

OptiMathSAT: A Tool for Optimization Modulo Theories ^{*}

Roberto Sebastiani and Patrick Trentin

DISI, University of Trento, Italy

Abstract. Many SMT problems of interest may require the capability of finding models that are *optimal* wrt. some objective functions. These problems are grouped under the umbrella term of *Optimization Modulo Theories – OMT*. In this paper we present OPTIMATHSAT, an OMT tool extending the MATHSAT5 SMT solver. OPTIMATHSAT allows for solving a list of optimization problems on SMT formulas with linear objective functions –on the Boolean, the rational and the integer domains, and on their combination thereof– including partial weighted MaxSMT. Multiple objective functions can be combined together and handled either independently, or lexicographically, or in a min-max/max-min fashion.

OPTIMATHSAT ships with an extended SMT-LIBV2 input syntax and C API bindings, and it preserves the incremental attitude of its underlying SMT solver.

1 Introduction

SMT solvers are currently used as backend engines in many formal verification (FV) tools for Hardware, Software and Hybrid Systems. Many SMT problems of interest for FV or for other disciplines, however, require the capability of finding models that are *optimal* wrt. some objective functions [16, 8, 18, 11, 9, 15, 14, 19, 6, 13, 20]. These problems are grouped under the umbrella term of *Optimization Modulo Theories – OMT*.

For instance, in SMT-based model checking with timed or hybrid systems, you may want to find executions which optimize the value of some parameter while fulfilling/violating some property –e.g., to find the minimum opening time interval for a railcrossing causing a safety violation. (See e.g. [19] for some examples.) Also, a recent application of OMT is the SMT-based computation of the worst-case execution time (WCET) of loop-free programs [12], which finds tighter over-approximations of the WCET than other state-of-the-art approaches. A longer list of OMT applications in formal verification and in other disciplines can be found in [16, 8, 11, 15, 14, 19].

In this paper we present OPTIMATHSAT, an OMT tool extending the MATHSAT5 SMT solver [3, 10], implementing the OMT procedures described in [18–20]. OPTIMATHSAT allows for solving a list of optimization problems on SMT formulas with linear objective functions –on the Boolean, the rational and the integer domains, and on

^{*} This work is supported by Semiconductor Research Corporation (SRC) under GRC Research Project 2012-TJ-2266 WOLF. We thank Alberto Griggio for support with MATHSAT5 code.

their combination thereof— including partial weighted MAXSMT. Multiple objective functions can be combined together and handled either independently, or lexicographically, or in a min-max/max-min fashion. Like MATHSAT5, it is freely available for research and evaluation purposes [4], and it is currently used in some innovative projects (see §5).

Related Tools. Currently few other OMT tools exist. Closest to OPTIMATHSAT are SYMBA [14] and the very-recent νZ [6, 7], which are both built on top of Z3. [14] considered the problem of optimizing multiple rational cost functions at the same time. SYMBA uses the underlying SMT solver as black-box, and it features additional ad hoc techniques for detecting unbounded costs and optimization. Currently SYMBA has no incremental interface, it handles neither strict inequalities nor combined theories. νZ supports both single-objective linear optimization —over a real, integer or bit-vector term — and multi-objective optimization in either boxed, lexicographic or Pareto-optimization mode. It ships with several specialized engines for MaxSMT and with pre-processing techniques that re-encode the 0-1 integer variables of the input formula into Pseudo-Boolean or MaxSMT constraints. We refer the reader to the related work section of [19] for a more-detailed analysis of other OMT-related approaches and tools.

Content. This paper is structured as follows. §2 provides a brief outline of OPTIMATHSAT architecture, followed by a description of its optimization functionalities and interfaces in §3. §4 presents a short example, and §5 reviews some recent interesting applications of OPTIMATHSAT. §6 concludes the paper with hints of some future developments. An extended version of this paper, containing a performance evaluation and some more details, is available from OPTIMATHSAT web page [4].

2 Architecture

OPTIMATHSAT is written in C++ and it is built as an extension of MATHSAT 5, which implements the standard lazy SMT paradigm (see [5]). Unlike the OMT algorithms in [14, 6], which are based on an *offline* architecture —in which the SMT Solver is incrementally called multiple times as a black-box— OPTIMATHSAT is based on an *inline architecture* —in which the SMT solver is run only once and its internal SAT solver is modified to handle the search for the optima [18–20]. Although harder to implement, the *inline* architecture has showed better performance for OPTIMATHSAT than the *offline* one [18, 19]. (We refer the reader to [19] for a comparison of the two architectures.)

