



From MINIZINC to Optimization Modulo Theories, and Back

Francesco Contaldo^(✉), Patrick Trentin^(✉), and Roberto Sebastiani^(✉)

DISI, University of Trento, Trento, Italy
francesco.contaldo@alumni.unitn.it,
{patrick.trentin, roberto.sebastiani}@unitn.it

Abstract. Optimization Modulo Theories (OMT) is an extension of SMT that allows for finding models that optimize objective functions. In this paper we aim at bridging the gap between Constraint Programming (CP) and OMT, *in both directions*. First, we have extended the OMT solver OPTIMATHSAT with a FLATZINC interface – which can also be used as FLATZINC-to-OMT encoder for other OMT solvers. This allows OMT tools to be used in combination with MZN2FZN on the large amount of CP problems coming from the MINIZINC community. Second, we have introduced a tool for translating SMT and OMT problems on the linear arithmetic and bit-vector theories into MINIZINC. This allows MINIZINC solvers to be used on a large amount of SMT/OMT problems.

We have discussed the main issues we had to cope with in either directions. We have performed an extensive empirical evaluation comparing three state-of-the-art OMT-based tools with many state-of-the-art CP tools on (i) CP problems coming from the MINIZINC challenge, and (ii) OMT problems coming mostly from formal verification. This analysis also allowed us to identify some criticalities, in terms of efficiency and correctness, one has to cope with when addressing CP problems with OMT tools, and vice versa.

1 Introduction

The last two decades have witnessed the rise of Satisfiability Modulo Theories (SMT) [12] as efficient tool for dealing with several applications of industrial interest, in particular in the contexts of Formal Verification (FV). SMT is the problem of finding value assignments satisfying some formula in first-order logic wrt. some background theory. Optimization Modulo Theories (OMT) [15, 33, 34, 36, 39, 42, 53, 56] is a more-recent extension of SMT searching for the *optimal* value assignment(s) w.r.t. some objective function(s), by means of a combination of SMT and optimization procedures. (Since OMT extends SMT, hereafter we often simply say “OMT” for both SMT and OMT.)

Several distinctive traits of OMT solvers –like, e.g., the efficient combination of Boolean and arithmetical reasoning, incrementality, the availability of decision procedures for infinite-precision arithmetic and the ability to produce conflict explanations– are a direct consequence of their tight relationship with the FV domain and its practical needs. On the whole, it appears that OMT can be a potentially interesting and efficient technology for dealing with Constraint Programming (CP) problems as well. At the

same time, modeling CP problems for OMT solvers requires a higher-level of expertise, because the same CP instance can have many possible alternative formulations, s.t. the performance of SMT solvers on each encoding are hardly predictable [26, 28, 29].

On the other hand, the availability, efficiency and expressiveness of CP tools makes them of potential interest as backend engines also for FV applications (e.g., [22, 23, 31]), in particular with SW verification, where currently SMT is the dominating backend technology, s.t. a large amount of SMT-encoded FV problems are available [11].

In this paper we aim at bridging the gap between CP and OMT, *in both directions*.

In the CP-to-OMT direction, we have extended the state-of-the-art OMT solver OPTIMATHSAT [58] with a FLATZINC interface (namely “FZN2OMT”). In combination with the standard MZN2FZN encoder [38], this new interface can be used to either (i) solve CP models with OPTIMATHSAT directly or (ii) generate OMT formulas encoded in the SMT-LIB [25] format with optimization extensions, to be fed to other OMT solvers, such as BCLT [16] and Z3 [15]. This allows state-of-the-art OMT technology to be used on MINIZINC problems coming from the CP community.

In the OMT-to-CP direction, we have introduced a tool for translating SMT and OMT problems on the theories of linear arithmetic over the integers and rational (\mathcal{LIRA}) and bit-vector (\mathcal{BV}) into MINIZINC models (hereafter “OMT2MZN”). This allows MINIZINC solvers to be used on OMT problems, giving them access to a large amount of OMT problems, mostly coming from formal verification.

With both directions, we first present and discuss the challenges we encountered and the solutions we adopted to address the differences between the two formalisms. Then we present an extensive empirical evaluation comparing three OMT tools with many state-of-the-art CP tools on (i) CP problems coming from the MINIZINC challenge, and (ii) OMT problems coming mostly from formal verification. This analysis allowed us to identify some criticalities, in terms of efficiency and correctness, one has to cope with when addressing CP problems with OMT tools, and vice versa.

Overall, our new compilers FZN2OMT and OMT2MZN in combination with the standard compiler MZN2FZN [38] provide a framework for translating problems encoded in the SMT-LIB or the MINIZINC format in either direction. This framework enables also for a comparison between OMT solvers and CP tools on problems that do not belong to their original application domain. To the best of our knowledge, this is the first time that such a framework has been proposed, and that the OMT and CP technologies have been extensively compared on problems coming from both fields.

Related Work. The tight connection between SMT and Constraint Programming (CP) has been known for a relatively long period of time [43] and it has previously been subject to investigation. Some works considered a direct encoding of CP [28, 29] and weighted CP [8] into SMT and MAXSMT, or an automatic framework for translating MINIZINC –a standard CP modeling language [40]– into SMT-LIB –the standard SMT format– [17, 18]. Other works explored the integration of typical SAT and SMT techniques within CP solvers [27, 45]. Nowadays, several MINIZINC solvers –like, e.g., HAIFACSP [61] and PICAT [62]– are at least partially based on SAT technology.

To this extent, our first contribution FZN2OMT also obviates the loss, due to obsolescence, of the FZN2SMT compiler proposed by Bofill et al. in [17, 18]. FZN2SMT is not

compatible with the changes that have been introduced to the MINIZINC and FLATZINC standards starting from version 2.0 of the MINIZINC distribution. Since some of these changes are not backward compatible, it is also not possible to use FZN2SMT in conjunction with an older version of the MZN2FZN compiler when dealing with recent MINIZINC models. Furthermore, FZN2SMT translates satisfaction problems into the Version 1 of the SMT-LIB standard and produces no SMT-LIB output in the case of optimization problems, that are solved directly. However, the optimization interface of modern OMT solvers is based on the Version 2 of the SMT-LIB standard. This makes it difficult to use it together with OMT solvers. Unfortunately, the FZN2SMT compiler is closed source, with only the binaries being freely distributed, and seemingly no longer maintained. This made it necessary to provide a new alternative solution to FZN2SMT. To this extent, our new FLATZINC interface of OPTIMATHSAT, FZN2OMT, translates both satisfaction and optimization problems in the Version 2 of the SMT-LIB standard enriched with the optimization extensions for OMT described in [58].

Content. The rest of the paper is organized as follows. In Sect. 2 we provide some background on OMT, MINIZINC and FLATZINC. In Sect. 3 we describe the process from MINIZINC to OMT. In Sect. 4 we describe the process from OMT to MINIZINC. In Sect. 5 we describe an empirical evaluation comparing a OMT-based tool with many state-of-the-art CP tools. Finally, in Sect. 6 we conclude and point out some further research directions.

A longer and more detailed version of this paper is publicly available as [24].

2 Background

Satisfiability Modulo Theories (SMT) is the problem of deciding the satisfiability of a first-order formula φ with respect to a combination of decidable first-order theories. Typical theories of SMT interest are (the theory of) linear arithmetic over the rationals (\mathcal{LRA}), the integers (\mathcal{LIA}) or their combination (\mathcal{LRA}), non-linear arithmetic over the rationals (\mathcal{NLRA}) or the integers (\mathcal{NLIA}), arrays (\mathcal{AR}), bit-vectors (\mathcal{BV}), floating-point arithmetic (\mathcal{FP}), and their combinations thereof. (See [12, 44, 52] for an overview.). The last two decades have witnessed the development of very efficient SMT solvers based on the so-called lazy-SMT schema [12, 52]. This has brought previously-intractable problems to the reach of state-of-the-art SMT solvers.

Optimization Modulo Theories (OMT), [15, 34, 36, 39, 42, 53, 56, 60], is an extension to SMT that allows for finding a model of a first-order formula φ that is *optimal* with respect to some objective function expressed in some background theory, by means of a combination of SMT and optimization procedures. State-of-the-art OMT tools allow optimization in a variety of theories, including linear arithmetic over the rationals (OMT(\mathcal{LRA})) [53] and the integers (OMT(\mathcal{LIA})) [15, 56], bit-vectors (OMT(\mathcal{BV})) [39] and floating-point numbers (OMT(\mathcal{FP})) [60].

