a, variable  $x_4$  has coefficient 5, which is represented in base  $\langle 4, 3 \rangle$  as 1 1, and therefore corresponds to the variables  $k_1^0$  and  $k_1^1$  in the lists  $K_0$  and  $K_1$  of node  $K$ .
- At a non-leaf node  $O$  with left and right children  $L$  and  $R$ , we add to list  $O_h$  the required new variables to represent any sum of the variables in  $L_h$  and  $R_h$ . First of all, we introduce a new variable to represent the carry digit whenever the sum of two weighted digits from  $L_h$  and  $R_h$  is greater or equal than  $\lambda_h$ . We name this variable  $\gamma_O^h$ , and it is only introduced if there exist two variables  $l_i^h, r_j^h$  such that  $i + j \geq \lambda_h$ , or such that  $i + j + 1 \geq \lambda_h$  if there exists the carry digit  $\gamma_O^{h-1}$ . Otherwise, the carry variable  $\gamma_O^h$  is assumed to be false. Then,  $O_h$  contains a variable  $o_\sigma^h$ , with  $\sigma = i + j \pmod{\lambda_h}$ , for any two variables  $l_i^h, r_j^h$ . Also, if  $\gamma_O^{h-1}$  exists,  $O_h$  contains variables  $o_\sigma^h$ , with  $\sigma = i + j + 1 \pmod{\lambda_h}$ .

Note that the GT encoding can be seen as a particular case of the MTO encoding where an empty base  $\Lambda = \langle \rangle$  is used.

The encoding contains, firstly, the clauses needed to propagate the sums of values from the leaf nodes to the root node. For every non-leaf node  $O$ , with children  $L$  and  $R$ , we add the following clauses.<sup>1</sup>

Sums when the carry-in does not exist. For all  $h \in 0..\beta - 1$ ,  $l_i^h \in L_h, r_j^h \in R_h$ ,

---

<sup>1</sup>The nomenclature of this formulation has been slightly changed w.r.t. the one in [8] without affecting the final result, for the sake of a unified notation in the current work.

where  $\sigma = i + j$ :

$$\overline{l_i^h} \vee \overline{r_j^h} \vee o_\sigma^h \vee \gamma_O^h \quad \text{if } \sigma < \lambda_h \quad (19)$$

$$\overline{l_i^h} \vee \overline{r_j^h} \vee \gamma_O^h \quad \text{if } \sigma \geq \lambda_h \quad (20)$$

$$\overline{l_i^h} \vee \overline{r_j^h} \vee o_\sigma^h \text{ mod } \lambda_h \quad \text{if } \sigma > \lambda_h \quad (21)$$

Sums when the carry-in does exist. For all  $h \in 1..\beta-1$ ,  $l_i^h \in L_h$ ,  $r_j^h \in R_h$ , where  $\sigma = i + j + 1$ :

$$\overline{\gamma_O^{h-1}} \vee \overline{l_i^h} \vee \overline{r_j^h} \vee o_\sigma^h \vee \gamma_O^h \quad \text{if } \sigma < \lambda_h \quad (22)$$

$$\overline{\gamma_O^{h-1}} \vee \overline{l_i^h} \vee \overline{r_j^h} \vee \gamma_O^h \quad \text{if } \sigma \geq \lambda_h \quad (23)$$

$$\overline{\gamma_O^{h-1}} \vee \overline{l_i^h} \vee \overline{r_j^h} \vee o_\sigma^h \text{ mod } \lambda_h \quad \text{if } \sigma > \lambda_h \quad (24)$$

Sums for the uppermost digits. For all  $l_i^\beta \in L_\beta$ ,  $r_j^\beta \in R_\beta$ , where  $\sigma = i + j$ :

$$\overline{l_i^\beta} \vee \overline{r_j^\beta} \vee o_\sigma^\beta \quad (25)$$

$$\overline{\gamma_O^{h-1}} \vee \overline{l_i^\beta} \vee \overline{r_j^\beta} \vee o_{\sigma+1}^\beta \quad (26)$$

Secondly, the encoding enforces that the value of the root node is not greater than the constant  $K$  of the PB constraint. To impose this constraint,  $K$  must also be represented in base  $\Lambda$ , and we refer to the  $d_h$  digit of this representation by  $K^h$ . The clauses to be added are the following, specified from the uppermost to the lowest digit, where  $O$  is the root node of the tree:

$$\overline{o_i^\beta} \quad \forall o_i^\beta \in O_\beta, i > K^\beta \quad (27)$$

$$\overline{o_{K^\beta}^\beta} \vee \overline{o_i^{\beta-1}} \quad \forall o_i^{\beta-1} \in O_{\beta-1}, i > K^{\beta-1} \quad (28)$$

$\vdots$

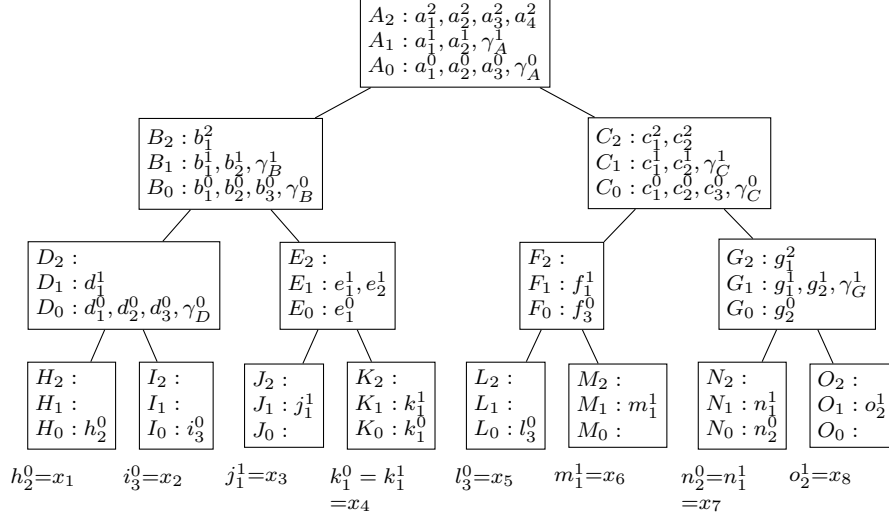
$$\overline{o_{K^\beta}^\beta} \vee \overline{o_{K^{\beta-1}}^{\beta-1}} \vee \dots \vee \overline{o_{K^2}^2} \vee \overline{o_i^1} \quad \forall o_i^1 \in O_1, i > K^1 \quad (29)$$

$$\overline{o_{K^\beta}^\beta} \vee \overline{o_{K^{\beta-1}}^{\beta-1}} \vee \dots \vee \overline{o_{K^1}^1} \vee \overline{o_i^0} \quad \forall o_i^0 \in O_0, i > K^0 \quad (30)$$


Note that this series of formulas can be stopped if we find an index  $h$  such that  $o_{K^h}^h \notin O_h$ , because  $\overline{o_{K^h}^h}$  is always true and is part of all the clauses from that point on.

### 7.1.1. Comparison with Decomposable Negation Normal Form

In Section 6.1.1 we showed that GT, GGT, RGT, and RGGT trees can be straightforwardly encoded into the knowledge compilation language Decomposable Negation Normal Form (DNNF) [24]. This is not the case for MTO: a straightforward translation into NNF does not have the *decomposable* property (that conjuncts do not share variables). Consider a non-leaf node  $O$  (with children  $L$  and  $R$ ), and digit  $h$ . The carry  $\gamma_O^h$  depends on values of  $L_h$  and  $R_h$ . The



(a) MTO.

(b) Coefficients in base  $(4, 3)$ .

(c) GMTO.

Figure 7: Example with base  $\Lambda = \langle 4, 3 \rangle$ ,  $P : 2x_1 + 3x_2 + 4x_3 + 5x_4 + 3x_5 + 4x_6 + 6x_7 + 8x_8 \leq 10$ . (a): binary tree of  $MTO(P)$ . (b): representation in base  $\Lambda$  of the coefficients of  $P$ . (c): binary tree of  $GMTO(P, \{\{x_1, x_2, x_3, x_4\}, \{x_5, x_6, x_7, x_8\}\})$ .

values of the digit  $O_h$  (represented in MTO with variables  $o_\sigma^h$ ) depend on  $L_h$  and  $R_h$  as well as the carry  $\gamma_O^h$  (see clause (22)) and as a result the NNF terms corresponding to variables  $o_\sigma^h$  may not be decomposable. As a concrete example, suppose  $\lambda_h = 3$  and there is no carry-in.  $o_2^h$  is true when  $l_1^h$  and  $r_1^h$  are both true and  $\gamma_O^h$  is false (conjunction A).  $\gamma_O^h$  is true when  $l_1^h$  and  $r_2^h$  (for example), so  $l_1^h$  is mentioned in two conjuncts of A and the NNF is not decomposable.

## 7.2. Generalized $n$ -Level Modulo Totalizer (GMTO)

The generalization from MTO to GMTO is analogous to the generalization from GT to GGT. We will instantiate the leaves of the totalizer in a way that each set of variables  $X_i \in \mathcal{X}$  is represented by a single leaf node. The leaf node  $O$  associated to set  $X_i$  will contain all variables of the form  $o_\sigma^h$  involved in the representation in the selected mixed radix base  $\Lambda$  of the coefficients  $q_j$

of variables  $x_j \in X_i$ . We denote as  $q_j^h$  the  $d_h$  digit of the representation of  $q_j$  in base  $\Lambda$ . For each set  $X_i \in \mathcal{X}$ , with leaf node name  $O$ , for each  $d_h$  digit with  $h \in 0..\beta$ , and for each value  $\sigma \in 1..\lambda_h$ :

- If there is no coefficient  $q_j$  such that  $x_j \in X_i$  and  $q_j^h = \sigma$ , then list  $O_h$  does not contain variable  $o_\sigma^h$ .
- If there is only one coefficient  $q_j$  such that  $x_j \in X_i$  and  $q_j^h = \sigma$ , then variable  $o_\sigma^h \in O_h$  is the same as variable  $x_j$ .
- If there is more than one coefficient  $q_j$  such that  $x_j \in X_i$  and  $q_j^h = \sigma$ , then variable  $o_\sigma^h \in O_h$  is a new variable, and we add the clauses  $\overline{x_j} \vee o_\sigma^h$  for any such  $q_j$ .

As in MTO,  $o_0^h$  is defined as the 1 constant and is always present for each  $h \in 0..\beta$  and for all nodes. An example is given in Figure 7c (the constants  $o_0^h$  are not represented).

Therefore, the GMTO encoding consists of the previous definition of the totalizer together with clauses (19)–(30). Note that assuming that an AMO constraint over each set  $X_i$  is satisfied, the value of the leaf node associated to  $X_i$  will be at least  $q_j$ , where  $x_j \in X_i$  is the variable that is set to true, and therefore the encoding correctly evaluates  $\sum_{i=1}^n q_i x_i \leq K$ .

The MTO encoding requires  $O(n\beta\lambda_0)$  auxiliary variables and  $O(n\beta\lambda_0^2)$  clauses (assuming  $\lambda_0 = \dots = \lambda_{\beta-1} = \lceil K^{\frac{1}{\beta}} \rceil$ ) [8], while the GMTO encoding requires  $O(N\beta\lambda_0)$  auxiliary variables and  $O(N\beta\lambda_0^2)$  clauses.

### 7.3. Construction of (G)MTO

In our experiments, in order to build the n-Level Modulo Totalizer, we follow Algorithm 5 from [8] which produces a balanced binary tree, with very small sizes and good performance in our experiments. Regarding the selection of a mixed radix base, we have implemented a greedy heuristic based on the description given in Section 4.3 of [8]. We keep adding values to  $\Lambda$  until  $\prod_{\lambda_h \in \Lambda} \lambda_h > K$ . To select the new value  $\lambda_h$  to add to  $\Lambda$ , we choose the number greater than 1 which is a divisor of the largest number of coefficients of the PB constraint and, in case of a tie, we choose the highest value. Each time we add a new value  $\lambda_h$  to  $\Lambda$ , all coefficients  $q_i$  are updated as  $q_i = \lfloor q_i / \lambda_h \rfloor$  in order to find the best divisor (base) for the next digit.

## 8. Global Polynomial Watchdog Encoding

In this section we first revisit the Global Polynomial Watchdog encoding for PB constraints from [9] and in Subsection 8.2 we provide its generalization to encode PB(AMO) constraints.



### 8.1. Global Polynomial Watchdog

The *Global Polynomial Watchdog* (GPW) encoding was introduced by Bailleux et al [9]. It uses as its basis a *polynomial watchdog* formula, denoted  $PW(P)$ , which is associated with a PB constraint  $P$ , and contains a variable named the *output* variable, denoted  $w$ . The  $PW(P)$  formula satisfies the following property:

*Lemma 2* (Lemma 1 in [9]). For any partial assignment to the variables of  $P$ , unit propagation on  $PW(P)$  assigns 1 to  $w$  if and only if this partial assignment cannot be extended to a model of  $P$ .

