Journal of Artificial Intelligence Research 1 (2008) 1-45 Submitted ?/??; published ?/??

Automated Reasoning in Modal and Description Logics via SAT Encoding: the Case Study of K(m)/$\mathcal{ALC}$-Satisfiability

Roberto Sebastiani ROBERTO.SEBASTIANI@DISI.UNITN.IT
Michele Vescovi MICHELE.VESCOVI@DISI.UNITN.IT

DISI, Università di Trento,
Via Sommarive 14, I-38050, Povo, Trento, Italy

Abstract

In the last two decades, modal and description logics have been applied to numerous areas of computer science, including knowledge representation, formal verification, database theory, distributed computing and, more recently, semantic web and ontologies. For this reason, the problem of automated reasoning in modal and description logics has been thoroughly investigated. In particular, many approaches have been proposed for efficiently handling the satisfiability of the core normal modal logic K_m, and of its notational variant, the description logic $\mathcal{ALC}$. Although simple in structure, $K_m/\mathcal{ALC}$ is computationally very hard to reason on, its satisfiability being PSPACE-complete.

In this paper we start exploring the idea of performing automated reasoning tasks in modal and description logics by encoding them into SAT, so that to be handled by state-of-the-art SAT tools; as with most previous approaches, we begin our investigation from the satisfiability in K_m. We propose an efficient encoding, and we test it on an extensive set of benchmarks, comparing the approach with the main state-of-the-art tools available. Although the encoding is necessarily worst-case exponential, from our experiments we notice that, in practice, this approach can handle most or all the problems which are at the reach of the other approaches, with performances which are comparable with, or even better than, those of the current state-of-the-art tools.

1. Motivations and Goals

In the last two decades, modal and description logics have provided an essential framework for many applications in numerous areas of computer science, including artificial intelligence, formal verification, database theory, distributed computing and, more recently, semantic web and ontologies. For this reason, the problem of automated reasoning in modal and description logics has been thoroughly investigated (e.g., Fitting, 1983; Ladner, 1977; Baader & Hollunder, 1991; Halpern & Moses, 1992; Baader, Franconi, Hollunder, Nebel, & Profitlich, 1994; Massacci, 2000). In particular, the research in modal and description logics has followed two parallel routes until the seminal work of Schild (1991), which proved that the core modal logic K_m and the core description logic $\mathcal{ALC}$ are one a notational variant of the other. Since then, analogous results have been produced for a bunch of other logics, so that, nowadays the two research lines have mostly merged into one research flow.

Many approaches have been proposed for efficiently reasoning in modal and description logics, starting from the problem of checking the satisfiability in the core normal modal logic K_m and in its notational variant, the description logic $\mathcal{ALC}$ (hereafter simply "K_m"). We classify them as follows.

©2008 AI Access Foundation. All rights reserved.

- The "classic" *tableau-based* approach (Fitting, 1983; Baader & Hollunder, 1991; Massacci, 2000) is based on the construction of propositional tableau branches, which are recursively expanded on demand by generating successor nodes in a candidate Kripke model. Kris (Baader & Hollunder, 1991; Baader et al., 1994), Crack (Franconi, 1998), LWB (Balsiger, Heuerding, & Schwendimann, 1998) were among the main representative tools of this approach.

- The *DPLL-based* approach (Giunchiglia & Sebastiani, 1996, 2000) differs from the previous one mostly in the fact that a DPLL procedure, which treats the modal subformulas as propositions, is used instead of the classic propositional tableaux procedure at each nesting level of the modal operators. KSAT (Giunchiglia & Sebastiani, 1996), ESAT (Giunchiglia, Giunchiglia, & Tacchella, 2001) *SAT (Tacchella, 1999), are the Representative tools of this approach.

These two approaches merged into the "modern" tableaux-based approach, which has been extended to work with more expressive description logics and to provide more sophisticate reasoning functions. Among the tools employing this approach, we recall FaCT/FaCT++ and DLP (Horrocks & Patel-Schneider, 1999), and Racer (V. Haarslev, 2001).[1]

- In the *translational* approach (Hustadt & Schmidt, 1999; Areces, Gennari, Heguiabehere, & de Rijke, 2000) the modal formula is encoded into first-order logic (FOL), and the encoded formula can be decided efficiently by a FOL theorem prover (Areces et al., 2000). Mspass (Hustadt, Schmidt, & Weidenbach, 1999) is the most representative tool of this approach.