A relevant strict subcase of OMT(\mathcal{LRA}) is OMT with Pseudo-Boolean objective functions (OMT(\mathcal{PB})) in the form $\sum_i w_i A_i$ s.t. w_i are rational values and A_i are Boolean variables whose values are interpreted as $\{0, 1\}$. Notice that OMT(\mathcal{PB}) is also equivalent to (partial weighted) MAXSMT, the SMT extension of MAXSAT, and that

OMT(\mathcal{PB}) and MAXSMT can be encoded into OMT(\mathcal{LRA}) but not vice versa [54]. Encoding OMT(\mathcal{PB})/MAXSMT into OMT(\mathcal{LRA}), however, is not the most efficient way to solve them, so that modern OMT solvers such as BCLT [16], OPTIMATHSAT [58] and Z3 [15] implement specialized OMT(\mathcal{PB})/MAXSMT procedures which are much more efficient than general-purpose OMT(\mathcal{LRA}) ones [15, 57, 58].

We stress the fact that —unlike with purely-combinatorial problems, which are encoded into SAT or MAXSAT and are thus solved by purely-Boolean search— typically OMT problems involve the interleaving of *both* Boolean and arithmetical search: search not only for the best truth-value assignment to the atomic subformulae, but also for the best values to the numerical variables compatible with such truth-value assignment [54].

To this date, few OMT solvers exist, namely BCLT [16], CEGIO [9], HAZEL [39], OPTIMATHSAT [58], PULI [33], SYMBA [36] and Z3 [15]. To this aim, we observe that (i) some of these solvers are quite recent, (ii) most of these solvers focus on different, partially overlapping, niche subsets of Optimization Modulo Theories, and (iii) the lack of an official Input/Output interface for OMT makes it hard to compare some of these tools with one another. OMT finds applications in the context of static analysis [19, 32], formal verification and model checking [37, 48], scheduling and planning with resources [33, 35, 46, 50], software security and requirements engineering [41], workflow analysis [13], machine learning [59], and quantum computing [14].

A distinctive trait of SMT (and OMT) solvers is the trade-off of speed against the ability to certify the correctness of the result of any computation, which is particularly important in the contexts of Formal Verification (FV) and Model Checking (MC). When dealing with linear arithmetic in particular, SMT solvers employ *infinite-precision arithmetic* software libraries to avoid numerical errors and overflows.

SMT-LIB [25] is the standard input format by SMT solvers, it provides a standardized definition of the most prominent theories supported by SMT solvers and the corresponding language primitives to use these features. At present, there is no standard input format for modeling optimization problems targeting OMT solvers, although there exist only minor syntactical differences between the major OMT solvers. The tools presented in this paper conform to the *extended* SMT-LIB format for OMT presented in [58], that includes language primitives for modeling objectives.

OPTIMATHSAT [53–58] is a state-of-the-art OMT solver based on the MATHSAT5 SMT solver [3, 21]. OPTIMATHSAT features both single- and multi-objective optimization over arbitrary sets of \mathcal{LRA} , \mathcal{LIA} , \mathcal{LIRA} , \mathcal{BV} , \mathcal{FP} , Pseudo-Boolean (\mathcal{PB}) and MAXSMT cost functions. Multiple objective functions can be combined with one another into a Lexicographic or a Pareto optimization problem, or independently solved in a single run (for the best efficiency).

MINIZINC [38, 40] is a widely adopted high-level declarative language for modeling Constraint Satisfaction Problems (CSP) and Constraint Optimization Problems (COP). The MINIZINC format defines three scalar types (bool, int and float) and two compound types (sets and fixed-size arrays of some scalar type). The standard provides an extensive list of predefined *global constraints*, a class of high-level language primitives that allows one to encode complex constraints in a compact way.




Fig. 1. Circular translation schema from MINIZINC to SMT-LIB and back, resulting from the composition of MZN2FZN, OPTIMATHSAT and OMT2MZN. In this picture, OPTIMATHSAT acts both as a FLATZINC/OMT solver, and also as a FLATZINC to SMT-LIB compiler.

FLATZINC is a lower-level language whose purpose is to bridge the gap between the high-level modeling in MINIZINC and the need for a fixed, and easy-to-parse, format that can simplify the implementation of the input interface of MINIZINC solvers. A MINIZINC model is typically flattened into a FLATZINC instance using the MZN2FZN compiler [38], and then solved with some MINIZINC tool.

3 From MINIZINC to OMT

We consider the problem of translating MINIZINC models into OMT problems first. Similarly to other MINIZINC solvers, we assume that the MINIZINC model is first translated into FLATZINC using the MZN2FZN standard compiler, as depicted in Fig. 1. We describe the main aspects of FZN2OMT, focusing on the challenges we have encountered and on the solutions we have adopted.

FLATZINC data-types. The first challenge is to find a suitable representation of the data-types supported by FLATZINC in SMT-LIB.

One possible choice for modeling the three basic scalar types of FLATZINC –namely `bool`, `int` and `float`– with SMT-LIB are the Boolean, bit-vector and floating-point theories respectively. However, the decision procedures for the bit-vector and floating-point numbers theories can be significantly more resource demanding than the decision procedure for the linear arithmetic theory (*LCRA*), in particular when dealing with a substantial amount of arithmetic computations. For this reason, we have opted to model FLATZINC `int` and `float` data-types with the SMT-LIB integer and rational types respectively, by default. For the case in which no substantial linear arithmetic computation is performed, we also optionally allow for encoding the FLATZINC `int` data-type as a SMT-LIB bit-vector.

For what concerns the two compound types of FLATZINC, that is the `set` and `array` data-types, we have chosen to proceed as follows. Given that OPTIMATHSAT lacks a decision procedure for the theory of finite sets [7], we model a set using the Boolean and integer theories, similarly to what has been done in [17]. The basic idea is to introduce a fresh Boolean variable for each element in the domain of a set, and use such variable as a placeholder for the membership of an integer element to the set instance. Differently from [17], we make an extensive use of cardinality networks [10] to encode constraints over the sets because they are handled more efficiently, for their nice arc-consistency properties. No action is required to encode a FLATZINC `array` into SMT-LIB, because it is used only as a container for other variables.

Floating-Point Precision. A consequence of encoding the FLATZINC `int` and `float` data-types with the linear arithmetic theory is that all of our computation is performed with infinite-precision arithmetic. This can result in a performance disadvantage wrt. other MINIZINC solvers using finite-precision arithmetic, due to the increased cost of each operation, but it has the benefit of guaranteeing the correctness of the final result of the computation.

Currently, the MINIZINC language does not allow one to express a certain quantity as an infinite-precision fraction between two constant numbers. Instead, the MZN2FZN compiler computes on-the-fly the result of any division operation between two constant integers or floating-point numbers applying the rules corresponding to the type of the operands. However, there are some instances in which we really need to be able to both express quantities and perform computation with infinite-precision arithmetic. One of such situations is to double-check the correctness of the MINIZINC models generated by the OMT2MZN compiler described in Sect. 4 (we have done this for the experimental evaluations in Sect. 5.2). In order to get around this limitation we developed a simple wrapper around the MZN2FZN compiler, called EMZN2FZN [5], that replaces any fraction among two constant `floats` with a fresh variable and, after the basic MZN2FZN compiler generated the FLATZINC model, the EMZN2FZN wrapper restores the original fractional values using the FLATZINC constraint `float.div()`.

FLATZINC Constraints. The SMT-LIB encoding of the majority of FLATZINC constraints in OPTIMATHSAT follows their definition in the FLATZINC Standard, with the exception of Pseudo-Boolean constraints, which we examine in detail later on. Several global constraints are also supported in the same way, because the OMT-solver currently lacks ad hoc and efficient decision procedures for dealing with them. Constant values and alias variables (e.g. those arising from the definition of some arrays) are propagated through the formula, so as to keep the set of problem variables as compact as possible. Those constraints requiring non-linear arithmetic –like, e.g., trigonometric, logarithmic and exponential functions– are currently not supported; this situation may change soon due to the recent extension of MATHSAT5 with a procedure for it [20].