We first summarise how to construct the formula  $PW(P)$  and then complete the definition of the GPW encoding. The first step is to rewrite the constraint into the form  $T + \sum_{i=1}^n q_i x_i < m \cdot 2^p$ , with a strict inequality, where  $p$ ,  $T$  and  $m$  are defined as follows:  $p = \lfloor \log_2(\max_{i=1..n}(q_i)) \rfloor$  is the index of the most significant bit in the binary representation of the largest coefficient  $q_i$ , where 0 is the index of the least significant bit. In other words,  $p + 1$  is the number of bits needed to represent  $q_i$  in binary notation;  $T$  is the smallest non-negative integer such that  $K + 1 + T$  is a multiple of  $2^p$ ;  $m = (K + 1 + T)/2^p$ .

Once the constraint is rewritten to this form, a set  $B_r$  of variables of  $P$  (called *bucket*) is computed for each bit  $0 \leq r \leq p$ . We denote by  $b_r(q_i)$  the  $r$ -th bit of the binary representation of the integer  $q_i$ . Bucket  $B_r$  contains all the variables  $x_i$  such that  $b_r(q_i) = 1$ . Bucket  $B_r$  also contains a 1 constant if  $b_r(T) = 1$ .

*Example 2.* The following is the transformation to apply to the PB constraint  $2x_1 + 3x_2 + 4x_3 + 7x_4 \leq 8$ . We have  $p = 2$ , and  $T = 3$  is the smallest integer such that  $K + 1 + T = 12$  is a multiple of  $2^p$ , with  $m = 3$ . Therefore, the constraint is expressed as  $3 + 2x_1 + 3x_2 + 4x_3 + 7x_4 < 12$ . The content of buckets  $B_0$ ,  $B_1$  and  $B_2$  is illustrated in Figure 8.

The idea is to decompose each coefficient in its binary representation and sum each bit having the same weight.

The formula  $PW(P)$  can be represented as a circuit, as can be seen in Figure 8 corresponding to Example 2. We denote by  $\langle B_r \rangle$  a vector with an arbitrary order containing the elements of bucket  $B_r$ . The formula  $PW(P)$  uses two main components: the formulas  $\phi(V)$  and  $\psi(V_1, V_2)$ . The formula  $\phi(V)$  has as input a vector of Boolean variables  $V$ , and has as output a vector of  $|V|$  variables named  $U(V)$ . The formula  $\phi(V)$  enforces that  $U(V)$  is the unary representation of the sum of the input variables. The formula  $\psi(V_1, V_2)$ , has as input two vectors of variables  $V_1$  and  $V_2$ , which are the unary representation of two integers, and has as output a vector of  $|V_1| + |V_2|$  variables named  $S$ . The formula  $\psi(V_1, V_2)$  enforces that  $S$  is the unary representation of  $V_1 + V_2$ . In the definition of  $PW(P)$ , we denote by  $S_r$  the output of the  $\psi$  formula related with bucket  $B_r$ , for  $1 \leq r \leq p$ , and we define  $S_0 = U(\langle B_0 \rangle)$ . Half of the value of  $S_k$  for a weight  $2^k$ , denoted as  $S_k^{1/2}$  is integrated in the sum for weight  $2^{k+1}$ .





Figure 8: (a): circuit representation of  $PW(2x_1 + 3x_2 + 4x_3 + 7x_4 \leq 8)$ . (b): circuit representation of  $PW(2x_1 + 3x_2 + 4x_3 + 7x_4 \leq 8, \{\{x_1, x_2\}, \{x_3, x_4\}\})$ .

Then, the formula  $PW(P)$  is defined as the conjunction of these two formulas:

$$\phi(\langle B_r \rangle) \quad 0 \leq r \leq p \quad (31)$$

$$\psi(U(\langle B_r \rangle), S_{r-1}^{1/2}) \quad 1 \leq r \leq p \quad (32)$$

The GPW encoding is defined as:

$$PW(P) \quad (33)$$

$$\bar{w} \quad (34)$$

The basic idea is that the  $m$ -th bit of  $S_p$ , represented with variable  $w$ , is set to 1 by UP if the sum of the constraint is greater or equal than  $m \cdot 2^p = K + 1 + T$ . If  $w$  is set to 1 the formula is not satisfied. We build formulas  $\phi$  and  $\psi$  as in Bailleux et al [9], where  $\phi$  is encoded with a totalizer, and  $\psi$  with an adder of unary numbers.

## 8.2. Generalized Global Polynomial Watchdog (GGPW)

We define GGPW by using a *generalized polynomial watchdog formula*  $PW(P, \mathcal{X})$  instead of the original polynomial watchdog formula. Again,  $P$  is normalised to

the form  $T + \sum_{i=1}^n q_i x_i < m \cdot 2^p$  in the same way as in  $PW(P)$ . For each set  $X_i$ ,  $PW(P, \mathcal{X})$  will contain a vector of variables  $Y_i = \langle y_{i,p}, y_{i,p-1}, \dots, y_{i,0} \rangle$ .

$Y_i$  is interpreted as a binary number, where for all  $x_l \in X_i$  such that  $x_l$  is true, at least the bits corresponding to the binary representation of  $q_l$  are set to one. Therefore, when exactly one  $x_l$  is true,  $Y_i$  will be greater than or equal to  $q_l$ . The following clauses define the variables  $Y_i$ :

$$\overline{x_l} \vee y_{i,r} \quad 0 \leq r \leq p, 1 \leq i \leq N, x_l \in X_i, b_r(q_l) = 1 \quad (35)$$

In this case bucket  $B_r$ , for each bit  $0 \leq r \leq p$ , will contain variables  $y_{1,r}, y_{2,r}, \dots, y_{N,r}$ . Bucket  $B_r$  will also contain a 1 constant if  $b_r(T) = 1$ .

The formula  $PW(P, \mathcal{X})$  is defined as the conjunction of (31), (32) and (35). Some considerations can be taken into account on Clauses (35) in order to optimise the encoding:

- If there is no  $x_l \in X_i$  such that  $b_r(q_l) = 1$ , and therefore variable  $y_{i,r}$  does not appear in any clause of (35), then this variable is not created nor included in any bucket.
- If there is only one variable  $x_l \in X_i$  such that  $b_r(q_l) = 1$ , then variable  $y_{i,r}$  is the variable  $x_l$  itself, and Clause (35) is not added for  $y_{i,r}$ .
- Otherwise,  $y_{i,r}$  is indeed a new variable and Clause (35) is added.

Figure 8 contains a circuit representation of  $PW(P, \mathcal{X})$ .

The GGPW encoding is defined by:

$$PW(P, \mathcal{X}) \quad (36)$$

$$\overline{w} \quad (37)$$

Just as with the other newly introduced encodings, given an assignment that satisfies an AMO constraint over each  $X_i \in \mathcal{X}$ , this encoding represents the PB constraint  $\sum_{i=1}^n q_i x_i \leq K$  in a more compact way.

The GPW encoding introduces  $O(n \log(n) \log(q_{max}))$  auxiliary variables and  $O(n^2 \log(n) \log(q_{max}))$  clauses, while the GGPW introduces  $O(N \log(N) \log(q_{max}))$  auxiliary variables and  $O(N^2 \log(N) \log(q_{max}))$  clauses, where  $q_{max} = \max_{i=1}^n q_i$ . This follows from the fact that a totalizer  $\phi$  with  $n$  input variables requires  $O(n \log(n))$  auxiliary variables and  $O(n^2 \log(n))$  clauses, and an adder  $\psi$  of unary numbers with  $n$  input variables requires  $O(n)$  auxiliary variables and  $O(n^2)$  clauses; see [9].

## 9. Local Polynomial Watchdog Encoding

In this section we briefly revisit the Local Polynomial Watchdog encoding for PB constraints from [9] and in Subsection 9.2 we provide its generalization to encode PB(AMO) constraints. For this generalization, we provide an efficient implementation that reuses many variables and clauses to obtain smaller formulas.

### 9.1. Local Polynomial Watchdog

In Section 10 we give details of the propagation strength of each encoding, and note that GPW does not have the same propagation strength as BDD, GT, RGT, and SWC. The Local Polynomial Watchdog (LPW) does have the same propagation strength as the others mentioned, at the cost of additional variables and clauses (specifically by including a different PW formula for each variable of the PB constraint). The definition of  $PW(P)$  is the same as in GPW, that is Constraints (31) and (32). Then, LPW is defined as:

$$PW(P[x_i]) \quad 1 \leq i \leq n \quad (38)$$

$$\overline{w(P[x_i])} \vee \overline{x_i} \quad 1 \leq i \leq n \quad (39)$$

where  $P[x_i]$  is the resulting PB constraint of setting  $x_i$  to 1 in  $P$ , and  $w(P[x_i])$  is the output variable of  $PW(P[x_i])$ .

### 9.2. Generalized Local Polynomial Watchdog (GLPW)

Similarly to LPW, we can define the GLPW to be a GAC encoding for PB(AMO) constraints. We can do that by simply modifying Formulas (38) and (39) as follows:

$$PW(P[x_l, X_i], \mathcal{X} \setminus \{X_i\}) \quad X_i \in \mathcal{X}, x_l \in X_i \quad (40)$$

$$\overline{w(P[x_l, X_i])} \vee \overline{x_l} \quad X_i \in \mathcal{X}, x_l \in X_i \quad (41)$$

where  $P[x_l, X_i]$  is the PB constraint resulting from setting  $x_l$  to 1 in  $P$  and setting any other variable  $x_{l'} \in X_i$  to 0, and  $w(P[x_l, X_i])$  is the output variable of  $PW(P[x_l, X_i], \mathcal{X} \setminus \{X_i\})$ .

The LPW introduces  $O(n^2 \log(n) \log(q_{max}))$  auxiliary variables and  $O(n^3 \log(n) \log(q_{max}))$  clauses, while the GLPW encoding introduces  $O(nN \log(N) \log(q_{max}))$  auxiliary variables and  $O(nN^2 \log(N) \log(q_{max}))$  clauses, where  $q_{max} = \max_{i=1}^n q_i$ . Basically, the sizes are multiplied by  $n$  w.r.t. those of GPW and GGPW, since we encode  $n$  polynomial watchdog formulas, one for each variable in  $P$ .

Bailleux et al [9] stated that the LPW encoding can be compacted by reusing some of the components of the different polynomial watchdogs introduced by Constraints (38), and some hints were given, but no method was detailed nor evaluated for this purpose. Here we provide implementation techniques to obtain small GLPW encodings.

Firstly, we build the totalizers for formulas  $\phi$  in a way that (i) any two formulas  $PW(P[x_l, X_i], \mathcal{X} \setminus \{X_i\})$ ,  $PW(P[x_{l'}, X_i], \mathcal{X} \setminus \{X_i\})$  contain exactly the same set of totalizers, and (ii) the totalizers of any two formulas  $PW(P[x_l, X_i])$ ,  $PW(P[x_{l'}, X_j])$ , with  $i \neq j$ , share most of the nodes. This is explained in Section 9.2.1.

Secondly, we reuse auxiliary variables and clauses while constructing formulas (40). This can help not only in GLPW but also to encode PB constraints without AMOs, i.e., in LPW encodings. Variables and clauses are reused globally across the different generalized local polynomial watchdog formulas for all  $X_i \in \mathcal{X}$ ,  $x_l \in X_i$ . This is explained in Subsection 9.2.2.

### 9.2.1. Structure of Totalizers

The difference between the two formulas  $PW(P[x_l, X_i], \mathcal{X} \setminus \{X_i\})$  and  $PW(P[x_{l'}, X_i], \mathcal{X} \setminus \{X_i\})$  is just the value of the right hand side  $K$  in the PB constraint.

*Example 3.* Consider the PB constraint and partition:

$$(P, \mathcal{X}) = (2x_1 + 3x_2 + 4x_3 + 7x_4 \leq 8, \{\{x_1, x_2\}, \{x_3, x_4\}\})$$

$$\begin{aligned} \text{With } x_1 = 1: & \quad PW(P[x_1, X_1], \mathcal{X} \setminus \{X_1\}) \equiv PW(4x_3 + 7x_4 \leq 6, \{\{x_3, x_4\}\}) \\ \text{With } x_2 = 1: & \quad PW(P[x_2, X_1], \mathcal{X} \setminus \{X_1\}) \equiv PW(4x_3 + 7x_4 \leq 5, \{\{x_3, x_4\}\}) \\ \text{With } x_3 = 1: & \quad PW(P[x_3, X_2], \mathcal{X} \setminus \{X_2\}) \equiv PW(2x_1 + 3x_2 \leq 4, \{\{x_1, x_2\}\}) \\ \text{With } x_4 = 1: & \quad PW(P[x_4, X_2], \mathcal{X} \setminus \{X_2\}) \equiv PW(2x_1 + 3x_2 \leq 1, \{\{x_1, x_2\}\}) \end{aligned}$$

The only difference between  $P[x_1, X_1]$  and  $P[x_2, X_1]$  is the right hand side constant in  $P[x_1, X_1]$  and  $P[x_2, X_1]$ . The same happens with  $P[x_3, X_2]$  and  $P[x_4, X_2]$ .