The optimization algorithm can explore the search space either in *linear-search* mode, by pruning one intermediate solution at a time, or in *binary-search* mode, by introducing cuts bisecting the search space, or in *adaptive-search* mode, which uses load-balancing adaptive heuristics to choose among the linear- and binary-search modes at each search step.

Some functionalities, such as the control loop for *lexicographic* optimization and the assertion of soft clauses, are handled at a higher level of abstraction by means of a combination of MATHSAT5 and OPTIMATHSAT API calls.

3 Optimization Functionalities

OPTIMATHSAT is mainly a tool for (single- and multiple-objective) OMT with linear objective functions $\text{OMT}(\mathcal{L}\mathcal{A} \cup \mathcal{T})$, s.t. “ $\mathcal{L}\mathcal{A}$ ” denotes linear arithmetic over either the rationals ($\mathcal{L}\mathcal{R}\mathcal{A}$), or the integers ($\mathcal{L}\mathcal{I}\mathcal{A}$) or their combination ($\mathcal{L}\mathcal{R}\mathcal{I}\mathcal{A}$), and \mathcal{T} denotes any other Nelson-Open theory supported by MATHSAT5. For each objective it is possible to specify both global and local bounds, if known.¹ OPTIMATHSAT can use this information to explore the search space in binary or in adaptive search mode, which might improve the overall performance of the solver. We support objective functions over the rational, integer and Boolean domains², or their combinations.

Here we provide a brief list of OPTIMATHSAT optimization functionalities, omitting the functionalities inherited from MATHSAT5 [10]. A detailed description of the implemented algorithms is presented in [18–20].

3.1 Single-Objective Optimization

We discuss first the case in which we have only one objective function, namely *obj*.

Linear Arithmetic Optimization over $\mathcal{L}\mathcal{R}\mathcal{A}$, $\mathcal{L}\mathcal{I}\mathcal{A}$ and $\mathcal{L}\mathcal{R}\mathcal{I}\mathcal{A}$. Given some term *obj* on $\mathcal{L}\mathcal{A}$, OPTIMATHSAT finds a solution (if any) which makes the term *obj* minimum/maximum. This is based on a combination of SMT and linear [integer] programming techniques.

Partial Weighted MAXSMT and SMT with Pseudo-Boolean Objectives (PB-SMT).

Given an input formula $\varphi_h \wedge \varphi_s$, where φ_h contains hard constraints and φ_s contains soft constraints with positive weights, the goal of *partial weighted* MAXSMT [16, 9] is to find a model M s.t. $M \models \varphi_h \wedge \varphi_s^M$ and φ_s^M is a subset of φ_s in which the soft-constraints have the largest cumulative weight possible. Similarly, OPTIMATHSAT allows also for defining Pseudo-Boolean objective functions in the form $\sum_i w_i \psi_i$, where w_i are numerical constants and ψ_i are sub-formulas.

Unlike with the procedures in [16, 8, 9, 6], which use specialized algorithms for MAXSMT/PB-SMT, OPTIMATHSAT works by encoding the problem into the optimization of an $\mathcal{L}\mathcal{R}\mathcal{A}$ term, as described in [19]. This allows for combining the MaxSMT terms with other objectives, as we describe in §3.2.

Notice that it is possible to interrupt the search of OPTIMATHSAT (e.g., by setting a timeout) and to still have access to the current sub-optimal solutions and its model.

3.2 Multi-Objective Combination

The interface of OPTIMATHSAT allows for combining multiple objective functions obj_1, \dots, obj_N in various ways.

¹ Local bounds have a special use in boxed multi-objective optimization (see §3.2). In single-objective and lexicographic optimization, they coincide with global bounds.

² i.e. MaxSMT and SMT with Pseudo-Boolean objective functions.

Multiple Independent Objectives [14]. (Aka Boxed Optimization [6].) OPTIMATHSAT can solve simultaneously N independent optimization problems $\langle \varphi, obj_1 \rangle, \dots, \langle \varphi, obj_N \rangle$, optionally building the corresponding optimum models M_0, \dots, M_N .³ (In the empirical evaluation presented in [20], we showed that using this optimization strategy can be considerably more efficient than solving N single-objective optimization problems.) This option is the default configuration.