Pseudo-Boolean Constraints. When dealing with Pseudo-Boolean sums of the form $\sum_{i=1}^{i=N} A_i \cdot w_i$, where A_i is a Boolean variable and w_i is a numerical weight, the MZN2FZN compiler associates a fresh 0/1 integer variable x_i to each A_i , and encodes the sum as $\sum_{i=1}^{i=N} x_i \cdot w_i$. Notice that the original A_i s may not be eliminated from the FLATZINC model, because they typically occur elsewhere in the problem, i.e. as part of a Boolean formula. From our own experience, this situation arises frequently, because Pseudo-Boolean sums are typically used to express cardinality constraints that have a variety of uses. As described in [57], one limitation of this naive approach is that SMT and OMT solvers do not typically handle this encoding efficiently. The main reason is that the pruning power of the conflict clause resulting from a conflicting assignment is typically limited to one specific Boolean assignment at a time, meaning that a large number of conflict clauses (possibly exponential) has to be generated along the search. As shown in [57], SMT and OMT solvers can benefit from encoding Pseudo-Boolean constraints with cardinality networks.

FZN2OMT goes through some effort in order to recognize Pseudo-Boolean sums over the integers, and replace the naive encoding with one based on cardinality networks. We note that using this technique generally results in a trade-off between solving time and the overhead of generating cardinality networks prior to starting the search, especially when dealing with a large number of variables.

Multi-objective Optimization. FZN2OMT allows for multiple optimization goals, of heterogeneous type, being defined within the same FLATZINC model. This is a non-standard extension to the FLATZINC format. Multiple objectives can be solved independently from one another, or combined into a Lexicographic or Pareto optimization goal. We refer the reader to [58] for details on the input encoding and the solver configuration.

Functionality. Given a satisfiability or optimization problem encoded in the FLATZINC format, OPTIMATHSAT can be used in the following ways (Fig. 1):

- to directly solve the problem, optionally enumerating any sub-optimal solution found during the search or all possible solutions with the same optimal value;
- to produce an OMTproblem encoded with the extended SMT-LIB format described in [58]. This problem can be directly solved with OPTIMATHSAT or, with minor transformations¹, fed as input to other OMT solvers such as BCLT and Z3.

4 From OMT to MINIZINC

In this section, we consider the problem of translating OMT formulas, encoded in the optimization-extended SMT-LIB format of [58], into MINIZINC models. Hereafter, we describe the main challenges we have faced and the solutions we have adopted. Further details about this conversion are available in [4].

General Translation Approach. The main challenge is to design an encoding from OMT to MINIZINC that is *correct* (i.e., it preserves in full the semantics of the input OMT problems), *effective* (i.e., it produces as output MINIZINC models which are as compact and easy-to-solve as possible), and *efficient* (i.e. it does it with the least consumption of time and memory). To this extent, one critical design choice is the way in which the internal representation of the input OMT formula is organized and converted in terms of MINIZINC primitives. After a preliminar experimental evaluation we determined that the sweet-spot, in terms of compactness and easiness to solve of the resulting MINIZINC model, is to adopt what we call “ ≥ 2 -father DAG-ification”: a Directed-Acyclic-Graph (DAG) internal representation of the formula where a fresh label is associated to all and only DAG nodes with at least two fathers, inlining all other nodes (see [24] for details).

¹ To make this step as easy as possible, we collected our scripts into a public repository [1].

Theories Restriction. The SMT-LIB standard describes a wide number of SMT theories, most of which have no direct counterpart in MINIZINC due to the few data-types supported (see Sect. 2)). Hence, of linear rational and integer arithmetic, and their combination. On this regard, we note that even though OMT2MZN can also handle the theory of bit-vectors, we do not cover it here because it is not used in the experimental evaluation in Sect. 5 (We cover it in the long version of this paper [24]). We leave the handling of other SMT theories to future work.

Linear Arithmetic Theory. On the surface, encoding linear arithmetic constraints coming from OMT in MINIZINC, using the *int* and *float* data-types, looks like a trivial task. In reality, this poses several challenges and it is subject to several limitations, due to a couple of facts.

First, in SMT-LIB the linear arithmetic theory requires the capability to perform *infinite-precision* computations. Unfortunately, to the best of our knowledge, no MINIZINC solver provides infinite-precision arithmetic reasoning, and the MZN2FZN compiler itself prevents representing arbitrarily-large and arbitrarily-precise quantities (e.g. the fine-grained decimal weights of the machine learning application in [59]).

Second, in OMT linear arithmetic variables are not required to be bounded and have quite often no explicit domain (i.e. they lack a lower-bound, an upper-bound or both), because it is not necessary for the problem at hand or it is implied by other constraints. This is in contrast with MINIZINC, whereby linear arithmetic variables are expected to have a finite domain and, when they lack one, their domain appears to be capped with some solver-dependent pair of values.

These restrictions are currently part of the MINIZINC language and the target application domain, and we do not see any obvious work-around solving them. We note that although there exist methods for bounding all variables in a given LP problem (e.g. [47]), these have been deemed too impractical at this stage of our investigation. Nonetheless, we have chosen to translate SMT-LIB linear arithmetic constraints with a corresponding MINIZINC encoding based on the *int* and *float* data-types. Although the encoding is not always applicable, it does still allow one to correctly translate a number of interesting OMT problems into MINIZINC, as witnessed by our experimental evaluation in Sect. 5.2. More in detail, the translation is done as follows. We declare each integer variable as unbounded, and then extend the MINIZINC model with the appropriate constraints bounding its domain when the input OMT formula contains any such information. Our empirical observation is that MINIZINC models generated in this way are correctly handed by all MINIZINC solvers which we have tried, with the exception of GUROBI, which returns an “unsupported” message. Floating-Point variables, instead, are always declared with a user-defined domain. This is because all of the MINIZINC solvers we have tried, among those that can handle floating-point constraints, require such information.

Other OMT Functionalities. Several problems of OMT interest require the capability of dealing with soft-constraints (i.e. Weighted MAXSMT) and also with multiple objectives, that are either considered independent goals or combined in a Lexicographic or Pareto-like fashion. To the best of our knowledge, the MINIZINC standard does not

allow for an explicit encoding of soft-constraints, nor to deal with more than one objective function at a time.

We encode (weighted) MAXSMT problems using a standard Pseudo-Boolean encoding, such as the one used in [53]. When dealing with OMT problems that contain N goals $\text{obj}_1, \dots, \text{obj}_N$, for $N > 1$, we use the following approach. If these objectives are independent targets, we generate N MINIZINC models, each with a different goal obj_i , and separately solve each model. If instead the multiple objectives belong to a Lexicographic OMT problem, then we generate a unique MINIZINC model that leverages the lexicographic-optimization functionality provided by MINISEARCH [49]. (In all other cases, MINISEARCH is not used). We do not have any encoding for dealing with Pareto-optimization, yet.

5 Experimental Evaluations

In this section we present an extensive empirical evaluation comparing OMT tools with many state-of-the-art CP tools on CP problems coming from the MINIZINC challenge (Sect. 5.1), and on OMT problems coming mostly from formal verification (Sect. 5.2).

The OMT solvers under evaluation are BCLT, OPTIMATHSAT (v. 1.6.0) and Z3 (v. 4.8.5). These are compared with some of the top-scoring solvers that participated at recent editions of the MINIZINC challenge, including CHOCO (v. 4.0.4), CHUFFED, G12(FD) (v. 1.6.0), GECODE (v. 6.0.1), GUROBI (v. 8.0.1), HAIFACSP (v. 1.3.0), JACOP (v. 4.5.0), IZPLUS (v. 3.5.0), OR-TOOLS (v. 6.7.4981) and PICAT (v. 2.4).

Remark 1. We could not include FZN2SMT [17, 18] in our experimental evaluation because it is not compatible with the features of MINIZINC that have been added since version 2.0.

We run all these experimental evaluations on two identical *8-core 2.20 Ghz Xeon* machines with 64 GB of RAM and running *Ubuntu Linux*. All the benchmark-sets, the tools and the scripts used to run these experiments, and some of the plots for the results in Tables 1, 2 and 3 which could not fit into this paper, can be downloaded from [2].

We stress the fact that the goal of these experiments is not to establish a winner among OMT and MINIZINC tools; rather, it is to assess the correctness, effectiveness and efficiency of our OMT-to-CP and CP-to-OMT encoders and, more generally, to investigate the feasibility of solving MINIZINC problems with OMT tools and vice versa, and to identify the criticalities in terms of efficiency and correctness in these processes.

5.1 Evaluation on MINIZINC Benchmark Sets

We consider the benchmark-sets used at the MINIZINC Challenge of 2016 (MC16) and 2019 (MC19), each comprised by 100 instances. For compatibility reasons, the version of MZN2FZN used to convert the problems to the FLATZINC format differs between the two benchmark-sets. We use version 2.2.1 and 2.3.2 (with patches) for the problems in MC16 and MC19 respectively. Due to recent changes in the FLATZINC format that affect the benchmarks in MC19, the version of some MINIZINC tools differs from what

described in Sect. 5 (see Table 1). In some cases, we had to download and compile the latest source available for the tool, i.e. the “nightly” version.