The only thing that prevents the contents of the buckets of  $PW(P[x_l, X_i], \mathcal{X} \setminus \{X_i\})$ ,  $PW(P[x_{l'}, X_i], \mathcal{X} \setminus \{X_i\})$ , and hence of the  $\phi$  formulas, to be the same, are the possible 1 constants introduced by  $T$  in the normalisation step. However, since formulas  $\phi$  sort the buckets, instead of putting a 1 constant in the input of  $\phi$  we can append a 1 constant directly to the first position of the output of  $\phi$  when required, i.e., the 1 constant now goes directly into formula  $\psi$  (see first step of Figure 9a). This way, the content of the buckets for  $PW(P[x_l, X_i], \mathcal{X} \setminus \{X_i\})$ ,  $PW(P[x_{l'}, X_i], \mathcal{X} \setminus \{X_i\})$  is exactly the same regardless of whether there is an input 1-constant. Therefore, we will only introduce one formula  $\phi$ , encoded as a totalizer, for each bit  $r$  and for each  $X_i \in \mathcal{X}$ , instead of one totalizer for each bit  $r$  and variable  $x_l$ . In fact we can go one step further, and since formulas  $\psi$  behave as mergers of two sorted lists, we can move the 1 constant directly to the output of formula  $\psi$  (see second step of Figure 9a). This second move of the 1 constant reduces to a small extent the sizes of formulas  $\psi$ , and most importantly, makes the content of  $S_r$  independent of any 1 constant in bucket  $B_r$ . This is illustrated in Figure 9, where  $S_r$  is exactly the same when  $b_r(T) = 1$  and when  $b_r(T) = 0$ . This lets us reuse the output of formulas  $\psi$  as will be explained in Section 9.2.2, thus saving variables and clauses.

Once the constants are moved to the output of  $\psi$ , for any two formulas  $PW(P[x_l, X_i], \mathcal{X} \setminus \{X_i\})$ ,  $PW(P[x_{l'}, X_j], \mathcal{X} \setminus \{X_j\})$  with  $X_i \neq X_j$ , their corresponding buckets  $B_r$  for a bit  $r$  only differ in one variable, that is  $y_{j,r}$  in the first formula and  $y_{i,r}$  in the second. From the perspective of a totalizer, this means that only one leaf node changes between the trees of the two totalizers for  $PW(P[x_l, X_i], \mathcal{X} \setminus \{X_i\})$  and  $PW(P[x_{l'}, X_j], \mathcal{X} \setminus \{X_j\})$ . Therefore, we can reuse most of the nodes of the totalizers, and their associated auxiliary variables, when constructing the formulas. This is illustrated in Figure 10, where we consider a PB(AMO) with  $N = 8$ , and hence GLPW requires introducing 8 totalizers of 7 leaf nodes. We can see that the proposed implementation requires 30 distinct totalizer nodes in total, while a naïve implementation without reusing nodes requires 104 distinct nodes.

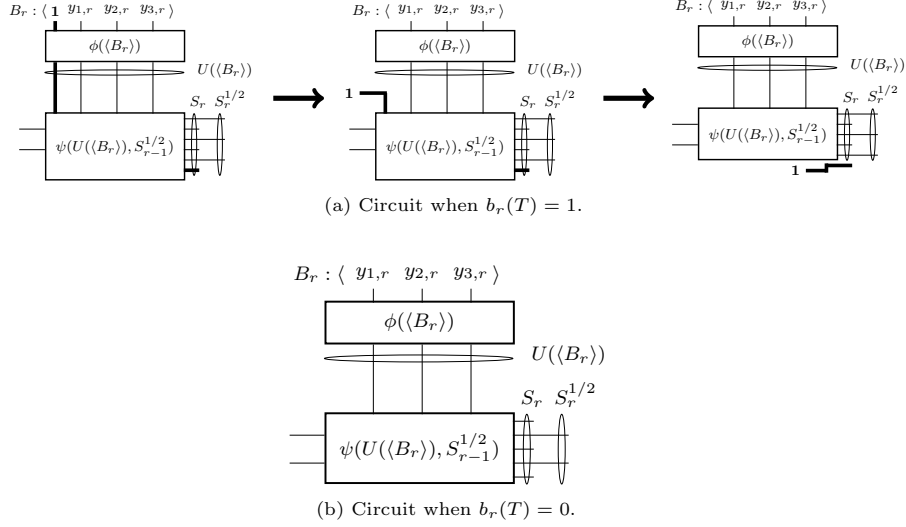


Figure 9: (a): example of moving the 1 constants introduced by  $T$  from the bucket to the output of formulas  $\psi$ . (b): representation of the corresponding formulas  $\phi$  and  $\psi$  when no 1 constant is introduced by  $T$ .




Figure 10: Example of reusing of totalizers of formula  $\phi$  for some bucket  $B_r$ , when only one leaf changes every time. In this example,  $N = 8$  and we assume that all variables  $y_{i,r}$  are present in  $B_r$ , for  $i \in 1..8$ . The tree on the top is never created but illustrates the virtual totalizer where no leaf node is missing. The nodes (variables) and edges (clauses) in boldface appear only in one totalizer, and the others are reused in most totalizers.

With the proposed implementation of GLPW, the number of required auxiliary variables is  $O(nN \log(q_{max}))$ , and the number of clauses is  $O(nN^2 \log(q_{max}))$ . This follows from the fact that we are only introducing one formula  $\phi$  for each  $i \in 1..N$  and for each bit  $r \in 0..p$ . This reduces the total asymptotic size of formulas  $\phi$ , but not the one of formulas  $\psi$  which become the dominating components of this size.

### 9.2.2. Reuse of Auxiliary Variables

Now we describe how to reuse auxiliary variables and clauses among the different generalized polynomial watchdog formulas (40). The degree of reusability is especially high in the polynomial watchdog formulas generated for variables belonging to the same group  $X_i$ . To a minor degree, we can also reuse parts of generalized polynomial watchdog formulas related to variables belonging to distinct groups. We maintain three maps ( $MY$ ,  $MT$ ,  $MM$ ) that store variables that can be reused if needed.

*Map MY.* This map is used to store the variables  $y_{i,r}$  that are defined by clauses (35), but avoiding having to introduce equivalent auxiliary variables. The key to an entry of  $MY$  is the set of variables  $V$  that logically imply  $MY(V)$  according to clauses (35). The following example shows a case in which variables  $y_{i,r}$  are reused.

*Example 4.* Consider a set  $X_i \in \mathcal{X}$  such that  $X_i = \{x_1, x_2, x_3\}$  and the associated coefficients are  $q_1 = 3, q_2 = 7, q_3 = 11$ . The binary representations of the coefficients are 0011, 0111, and 1011, respectively. We need at least 4 bits to represent these numbers, and therefore we need 4 buckets. Looking at the bit of most weight ( $r = 3$ ), the set of variables which have  $b_r(q) = 1$  is just  $\{x_3\}$ . In this case, since it is only one variable, no new variable is added nor clauses (35), and we set  $MY(\{x_3\}) = x_3$ . Similarly, looking at the following bit  $r = 2$ , we set  $MY(\{x_2\}) = x_2$ . For  $r = 1$  we introduce an auxiliary variable  $y_{i,1}$ , add clauses  $\overline{x_1} \vee y_{i,1}$ ,  $\overline{x_2} \vee y_{i,1}$  and  $\overline{x_3} \vee y_{i,1}$ , and define  $MY(\{x_1, x_2, x_3\}) = y_{i,1}$ . For the least significant bit  $r = 0$ , we can reuse variable  $MY(\{x_1, x_2, x_3\})$  instead of introducing variable  $y_{i,0}$  and the corresponding extra clauses.

*Map MT.* This map is used to implement the totalizers of formulas  $\phi$ , avoiding building a subtree twice with the same leaf nodes. Given as key a set of variables corresponding to leaves of a totalizer, it returns the variables of the root of the totalizer.

*Map MM.* This map is used to store the output formulas  $\psi$ , i.e., the vectors  $S_r$ , and reuse them when possible. The key of an entry is a triplet  $(S_{r-1}, c, U(\langle B_r \rangle))$ , where:

- $S_{r-1}$  is the vector of output variables of formula  $\psi$  for the previous bit.
- $c = 1$  if  $S_{r-1}^{1/2}$  was defined considering  $b_{r-1}(T) = 1$ , and 0 otherwise (see Figure 9). That is,  $c = 1$  if the left input of  $\psi$  are the even positions of  $S_{r-1}$ , and  $c = 0$  if the odd positions must be taken. See for instance that circuits  $\phi$  and  $\psi$  are identical in Figure 9a (right) and Figure 9b. The presence or not of a 1 constant (i.e., of  $c$ ) determines which bits  $S_r^{1/2}$  (either odd or even) are taken as input of the next formula  $\psi$  for bit  $r + 1$ .
- $U(\langle B_r \rangle)$  is the output of  $\phi(\langle B_r \rangle)$ , without including the 1 constant if present.

Note that an entry of  $MM$  contains all the bits of  $S_r$ . The choice of even or odd bits from  $S_r$  used as input of  $\phi$  for bucket  $r + 1$  can be different at each generalized local polynomial watchdog formula.

It is natural to ask whether preprocessing could achieve the same or similar reuse of variables and clauses when given a naïve GLPW encoding. None of the common subexpression elimination algorithms implemented in Savile Row [15], which is the constraint reformulation tool that we use in our experiments, would extract the common subtrees because the auxiliary variables have distinct names in each subtree. In SAT preprocessing, Equivalent Literal Substitution (ELS) [25] examines the set of binary clauses and detects sets of literals that take the same truth value in all solutions. The variables reused through maps  $MY$ ,  $MT$ , and  $MM$  do not necessarily meet the definition of equivalence. Consider the variables  $y_{i,1}$  and  $y_{i,0}$  in Example 4: when all  $x$  variables are false,  $y_{i,1}$  and  $y_{i,0}$  are free and so can take different values, therefore they are not equivalent literals.

## 10. Propagation Properties

In this section we review the principal propagation properties that SAT encodings can have, and we describe the propagation characteristics of the new generalized encodings. In summary, all the presented PB(AMO) encodings preserve the propagation strength of their counterpart PB encodings.

*Unit propagation* (UP) is the main propagation mechanism used in modern SAT solvers. It is based on the principle that if a clause contains a single literal (i.e., under a given assignment, all literals but one are false), then every model must make that literal true. Hence, the assignment can be extended with this literal. The principal consistency notions that a SAT encoding  $E$  of a constraint  $C$  can achieve with UP are the following:

- $E$  is said to be *consistency checker* (CC) when: given any partial assignment  $A$ , if  $A$  cannot be extended to a model of  $C$ , then unit propagating  $A$  on  $E$  will falsify some clause.
- $E$  is said to be *generalized arc consistent* (GAC) when: given any partial assignment  $A$ , if a variable  $x$  in  $C$  is true (respectively false) in every extension of  $A$  satisfying  $C$ , then unit propagating  $A$  on  $E$  will extend  $A$  to  $A \cup \{x\}$  (respectively  $A \cup \{\bar{x}\}$ ).

Although we do not consider them in this work, it is worth mentioning that there also exist stronger consistency notions that do not relate to the encoded constraint  $C$  but are defined on all variables of  $E$ , including auxiliary variables:

- A formula  $E$  is said to be *unit refutation complete* (URC) when: given any partial assignment  $A$  that cannot be extended to a model of  $E$ , then unit propagating  $A$  on  $E$  will falsify some clause.
- A formula  $E$  is said to be *propagation complete* (PC) when: given any partial assignment  $A$ , if a variable  $x$  in  $E$  is true (respectively false) in



Table 1: Propagation strength of the different encodings.

<b>BDD</b>	<b>SWC</b>	<b>GT/RGT</b>	<b>MTO</b>	<b>GPW</b>	<b>LPW</b>
<b>MDD</b>	<b>GSWC</b>	<b>GGT/RGGT</b>	<b>GMTO</b>	<b>GGPW</b>	<b>GLPW</b>
GAC	GAC	GAC	-	CC	GAC

every extension of  $A$  satisfying  $E$ , then unit propagating  $A$  on  $E$  will extend  $A$  to  $A \cup \{x\}$  (respectively  $A \cup \{\bar{x}\}$ ).

Note that any GAC encoding is also CC, any PC encoding is both URC and GAC, and any URC encoding is also CC.

When clear from the context, these properties are usually attributed to the encoding methods. For instance, we say that the MDD encoding is GAC, meaning that all formulas obtained with the MDD encoding method are GAC, and we say that GMTO is not CC, meaning that there exist PB(AMO) constraints whose encoding obtained with the GMTO encoding method is not CC.

The propagation strength of the encodings considered in this paper is summarised in Table 1. Previous works have already proved the propagation strength for encoding PB constraints of BDD [5], SWC [6], GT [7], and GPW and LPW [9]. Regarding RGT, it can be easily proved that it is also GAC using the same reasoning that was applied in the GAC proof of GT in [7]. The key idea of the proof is that UP propagates to a parent node the sum of the values reached by the two children given the current assignment. This also happens in RGT, where Property 8 ensures that any sum value is always represented by a corresponding interval.

**Theorem 1.** *RGT is a GAC encoding of PB constraints.*

We are however interested in the propagation strength of encodings of PB(AMO) constraints. There are many GAC encodings of AMO constraints, for instance the pairwise encoding [26]. Since we consider monotonic PB(AMO) constraints, it is straightforward to see that the previous PB encodings can be used in conjunction with encodings of AMO constraints to obtain PB(AMO) encodings with the same propagation strength.