- The *CSP-based approach* (Brand, Gennari, & de Rijke, 2003) differs from the tableaux-based and DPLL-based ones mostly in the fact that a CSP engine is used instead of tableaux/DPLL. KCSP is the only representative tool of this approach.

- In the *Inverse-method* approach (Voronkov, 1999, 2001), a search procedure is based on the "inverted" version of a sequent calculus (which can be seen as a modalized version of propositional resolution). KK (Voronkov, 1999) is the only representative tool of this approach.

- In the *Automata-theoretic* approach, (a BDD-based symbolic representation of) a tree automaton accepting all the tree models of the input formula is implicitly built and checked for emptiness (Pan, Sattler, & Vardi, 2002; Pan & Vardi, 2003). KBDD (Pan & Vardi, 2003) is the only representative tool of this approach.

- Pan and Vardi (2003) presented also an encoding of K-satisfiability into QBF-satisfiability (which is PSpace-complete too), combined with the use of a state-of-the-art QBF solver. We call this approach *QBF-encoding* approach.

1. Notice that there is not an universal agreement on the terminology "tableaux-based" and "DPLL-based". E.g., tools like FaCT, DLP, and Racer are most often called "tableau-based", although they use a DPLL-like algorithm instead of propositional tableaux for handling the propositional component of reasoning (Horrocks, 1998; Patel-Schneider, 1998; Horrocks & Patel-Schneider, 1999; V. Haarslev, 2001).

To the best of our knowledge, the last four approaches so far are restricted to the satisfiability in K_m only.

A significant amount of benchmarks formulas have been produced for testing the effectiveness of the different techniques (Halpern & Moses, 1992; Giunchiglia, Roveri, & Sebastiani, 1996; Heuerding & Schwendimann, 1996; Horrocks, Patel-Schneider, & Sebastiani, 2000; Massacci, 1999; Patel-Schneider & Sebastiani, 2001, 2003).

In the last two decades we have also witnessed an impressive advance in the efficiency of propositional satisfiability techniques (SAT), which has brought large and previously-intractable problems at the reach of state-of-the-art SAT solvers. Most of the success of SAT technologies is motivated by the impressive efficiency reached by current implementations of the Davis-Putnam-Logemann-Loveland procedure, DPLL (Davis & Putnam, 1960; Davis, Longemann, & Loveland, 1962), in its most-modern variants (Silva & Sakallah, 1996; Moskewicz, Madigan, Zhao, Zhang, & Malik, 2001; Eén & Sörensson, 2004). Current implementations can handle formulas in the order of 10^7 variables and clauses.

As a consequence, many hard real-world problems have been successfully solved by encoding into SAT (including, e.g., circuit verification and synthesis, scheduling, planning, model checking, ATPG, cryptanalysis, gene mapping). Effective encodings into SAT have been proposed also for the satisfiability problems in quantifier-free FOL theories which are of interest for formal verification (Strichman, Seshia, & Bryant, 2002; Seshia, Lahiri, & Bryant, 2003; Strichman, 2002). Notably, successful SAT encodings include also PSpace-complete problems, like planning (Kautz, McAllester, & Selman, 1996) and model checking (Biere, Cimatti, Clarke, & Zhu, 1999).

In this paper we start exploring the idea of performing automated reasoning tasks in modal and description logics by encoding them into SAT, so that to be handled by state-of-the-art SAT tools; as with most previous approaches, we begin our investigation from the satisfiability in K_m.

In theory, the task may look hopeless because of worst-case complexity issues: in fact, with few exceptions, the satisfiability problem in most modal and description logics is not in NP, typically being PSpace-complete or even harder —PSpace-complete for K_m (Ladner, 1977; Halpern & Moses, 1992)— so that the encoding is in worst-case exponential or larger.[2]

In practice, however, a few considerations allow for conjecturing that this approach may be competitive with the state-of-the-art approaches. First, the non-polynomial bounds above are *worst-case* bounds, and formulas may have different behaviors from that of the pathological formulas which can be found in textbooks. (E.g., notice that the exponentiality is based on the hypothesis of unboundedness of some parameter like the modal depth; Halpern & Moses, 1992; Halpern, 1995.) Second, some tricks in the encoding may allow for reducing the size of the encoded formula significantly. Third, as the amount of RAM memory in current computers is in the order of the GBytes and current SAT solvers can successfully handle huge formulas, the encoding of many modal formulas (at least of those which are not too hard to solve also for the competitors) may be at the reach of a SAT solver. Finally, even for PSpace-complete logics like K_m, also other state-of-the-art approaches are not guaranteed to use polynomial memory.