We run each MINIZINC solver with the corresponding directory of global constraints, and we run each MINIZINC and OMT tool with the default options. We consider two OMT encodings of the original FLATZINC problems, LA and BV. The first encodes the FLATZINC *int* type with the theory of *linear integer arithmetic*, whereas the second is based on the theory of *bit-vectors*. We evaluate each OMT solver on both SMT-LIB encodings, except for BCLT that has no support for bit-vector optimization. For uniformity reasons with the other OMT solvers, we evaluate OPTIMATHSAT using its SMT-LIB interface only, using thus its FZN2OMT interface as an external tool, like with the other OMT solvers. We note that the solving time for all OMT solvers includes the time required for translating the formula from the FLATZINC to the SMT-LIB format. Each solver, either OMT or MINIZINC, is given up to 1200s. to solve each problem, not including the time taken by MZN2FZN to flatten it.

We verify the correctness of the results by automatically checking that all terminating solvers agree on the (possibly optimal) solution and, when this is not the case, we manually investigate the inconsistency.

Experiment Results. The results of this experiment are shown in Table 1, with separate numbers for satisfiability (s) and optimization (o) instances in each benchmark-set. Using the experimental data, we separately computed the virtual best configuration among all MINIZINC solvers (i.e. VIRTUAL BEST(MINIZINC)), all OMT solvers (i.e. VIRTUAL BEST(OMT)), and also the virtual best among all tools considered in the experiment (i.e. VIRTUAL BEST(ALL)). The last two columns in the table list the number of problems solved by the given configuration in the same amount of time as the VIRTUAL BEST() of each group (col. *BT1*) and as the VIRTUAL BEST(all) (col. *BT2*).

We start by looking at the MINIZINC solvers in Table 1. The performance ladder is dominated by OR-TOOLS(SAT) and PICAT(SAT), closely followed by GUROBI, HAIFACSP and CHUFFED (in MC19). By looking at column *BT1*, we observe that the top-performing MINIZINC solvers tend to dominate over all the others. Looking at the results of the MC19 experiment, we notice a significant increase in the number of errors with respect to the benchmark-set of the MC16 edition, as well as a handful of problems solved incorrectly. In the case of GUROBI and PICAT(SAT), the MZN2FZN compiler encountered an error over a few instances. As a consequence, the total number of problems is smaller than 100 for both tools. After taking a closer look, we ascribe this phenomenon to the recent changes in the MINIZINC/FLATZINC format, that has created some minor issues with some tools that have not been adequately updated.

Looking at the OMT tools only, we observe that Z3 has leading performance over the other solvers. When compared to the MINIZINC solvers, the OMT solvers place themselves in the middle of the rank on both benchmark-sets. Given the fact that none of the OMT solvers has specialized procedures or encodings for dealing with global constraints, we consider this an interesting result.

5.2 Evaluation on OMT Benchmark Sets

In this experimental evaluation we use OMT formulas taken from well-known, publicly available, repositories. We characterize these benchmark-sets as follows:

Table 1. MINIZINC Challenge formulas. The columns list the total number of instances (inst.), of timeouts (timeout), of run-time errors (error), of unsupported problems (unsup.), of incorrectly solved instances (incor.), of correctly solved instances (correct), the total solving time for all solved instances (time), the number of instances solved in the shortest time within the same category (BT1) and those solved in the shortest time considering all tools (BT2).

| tool, configuration & encoding | MINIZINC Challenge 2016 | | | | | | | | | | | | | | | | | |
|--------------------------------|-------------------------|----|---------|----|-------|----|--------|----|--------|---|---------|----|-----------|-------|-----|----|-----|----|
| | inst. | | timeout | | error | | unsup. | | incor. | | correct | | time (s.) | | BT1 | | BT2 | |
| | s | o | s | o | s | o | s | o | s | o | s | o | s | o | s | o | s | o |
| MINIZINC Challenge 2016 | | | | | | | | | | | | | | | | | | |
| PICAT(CP) | 15 | 85 | 9 | 70 | 0 | 0 | 0 | 0 | 0 | 0 | 6 | 15 | 2281 | 6043 | 0 | 0 | 0 | 0 |
| G12(FD) | 15 | 85 | 4 | 71 | 1 | 3 | 0 | 0 | 0 | 0 | 10 | 11 | 4436 | 4220 | 0 | 0 | 0 | 0 |
| CHOCO() | 15 | 85 | 3 | 50 | 0 | 0 | 0 | 0 | 0 | 0 | 12 | 35 | 4256 | 11423 | 1 | 0 | 1 | 0 |
| IZPLUS() | 15 | 85 | 6 | 44 | 0 | 0 | 0 | 0 | 0 | 0 | 9 | 41 | 999 | 5492 | 3 | 4 | 3 | 4 |
| CHUFFED() | 15 | 85 | 2 | 40 | 0 | 0 | 5 | 0 | 0 | 0 | 8 | 45 | 635 | 4187 | 0 | 5 | 0 | 5 |
| JACoP() | 15 | 85 | 3 | 39 | 0 | 0 | 0 | 0 | 0 | 0 | 12 | 46 | 3411 | 12825 | 0 | 0 | 0 | 0 |
| GUROBI() | 15 | 85 | 6 | 22 | 0 | 0 | 0 | 0 | 0 | 0 | 9 | 63 | 2346 | 3037 | 0 | 15 | 0 | 15 |
| HAIFACSP() | 15 | 85 | 4 | 23 | 0 | 0 | 0 | 0 | 0 | 0 | 11 | 62 | 591 | 4444 | 0 | 11 | 0 | 11 |
| PICAT(SAT) | 15 | 85 | 1 | 26 | 0 | 0 | 0 | 0 | 0 | 0 | 14 | 59 | 151 | 7293 | 10 | 1 | 10 | 1 |
| OR-TOOLS(SAT) | 15 | 85 | 1 | 15 | 0 | 0 | 0 | 0 | 0 | 0 | 14 | 70 | 555 | 1338 | 1 | 45 | 1 | 45 |
| VIRTUAL BEST(MINIZINC) | 15 | 85 | 0 | 7 | 0 | 0 | 0 | 0 | 0 | 0 | 15 | 78 | 146 | 3514 | - | - | - | - |
| MINIZINC Challenge 2019 | | | | | | | | | | | | | | | | | | |
| OPTIMATHSAT(INT) | 15 | 85 | 10 | 38 | 0 | 0 | 0 | 0 | 0 | 0 | 5 | 47 | 604 | 4856 | 1 | 20 | 0 | 0 |
| OPTIMATHSAT(BV) | 15 | 85 | 2 | 42 | 0 | 0 | 0 | 0 | 0 | 0 | 13 | 43 | 3664 | 8561 | 11 | 2 | 0 | 0 |
| BCLT(INT) | 15 | 85 | 10 | 33 | 0 | 0 | 0 | 0 | 0 | 0 | 5 | 52 | 1117 | 5998 | 0 | 15 | 0 | 2 |
| Z3(INT) | 15 | 85 | 10 | 32 | 0 | 0 | 0 | 0 | 0 | 0 | 5 | 53 | 676 | 10424 | 0 | 11 | 0 | 0 |
| Z3(BV) | 15 | 85 | 5 | 28 | 0 | 0 | 0 | 0 | 0 | 0 | 10 | 57 | 2938 | 11113 | 2 | 19 | 0 | 0 |
| VIRTUAL BEST(OMT) | 15 | 85 | 1 | 21 | 0 | 0 | 0 | 0 | 0 | 0 | 14 | 64 | 3842 | 6432 | - | - | - | - |
| VIRTUAL BEST(ALL) | 15 | 85 | 0 | 7 | 0 | 0 | 0 | 0 | 0 | 0 | 15 | 78 | 146 | 3514 | - | - | - | - |
| MINIZINC Challenge 2019 | | | | | | | | | | | | | | | | | | |
| PICAT(CP) [2.7B12] | 10 | 90 | 8 | 67 | 0 | 11 | 0 | 5 | 0 | 0 | 2 | 7 | 54 | 1440 | 1 | 0 | 1 | 0 |
| IZPLUS() | 10 | 90 | 5 | 71 | 0 | 4 | 0 | 0 | 0 | 0 | 5 | 15 | 14 | 3077 | 1 | 3 | 1 | 3 |
| G12(FD) | 10 | 90 | 5 | 64 | 0 | 10 | 0 | 0 | 0 | 0 | 5 | 16 | 323 | 4010 | 0 | 0 | 0 | 0 |
| CHOCO(STD) | 10 | 90 | 4 | 63 | 0 | 5 | 0 | 0 | 0 | 0 | 6 | 22 | 415 | 4312 | 0 | 0 | 0 | 0 |
| GECODE() [6.2.0] | 10 | 90 | 4 | 63 | 0 | 0 | 0 | 0 | 0 | 0 | 6 | 27 | 420 | 5094 | 0 | 6 | 0 | 6 |
| JACoP() [4.8] | 10 | 90 | 4 | 55 | 0 | 6 | 0 | 0 | 0 | 0 | 6 | 29 | 260 | 5467 | 0 | 1 | 0 | 1 |
| HAIFACSP() | 10 | 90 | 0 | 47 | 0 | 10 | 0 | 2 | 2 | 8 | 31 | 2 | 6408 | 4 | 7 | 4 | 4 | |
| CHUFFED() [NIGHTLY] | 10 | 90 | 0 | 43 | 0 | 0 | 5 | 10 | 0 | 0 | 5 | 37 | 1 | 4886 | 3 | 19 | 3 | 19 |
| GUROBI() [8.1.1] | 10 | 80 | 0 | 48 | 0 | 0 | 0 | 0 | 0 | 0 | 10 | 32 | 705 | 2895 | 2 | 6 | 2 | 4 |
| PICAT(SAT) [2.7B12] | 10 | 90 | 0 | 45 | 0 | 5 | 0 | 0 | 0 | 1 | 10 | 39 | 275 | 9894 | 0 | 7 | 0 | 5 |
| OR-TOOLS(SAT) [NIGHTLY] | 10 | 90 | 5 | 42 | 0 | 3 | 0 | 0 | 0 | 0 | 5 | 45 | 8 | 7239 | 0 | 13 | 0 | 11 |
| VIRTUAL BEST(MINIZINC) | 10 | 90 | 0 | 29 | 0 | 0 | 0 | 0 | 0 | 0 | 10 | 61 | 9 | 5247 | - | - | - | - |
| MINIZINC Challenge 2019 | | | | | | | | | | | | | | | | | | |
| OPTIMATHSAT(INT) [1.6.4.1] | 10 | 90 | 5 | 62 | 0 | 0 | 0 | 0 | 0 | 0 | 5 | 28 | 4 | 3650 | 2 | 10 | 0 | 0 |
| OPTIMATHSAT(BV) [1.6.4.1] | 10 | 90 | 4 | 59 | 0 | 5 | 0 | 0 | 0 | 0 | 6 | 26 | 484 | 7271 | 0 | 1 | 0 | 0 |
| BCLT(INT) | 10 | 90 | 5 | 60 | 0 | 0 | 0 | 0 | 0 | 0 | 5 | 30 | 6 | 3369 | 0 | 6 | 0 | 5 |
| Z3(INT) | 10 | 90 | 5 | 64 | 0 | 0 | 0 | 0 | 0 | 0 | 5 | 26 | 4 | 5358 | 3 | 6 | 0 | 1 |
| Z3(BV) | 10 | 90 | 0 | 55 | 0 | 2 | 0 | 0 | 0 | 0 | 10 | 33 | 1629 | 7550 | 5 | 17 | 0 | 3 |
| VIRTUAL BEST(OMT) | 10 | 90 | 0 | 48 | 0 | 2 | 0 | 0 | 0 | 0 | 10 | 40 | 1624 | 5179 | - | - | - | - |
| VIRTUAL BEST(ALL) | 10 | 90 | 0 | 29 | 0 | 0 | 0 | 0 | 0 | 0 | 10 | 61 | 9 | 4919 | - | - | - | - |