*Lemma 3.* Let  $\mathcal{P}$  be a PB(AMO) constraint of the form  $P \wedge M_1 \wedge \dots \wedge M_N$ . Let  $E$  be any GAC (respectively CC) encoding of  $P$ . Then the conjunction of  $E$  with a GAC (resp. CC) encoding of  $M_1 \wedge \dots \wedge M_N$  is a GAC (resp. CC) encoding of  $\mathcal{P}$ .

The generalized PB(AMO) encodings considered in this paper maintain the propagation properties of their counterpart non-generalized encodings. This is proved in Theorem 2. The intuition of its proof is that each one of the formulas obtained with MDD, GSWC, GGT, RGGT, GGPW and GLPW for an input  $(P, \mathcal{X})$  can be roughly seen as a union of many encodings obtained respectively with BDD, SWC, GT, RGT, GPW and LPW. In particular, they encode all possible PB constraints obtained by keeping in  $P$  just one variable for each  $X_i$ . For instance,  $GSWC(q_1x_1 + q_2x_2 + q_3x_3 + q_4x_4 \leq K, \{\{x_1, x_2\}, \{x_3, x_4\}\})$  is

roughly an efficient union of  $SWC(q_1x_1 + q_3x_3 \leq K)$ ,  $SWC(q_1x_1 + q_4x_4 \leq K)$ ,  $SWC(q_2x_2 + q_3x_3 \leq K)$  and  $SWC(q_2x_2 + q_4x_4 \leq K)$ , that shares auxiliary variables and clauses. Therefore, GSWC contains all the required clauses to enforce GAC on the possible PB constraints that satisfy the AMOs.

**Theorem 2.** *Let  $\mathcal{P}$  be a PB(AMO) constraint of the form  $P \wedge M_1 \wedge \dots \wedge M_N$ , where  $\mathcal{X} = \{X_1, \dots, X_N\}$  is a partition of the variables in  $P$  such that  $X_i = \text{scope}(M_i)$ . The following hold:*

- *Let  $E$  be any encoding among MDD, GSWC, GGT, RGGT and GLPW. Then the conjunction of  $E(P, \mathcal{X})$  with a GAC encoding of  $M_1 \wedge \dots \wedge M_N$  is a GAC encoding of  $\mathcal{P}$ .*
- *The conjunction of GGPW( $P, \mathcal{X}$ ) with a CC encoding of  $M_1 \wedge \dots \wedge M_N$  is a CC encoding of  $\mathcal{P}$ .*

*Proof.* We prove the theorem for GSWC, but the proof for the other encodings is analogous. Let  $S$  denote the conjunction of  $GSWC(P, \mathcal{X})$  with a GAC encoding of  $M_1 \wedge \dots \wedge M_N$ . Let  $A$  be a partial assignment to the variables of  $S$  which is extendible to a satisfying assignment of  $\mathcal{P}$ . Therefore, no AMO constraint  $M_i$  is violated under  $A$ . We need to show that for every variable  $x$  of  $\mathcal{P}$  such that  $x$  is not assigned in  $A$ , if  $A \cup \{x\}$  cannot be extended to a satisfying assignment of  $\mathcal{P}$ , then  $x$  is set to false by unit propagating  $A$  on  $S$  (note that  $A \cup \{\bar{x}\}$  can always be extended to a satisfying assignment due to decreasing monotonicity, so we don't need to consider this case). W.l.o.g., assume that  $x_1 \in X_1$  is such variable. If  $A \cup \{x_1\}$  cannot be extended to a satisfying assignment of  $M_1 \wedge \dots \wedge M_N$  then, by the assumption that  $S$  contains a GAC encoding of  $M_1 \wedge \dots \wedge M_N$ , we have that  $x_1$  is set to false by unit propagation. Assume now the contrary, i.e., that  $A \cup \{x_1\}$  can be extended to an assignment satisfying the AMOs. In this case, the reason why UP should set  $x_1$  to false is that  $A \cup \{x_1\}$  cannot be extended to satisfy  $P$ . Since  $A \cup \{x_1\}$  does not violate  $M_1 \wedge \dots \wedge M_N$ , at most one variable in  $X_i$  is true in  $A$ , for  $2 \leq i \leq N$ , and no variable in  $X_1$  is true in  $A$ . Let us construct a PB constraint  $P'$  from  $P$  by picking one variable  $x_{j_i}$  from each set  $X_i$ ,  $2 \leq i \leq N$ , as follows: if  $X_i$  contains a variable which is true in  $A$ , then this is the variable to be picked up from  $X_i$ , otherwise pick up any variable. We define  $P' : q_1x_1 + \sum_{i=2}^N q_{j_i}x_{j_i} \leq K$ . Since  $P'$  contains all variables of  $P$  which are true in  $A$ , and due to the monotonicity of  $P$ , we have that  $q_1x_1 + \sum_{i=2}^N q_{j_i}x_{j_i} \leq K$  is equisatisfiable to  $\sum_{i=1}^n q_i x_i \leq K$  under the assignment  $A \cup \{x_1\}$ , and therefore  $A \cup \{x_1\}$  cannot be extended to a model of  $P'$ . It is not hard to see that  $GSWC(P, \mathcal{X})$  contains all clauses of  $SWC(P')$ . Since the SWC encoding is GAC,  $S$  contains all the clauses required to set  $x_1$  to false by UP.  $\square$

An alternative proof that MDD is GAC was presented in [12]. Note that LPW as well as GLPW are the only GAC encodings of polynomial size considered in this paper. The other ones have pseudo-polynomial size complexity.

It was stated in [9] that the GPW encoding is not GAC. Since GGPW is a generalization of GPW, it is also not GAC. For instance, one could consider

a PB(AMO) constraint where all AMO constraints have size 1. In this case, GPW and GGPW are identical. The same happens with MTO, which is also not GAC as stated in [8], and therefore GMTO is not GAC either. In fact, MTO is not even a CC encoding of PB constraints, and therefore GMTO is not a CC encoding of PB(AMO) constraints.

**Theorem 3.** *MTO is not a CC encoding of PB constraints.*

*Proof.* Consider the PB constraint and MTO encoding of Figure 7a. Let partial assignment  $A = \{x_6, x_8\}$  (i.e.  $A = \{m_1^1, o_2^1\}$ ). Clearly  $A$  cannot be extended to a model of  $P$  since  $4x_6 + 8x_8 > 10$ . The only generated clauses containing  $m_1^1$  and  $o_2^1$  are:

$$(19) : \overline{l_0^1} \vee \overline{m_1^1} \vee f_1^1 \vee \gamma_F^1 \quad (19) : \overline{n_0^1} \vee \overline{o_2^1} \vee g_2^1 \vee \gamma_G^1 \quad (20) : \overline{n_1^1} \vee \overline{o_2^1} \vee \gamma_G^1$$

where we paint in red the literals of type  $\gamma_O^h$ ,  $l_0^h$  and  $r_0^h$  which are trivially false and are not included in these clauses by construction. Similarly we do not include clauses of type (22) to (24) which are satisfied due to  $\gamma_O^{h-1}$  being trivially false.

Unit propagation only assigns literal  $f_1^1$ , due to clause  $\overline{m_1^1} \vee f_1^1$ . The only remaining clauses containing variable  $f_1^1$  are:

$$(19) : \overline{f_1^1} \vee \overline{g_0^1} \vee c_1^1 \vee \gamma_C^1 \quad (22) : \overline{\gamma_C^0} \vee \overline{f_1^1} \vee \overline{g_0^1} \vee c_2^1 \vee \gamma_C^1 \\ (19) : \overline{f_1^1} \vee \overline{g_1^1} \vee c_2^1 \vee \gamma_C^1 \quad (23) : \overline{\gamma_C^0} \vee \overline{f_1^1} \vee \overline{g_1^1} \vee \gamma_C^1 \\ (20) : \overline{f_1^1} \vee \overline{g_2^1} \vee \gamma_C^1 \quad (23) : \overline{\gamma_C^0} \vee \overline{f_1^1} \vee \overline{g_2^1} \vee \gamma_C^1 \quad (24) : \overline{\gamma_C^0} \vee \overline{f_1^1} \vee \overline{g_2^1} \vee c_1^1$$

No clause is falsified and no other literal is unit-propagated.  $\square$

*Corollary 1.* The conjunction of GMTO with an encoding of the AMO constraints is not a CC encoding of PB(AMO) constraints.

Although GMTO is the encoding with the worst propagation properties among all the encodings considered in this paper, in Section 11 it can be observed that it produces the smallest formulas by far in the selected benchmark sets. As a result, the GMTO encoding ends up providing the best performance for some benchmarks.

## 11. Experiments

In this section we report on a comparison between the different encodings for PB(AMO) constraints, and also between those and the classical encodings for PB constraints. For this purpose, we solve problems containing PB constraints as well as sources of incompatibility between their variables, i.e. mutexes. We provide empirical evidence of the usefulness of taking into account existing AMO constraints when encoding PB constraints. We show that all new PB(AMO) encodings perform significantly better than their counterpart PB encodings, based on executions of two different SAT solvers which are representative of the state-of-the-art. We also show the good performance of the reduced generalized totalizer encodings (RGT / RGGT). We study in detail the impact of the RGGT reduction algorithm on the size of the generalized totalizers.

### 11.1. Experimental Setting

On one hand we consider two problems consisting essentially of conjunctions of PB constraints and AMO constraints: the Multi-Choice Multidimensional Knapsack Problem (MMKP) and Combinatorial Auctions (CA). For these problems, the AMO constraints between Boolean variables are explicitly stated in the problem definition, and hence we can directly define the PB(AMO) constraints, as we describe in Subsections 11.1.1 and 11.1.2.

On the other hand we consider three challenging problems that are not essentially a set of PB(AMO) constraints but where PB constraints play an important role as well. Namely, we consider two extensions of the highly studied Resource-Constrained Project Scheduling Problem (RCPSP): Multi-mode RCPSP (MRCPSP) [27] and RCPSP with Time-Dependent Resource Capacities and Requests (RCPSP/t) [28]. We also consider the Nurse Scheduling Problem (NSP) [29]. These three problems have been modelled with the constraint programming modelling language Essence Prime [30]. We automatically detect AMO constraints and generate the SAT formulas using Savile Row [31], as described in [14]. In brief, Savile Row detects pairs of Boolean variables that cannot be true at the same time (mutexes) by using constraint programming propagation schemes, and then builds disjoint AMO constraints from those mutexes. The detected AMO constraints are translated to SAT with the 2-product encoding [32]. In Subsections 11.1.3, 11.1.4 and 11.1.5 we provide a small description of such problems together with a discussion on the source of the mutexes, which in most cases are implicit, i.e. not stated by an explicit constraint in the model.

We have chosen MMKP, which is essentially a set of PB(AMO) constraints, in order to craft three different benchmark sets with different parameters, with the aim of showing which encodings are better suited for different kinds of PB(AMO) constraints (e.g. with different numbers of variables, coefficient values or AMO sizes). For the other problems, we have used representative benchmarks from the literature. We have considered the decision version of CA, MRCPSP, RCPSP/t and NSP, which are optimisation problems. This means that for maximisation problems (CA), and respectively minimisation problems (the others), we set a lower bound (resp. upper bound) on the objective function. In order to obtain both satisfiable and unsatisfiable instances, we consider each instance twice with two different bounds: a bound equal to the best known objective (in most cases the optimum), and a bound equal to the best plus one (in CA) or minus one (in the others). Table 2 summarises the properties of each benchmark set.

Overall, we consider all the problems that have been studied in the previous works related to PB(AMO) constraints [14, 12]. All of them are NP-hard. In each case the AMO constraints are encoded using only binary clauses and a small number of additional variables. The AMO encodings are invariant when comparing the various PB and PB(AMO) encodings, so the clauses and variables of the AMO encodings are not included in the reported formula sizes.

Table 2: Summary of each set, containing in this order: number of instances; approximate number of PB constraints in each instance; approximate average number or range of number of variables in a PB constraint; average size of the AMOs; range of values of the coefficients.

	set	PB count	PB size	AMO size	coef.
<b>MMKP1</b>	500	10	150	10	[1,1000]
<b>MMKP2</b>	500	10	150	10	[1,60]
<b>MMKP3</b>	400	50	75	5	[1,10]
<b>CA</b>	340	1	[70,200]	8	[1,2500]
<b>MRCPSpj30</b>	1004	68	29	4	[1,10]
<b>RCPSP/Tj120</b>	720	624	158	9	[1,10]
<b>NSP</b>	400	1	700	4	[1,4]

### 11.1.1. MMKP Instances

These have been generated using the MMKP instance generator from [33]. Each instance is defined by four parameters:  $L$  is the number of PB constraints,  $N$  is the number of AMO constraints,  $M$  is the number of Boolean variables in each AMO constraint, and  $Q$  is the maximum coefficient of a variable in a PB constraint. The variables of the AMO constraints are disjoint, so there are a total of  $n = N \cdot M$  Boolean variables in each instance. The PB constraints contain all  $n$  variables. The  $j$ -th variable in the  $i$ -th AMO constraint is named  $x_{i,j}$ . The coefficients in the PB constraints are generated uniformly and independently at random in the range  $[1, Q]$ . The resulting instance has the following constraints:

$$\sum_{i=1}^N \sum_{j=1}^M q_{i,j,k} \cdot x_{i,j} \leq K_k \quad 1 \leq k \leq L \quad (42)$$

$$\sum_{j=1}^M x_{i,j} \leq 1 \quad 1 \leq i \leq N \quad (43)$$

$$\sum_{j=1}^M x_{i,j} \geq 1 \quad 1 \leq i \leq N \quad (44)$$

The conjunction of PB and AMO constraints (42) and (43) is not a hard problem, since a trivial solution is to set all variables  $x_{i,j}$  to 0. For this reason we add *at-least-one* constraints (44), requiring that at least one variable in each AMO group is set to true. This way, (43) together with (44) form an *exactly-one* constraint. When generating the SAT formulas, constraint (43) is encoded with the Regular encoding [34].