Lexicographic Optimization. OPTIMATHSAT optimizes lexicographically the objectives obj_1, \dots, obj_N by decreasing level of priority, which corresponds to the input order of the objectives. If any objective obj_i is unsatisfiable or unbounded, the search returns.

Min-max & Max-min. The goal of a min-max problem is to find the maximum value of an obj s.t. $\bigwedge_{i=0}^N (obj \leq obj_i) \wedge \bigvee_{i=0}^N (obj_i = obj)$, obj being a fresh variable.⁴ Max-min is dual. OPTIMATHSAT provides syntactic-sugar extensions to SMT-LIBV2 that allow for encoding this type of objectives.

Linear Combination. Obviously, one can also create objectives that are a linear combination of other objectives obj_1, \dots, obj_N , i.e., $obj = \sum_{i=1}^N w_i \cdot obj_i$.

We remark that all the above combinations hold for obj_i cost functions over every domain, including Boolean. For instance, you can combine together MaxSMT with OMT optimization over Integer or Real objectives.

3.3 Interfaces

Input Language. OPTIMATHSAT functions are accessible through a list of commands, extending the SMT-LIBV2 syntax, which is shown in a concise description in Figure 1. Notice that, differently from νZ [7], in case of a MAXSMT problem we require the user to build *explicitly* a minimization objective using the ID associated with the asserted soft clauses, i.e., by writing “(minimize ID)”. The advantage of this requirement is that we allow for arbitrary composition of the MAXSMT objective with other linear arithmetic functions, which can be useful in particular contexts (for instance, to build obj functions on mixed Boolean/numeric domains, as with Linear Generalized Disjunctive Programming (LGDP) problems [17]).

C API. The optimization functions of OPTIMATHSAT are also available through its C API, which extends that of MATHSAT5 [3]. A detailed documentation of the C API, the SMT-LIBV2 language extensions and some usage examples are accessible on OPTIMATHSAT website [4].

³ Since the N input problems are independent to one another, the local bounds of each objective obj_i do not have any side effect on the feasible solutions of all other objectives obj_j , as if the N problems were solved separately.

⁴ Notice that in the actual encoding we drop the “ \vee ” part of the formula, since it is unnecessary and may cause extra Boolean search.

```

; assert soft clauses for MaxSMT (weight/dweight
;   defaults to 1, id defaults to 'I')
(assert-soft <term> [:weight <numeral>] [:id <string>])
(assert-soft <term> [:dweight <decimal>] [:id <string>])
; push new objectives on the stack
(minimize|maximize <term> [:local-lb <decimal>]
[:local-ub <decimal>])
(minmax|maxmin <term>...<term> [:local-lb <decimal>]
[:local-ub <decimal>])
; set multi-objective optimization mode
(set-option :opt.priority box|lex)
; optimize objective(s) in the stack
(check-sat)
; load model of objective with stack index <numeral>
(set-model <numeral>)
; retrieve model values
(get-model)
(get-value <VAR>)

```

Fig. 1. SMT-LIBv2 Optimization Extensions, square brackets corresponds to optional parameters, whereas “|” stands for alternative choices.

Incremental Interface. Like MATHSAT5, OPTIMATHSAT provides a PUSH/POP interface for adding and removing objectives and pieces of formulas from the formula stack, which allows for reusing information from one optimization search to another to improve the global performance of the search [20].

4 Example

In Figure 2 we present a toy example that illustrates how to encode a problem into the extended SMT-LIBv2 language of OPTIMATHSAT. A small company urgently needs 250 units of some goods. Suppliers s_1, s_2, s_3, s_4 offer to supply up to 250, 150, 100, 100 units of goods starting from the minimum quantity of 50, 100, 100, 50 units respectively. Their prices are 23\$, 21\$, 20\$, 10\$ per unit respectively. Our goal is (A) to minimize the overall purchase cost and, at cost tie, (B) to maximize the number of suppliers with whom we seal the deal, so that to minimize the risk of delivery delays and spread the payment in smaller chunks.