- *SAL* [integers]: 66 SMT-based Bounded Model Checking and K-Induction parametric problems created with the *SAL* model checker [6];
- *SAL* [rationals]: as above, with problems on the rationals;
- *Symba* [rationals]: 2632 *bounded*² software verification instances derived from a set of C programs used in the Software Verification Competition of 2013 [36];
- *Jobshop and Strip Packing* [rationals]: 190 problems taken from [51,53];
- *Machine Learning* [rationals]: 510 OMTinstances generated with the *PYLMT* tool based on Machine Learning Modulo Theories [59].

² We discarded any *unbounded* instance in the original benchmark-set in [36].

The first benchmark-set is on the integers, whereas the other four are on the rationals. We stress the fact that all formulas contained in all benchmark-sets are *satisfiable*.

Remark 2. Although there exists a repository of multi-objective OMT formulas (e.g. [36,56]), we have chosen to not include these in our experimental evaluation. The reason for this is twofold. First, such comparison would likely be unfair wrt. CP tools because that the workaround for dealing with multi-independent OMT formulas described in Sect. 4 is not competitive with the integrated optimization schema provided by OMT solvers [36,56]. In fact, the experimental evidence in [36,56] collected on a group of OMT solvers indicates that the latter approach can be an order of magnitude faster than the former one. Second, the workaround for dealing with lexicographic-optimization is limited by the fact that MINISEARCH is not fully compatible with recent versions of MINIZINC, and it only works with a restricted set of tools.

We have used the OMT2MZN tool described in Sect. 4 to translate each OMT formula to the MINIZINC format. OMT2MZN is written in Python and it is built on top of PYSMT [30], a general-purpose Python library for solving SMT problems, and it is available at [4]. During this step, it has been necessary to impose a finite domain to any unconstrained SMT-LIB rational variable, because otherwise none of the MINIZINC solvers would have been able to deal with them. We have experimented with two different domains: the largest feasible domain for floating-point variables of 32 bits (i.e. $\pm 3.402823e + 38$) for the first two benchmark-sets, and the largest feasible domain for integer variables (i.e. $\pm 2^{31}$) for the last two.

We consider two OPTIMATHSAT configurations: OPTIMATHSAT(SMT), solving the original OMTformulas, and OPTIMATHSAT(FZN), executed on the generated MINIZINC instances. The benefits of this choice is two-fold. First, we can double-check the correctness of such encoding, by comparing the optimum models generated in the two cases. Second, we can verify whether there is any performance loss caused by the encoding of the formula.

Only four of the MINIZINC solvers listed in Sect. 5 support floating-point reasoning. This limited the number of tools that could be used with some OMT benchmark-sets. The running-time of each MINIZINC solver reported in these experiments (including OPTIMATHSAT(FZN)) is comprehensive of the time taken by the MZN2FZN compiler, because the latter can sometime solve the input formulas on its own. The overall timeout is set to 600 s.

Notice that the optimal solutions found by OPTIMATHSAT(SMT) have been previously independently verified with a third-party SMT tool as reported in previous publications [55–57].³ Therefore, we verify the correctness of the results found by any other configuration by comparing them with those found by OPTIMATHSAT(SMT), and otherwise mark the result as “unverified”.

Experimental Results over the Integers. In this experiment, we evaluate the SAL (*over integers*) benchmark-set. The results are collected in Table 2.

³ For every OMT problem $\langle \varphi, \text{obj} \rangle$ s.t. OPTIMATHSAT(SMT) returns a minimum value min for obj on the formula φ , we say min is correct iff $\varphi \wedge (\text{obj} = min)$ is satisfiable and $\varphi \wedge (\text{obj} < min)$ is unsatisfiable. (Dual for maximization.).