We provide three different benchmark sets with different parameters. The instances in a benchmark set are distributed in families, and every family has values of  $K_k$  distributed uniformly at random around a different mean in the range  $[1, M \cdot Q]$ . The values of  $K_k$  are proportional to the values of the coefficients in order to avoid introducing trivially satisfiable PB constraints. We choose different values of  $K_k$  to ensure that in the benchmark sets there are

instances of different hardness, and that approximately half of the instances are satisfiable.

**MMKP1** 100 families of 5 instances, with  $L = 10$ ,  $N = 15$ ,  $M = 10$ ,  $Q = 1000$ .  
The families have linearly increasing  $K_k$  values from family 1 (capacities of about 1000) to family 100 (capacities of about 14000).

**MMKP2** 100 families of 5 instances, with  $L = 10$ ,  $N = 15$ ,  $M = 10$ ,  $Q = 60$ .  
The families have linearly increasing  $K_k$  values from family 1 (capacities of about 100) to family 100 (capacities of about 800).

**MMKP3** 20 families of 20 instances, with  $L = 50$ ,  $N = 15$ ,  $M = 5$ ,  $Q = 10$ .  
The values of  $K_k$  linearly increase in each family, ranging between 65 and 100.

As can be seen in Table 3, the three benchmark sets are diverse regarding the sizes of the formulas generated by all encodings. Namely, formulas obtained from MMKP1 are approximately one order of magnitude larger than those obtained from MMKP2, and these are approximately one order of magnitude larger than the formulas obtained from MMKP3.

#### 11.1.2. CA Instances

In the CA problem, there is a number of sets of items which are demanded (bids). Each bid has an associated profit, and some bids are incompatible, i.e. they cannot be selected together because they contain a same item. The problem consists of selecting a subset of bids to maximise the obtained profit. This problem can be naturally modelled with a Boolean variable for each bid: a variable is set to true if and only if its associated bid is selected; AMO constraints appear when setting incompatibilities between bids; finally, the maximisation of the profit is modelled as a pseudo-Boolean objective function. The AMO groups are constructed using the following greedy procedure: we start with an empty list that will hold disjoint sets of bids and we process each bid  $b$  in turn;  $b$  is added to the first set in which all existing bids are pairwise incompatible with  $b$ ; if no such set exists, a new set containing  $b$  is added to the list.

We consider the set of instances presented by Bofill et al [35], which were generated with the Combinatorial Auctions Test Suite [36]. There are a total of 170 optimisation instances, each one containing a set of between 70 and 200 bids, and a subset of these bids must be selected. As explained before, from each instance we have defined two instances of the decision problem with two different lower bounds, and hence there is a total of 340 instances.

#### 11.1.3. MRCPSP instances

The MRCPSP consists of deciding a start time and an execution mode for each of the activities of a project. These activities have demands of arbitrary quantities on shared resources of limited capacity, and it must be ensured that those capacities are never surpassed. These constraints can be naturally modelled with PB constraints. Also, there are predefined end-start precedence relations between activities that must be respected. Finally, the duration and

demands on resources of each activity depends on the selected execution mode, and just one mode must be chosen. An optimal solution must minimise the total duration of the project. We consider the time-indexed model from [14], where Boolean variables express whether an activity is running in a particular mode at a particular time instant. There are mainly three reasons of pair-wise incompatibility between these variables: precedences between activities, the requirement of single-mode selection, and the limited capacity of the resources.

For this problem, we consider the 552 feasible instances of the j30 set from PSPLib [37], each of them with two upper bounds. We refer to this set as MRCPSPj30.

#### 11.1.4. RCPSP/t instances

RCPSP/t is another extension of RCPSP. Unlike MRCPSP, there is just one execution mode per activity. However, the demands of an activity over each resource can change during its execution, i.e. the demand depends on how many time units have passed since the activity started. A natural way to deal with this characteristic is by introducing Boolean variables stating whether an activity has started at a particular time instant, and express PB resource constraints in terms of these variables [38]. Therefore there is a new source of pair-wise incompatibilities in addition to precedences and limited resource capacities, that is the fact that an activity cannot start at two different time instants.

For this problem, we consider a representative subset of 360 instances out of the 3600 instances of set j120 from [28]. The J120 set is composed of 360 families of 10 instances with similar characteristics, and from every family we have chosen the first one. Again, each instance is considered with two upper bounds. We refer to this set as RCPSP/Tj120.

#### 11.1.5. NSP instances

The NSP is the problem of deciding the daily shifts of nurses according to their preferences. Many variants of this problem have been proposed [39, 29]. Here we consider the basic definition that was also used in [14] as well as the same Essence Prime model and set of instances. In particular, the considered problem consists in assigning a shift for each nurse and day, while ensuring that each shift has enough nurses and that nurses do not work too many days in a week. Also, the nurses have some preferences regarding their shift assignments, and there is a penalisation for each preference that is not satisfied. The sum of such penalisations must be minimised. For the decision version of the problem, this objective function becomes an upper bound on the total penalisation, which can be modelled with a PB constraint. AMO constraints appear since each nurse can only work in one shift each day.

The instances we solve are a set of 200 instances taken at random from the N25 set from NSPLib [29]. Again, each instance is solved with two upper bounds.

### 11.2. Comparison Between PB and PB(AMO) Encodings

All instances of all sets have been encoded into SAT using the PB(AMO) encodings introduced in this paper, as well as using the corresponding original PB encodings. For completeness we also report results on the BDD- and MDD-based encodings from [5] and [12] summarised in Section 4.

The generated SAT formulas have been solved with a timeout setting of 600 seconds using two different SAT solvers. The first solver is CaDiCaL [1], which was the system that solved more instances than any other competitor in the SAT Race 2019 [40]. The second solver is Glucose version 4.1 [2], which has also obtained top positions in previous SAT competitions and is currently the core of many other state-of-the-art solvers such as MapleLCMD. All formulas have been solved on the same machine: an 8GB, 3.10GHz Intel<sup>®</sup> Xeon<sup>®</sup> E3-1220v2.

#### 11.2.1. Formula Size Comparison

Table 3 contains the sizes of the encodings of PB and PB(AMO) constraints. For each benchmark set and for each encoding we report, in thousands, the mean number of variables (*vars.*) and clauses (*cl.*) required to encode one PB or PB(AMO) constraint. We omit the number of variables and clauses required to encode the AMO constraints because it is the same for each encoding and negligible in magnitude. Column *g.t.* contains the mean computation time required to generate the SAT encoding of the instances of the benchmark set. A long dash (—) means that the encoding has been discarded for that benchmark set, because the formulas are too large and the generation of most instances either ran out of memory or did not finish in less than 600 seconds. We highlight in boldface the smallest number of variables and clauses among all encodings for each benchmark set and also distinguishing between PB and PB(AMO) encodings.

In all sets, using PB(AMO) encodings lets us reduce both the number of clauses and the number of variables w.r.t. their counterpart PB encodings. The reduction rate ranges from one half to three orders of magnitude. It is noticeable that we obtain a high reduction even in the sets with the smallest AMO constraints, with only 4 and 5 variables per AMO. The decrease in size also clearly affects positively the generation time in MMKP and CA sets, which is at least halved in most encodings. However the sets where formulas are generated with Savile Row have generally higher generation times, and the difference between PB and PB(AMO) encodings is not always significant. This is because the automatic AMO detection requires more time than the ad-hoc methods used in CA and MMKP. The generation times are particularly large in RCPSP/Tj120, where a huge number of mutexes are detected.

The use of PB(AMO) encodings is crucial in some cases. For instance, with GSWC and GLPW in MMKP1 we are able to generate instances that run out of memory with their counterpart encodings SWC and LPW. Similarly, with RGGT in MMKP1 we obtain a reasonable generation time compared to the one of RGT.

We also observe that GGPW and GMT0, which digit-wise decompose the coefficients of the PB, produce dramatically smaller sizes compared to the other



encodings. This difference is of approximately two orders of magnitude in the number of clauses and variables in MMKP1 and CA. This is because these benchmark sets have large coefficients (and also large  $K$ ) and this fact penalises the encodings with a size proportional to the value of  $K$ . Although LPW and GLPW also use a digit-wise decomposition, they produce significantly larger encodings than GGPW and GMTO. This is particularly noticeable in RCPSP/Tj120 and NSP, where the generated formulas are many orders of magnitude larger than those obtained with other encodings. This is because GLPW needs to encode many polynomial watchdog formulas. Nevertheless, GLPW generates smaller formulas than the other GAC encodings in CA and MMKP1, which have large coefficients. Recall that GLPW is the only GAC encoding that generates polynomial size formulas w.r.t. the size of the PB. We have observed that the refinements of GLPW proposed in Section 9.2.1 and Section 9.2.2 are crucial, since a naïve implementation generates huge formulas in all sets, as is the case with LPW.

Regarding the reduction process applied in RGT/RGGT, we observe that it generally reduces the sizes of the formulas w.r.t. GT/GGT encodings, in some cases halving the number of variables and clauses. Also, comparing the minRatio heuristic to the default with GT and GGT encodings, we observe that the minRatio heuristic produces smaller formulas than the default (i.e. GT is smaller than GTd, and GGT is smaller than GGTd) in terms of the number of clauses. For MMKP1 and CA, GT and GGT are able to generate formulas but GTd and GGTd are not. However, the minRatio encodings generally produce more variables.

### 11.2.2. Solving Time Comparison

Tables 4 and 5 contain statistics regarding the solving times using the SAT solvers CaDiCaL and Glucose respectively. The tables contain, for each benchmark set and encoding: first quartile ( $Q1$ ), median ( $med$ ) and third quartile ( $Q3$ ) of solving time in seconds, where t.o. means execution aborted at 600 seconds; and the number of instances that timed out before being solved ( $t.o.$ ). We highlight in boldface the best encoding regarding number of timeouts, breaking ties with  $Q3$ , for each solver, for each benchmark set, and also distinguishing between PB and PB(AMO) encodings. The globally best values for each dataset considering both solvers are underlined.

Even though both are CDCL solvers, CaDiCaL incorporates many inprocessing techniques that modify the formula on the fly, as well as local search. These are not included in Glucose, and this might explain that in some cases they perform quite differently. In any case, the results clearly show that using PB(AMO)s substantially improves the solving times. This solving time reduction is consistent with the observed reduction in the sizes of the formulas. In most cases the solving times are reduced by approximately one order of magnitude. Also the number of timeouts is significantly reduced, sometimes more than halved. There are cases where this improvement is even more evident, as in RCPSP/Tj120 with CaDiCaL, and especially in NSP where almost no solutions are found with original PB encodings with any solver.

Table 3: Number of variables and clauses in thousands and generation time in seconds, for each set and PB(AMO) encoding.