A simple OMT encoding of the problem is shown in Figure 2. In this example there are two combinations of suppliers $-s_2, s_4$ and s_1, s_3, s_4 from which we can purchase the goods at the minimum cost of 4150\$. Therefore, the tie is broken by means of our secondary goal (B), which imposes our preference on the second solution. The optimum model of a lexicographic optimization is always associated with the top-most objective on the internal stack. Since in this example there are only two objectives, this objective can be selected by passing 1 to the `set-model` command. Alternatively the user can also use negative indexing, and select the top-most objective by typing `(set-model -1)`. As mentioned in §3.3, notice that we explicitly ask for `ignored-suppliers`,

```

; set goods quantity
(assert (= 250 (+ q1 q2 q3 q4)))
; set goods offered by each supplier
(assert (and (or (= q1 0) (and (<= 50 q1) (<= q1 250)))
             (or (= q2 0) (and (<= 100 q2) (<= q2 150)))
             (or (= q3 0) (and (<= 100 q3) (<= q3 100)))
             (or (= q4 0) (and (<= 50 q4) (<= q4 100)))))
; a supplier is 'used' if sends more than zero items
(assert (and (=> s1 (not (= q1 0))) (=> s2 (not (= q2 0)))
             (=> s3 (not (= q3 0))) (=> s4 (not (= q4 0)))))
; supply from the largest number of suppliers
(assert-soft s1 :id ignored_suppliers)
(assert-soft s2 :id ignored_suppliers)
(assert-soft s3 :id ignored_suppliers)
(assert-soft s4 :id ignored_suppliers)
; set goal (A)
(minimize (+ (* q1 23) (* q2 21) (* q3 20) (* q4 10)))
; set goal (B)
(minimize ignored_suppliers)
; optimize lexicographically
(set-option :opt.priority lex)
(check-sat)
; print model
(set-model 1)
(get-model)

```

Fig. 2. SMT-LIBv2 encoding of the problem.

the label of the MAXSMT constraints, to be minimized. OPTIMATHSAT solved the problem in 10ms.

5 Applications

We briefly mention two examples of recent applications –which are very innovative in their respective domains– that have been technologically enabled by OMT and use OPTIMATHSAT as backend automated-reasoning engine.

Structured Learning Modulo Theories. In Machine Learning applications, performing inference and learning in hybrid domains –characterized by both continuous and Boolean/discrete variables– is a particularly daunting task. *Structured Learning Modulo Theories (SLMT)* [21] addresses the problem by combining (Structured-Output) Support Vector Machines (SVNs) with OMT, so that the latter plays the role of inference and separation oracle for the former. The tool LMT implementing the SLMT method [2] uses OPTIMATHSAT as backend OMT engine.

Automated Reasoning on Constrained Goal Models. Goal Models (GM) are widely used in Requirements Engineering to represent software requirements, objectives,

and design qualities [22]. *Constrained Goal Models (CGM)* are a novel, formal version of GM which are enriched with constraints so that to handle preferences, numerical attributes and resources (e.g., scores, financial cost, workforce, etc.). OPTIMATHSAT is used as a backend reasoning engine of CGM-TOOL [1], a tool for building and reasoning on *CGMs*, allowing for automatically verifying the realizability of a CGM and for finding optimal realizations according to some specified criterion.

6 Future Developments

We plan to extend OPTIMATHSAT capabilities along several directions. For instance, we are interested into generalizing our implementation to support objective functions extended on other theories, i.e. bit-vector. We are also considering to add the possibility of combining multiple objectives for Pareto-optimization. Finally, we plan to parallelize OMT so that to exploit the multi-core architectures of modern CPUs.