Table 2. SAL over integers. A SAT result is marked as *correct* when the objective value matches the reference solution provided by OPTIMATHSAT(SMT) (when run without a timeout), as *incorrect* otherwise.

| tool & configuration | inst. | timeout | tool-er. | unsupp. | terminated | | | | |
|-------------------------|-------|---------|----------|---------|------------|---------|----------------|----------------|----------------|
| | | | | | incor. | correct | tot. time (s.) | avg. time (s.) | med. time (s.) |
| GUROBI() | 66 | 0 | 0 | 66 | 0 | 0 | 0 | 0.00 | 0.00 |
| G12(FD) | 66 | 0 | 66 | 0 | 0 | 0 | 0 | 0.00 | 0.00 |
| IZPLUS() | 66 | 0 | 66 | 0 | 0 | 0 | 0 | 0.00 | 0.00 |
| JACOPO() | 66 | 0 | 66 | 0 | 0 | 0 | 0 | 0.00 | 0.00 |
| CHUFFED() | 66 | 19 | 47 | 0 | 0 | 0 | 0 | 0.00 | 0.00 |
| OR-TOOLS(SAT) | 66 | 57 | 9 | 0 | 0 | 0 | 0 | 0.00 | 0.00 |
| CHOCO() | 66 | 66 | 0 | 0 | 0 | 0 | 0 | 0.00 | 0.00 |
| HAIFACSP() | 66 | 66 | 0 | 0 | 0 | 0 | 0 | 0.00 | 0.00 |
| PICAT(CP) | 66 | 66 | 0 | 0 | 0 | 0 | 0 | 0.00 | 0.00 |
| GECODE() | 66 | 66 | 0 | 0 | 0 | 0 | 0 | 0.00 | 0.00 |
| GUROBI(L) | 66 | 63 | 0 | 0 | 0 | 3 | 166 | 55.49 | 52.44 |
| PICAT(SAT) | 66 | 62 | 0 | 0 | 0 | 4 | 1667 | 416.85 | 467.09 |
| VIRTUAL BEST(MINI-ZINC) | 66 | 62 | 0 | 0 | 0 | 4 | 718 | 179.51 | 78.54 |
| OPTIMATHSAT(FZN) | 66 | 18 | 0 | 0 | 0 | 48 | 7113 | 148.20 | 70.52 |
| VIRTUAL BEST(FZN) | 66 | 18 | 0 | 0 | 0 | 48 | 7113 | 148.20 | 70.52 |
| OPTIMATHSAT(SMT) | 66 | 22 | 0 | 0 | 0 | 44 | 2657 | 60.41 | 18.72 |
| VIRTUAL BEST(ALL) | 66 | 16 | 0 | 0 | 0 | 50 | 5037 | 100.75 | 25.13 |

Table 3. OMT Problems defined over the rationals. A SAT result is marked as *correct* when the objective value matches the reference solution provided by OPTIMATHSAT(SMT) with an absolute error $\Delta < 10^{-6}$. A result is marked as *unverified* when we have no reference solution and *incorrect* if neither of the previous two conditions apply.

| tool & configuration | instances | timeout | tool-errors | terminated | | | | | incorrect results | | | | | | |
|---------------------------------------------|-----------|---------|-------------|------------|----------|------------|----------------|----------------|-------------------|-------|-------------------------|-------------------------|-------------------------|----------------------|----------------------|
| | | | | incorrect | verified | unverified | tot. time (s.) | avg. time (s.) | med. time (s.) | unsat | $\Delta \wedge 10^{-6}$ | $\Delta \wedge 10^{-3}$ | $\Delta \wedge 10^{-1}$ | $\Delta \wedge 10^0$ | $\Delta \wedge 10^1$ |
| SAL, Symba, Jobshop and Strippacking | | | | | | | | | | | | | | | |
| GECODE() | 2888 | 2733 | 0 | 0 | 155 | 0 | 10800 | 69.68 | 18.63 | 0 | 0 | 0 | 0 | 0 | |
| G12(MIP) | 2888 | 10 | 0 | 2855 | 0 | 23 | 317 | 13.79 | 11.21 | 2765 | 90 | 90 | 86 | 39 | 0 |
| GUROBI() | 2888 | 48 | 0 | 2728 | 104 | 8 | 3961 | 35.37 | 2.14 | 2684 | 44 | 32 | 32 | 1 | 0 |
| VIRTUAL BEST(MINI-ZINC) | 2888 | 0 | 0 | 2628 | 237 | 23 | 13801 | 53.08 | 6.97 | - | - | - | - | - | - |
| OPTIMATHSAT(FZN) | 2888 | 31 | 0 | 0 | 2854 | 3 | 22320 | 7.81 | 0.40 | 0 | 0 | 0 | 0 | 0 | 0 |
| VIRTUAL BEST(FZN) | 2888 | 0 | 0 | 11 | 2854 | 23 | 20674 | 7.19 | 0.40 | - | - | - | - | - | - |
| OPTIMATHSAT(SMT) | 2888 | 23 | 0 | 0 | 2865 | 0 | 15676 | 5.47 | 0.08 | 0 | - | - | - | - | - |
| VIRTUAL BEST(ALL) | 2888 | 0 | 0 | 0 | 2865 | 23 | 15183 | 5.26 | 0.08 | - | - | - | - | - | - |
| Machine Learning | | | | | | | | | | | | | | | |
| GECODE() | 510 | 322 | 0 | 164 | 24 | 0 | 11 | 0.44 | 0.43 | 147 | 17 | 17 | 2 | 0 | 0 |
| G12(MIP) | 510 | 108 | 0 | 400 | 2 | 0 | 225 | 112.47 | 112.47 | 400 | 0 | 0 | 0 | 0 | 0 |
| GUROBI() | 510 | 9 | 0 | 472 | 28 | 1 | 201 | 6.92 | 3.17 | 468 | 4 | 4 | 2 | 0 | 0 |
| VIRTUAL BEST(MINI-ZINC) | 510 | 9 | 0 | 464 | 36 | 1 | 383 | 10.34 | 0.46 | - | - | - | - | - | - |
| OPTIMATHSAT(FZN) | 510 | 7 | 0 | 237 | 263 | 3 | 2797 | 10.52 | 2.21 | 177 | 60 | 59 | 0 | 0 | 0 |
| OPTIMATHSAT(FZN+E) | 510 | 92 | 0 | 0 | 415 | 3 | 1197 | 2.86 | 2.03 | 0 | 0 | 0 | 0 | 0 | 0 |
| VIRTUAL BEST(FZN) | 510 | 7 | 0 | 83 | 417 | 3 | 1366 | 3.25 | 2.03 | - | - | - | - | - | - |
| OPTIMATHSAT(SMT) | 510 | 10 | 0 | 0 | 500 | 0 | 5766 | 11.53 | 12.15 | 0 | - | - | - | - | - |
| VIRTUAL BEST(ALL) | 510 | 7 | 0 | 0 | 500 | 3 | 2290 | 4.55 | 2.05 | - | - | - | - | - | - |

We notice first that OPTIMATHSAT(FZN) always produces correct results and it shows comparable performances in terms of number of problems solved wrt. the base-

line OPTIMATHSAT(SMT), solving even 4 problems more. (We conjecture that the latter fact should be attributed to the limited, but effective, deduction capabilities of the MZN2FZN compiler, that may have helped OPTIMATHSAT in solving the input formulas.) This suggests that, at least on problems on the integers, OMT2MZN is efficient and effective and does not affect correctness.

In general, MINIZINC solvers do not seem to deal efficiently with this benchmark-set. Some tools have experienced some internal error (e.g. dumped-core, segmentation fault), some others have been killed to a high memory consumption (over 32GB), whereas the majority of the remaining tools had a timeout.

We explain this behavior with the fact that the given benchmark set is characterized by the presence of a heavy Boolean structure combined with arithmetical constraints, which requires the efficient combination of strong Boolean-reasoning capabilities (e.g., efficiently handling chains of unit propagations) with strong arithmetical-solving&optimization capabilities, which is a typical feature of OMT solvers.

None of the input formulas was initially supported by GUROBI. After restricting the bound of every integer variable to $\pm 10^6$, GUROBI(L) was able to solve 3 instances within the timeout. Among the MINIZINC solvers, the best result is obtained by PICAT(SAT), that solved 4 problems out of 66.

Experimental Results Over the Rationals. We consider first the first three benchmark-sets over the rationals: SAL over rationals, Symba, JobShop&Strip-Packing. (Separate tables for the four benchmarks are reported in the extended version of this paper [24].) Of all MINIZINC solvers we have tried, only three are able to deal with floating-point constraints. The results are shown in Table 3. Since each of the input formulas is satisfiable, we consider a result incorrect either when it is equal to UNSAT, or when the relative error Δ exceeds 10^{-6} , s.t.: $\Delta \stackrel{\text{def}}{=} \frac{|o_{smt} - o_{fzn}|}{|o_{smt}|}$, o_{smt} and o_{fzn} being the optimal value found by OPTIMATHSAT(SMT) and the optimal value found by the MINIZINC solver under test respectively. (Recall that the former was previously checked to be correct.)

Similarly to the previous experiment on the integers, OPTIMATHSAT(FZN) always produces correct results, and display comparable performance wrt. OPTIMATHSAT(OMT) in terms of number of instances being solved, solving somewhat fewer problems. This is not the case of the other three MINIZINC solvers. Among these, GECODE experienced a timeout on the majority of the formulas being considered, G12(MIP) returned mostly incorrect answers, whereas GUROBI seems to have the best performance, in particular on the third benchmark-set.