		PB				PB(AMO)			
		enc.	vars.	cl.	g.t.	enc.	vars.	cl.	g.t.
MMKP1	BDD	596.60	1193.20	19.84	MDD	25.09	263.37	2.68	
	SWC	—	—	—	GSWC	104.84	1072.04	5.97	
	GTd	—	—	—	GGTd	—	—	—	
	GT	831.49	1805.75	88.71	GGT	61.91	676.54	6.52	
	RGT	634.81	1411.53	116.13	RGGT	25.00	275.01	4.53	
	MTO	<b>3.56</b>	<b>10.30</b>	0.17	GMTO	<b>0.49</b>	<b>1.95</b>	0.09	
	GPW	5.90	76.93	0.48	GGPW	0.99	4.43	0.04	
	LPW	—	—	—	GLPW	30.52	220.56	2.09	
MMKP2	BDD	40.60	81.20	1.10	MDD	2.04	19.59	0.19	
	SWC	68.11	135.49	0.82	GSWC	6.41	61.64	0.36	
	GTd	10.00	1639.65	10.47	GGTd	1.93	120.67	0.82	
	GT	47.85	101.03	4.01	GGT	4.15	41.65	0.36	
	RGT	37.34	80.00	3.99	RGGT	2.01	20.94	0.24	
	MTO	<b>2.40</b>	<b>7.15</b>	0.08	GMTO	<b>0.33</b>	<b>1.19</b>	0.02	
	GPW	3.48	42.33	0.26	GGPW	0.59	2.46	0.03	
	LPW	—	—	—	GLPW	13.26	90.08	0.80	
MMKP3	BDD	3.27	6.55	0.37	MDD	0.46	2.18	0.11	
	SWC	6.05	12.02	0.40	GSWC	1.16	5.52	0.18	
	GTd	1.31	30.90	1.03	GGTd	0.38	4.58	0.17	
	GT	4.27	8.78	1.10	GGT	0.80	3.89	0.18	
	RGT	3.15	6.54	1.04	RGGT	0.45	2.32	0.13	
	MTO	0.83	<b>1.78</b>	0.10	GMTO	<b>0.21</b>	<b>0.58</b>	0.04	
	GPW	<b>0.82</b>	5.11	0.17	GGPW	0.33	1.17	0.05	
	LPW	19.38	213.59	7.45	GLPW	4.38	24.04	1.04	
CA	BDD	155.96	312.06	0.44	MDD	38.60	160.71	0.22	
	SWC	376.91	747.11	0.45	GSWC	109.55	381.73	0.23	
	GTd	—	—	—	GGTd	—	—	—	
	GT	166.92	458.77	1.07	GGT	78.94	247.68	0.35	
	RGT	116.17	349.09	1.62	RGGT	34.90	150.98	0.52	
	MTO	<b>1.21</b>	<b>6.04</b>	0.02	GMTO	<b>0.41</b>	<b>1.35</b>	0.01	
	GPW	1.86	17.96	0.02	GGPW	0.72	3.81	0.01	
	LPW	67.03	1327.85	0.84	GLPW	16.60	154.37	0.12	
MRCSP30	BDD	0.42	0.81	5.37	MDD	0.06	0.21	5.33	
	SWC	0.85	1.68	5.35	GSWC	0.15	0.50	5.30	
	GTd	0.18	1.60	5.46	GGTd	0.06	0.41	5.35	
	GT	0.50	1.03	5.70	GGT	0.11	0.33	5.37	
	RGT	0.34	0.73	5.68	RGGT	0.05	0.21	5.36	
	MTO	0.26	0.85	5.37	GMTO	0.06	0.16	5.30	
	GPW	0.28	1.32	5.35	GGPW	0.08	0.21	5.30	
	LPW	4.77	42.24	5.64	GLPW	0.72	2.57	5.40	
RCSP/Tj120	BDD	2.61	5.19	35.52	MDD	0.30	1.03	33.27	
	SWC	3.92	7.97	34.88	GSWC	0.42	1.39	32.92	
	GTd	0.78	4.77	35.15	GGTd	0.21	2.09	33.55	
	GT	2.44	4.99	81.93	GGT	0.44	1.37	34.25	
	RGT	2.24	4.59	78.65	RGGT	0.34	1.20	34.12	
	MTO	1.79	22.30	36.86	GMTO	0.19	0.59	33.29	
	GPW	2.27	43.82	34.96	GGPW	0.26	1.12	33.16	
	LPW	190.51	11269.63	60.49	GLPW	5.51	46.33	35.41	
NSP	BDD	113.28	226.31	2.03	MDD	4.22	12.38	1.70	
	SWC	175.64	351.48	1.85	GSWC	5.06	14.74	1.72	
	GTd	6.59	249.81	2.09	GGTd	1.28	15.96	1.78	
	GT	143.97	288.45	14.30	GGT	4.80	13.77	2.34	
	RGT	112.45	225.41	14.36	RGGT	4.23	12.55	2.36	
	MTO	7.72	14.12	1.85	GMTO	1.34	7.19	1.74	
	GPW	7.93	242.72	1.87	GGPW	1.91	26.99	1.73	
	LPW	1180.69	101527.19	4.47	GLPW	109.52	3961.55	2.35	

Regardless of which solver is used, the GMTO encoding is clearly the best in MMKP1 and MMKP2, and is very close to the best encoding for CA, namely GGPW. As mentioned before, all encodings except GMTO and GGPW generate huge formulas for MMKP1 and CA, and large formulas for MMKP2. These large sizes adversely affect the solving times. Among GAC encodings, the best one in MMKP1 and CA is GLPW, which is among the smallest GAC encodings for those benchmark sets. For datasets with smaller coefficient values, i.e. MMKP3, MRCPSpj30, RCPSP/Tj120 and NSP, there is not a clear winner, although GGPW, RGGT and GMTO perform the best in many cases.

In all of our benchmark sets, GGPW is better than GLPW even though the latter is GAC. This is possibly because GLPW always generates significantly larger formulas.

Looking at the crafted MMKP datasets, instances in MMKP3 contain more PB constraints than the others, and the values of  $K$  are distributed around the transition value from unsatisfiable instances to satisfiable instances. We have observed empirically that it is in this transition where the instances become harder. It is precisely in MMKP3 where we can observe that non-GAC encodings worsen significantly their performance with respect to MMKP1 in comparison to GAC encodings, such as RGGT.

The picture regarding the minRatio heuristic is somewhat complicated. For MMKP1 and CA, it was not even possible to generate the formulas for GTd and GGTd. However, for MMKP2, MMKP3 and MRCPSpj30 with the PB encodings (comparing GT to GTd), the default heuristic is superior with both solvers. In this case minRatio causes GT to generate many more variables than GTd. In other datasets, there is no clear best option. When using the PB(AMO) encodings, the sizes of GGT and GGTd are more similar than the sizes of GT and GTd. Overall it seems that the minRatio heuristic avoids catastrophic worst-case behaviour when the range of coefficients is large, but otherwise is not clearly better or worse than the default heuristic for the PB(AMO) encodings.

Finally, we observe that the reduction technique introduced in the RGT and RGGT encodings improves the solving times in all MMKP and CA sets and with both solvers compared to GT and GGT respectively, and the number of timeouts is also reduced in all cases.

### 11.2.3. RGT and RGGT

Figure 11 illustrates how the reduction step of RGGT affects the number of variables required to encode the PB(AMO) constraint (in comparison with GGT). We omit the root node because GT, GGT, RGT, and RGGT all use only one SAT variable at the root. Figure 11a shows how much reduction happens on average at different depths in the RGGT tree. The reduction factor (defined as  $\frac{|A.vals|}{|A.intervals|}$  for each node  $A$ ) is largest close to the root. It is interesting that considerable reduction is still taking place beyond depth 10 for some benchmark sets. This may be the result of the minRatio heuristic which tends to produce unbalanced trees (as described in Section 6.4), spreading the leaf nodes over a wide range of depths. GGT leaf nodes typically have smaller *vals* sets than

Table 4: Solving times and number of timeouts using CaDiCaL.

		PB				PB(AMO)					
		enc.	Q1	med	Q3	t.o.	enc.	Q1	med	Q3	t.o.
MMKP1	BDD	27.75	47.62	t.o.	149	MDD	3.63	5.64	274.05	118	
	SWC	—	—	—	—	GSWC	21.32	37.32	292.34	107	
	GTd	—	—	—	—	GGTd	—	—	—	—	
	GT	48.86	65.47	t.o.	146	GGT	14.60	19.83	314.57	116	
	RGT	31.93	49.30	t.o.	146	RGGT	3.90	5.75	181.21	110	
	MTO	0.21	0.97	7.75	<b>43</b>	GMTO	0.03	0.07	1.03	<b>27</b>	
	GPW	1.79	2.08	24.87	68	GGPW	0.07	0.11	4.30	55	
	LPW	—	—	—	—	GLPW	5.64	6.38	90.49	89	
MMKP2	BDD	1.95	2.90	346.01	116	MDD	0.25	0.48	20.38	73	
	SWC	3.17	4.43	110.42	78	GSWC	1.27	1.91	11.14	51	
	GTd	39.46	58.30	192.44	91	GGTd	1.93	2.22	39.48	65	
	GT	2.35	3.88	487.41	120	GGT	0.89	1.16	16.69	69	
	RGT	2.06	2.86	445.05	118	RGGT	0.27	0.47	12.58	66	
	MTO	0.13	0.60	5.98	<b>43</b>	GMTO	0.02	5	0.78	<b>23</b>	
	GPW	0.83	1.10	18.74	71	GGPW	0.04	0.07	2.98	52	
	LPW	—	—	—	—	GLPW	1.92	2.32	35.59	71	
MMKP3	BDD	t.o.	t.o.	t.o.	318	MDD	84.74	248.11	t.o.	135	
	SWC	167.89	558.67	t.o.	189	GSWC	32.97	127.61	t.o.	101	
	GTd	213.55	t.o.	t.o.	223	GGTd	71.75	259.29	t.o.	142	
	GT	t.o.	t.o.	t.o.	330	GGT	48.68	208.53	t.o.	131	
	RGT	t.o.	t.o.	t.o.	329	RGGT	46.53	202.19	t.o.	126	
	MTO	81.57	447.52	t.o.	<b>170</b>	GMTO	22.91	100.24	551.00	<b>98</b>	
	GPW	114.80	507.51	t.o.	194	GGPW	78.22	374.97	t.o.	172	
	LPW	t.o.	t.o.	t.o.	349	GLPW	t.o.	t.o.	t.o.	311	
CA	BDD	0.01	0.05	9.77	28	MDD	0.01	0.04	5.33	14	
	SWC	0.03	0.20	27.05	37	GSWC	0.02	0.09	10.89	23	
	GTd	—	—	—	—	GGTd	—	—	—	—	
	GT	0.02	0.11	13.77	18	GGT	0.02	0.08	6.32	17	
	RGT	0.02	0.08	8.65	18	RGGT	0.02	0.07	3.18	13	
	MTO	0.02	0.05	0.35	0	GMTO	0.01	0.03	0.22	0	
	GPW	0.01	0.04	<b>0.23</b>	<b>0</b>	GGPW	0.01	0.03	<b>0.14</b>	<b>0</b>	
	LPW	0.03	0.57	6.02	0	GLPW	0.02	0.08	1.53	0	
MRCFSPj30	BDD	0.03	0.07	2.17	20	MDD	0.02	0.04	0.11	8	
	SWC	0.04	0.11	3.03	24	GSWC	0.03	0.06	0.17	7	
	GTd	0.04	0.09	2.57	40	GGTd	0.02	0.05	0.14	8	
	GT	0.03	0.07	1.95	<b>19</b>	GGT	0.02	0.05	0.14	8	
	RGT	0.02	0.06	1.55	21	RGGT	0.02	0.05	<b>0.12</b>	<b>7</b>	
	MTO	0.04	0.09	3.19	41	GMTO	0.02	0.05	0.16	7	
	GPW	0.03	0.08	2.04	25	GGPW	0.02	0.05	0.12	8	
	LPW	0.06	0.48	20.67	49	GLPW	0.04	0.09	0.38	10	
RCFSP/Tj120	BDD	0.21	4.87	257.43	117	MDD	0.16	1.03	70.42	41	
	SWC	0.30	8.44	479.54	162	GSWC	0.19	1.70	85.12	41	
	GTd	0.26	3.33	177.85	<b>75</b>	GGTd	0.16	1.79	76.00	35	
	GT	0.25	4.83	311.58	142	GGT	0.19	1.52	76.98	38	
	RGT	0.23	4.56	346.34	154	RGGT	0.16	1.28	83.68	38	
	MTO	0.39	25.64	t.o.	223	GMTO	0.15	6.39	209.17	66	
	GPW	0.32	7.91	t.o.	187	GGPW	0.18	1.10	73.26	<b>26</b>	
	LPW	10.48	t.o.	t.o.	467	GLPW	0.32	12.53	599.24	180	
NSP	BDD	t.o.	t.o.	t.o.	394	MDD	0.03	0.18	1.89	15	
	SWC	t.o.	t.o.	t.o.	374	GSWC	0.06	0.34	1.76	17	
	GTd	t.o.	t.o.	t.o.	<b>333</b>	GGTd	0.04	0.15	1.56	17	
	GT	t.o.	t.o.	t.o.	389	GGT	0.05	0.33	3.36	29	
	RGT	t.o.	t.o.	t.o.	391	RGGT	0.03	0.23	3.31	27	
	MTO	t.o.	t.o.	t.o.	375	GMTO	0.08	0.32	1.92	23	
	GPW	t.o.	t.o.	t.o.	336	GGPW	0.05	0.16	<b>1.32</b>	<b>15</b>	
	LPW	t.o.	t.o.	t.o.	400	GLPW	0.73	5.71	27.49	22	

Table 5: Solving times and number of timeouts using Glucose.