2. We implicitly make the assumption NP $\neq$ PSpace.

In this paper we show that, at least for the satisfiability K_m, by exploiting some smart optimizations in the encoding the SAT-encoding approach becomes competitive in practice with previous approaches. To this extent, the contributions of this paper are manyfold.

- We propose a basic encoding of K_m formulas into purely-propositional ones, and prove that the encoding is satisfiability-preserving.

- We describe some optimizations of the encoding, both in form of preprocessing and of on-the-fly simplification. These techniques allow for significant (and in some cases dramatic) reductions in the size of the resulting Boolean formulas, and in performances of the SAT solver thereafter.

- We perform a very extensive empirical comparison against the main state-of-the-art tools available. We show that, despite the NP-vs.-PSpace issue, this approach can handle most or all the problems which are at the reach of the other approaches, with performances which are comparable with, and sometimes even better than, those of the current state-of-the-art tools. In our perspective, this is the most surprising contribution of the paper.

- As a byproduct of our work, we obtain an empirical evaluation of current tools for K_m-satisfiability available, which is very extensive in terms of both amount and variety of benchmarks and of number and representativeness of the tools evaluated. We are not aware of any other such evaluation in the recent literature.

We also stress the fact that with our approach the encoder can be interfaced with every SAT solver in a plug-and-play manner, so that to benefit for free of every improvement in the technology of SAT solvers which has been or will be made available.

Content. The paper is structured as follows. In Section 2 we provide the necessary background notions on modal logics and SAT. In Section 3 we describe the basic encoding from K_m to SAT. In Section 4 we describe and discuss the main optimizations, and provide many examples. In Section 5 we present the empirical evaluation, and discuss the results. In Section 6 we present some related work and current research trends. In Section 7 we conclude, and describe some possible future evolutions.

Note to reviewers. A six-page preliminary version of this paper, containing some of the basic ideas presented here, was presented at SAT'06 conference (Sebastiani & Vescovi, 2006). For both the reviewers' and the readers' convenience, an online appendix is provided, containing all plots of Section 5 in full size. Moreover, in order to make the results reproducible, the encoder, the benchmarks and the random generators with the seeds used are also available in the online appendix.

2. Background

In this section we provide the necessary background on the modal logic K_m (Section 2.1) and on SAT and the DPLL procedure (Section 2.2).

2.1 The Modal Logic K_m

We recall some basic definitions and properties of K_m. Given a non-empty set of primitive propositions $\mathcal{A} = \{A_1, A_2, \ldots\}$ and a set of m modal operators $\mathcal{B} = \{\Box_1, \ldots, \Box_m\}$, the language of K_m is the least set of formulas containing $\mathcal{A}$, closed under the set of propositional connectives $\{\neg, \wedge\}$ and the set of modal operators in $\mathcal{B}$. Notationally, we use the Greek letters $\alpha, \beta, \varphi, \psi, \nu, \pi$ to denote formulas in the language of K_m (K_m-formulas hereafter). We use the standard abbreviations, that is: "$\Diamond_r\varphi$" for "$\neg\Box_r\neg\varphi$", "$\varphi_1 \vee \varphi_2$" for "$\neg(\neg\varphi_1 \wedge \neg\varphi_2)$", "$\varphi_1 \rightarrow \varphi_2$" for "$\neg(\varphi_1 \wedge \neg\varphi_2)$", "$\varphi_1 \leftrightarrow \varphi_2$" for "$\neg(\varphi_1 \wedge \neg\varphi_2) \wedge \neg(\varphi_2 \wedge \neg\varphi_1)$", "$\top$" and "$\bot$" for the constants "*True*" and "*False*". (Hereafter formulas like $\neg\neg\psi$ are implicitly assumed to be simplified into ψ, so that, if ψ is $\neg\phi$, then by "$\neg\psi$" we mean "ϕ".) Notationally, we often write "$(\bigwedge_i l_i) \rightarrow \bigvee_j l_j$" for the clause "$\bigvee_j \neg l_i \vee \bigvee_j l_j$", and "$(\bigwedge_i l_i) \rightarrow (\bigwedge_j l_j)$" for the conjunction of clauses "$\bigwedge_j(\bigvee_i \neg l_i \vee l_j)$". We denote by $\Box_r^i$ the nested application of the $\Box_r$ operator i times: $\Box_r^0\psi := \psi$ and $\Box_r^{i+1}\psi := \Box_r\Box_r^i\psi$. We call *depth* of φ, written *depth(φ)*, the maximum number of nested modal operators in φ. We call a *propositional atom* every primitive proposition in $\mathcal{A}$, and a *propositional literal* every propositional atom (*positive literal*) or its negation (*negative literal*).