We attribute the large number of incorrect results returned by all three MINIZINC solvers to the fact that these tools use finite-precision floating-point arithmetic internally. The incorrect behavior of some of these solvers (e.g. GUROBI) can also be partially explained with the large domain of floating-point variables in these problems. However, given the nature of these input instances, it was not possible for us to assign a smaller domain to each variable in the problem *a priori*.

We analyze separately the results for the last benchmark-set reported in Table 3. The peculiar aspect of the Machine Learning benchmark-set [59] is that it is characterized by Pseudo-Boolean sums over rational weights, and by very fine-grained rational values⁴.

⁴ For example, $\frac{1799972218749879}{2251799813685248}$ is a sample weight value from problems in [59].

Unfortunately, these fine-grained rational values are rounded by the standard `MZN2FZN` compiler, which causes the incorrect results even of `OPTIMATHSAT(FZN)` in Table 3, despite the fact that `OPTIMATHSAT` uses infinite-precision arithmetic.

In order to overcome this issue, we leverage the `EMZN2FZN` compiler described in Sect. 3 so that the original fractional values are preserved in the resulting `FLATZ-INCmodel`, and show that with this approach `OPTIMATHSAT` does not produce incorrect results any longer (configuration `OPTIMATHSAT(FZN+E)` in Table 3), solving correctly 152 problems more than `OPTIMATHSAT(FZN)`.

Overall, since there are at least 237 formulas affected by the above issue with the `MZN2FZN` compiler, we avoid an in-depth discussion of the results obtained by the other `MINIZINC` solvers. However, at a first glance the situation does not seem to differ from the other benchmark-sets over the rationals.

5.3 Discussion

On the whole, from our experiments, OMT tools appear to be still at some disadvantage when dealing with `MINIZINC` problems wrt. specific tools, and vice versa.

On the one hand, OMT solvers seem to be penalized by their lack of efficient ad hoc decision procedures for dealing with global constraints. Moreover, the approach taken by the `MZN2FZN` compiler, that creates lots of alias Boolean, integer and floating-point variables for dealing with Pseudo-Boolean constraints, is particularly challenging to deal with efficiently by an OMT solver.

On the other hand, `MINIZINC` solvers seem to suffer with problems needing an arithmetic-reasoning component combined with heavy Boolean-reasoning component. Even more importantly, the lack of infinite-precision linear arithmetic procedures causes a number of incorrect results when dealing with OMT problems over the rationals. Both of these points need to be addressed in order to deal with the vast number of Formal Verification and Model Checking applications in the SMT/OMT domain.

6 Conclusions and Future Work

In this paper we have taken a first step forward towards bridging the `MINIZINC` and the OMT communities. The ultimate goal is to obtain a correct, effective and efficient fully-automated system for translating problems from one community to the other, so as to extend the application domain of both communities. With our experimental evaluation, we have identified some criticalities that need to be addressed by each community in order to solidify this union.

We plan to push this investigation forward as follows. In the short term, we plan to address the inefficient handling of Pseudo-Boolean constraints over the rationals revealed by the experimental evaluation in Sect. 5.2. In order to deal with those `FLATZ-INC` constraints that require non-linear arithmetic, we envisage an opportunity to either extend `OPTIMATHSAT` with proper handling of the non-linear arithmetic theory [20] or to experiment with an encoding based on the floating-point theory [60]. This objective goes hand in hand with the extension of `OMT2MZN` to deal with other SMT theories. In the long term, OMT solving may also benefit from adopting efficient ad hoc decision

procedures for frequently used global constraints. Finally, we plan to broaden the scope of our investigation and include other OMT solvers in our study.

References

1. FZN2OMT. <https://github.com/PatrickTrentin88/fzn2omt>
2. Benchmarks, Tools and Data. <http://disi.unitn.it/trentin/resources/cpaior2020.tar.xz>
3. MathSAT 5. <http://mathsat.fbk.eu/>
4. OMT2MZN. <https://github.com/cespio/omt2mzn>
5. EMZN2FZN Repository. <https://github.com/PatrickTrentin88/emzn2fzn>
6. Sal, symbolic analysis laboratory. <http://sal.csl.sri.com>
7. SMT-LIB Format for Finite Lists, Sets and Maps. <https://www.cprover.org/SMT-LIB-LSM/>
8. Ansótegui, C., Bofill, M., Palahí, M., Suy, J., Villaret, M.: Solving weighted CSPs with meta-constraints by reformulation into satisfiability modulo theories. *Constraints* **18**(2), 236–268 (2013). <https://doi.org/10.1007/s10601-012-9131-1>
9. Araújo, R., Bessa, I., Cordeiro, L.C., Filho, J.E.C.: SMT-based verification applied to non-convex optimization problems. In: 2016 VI Brazilian Symposium on Computing Systems Engineering (SBESC), November 2016
10. Asín, R., Nieuwenhuis, R., Oliveras, A., Rodríguez-Carbonell, E.: Cardinality networks: a theoretical and empirical study. *Constraints* **16**(2), 195–221 (2011). <https://doi.org/10.1007/s10601-010-9105-0>
11. Barrett, C., Ranise, S., Stump, A., Tinelli, C.: The satisfiability modulo theories library (SMT-LIB) (2010). <http://www.smtlib.org>
12. Barrett, C., Sebastiani, R., Seshia, S.A., Tinelli, C.: *Satisfiability Modulo Theories*, vol. 185, chap. 26, pp. 825–885. IOS Press, February 2009
13. Bertolissi, C., dos Santos, D.R., Ranise, S.: Solving multi-objective workflow satisfiability problems with optimization modulo theories techniques. In: SACMAT. ACM (2018)
14. Bian, Z., Chudak, F., Macready, W., Roy, A., Sebastiani, R., Varotti, S.: Solving SAT and MaxSAT with a quantum annealer: foundations and a preliminary report. In: Dixon, C., Finger, M. (eds.) *FroCoS 2017*. LNCS (LNAI), vol. 10483, pp. 153–171. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66167-4_9
15. Bjørner, N., Phan, A.-D., Fleckenstein, L.: vZ - an optimizing SMT solver. In: Baier, C., Tinelli, C. (eds.) *TACAS 2015*. LNCS, vol. 9035, pp. 194–199. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46681-0_14
16. Bofill, M., Nieuwenhuis, R., Oliveras, A., Rodríguez-Carbonell, E., Rubio, A.: The barcelogic SMT solver. In: Gupta, A., Malik, S. (eds.) *CAV 2008*. LNCS, vol. 5123, pp. 294–298. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-70545-1_27
17. Bofill, M., Palahí, M., Suy, J., Villaret, M.: Solving constraint satisfaction problems with SAT modulo theories. *Constraints* **17**(3), 273–303 (2012). <https://doi.org/10.1007/s10601-012-9123-1>
18. Bofill, M., Suy, J., Villaret, M.: A system for solving constraint satisfaction problems with SMT. In: Strichman, O., Szeider, S. (eds.) *SAT 2010*. LNCS, vol. 6175, pp. 300–305. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14186-7_25
19. Candeago, L., Larraz, D., Oliveras, A., Rodríguez-Carbonell, E., Rubio, A.: Speeding up the constraint-based method in difference logic. In: Creignou, N., Le Berre, D. (eds.) *SAT 2016*. LNCS, vol. 9710, pp. 284–301. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-40970-2_18
20. Cimatti, A., Griggio, A., Irfan, A., Roveri, M., Sebastiani, R.: Incremental linearization for satisfiability and verification modulo nonlinear arithmetic and transcendental functions. *ACM Trans. Comput. Logic* **19**(3), 1–52 (2018)