		PB					PB(AMO)				
		enc.	Q1	med	Q3	t.o.	enc.	Q1	med	Q3	t.o.
MMKP1	BDD	9.80	12.18	t.o.	138	MDD	3.60	8.66	144.55	104	
	SWC	—	—	—	—	GSWC	4.42	5.83	249.15	112	
	GTd	—	—	—	—	GGTd	—	—	—	—	
	GT	12.06	15.74	t.o.	137	GGT	2.57	2.77	97.69	103	
	RGT	8.95	11.43	t.o.	136	RGGT	2.49	7.14	110.08	99	
	MTO	0.23	0.93	24.53	<b>68</b>	GMTO	0.02	0.03	0.99	<b>37</b>	
	GPW	0.92	0.96	16.01	81	GGPW	0.04	0.04	6.39	69	
	LPW	—	—	—	—	GLPW	2.06	2.16	66.93	90	
	MMKP2	BDD	2.94	3.52	403.91	119	MDD	0.19	0.29	15.79	69
SWC		4.08	5.53	145.60	102	GSWC	0.49	0.57	6.15	58	
GTd		5.33	6.98	244.01	116	GGTd	2.49	8.78	71.88	99	
GT		2.79	3.47	491.06	122	GGT	0.32	0.37	18.47	73	
RGT		2.36	2.92	355.14	120	RGGT	0.16	0.25	11.83	69	
MTO		0.14	0.58	13.52	<b>62</b>	GMTO	0.02	0.02	0.56	<b>38</b>	
GPW		0.46	0.48	14.08	80	GGPW	0.02	0.03	3.65	67	
LPW		—	—	—	—	GLPW	0.90	0.98	21.37	74	
MMKP3		BDD	t.o.	t.o.	t.o.	334	MDD	46.49	155.70	t.o.	104
	SWC	228.59	t.o.	t.o.	215	GSWC	33.69	140.21	584.10	100	
	GTd	219.36	t.o.	t.o.	215	GGTd	72.66	263.05	t.o.	146	
	GT	t.o.	t.o.	t.o.	333	GGT	47.52	153.11	t.o.	109	
	RGT	t.o.	t.o.	t.o.	332	RGGT	40.92	135.16	548.87	<b>94</b>	
	MTO	118.02	521.95	t.o.	<b>195</b>	GMTO	74.93	411.42	t.o.	172	
	GPW	144.17	t.o.	t.o.	214	GGPW	128.82	t.o.	t.o.	226	
	LPW	t.o.	t.o.	t.o.	349	GLPW	483.77	t.o.	t.o.	286	
	CA	BDD	0.01	0.03	5.16	20	MDD	0.01	0.02	1.88	12
SWC		0.01	0.35	14.47	34	GSWC	0.01	0.08	3.56	11	
GTd		—	—	—	—	GGTd	—	—	—	—	
GT		0.01	0.09	5.23	15	GGT	0.01	0.05	2.05	9	
RGT		0.01	0.04	4.36	11	RGGT	0.01	0.02	1.26	4	
MTO		0.01	0.02	0.26	0	GMTO	0.01	0.01	0.13	0	
GPW		0.01	0.02	<b>0.09</b>	<b>0</b>	GGPW	0.01	0.01	<b>0.05</b>	<b>0</b>	
LPW		0.00	0.34	3.01	0	GLPW	0.01	0.04	0.56	0	
MRCFSPj30		BDD	0.00	0.07	1.17	34	MDD	0.00	0.02	0.07	13
	SWC	0.00	0.13	1.44	31	GSWC	0.00	0.03	0.09	13	
	GTd	0.00	0.13	1.51	49	GGTd	0.00	0.03	0.11	13	
	GT	0.00	0.07	0.97	33	GGT	0.00	0.02	0.08	13	
	RGT	0.00	0.05	<b>0.86</b>	<b>31</b>	RGGT	0.00	0.02	0.07	<b>11</b>	
	MTO	0.00	0.04	2.51	56	GMTO	0.00	0.02	0.06	15	
	GPW	0.00	0.05	1.09	34	GGPW	0.00	0.02	0.05	13	
	LPW	0.00	0.66	7.07	66	GLPW	0.00	0.11	0.38	14	
	RCFSP/Tj120	BDD	0.34	7.62	t.o.	241	MDD	0.09	1.65	t.o.	201
SWC		0.78	10.15	t.o.	269	GSWC	0.13	1.93	t.o.	194	
GTd		0.34	6.16	t.o.	237	GGTd	0.13	2.91	t.o.	198	
GT		0.39	5.40	t.o.	<b>201</b>	GGT	0.11	1.68	t.o.	208	
RGT		0.37	6.00	t.o.	203	RGGT	0.09	1.53	t.o.	204	
MTO		0.44	67.49	t.o.	315	GMTO	0.12	8.24	t.o.	<b>189</b>	
GPW		0.48	5.19	t.o.	282	GGPW	0.11	1.18	t.o.	245	
LPW		3.70	t.o.	t.o.	460	GLPW	0.37	6.49	t.o.	268	
NSP		BDD	t.o.	t.o.	t.o.	398	MDD	0.01	0.07	0.82	19
	SWC	t.o.	t.o.	t.o.	400	GSWC	0.02	0.09	1.27	19	
	GTd	t.o.	t.o.	t.o.	398	GGTd	0.01	0.08	0.75	21	
	GT	t.o.	t.o.	t.o.	398	GGT	0.02	0.10	2.12	37	
	RGT	t.o.	t.o.	t.o.	398	RGGT	0.01	0.09	2.64	39	
	MTO	t.o.	t.o.	t.o.	396	GMTO	0.03	0.12	2.61	27	
	GPW	t.o.	t.o.	t.o.	<b>395</b>	GGPW	0.02	0.07	0.52	<b>13</b>	
	LPW	t.o.	t.o.	t.o.	400	GLPW	0.71	3.21	9.57	28	

internal nodes and are therefore more likely to permit non-trivial intervals to be created in their sibling node.

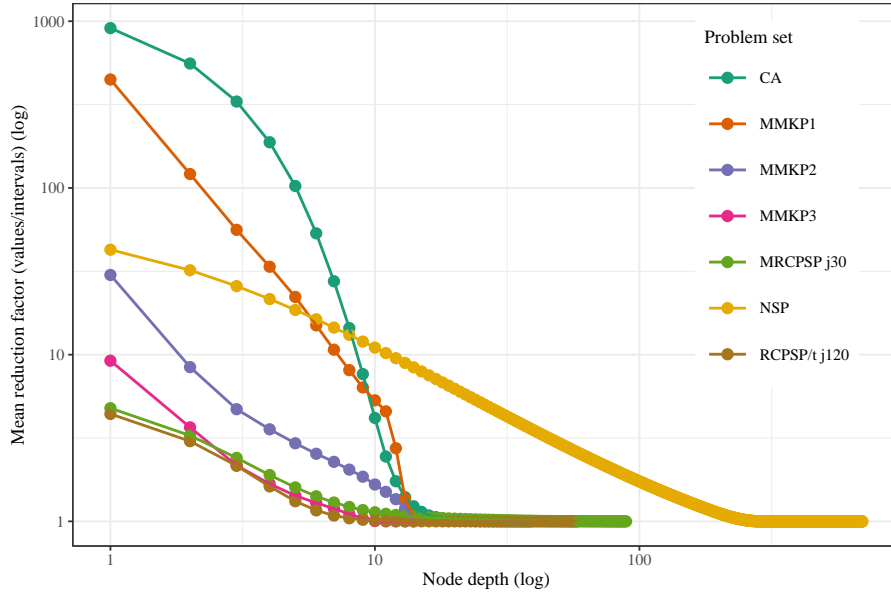
Both plots show that less reduction occurs as we progress through sets MMKP1, MMKP2, and MMKP3, just as in Table 3 the relative reduction in SAT variables and clauses from GGT to RGGT decreases across these three problem sets. Recall that the coefficients are being sampled with a maximum value of 1000, 60, and 10 respectively – this progression reduces the likelihood of selecting close but distinct coefficients which can be merged into an interval. The reduction appears to benefit all CA instances, the highest reduction factor being observed in the middle-sized instances. The reduction factors for MRCPSP and RCPSP/t are modest compared to most of the other problem classes. NSP has a significant average reduction factor even beyond depth 100, but the trees are very deep and the reduction for entire trees (Figure 11b) is modest.

## 12. Related Work

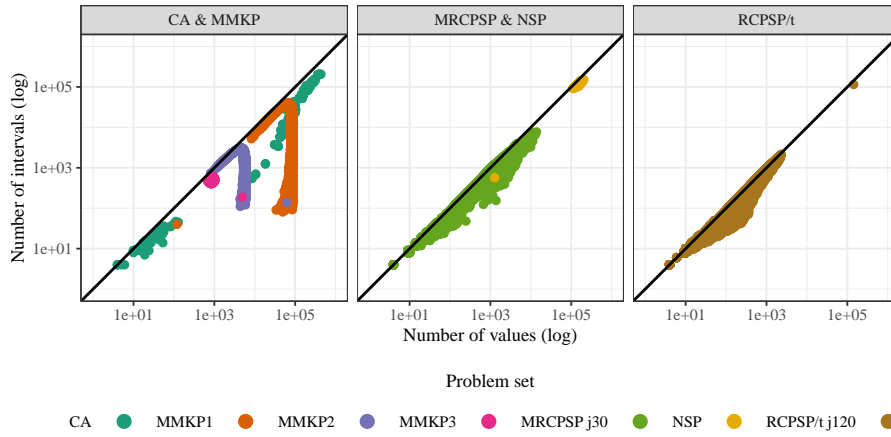
Encodings of PB constraints based on Multi-valued Decision Diagrams (MDDs) have been extensively studied. It was in this context that PB(AMO) constraints were first introduced by Bofill, Coll, Suy, and Villaret [11]. Originally PB(AMO) constraints were defined in a slightly different way and were referred to as AMO-PB constraints. An MDD-based SAT encoding of AMO-PB constraints was introduced and successfully applied to solve variants of the Resource-Constrained Project Scheduling Problem (RCPSP). The MDD encoding approach was later revisited by the same authors [12] within a more general framework named PB modulo  $\mathcal{C}$ , or  $\text{PB}(\mathcal{C})$ , where  $\mathcal{C}$  stands for any kind of collateral constraint over the variables of the PB constraint. PB(AMO) constraints were presented as a particular case of  $\text{PB}(\mathcal{C})$ . Other collateral constraints were also considered, such as *exactly-one* and *implication chains* (i.e. monotonically non-decreasing sequences) [12]. Earlier, Abío et al [41] studied decision diagram encodings of PB constraints in conjunction with implication chains, and also showed that implication chains can be used to represent AMO constraints. Bofill et al [12] compared their MDD encoding of PB(AMO) constraints to Abío et al’s encoding [41] (with AMOs represented as implication chains), and it was shown that the two MDD-based techniques had similar performance.

A number of alternatives to decision diagrams exist for encoding PB constraints into SAT. In this work we have reviewed a large sample of state-of-the-art SAT encodings for PB constraints, and the related work is presented in detail from Section 4 to Section 9: an encoding based on decision diagrams in Section 4; the Sequential Weight Counter (SWC) in Section 5; the Generalized Totalizer (GT) in Section 6; the n-Level Modulo Totalizer (MTO) in Section 7; the Global Polynomial Watchdog (GPW) in Section 8; and the Local Polynomial Watchdog (LPW) in Section 9. Each of these encodings have been generalized to encode PB(AMO) constraints, either here or elsewhere [11].

There are a number of other PB encodings that have not been generalized to PB(AMO) constraints, but most likely could be using techniques similar to



(a) The mean reduction factor per node as depth increases



(b) Comparison of # of values versus # of intervals for entire GGT resp. RGGT trees

Figure 11: Effects of reduction observed in RGGT (with MinRatio heuristic)

the ones used here. For instance, Manthey et al [10] proposed an alternative to GPW and LPW called Binary Merger, in which formulas  $\phi$  and  $\psi$  are built using sorting networks and odd-even merger circuits [42] respectively. The size bound of formulas obtained with Binary Merger is asymptotically smaller than that of GPW. However, in this work we have chosen the original GPW definition because in our preliminary experiments GPW has shown slightly better performance. Eén and Sörensson [4] presented a BDD-based encoding very similar to the one we use, and also two further SAT encodings for PB constraints. The first, named *Sorters*, consists of a sequence of digit-wise sums similar to GPW but also accepting mixed radix bases. A PB(AMO) version of this encoding could be easily defined by introducing auxiliary variables for each group  $X_i$  and bucket, as done in GGPW but considering mixed radix bases. The other one, named *Adders*, performs digit-wise sums using a circuit of full adders and hence introducing carry bits. Again, we could straightforwardly generalize *Adders* by introducing auxiliary variables for each bit and group  $X_i$ . In Eén and Sörensson’s own experiments the Adders encoding was inferior to their BDD-based encoding and the Sorters encoding [4].

Conjunctions of PB constraints and AMO constraints have also been considered in the context of Mixed Integer Linear Programming. Achterberg et al [43] describe a presolving step where an AMO constraint is used to replace a set of 0/1 variables (taken from the scope of the AMO constraint) with an integer variable, and to redefine a PB constraint to use the integer variable in place of the 0/1 variables.

Ansótegui et al. [14] integrated the MDD-based PB(AMO) encoding (described in Section 4) into the automatic reformulation pipeline of Savile Row [15]. Similarly to our work, the output of that reformulation process is a SAT formula. However, in that case the input is a CP model written in Essence Prime, which is an expressive constraint programming language that supports finite domain variables and global constraints, among others. The input CP model contains a set of linear constraints, and AMO constraints over their variables are automatically detected. Some AMO constraints are detected by means of a syntactic check. Other AMOs are retrieved by detecting cliques in a graph of mutexes between pairs of literals. Mutexes are detected by probing: a literal is set to true and constraint propagation is applied to detect other literals that are incompatible with it. The focus of Ansótegui et al. is on automatic detection of AMOs, therefore it is complementary to developing new PB(AMO) encodings.

### 13. Conclusions and Future Work

When solving a combinatorial problem with SAT, the size and properties of the encoding are of vital importance. Arithmetic can be challenging to encode into SAT, and there has been a great deal of work on encoding the PB constraint in particular. Our focus has been on PB(AMO) constraints, which are conjunctions of one PB constraint and any number of AMO constraints. We have defined five new encodings for PB(AMO) constraints by generalising existing state-of-the-art encodings of PB constraints. In each case, the size of



the PB(AMO) encoding is substantially reduced compared to its corresponding PB encoding. Moreover, the propagation properties of the original encodings are preserved in the new ones.