References

1. CGM-Tool. www.cgm-tool.eu.
2. LMT. <http://disi.unitn.it/~teso/lmt/lmt.tgz>.
3. MathSAT 5. <http://mathsat.fbk.eu/>.
4. OptiMathSAT. <http://optimathsat.disi.unitn.it/>.
5. C. W. Barrett, R. Sebastiani, S. A. Seshia, and C. Tinelli. Satisfiability Modulo Theories. In *Handbook of Satisfiability*, chapter 26, pages 825–885. IOS Press, 2009.
6. N. Bjorner and A.-D. Phan. νZ - Maximal Satisfaction with Z3. In *Proc SCSS. Invited presentation.*, Gammart, Tunisia, December 2014. EasyChair Proceedings in Computing (EPiC). <http://www.easychair.org/publications/?page=862275542>.
7. N. Bjorner, A.-D. Phan, and L. Fleckenstein. νZ - An Optimizing SMT Solver. In *Proc. TACAS*, volume 9035 of *LNCS*. Springer, 2015. To Appear.
8. A. Cimatti, A. Franzén, A. Griggio, R. Sebastiani, and C. Stenico. Satisfiability modulo the theory of costs: Foundations and applications. In *TACAS*, volume 6015 of *LNCS*. Springer, 2010.
9. A. Cimatti, A. Griggio, B. J. Schaafsma, and R. Sebastiani. A Modular Approach to MaxSAT Modulo Theories. In *SAT*, volume 7962 of *LNCS*, July 2013.
10. A. Cimatti, A. Griggio, B. J. Schaafsma, and R. Sebastiani. The MathSAT 5 SMT Solver. In *TACAS*, volume 7795 of *LNCS*. Springer, 2013.
11. I. Dillig, T. Dillig, K. L. McMillan, and A. Aiken. Minimum Satisfying Assignments for SMT. In *CAV*, pages 394–409, 2012.
12. J. Henry, M. Asavoa, D. Monniaux, and C. Maïza. How to compute worst-case execution time by optimization modulo theory and a clever encoding of program semantics. *SIGPLAN Not.*, 49(5):43–52, June 2014.
13. D. Larraz, A. Oliveras, E. Rodríguez-Carbonell, and A. Rubio. Minimal-Model-Guided Approaches to Solving Polynomial Constraints and Extensions. In *SAT*, 2014.
14. Y. Li, A. Albarghouthi, Z. Kincad, A. Gurfinkel, and M. Chechik. Symbolic Optimization with SMT Solvers. In *POPL*. ACM Press., 2014.
15. P. Manolios and V. Papavasileiou. Ilp modulo theories. In *CAV*, pages 662–677, 2013.

16. R. Nieuwenhuis and A. Oliveras. On SAT Modulo Theories and Optimization Problems. In *SAT*, volume 4121 of *LNCS*. Springer, 2006.
17. R. Raman and I. Grossmann. Modelling and computational techniques for logic based integer programming. *Computers and Chemical Engineering*, 18(7):563 – 578, 1994.
18. R. Sebastiani and S. Tomasi. Optimization in SMT with LA(Q) Cost Functions. In *IJCAR*, volume 7364 of *LNAI*, pages 484–498. Springer, July 2012.
19. R. Sebastiani and S. Tomasi. Optimization Modulo Theories with Linear Rational Costs. *ACM Transactions on Computational Logics*, 16(2), 2015.
20. R. Sebastiani and P. Trentin. Pushing the Envelope of Optimization Modulo Theories with Linear-Arithmetic Cost Functions. In *Proc. TACAS*, volume 9035 of *LNCS*. Springer, 2015.
21. S. Teso, R. Sebastiani, and A. Passerini. Structured Learning Modulo Theories. *Artificial Intelligence*, 2015. To appear. <http://disi.unitn.it/~rseba/publist.html>.
22. A. Van Lamsweerde. Goal-Oriented Requirements Engineering: A Guided Tour. In *Proceedings of the Fifth IEEE International Conference on Requirements Engineering*, RE'01, pages 249–. IEEE Computer Society, 2001.

A Appendix: An Empirical Evaluation

Note for the reviewers. *In compliance with CAV'15 call for papers, in this appendix we report some empirical results. Some of them are novel, some other are taken from [20]. Although we are aware that this is not guaranteed to be taken into account, we believe it may provide support to the tool description. If the paper is accepted, this appendix will be removed and added to some publicly-available extended version of the paper.*

In order to appreciate the performances of OPTIMATHSAT, we provide here an empirical evaluation wrt. the main OMT tools available, SYMBA and νZ , on both single-objective and boxed multi-objective problems.

All tests were executed with a timeout of 1200 seconds on two identical machines, both equipped with *8-core 2.20Ghz Xeon* processor and 64GB of RAM, on Ubuntu Linux 3.8-0-29 kernel. The results, the input OMT problems, the binaries of the tools used and full-size plots are publicly available to allow for reproducing our experiments.⁵

First, we consider the multi-independent-objective problems from [14], coming from Software Verification problems, and we report in Figure 3 some results from [20]. (We refer the reader to [14, 20] for a more detailed description of these benchmarks and experiments.) Since each multi-independent-objective problem encodes a set of independent single-objective problems, in Figure 3, we report a comparison among three different configurations of OPTIMATHSAT and νZ (SINGLE-OBJECTIVE, INCREMENTAL, MULTI-OBJECTIVE) and the two versions of SYMBA in [14] (both MULTI-OBJECTIVE).