21. Cimatti, A., Griggio, A., Schaafsma, B.J., Sebastiani, R.: The MathSAT5 SMT solver. In: Piterman, N., Smolka, S.A. (eds.) TACAS 2013. LNCS, vol. 7795, pp. 93–107. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36742-7_7
22. Collavizza, H., Rueher, M.: Exploration of the capabilities of constraint programming for software verification. In: Hermanns, H., Palsberg, J. (eds.) TACAS 2006. LNCS, vol. 3920, pp. 182–196. Springer, Heidelberg (2006). https://doi.org/10.1007/11691372_12
23. Collavizza, H., Rueher, M., Van Hentenryck, P.: CPBPV: a constraint-programming framework for bounded program verification. *Constraints* **15**(2), 238–264 (2010). <https://doi.org/10.1007/s10601-009-9089-9>
24. Contaldo, F., Trentin, P., Sebastiani, R.: From minizinc to optimization modulo theories, and back (extended version) (2019). <https://arxiv.org/abs/1912.01476>
25. SMTURL: SmtLibav2. www.smtlib.cs.uiowa.edu/
26. Elgabou, H.A., Frisch, A.M.: Encoding the lexicographic ordering constraint in SAT modulo theories. In: Proceedings of Thirteenth International Workshop on Constraint Modelling and Reformulation, September 2014
27. Feydy, T., Stuckey, P.J.: Lazy clause generation reengineered. In: Gent, I.P. (ed.) CP 2009. LNCS, vol. 5732, pp. 352–366. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-04244-7_29
28. Frisch, A.M., Giannaros, P.A.: SAT Encodings of the At-Most-k Constraint Some Old, Some New, Some Fast, Some Slow (2010)
29. Frisch, A.M., Palahí, M.: Anomalies in SMT solving: difficulties in modelling combinatorial problems. In: Proceedings of Thirteenth International Workshop on Constraint Modelling and Reformulation, September 2014
30. Gario, M., Micheli, A.: PySMT: a solver-agnostic library for fast prototyping of SMT-based algorithms. In: SMT Workshop 2015, 13th International Workshop on Satisfiability Modulo Theories (2015)
31. Grinchtein, O., Carlsson, M., Pearson, J.: A constraint optimisation model for analysis of telecommunication protocol logs. In: Blanchette, J.C., Kosmatov, N. (eds.) TAP 2015. LNCS, vol. 9154, pp. 137–154. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-21215-9_9
32. Karpenkov, G.E.: Finding inductive invariants using satisfiability modulo theories and convex optimization. Theses, Université Grenoble Alpes (2017)
33. Kovásznai, G., Erdélyi, B., Biró, C.: Investigations of graph properties in terms of wireless sensor network optimization. In: 2018 IEEE International Conference on Future IoT Technologies (Future IoT), January 2018
34. Larraz, D., Oliveras, A., Rodríguez-Carbonell, E., Rubio, A.: Minimal-model-guided approaches to solving polynomial constraints and extensions. In: Sinz, C., Egly, U. (eds.) SAT 2014. LNCS, vol. 8561, pp. 333–350. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-09284-3_25
35. Leofante, F., Abraham, E., Niemueller, T., Lakemeyer, G., Tacchella, A.: Integrated synthesis and execution of optimal plans for multi-robot systems in logistics. *Inf. Syst. Front.* **21**(1), 87–107 (2018)
36. Li, Y., Albarghouthi, A., Kincad, Z., Gurfinkel, A., Chechik, M.: Symbolic optimization with SMT solvers. In: POPL (2014)
37. Liu, T., Tyszbrowicz, S.S., Beckert, B., Taghdiri, M.: Computing exact loop bounds for bounded program verification. In: Larsen, K.G., Sokolsky, O., Wang, J. (eds.) SETTA 2017. LNCS, vol. 10606, pp. 147–163. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-69483-2_9
38. MiniZinc. www.minizinc.org

39. Nadel, A., Ryvchin, V.: Bit-vector optimization. In: Chechik, M., Raskin, J.-F. (eds.) TACAS 2016. LNCS, vol. 9636, pp. 851–867. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49674-9_53
40. Nethercote, N., Stuckey, P.J., Becket, R., Brand, S., Duck, G.J., Tack, G.: MiniZinc: towards a standard CP modelling language. In: Bessière, C. (ed.) CP 2007. LNCS, vol. 4741, pp. 529–543. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74970-7_38
41. Nguyen, C.M., Sebastiani, R., Giorgini, P., Mylopoulos, J.: Multi-objective reasoning with constrained goal models. *Requirements Eng.* **23**(2), 189–225 (2016). <https://doi.org/10.1007/s00766-016-0263-5>
42. Nieuwenhuis, R., Oliveras, A.: On SAT modulo theories and optimization problems. In: Biere, A., Gomes, C.P. (eds.) SAT 2006. LNCS, vol. 4121, pp. 156–169. Springer, Heidelberg (2006). https://doi.org/10.1007/11814948_18
43. Nieuwenhuis, R., Oliveras, A., Rodríguez-Carbonell, E., Rubio, A.: Challenges in satisfiability modulo theories. In: Baader, F. (ed.) RTA 2007. LNCS, vol. 4533, pp. 2–18. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-73449-9_2
44. Nieuwenhuis, R., Oliveras, A., Tinelli, C.: Solving SAT and SAT modulo theories: from an abstract Davis-Putnam-Logemann-Loveland procedure to DPLL(T). *J. ACM* **53**(6), 937–977 (2006)
45. Ohrimenko, O., Stuckey, P.J., Codish, M.: Propagation via lazy clause generation. *Constraints* **14**(3), 357–391 (2009). <https://doi.org/10.1007/s10601-008-9064-x>
46. Oliver, R.S., Craciunas, S.S., Steiner, W.: IEEE 802.1Qbv gate control list synthesis using array theory encoding. In: 2018 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS), April 2018
47. Papadimitriou, C.H.: On the complexity of integer programming. *J. ACM* **28**(4), 765–768 (1981)
48. Ratschan, S.: Simulation based computation of certificates for safety of dynamical systems. In: Abate, A., Geeraerts, G. (eds.) FORMATS 2017. LNCS, vol. 10419, pp. 303–317. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-65765-3_17
49. Rendl, A., Guns, T., Stuckey, P.J., Tack, G.: MiniSearch: a solver-independent meta-search language for MiniZinc. In: Pesant, G. (ed.) CP 2015. LNCS, vol. 9255, pp. 376–392. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-23219-5_27
50. Roselli, S.F., Bengtsson, K., Åkesson, K.: SMT solvers for job-shop scheduling problems: models comparison and performance evaluation. In: 2018 IEEE 14th International Conference on Automation Science and Engineering (CASE), August 2018
51. Sawaya, N.W., Grossmann, I.E.: A cutting plane method for solving linear generalized disjunctive programming problems. *Comput. Chem. Eng.* **29**(9), 1891–1913 (2005)
52. Sebastiani, R.: Lazy satisfiability modulo theories. *J. Satisf. Boolean Model. Comput. JSAT* **3**(3–4), 141–224 (2007)
53. Sebastiani, R., Tomasi, S.: Optimization in SMT with $\mathcal{L}A(Q)$ cost functions. In: Gramlich, B., Miller, D., Sattler, U. (eds.) IJCAR 2012. LNCS (LNAI), vol. 7364, pp. 484–498. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-31365-3_38
54. Sebastiani, R., Tomasi, S.: Optimization modulo theories with linear rational costs. *ACM Trans. Comput. Logics* **16**(2), 1–43 (2015)
55. Sebastiani, R., Trentin, P.: OptiMathSAT: a tool for optimization modulo theories. In: Kroening, D., Păsăreanu, C.S. (eds.) CAV 2015. LNCS, vol. 9206, pp. 447–454. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-21690-4_27
56. Sebastiani, R., Trentin, P.: Pushing the envelope of optimization modulo theories with linear-arithmetic cost functions. In: Baier, C., Tinelli, C. (eds.) TACAS 2015. LNCS, vol. 9035, pp. 335–349. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46681-0_27

57. Sebastiani, R., Trentin, P.: On optimization modulo theories, MaxSMT and sorting networks. In: Legay, A., Margaria, T. (eds.) TACAS 2017. LNCS, vol. 10206, pp. 231–248. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54580-5_14
58. Sebastiani, R., Trentin, P.: OptiMathSAT: a tool for optimization modulo theories. *J. Autom. Reason.* **64**(3), 423–460 (2018)
59. Teso, S., Sebastiani, R., Passerini, A.: Structured learning modulo theories. *Artif. Intell.* **244**, 166–187 (2017)
60. Trentin, P., Sebastiani, R.: Optimization modulo the theory of floating-point numbers. In: Fontaine, P. (ed.) CADE 2019. LNCS (LNAI), vol. 11716, pp. 550–567. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-29436-6_33
61. Veksler, M., Strichman, O.: Learning general constraints in CSP. *Artif. Intell.* **238**, 135–153 (2016)
62. Zhou, N.-F., Kjellerstrand, H.: Optimizing SAT encodings for arithmetic constraints. In: Beck, J.C. (ed.) CP 2017. LNCS, vol. 10416, pp. 671–686. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66158-2_43