We performed experiments with two recent CDCL SAT solvers (Glucose and CaDiCaL) using five problem classes: the Multi-Choice Multidimensional Knapsack Problem where we can control the parameters of the PB(AMO) constraints, the Combinatorial Auctions problem, the Multi-mode Resource-Constrained Project Scheduling Problem, the Resource-Constrained Project Scheduling Problem with Time-Dependent Resource Capacities and Requests, and the Nurse Scheduling Problem. The new PB(AMO) encodings are dramatically smaller and more efficient than their counterpart PB encodings. We have observed size reductions of an order of magnitude, and also solving time improvements of an order of magnitude in several cases (comparing median times). In almost every case, the PB(AMO) encoding solves more instances within the time limit than its corresponding PB encoding.

We have also shown that there is no single best encoding for PB(AMO) constraints, but it depends on the characteristics of the instances at hand. The benchmark instances that we consider expose some strengths and weaknesses of the different encodings. For example, the GGPW and GMTO encodings generate extremely small formulas, and they represent the best choices for some benchmark sets despite their poor propagation properties.

We also contribute a new encoding that improves the Generalized Totalizer by collecting equivalent values into intervals. When applied to PB constraints, the new encoding is named Reduced Generalized Totalizer (RGT), and for PB(AMO) constraints it is the Reduced Generalized Generalized Totalizer (RGGT). In terms of size, RGT and RGGT are never worse than GT and GGT respectively, and they are often significantly better. The improvement in formula size translates into faster solving times, and in fact RGGT is the overall best choice for two of the benchmark sets.

The success of the PB(AMO) encodings immediately suggests two avenues of future work. The first is to investigate whether other constraints could be exploited in addition to the set of AMO constraints to further reduce the size of the encodings. For example, in the Combinatorial Auctions problem the AMO constraints represent cliques in a mutex graph. Other mutexes (outside the cliques) are not presently used in any way. In an encoding based on Generalized Totalizer, the additional mutexes could potentially rule out values of internal nodes, saving both variables and clauses. Similarly, in the MDD encoding the additional mutexes could potentially rule out nodes and edges. The second avenue is to automatically select an appropriate PB(AMO) encoding based on properties of the PB(AMO) constraint, such as the number of variables, the magnitude of the coefficients, and the sizes and number of the cells in the AMO partition. Given that there is no single best encoding, and differences in performance are often substantial, an accurate encoding selection method would be valuable.

### Acknowledgements

Work partially supported by grant RTI2018-095609-B-I00 (MCIU/AEI/FEDER, UE). Jordi Coll is supported by grant *Ayudas para Contratos Predoctorales 2016* (grant number BES-2016-076867, funded by MINECO and co-funded by FSE), and partially funded by the French Agence Nationale de la Recherche, reference ANR-19-CHIA-0013-01, and by Archimedes institute, Aix-Marseille University. Felix Ulrich-Oltean is supported by grant EP/R513386/1 from the UK Engineering and Physical Sciences Research Council.

### References

- [1] A. Biere, CaDiCaL at the SAT Race 2019, in: M. Heule, M. Järvisalo, M. Suda (Eds.), Proc. of SAT Race 2019 – Solver and Benchmark Descriptions, Vol. B-2019-1 of Department of Computer Science Series of Publications B, University of Helsinki, 2019, pp. 8–9.
- [2] G. Audemard, L. Simon, On the glucose SAT solver, *International Journal on Artificial Intelligence Tools* 27 (1) (2018) 1–25.
- [3] T. Philipp, P. Steinke, PBLib - A Library for Encoding Pseudo-Boolean Constraints into CNF, in: Theory and Applications of Satisfiability Testing - SAT 2015, 18th International Conference, Vol. 9340 of LNCS, Springer, 2015, pp. 9–16.
- [4] N. Eén, N. Sorensson, Translating pseudo-Boolean constraints into SAT, *Journal on Satisfiability, Boolean Modeling and Computation* 2 (2006) 1–26.
- [5] I. Abío, R. Nieuwenhuis, A. Oliveras, E. Rodríguez-Carbonell, V. Mayer-Eichberger, A new look at BDDs for pseudo-Boolean constraints, *Journal of Artificial Intelligence Research* 45 (2012) 443–480.
- [6] S. Hölldobler, N. Manthey, P. Steinke, A compact encoding of pseudo-Boolean constraints into SAT, in: KI 2012: Advances in Artificial Intelligence - 35th Annual German Conference on Artificial Intelligence, Vol. 7526 of LNCS, Springer, 2012, pp. 107–118.
- [7] S. Joshi, R. Martins, V. M. Manquinho, Generalized Totalizer Encoding for Pseudo-Boolean Constraints, in: Principles and Practice of Constraint Programming - CP 2015, 21st International Conference, Vol. 9255 of LNCS, Springer, 2015, pp. 200–209.
- [8] A. Zha, M. Koshimura, H. Fujita, N-level modulo-based CNF encodings of pseudo-Boolean constraints for MaxSAT, *Constraints* 24 (2) (2019) 133–161.

- [9] O. Bailleux, Y. Boufkhad, O. Roussel, New Encodings of Pseudo-Boolean Constraints into CNF, in: *Theory and Applications of Satisfiability Testing - SAT 2009*, 12th International Conference, Vol. 5584 of LNCS, Springer, 2009, pp. 181–194.
- [10] N. Manthey, T. Philipp, P. Steinke, A More Compact Translation of Pseudo-Boolean Constraints into CNF such that Generalized Arc Consistency is Maintained, in: *KI 2014: Advances in Artificial Intelligence - 37th Annual German Conference on AI*, Vol. 8736 of LNCS, Springer, 2014, pp. 123–134.
- [11] M. Bofill, J. Coll, J. Suy, M. Villaret, Compact MDDs for Pseudo-Boolean Constraints with At-Most-One Relations in Resource-Constrained Scheduling Problems, in: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence - IJCAI 2017*, ijcai.org, 2017, pp. 555–562.
- [12] M. Bofill, J. Coll, J. Suy, M. Villaret, An mdd-based sat encoding for pseudo-boolean constraints with at-most-one relations, *Artificial Intelligence Review* (2020) 1–32.
- [13] M. Bofill, J. Coll, J. Suy, M. Villaret, An Efficient SMT Approach to Solve MRCPSP/max Instances with Tight Constraints on Resources, in: *Principles and Practice of Constraint Programming - CP 2017*, 23rd International Conference, Vol. 10416 of LNCS, Springer, 2017, pp. 71–79.
- [14] C. Ansótegui, M. Bofill, J. Coll, N. Dang, J. L. Esteban, I. Miguel, P. Nightingale, A. Z. Salamon, J. Suy, M. Villaret, Automatic detection of at-most-one and exactly-one relations for improved SAT encodings of pseudo-Boolean constraints, in: *Proceedings of 25th International Conference Principles and Practice of Constraint Programming - CP*, Vol. 11802 of Lecture Notes in Computer Science, Springer, 2019, pp. 20–36.  
doi:10.1007/978-3-030-30048-7\_2.  
URL [https://doi.org/10.1007/978-3-030-30048-7\\_2](https://doi.org/10.1007/978-3-030-30048-7_2)
- [15] P. Nightingale, Ö. Akgün, I. P. Gent, C. Jefferson, I. Miguel, P. Spracklen, Automatically improving constraint models in Savile Row, *Artificial Intelligence* 251 (2017) 35–61.
- [16] C. Basnet, J. Wilson, Heuristics for determining the number of warehouses for storing non-compatible products, *International transactions in operational research* 12 (5) (2005) 527–538.
- [17] P. R. Ma, E. Y. S. Lee, M. Tsuchiya, A Task Allocation Model for Distributed Computing Systems, *IEEE Trans. Computers* 31 (1) (1982) 41–47.
- [18] D. Pisinger, Budgeting with bounded multiple-choice constraints, *European Journal of Operational Research* 129 (3) (2001) 471–480.

- [19] R. Watson, Packet networks and optimal admission and upgrade of service level agreements: applying the utility model. ma sc, Ph.D. thesis, Department of ECE, University of Victoria (2001).
- [20] S. De Vries, R. V. Vohra, Combinatorial auctions: A survey, *INFORMS Journal on computing* 15 (3) (2003) 284–309.
- [21] C. E. Miller, A. W. Tucker, R. A. Zemlin, Integer programming formulation of traveling salesman problems, *Journal of the ACM* 7 (4) (1960) 326–329.
- [22] M. Bofill, J. Coll, J. Suy, M. Villaret, SAT encodings of pseudo-Boolean constraints with at-most-one relations, in: *Proceedings of the 16th International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research - CPAIOR*, Vol. 11494 of *Lecture Notes in Computer Science*, Springer, 2019, pp. 112–128. doi: 10.1007/978-3-030-19212-9\_8.  
URL [https://doi.org/10.1007/978-3-030-19212-9\\_8](https://doi.org/10.1007/978-3-030-19212-9_8)
- [23] O. Bailleux, Y. Boufkhad, Efficient CNF Encoding of Boolean Cardinality Constraints, in: *Principles and Practice of Constraint Programming - CP 2003*, 9th International Conference, Vol. 2833 of *LNCS*, Springer, 2003, pp. 108–122.
- [24] A. Darwiche, P. Marquis, A knowledge compilation map, *Journal of Artificial Intelligence Research* 17 (2002) 229–264.
- [25] A. Van Gelder, Toward leaner binary-clause reasoning in a satisfiability solver, *Annals of Mathematics and Artificial Intelligence* 43 (1) (2005) 239–253.
- [26] A. Biere, M. J. H. Heule, H. van Maaren, T. Walsh (Eds.), *Handbook of Satisfiability*, Vol. 185 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2009.
- [27] P. Brucker, A. Drexl, R. Möhring, K. Neumann, E. Pesch, Resource-Constrained Project Scheduling: Notation, Classification, Models, and Methods, *European Journal of Operational Research* 112 (1) (1999) 3 – 41.
- [28] S. Hartmann, Project scheduling with resource capacities and requests varying with time: a case study, *Flexible Services and Manufacturing Journal* 25 (1-2) (2013) 74–93.
- [29] M. Vanhoucke, B. Maenhout, NSPLib: a nurse scheduling problem library: a tool to evaluate (meta-)heuristic procedures, in: S. Brailsford, P. Harper (Eds.), *Operational research for health policy: making better decisions*, Peter Lang, 2007, pp. 151–165.

- [30] P. Nightingale, A. Rendl, Essence' description, arXiv:1601.02865 (2016).  
arXiv:1601.02865.  
URL <https://arxiv.org/abs/1601.02865>
- [31] P. Nightingale, Ö. Akgün, I. P. Gent, C. Jefferson, I. Miguel, P. Spracklen, Automatically improving constraint models in Savile Row, *Artificial Intelligence* 251 (2017) 35–61. doi:10.1016/j.artint.2017.07.001.
- [32] J. Chen, A new sat encoding of the at-most-one constraint, *Proc. Constraint Modelling and Reformulation*.
- [33] B. Han, J. Leblet, G. Simon, Hard multidimensional multiple choice knapsack problems, an empirical study, *Computers & Operations Research* 37 (1) (2010) 172 – 181.
- [34] C. Ansótegui, F. Manyá, Mapping problems with finite-domain variables to problems with boolean variables, in: *International conference on theory and applications of satisfiability testing*, Springer, 2004, pp. 1–15.
- [35] M. Boffill, M. Palahí, J. Suy, M. Villaret, Solving Intensional Weighted CSPs by Incremental Optimization with BDDs, in: *CP: Principles and Practice of Constraint Programming*, LNCS 8656, Springer, 2014, pp. 207–223. doi:10.1007/978-3-319-10428-7\\_17.
- [36] K. Leyton-Brown, Y. Shoham, A test suite for combinatorial auctions, in: *Combinatorial auctions*, The MIT Press, 2006, Ch. 18, pp. 451–478.
- [37] R. Kolisch, A. Sprecher, PSPLIB - A Project Scheduling Problem Library, *European Journal of Operational Research* 96 (1) (1997) 205–216.
- [38] M. Boffill, J. Coll, J. Suy, M. Villaret, Smt encodings for resource-constrained project scheduling problems, *Computers & Industrial Engineering* 149 (2020) 106777.
- [39] P. De Causmaecker, G. Vanden Berghe, A categorisation of nurse rostering problems, *Journal of Scheduling* 14 (1) (2011) 3–16. doi:10.1007/s10951-010-0211-z.
- [40] M. J. Heule, M. Jarvisalo, M. Suda, *Proceedings of sat race 2019*, SAT RACE 2019.
- [41] I. Abío, V. Mayer-Eichberger, P. J. Stuckey, Encoding linear constraints with implication chains to CNF, in: *International Conference on Principles and Practice of Constraint Programming(CP)*, Vol. 9255 of LNCS, Springer, 2015, pp. 3–11.
- [42] R. Asín, R. Nieuwenhuis, A. Oliveras, E. Rodríguez-Carbonell, Cardinality Networks: a theoretical and empirical study, *Constraints* 16 (2) (2011) 195–221. doi:10.1007/s10601-010-9105-0.

- [43] T. Achterberg, R. E. Bixby, Z. Gu, E. Rothberg, D. Weninger, Presolve reductions in mixed integer programming, *INFORMS Journal on Computing* 32 (2) (2020) 473–506.