We start by considering the first row of plots in Figure 3 and the last three rows in the table, where we compare the three above-mentioned configurations of OPTIMATHSAT. The plot on the left shows the speedup of using OPTIMATHSAT on single objectives incrementally, in comparison with restarting the tool on each new objective. The middle plot shows an even more substantial performance improvement by running OPTIMATHSAT in boxed MULTI-OBJECTIVE mode. Notice that, as shown in the right plot, the latter mode is to be preferred to the INCREMENTAL mode as it provides uniformly better performance.

In the table and in the second row of plots in Figure 3, we compare the multi-objective version of OPTIMATHSAT against SYMBA and νZ . We observe that multi-objective OPTIMATHSAT performs much better than the default configuration of SYMBA, and significantly better than both SYMBA(40)+OPT-Z3 and νZ -MULTI-OBJECTIVE.

Second, we tested the performances of OPTIMATHSAT vs. those of νZ over sets of single-objective problems from [18, 19], coming from LGDP (strip-packing, job-shop) and formal verification of parametric systems (SAL). The results are shown in Figure 4. They show that OPTIMATHSAT performs a little worse than νZ on strip-packing, significantly better on job-shop, and comparably on SAL (3 problems solved more, but little worse times on problems in which both terminate).

⁵ The TACAS'15 tests are available at <http://disi.unitn.it/%7Etrentin/resources/tacas15.tar.gz>; The novel tests are available at <http://disi.unitn.it/%7Etrentin/resources/cav15.tar.gz>; SYMBA is available at <https://bitbucket.org/arieg/symba/src>.

Tool:	#inst.	#solved	#timeout	time
SYMBA(100)	1103	1091	12	10917
SYMBA(40)+OPT-Z3	1103	1103	0	1128
νZ -multiobjective	1103	1090	13	1761
νZ -incremental	1103	1100	3	8683
νZ -singleobjective	1103	1101	2	10002
optimathsat-multiobjective	1103	1103	0	901
optimathsat-incremental	1103	1103	0	3477
optimathsat-singleobjective	1103	1103	0	16161

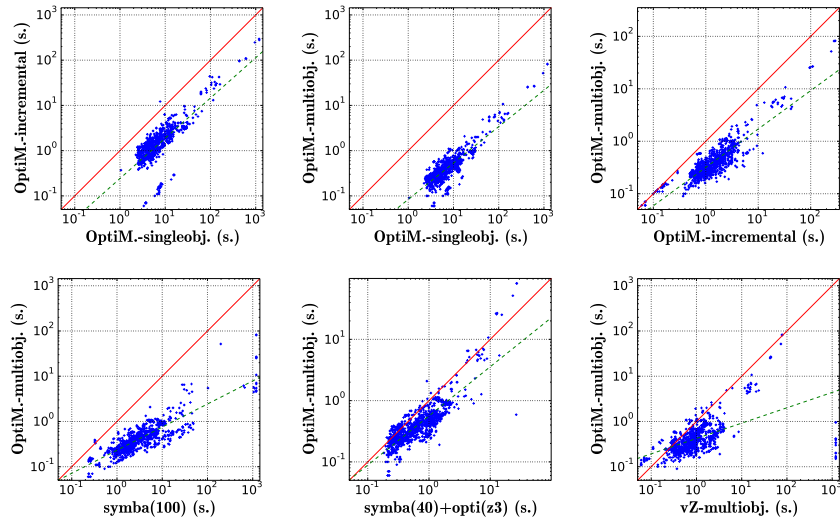


Fig. 3. Comparison of different versions of OPTIMATHSAT, SYMBA and νZ on the SW verification problems in [14]. Plots, first row: pairwise comparisons between different versions of OPTIMATHSAT; second row: pairwise comparisons between OPTIMATHSAT-MULTIOBJECTIVE, the two versions of SYMBA and νZ -MULTIOBJECTIVE.

Tool:	#inst.	#solved	#timeout	time
strip-packing				
OPTIMATHSAT	600	561	41	36209
νZ	600	563	37	47427
job-shop				
OPTIMATHSAT	600	586	14	63725
νZ	600	530	70	135230
SAL				
OPTIMATHSAT	544	544	0	88033
νZ	544	541	3	64837




Fig. 4. Comparison of OPTIMATHSAT and νZ on the single-objective problems of [19]: strip-packing, job-shop and SAL.