In order to make our presentation more uniform, and to avoid considering the polarity of subformulas, we adopt the representation of K_m-formulas (e.g., Fitting, 1983; Massacci, 2000) from the following table:

α	α_1	α_2	β	β_1	β_2	π^r	π_0^r	ν^r	ν_0^r
$(\varphi_1 \wedge \varphi_2)$	φ_1	φ_2	$(\varphi_1 \vee \varphi_2)$	φ_1	φ_2	$\Diamond_r\varphi_1$	φ_1	$\Box_r\varphi_1$	φ_1
$\neg(\varphi_1 \vee \varphi_2)$	$\neg\varphi_1$	$\neg\varphi_2$	$\neg(\varphi_1 \wedge \varphi_2)$	$\neg\varphi_1$	$\neg\varphi_2$	$\neg\Box_r\varphi_1$	$\neg\varphi_1$	$\neg\Diamond_r\varphi_1$	$\neg\varphi_1$
$\neg(\varphi_1 \rightarrow \varphi_2)$	φ_1	$\neg\varphi_2$	$(\varphi_1 \rightarrow \varphi_2)$	$\neg\varphi_1$	φ_2				

in which non-literal K_m-formulas are grouped into four categories: α's (conjunctive), β's (disjunctive), π's (existential), ν's (universal). Importantly, all such formulas occur in the main formula with positive polarity only. This allows for disregarding the issue of polarity of subformulas.

The semantic of modal logics is given by means of Kripke structures. A *Kripke structure* for K_m is a tuple $\mathcal{M} = \langle \mathcal{U}, \mathcal{L}, \mathcal{R}_1, \ldots, \mathcal{R}_m \rangle$, where $\mathcal{U}$ is a set of states, $\mathcal{L}$ is a function $\mathcal{L} : \mathcal{A} \times \mathcal{U} \longmapsto \{True, False\}$, and each $\mathcal{R}_r$ is a binary relation on the states of $\mathcal{U}$. With an abuse of notation we write "$u \in \mathcal{M}$" instead of "$u \in \mathcal{U}$". We call a *situation* any pair $\mathcal{M}, u$, $\mathcal{M}$ being a Kripke structure and $u \in \mathcal{M}$. The binary relation $\models$ between a modal formula φ and a situation $\mathcal{M}, u$ is defined as follows:

$\mathcal{M}, u \models \top$;
$\mathcal{M}, u \not\models \bot$;
$\mathcal{M}, u \models A_i,\ A_i \in \mathcal{A} \iff \mathcal{L}(A_i, u) = True$;
$\mathcal{M}, u \models \neg A_i,\ A_i \in \mathcal{A} \iff \mathcal{L}(A_i, u) = False$;
$\mathcal{M}, u \models \alpha \iff \mathcal{M}, u \models \alpha_1$ *and* $\mathcal{M}, u \models \alpha_2$;
$\mathcal{M}, u \models \beta \iff \mathcal{M}, u \models \beta_1$ *or* $\mathcal{M}, u \models \beta_2$;
$\mathcal{M}, u \models \pi^r \iff \mathcal{M}, w \models \pi_0^r$ for some $w \in \mathcal{U}$ s.t. $\mathcal{R}_r(u, w)$ holds in $\mathcal{M}$;
$\mathcal{M}, u \models \nu^r \iff \mathcal{M}, w \models \nu_0^r$ for every $w \in \mathcal{U}$ s.t. $\mathcal{R}_r(u, w)$ holds in $\mathcal{M}$.

"$\mathcal{M}, u \models \varphi$" should be read as "$\mathcal{M}, u$ *satisfy* φ in K_m" (alternatively, "$\mathcal{M}, u$ K_m-satisfies φ"). We say that a K_m-formula φ is satisfiable in K_m (K_m-satisfiable henceforth) if and only

```
1.      SatValue DPLL (formula φ, assignment μ) {
2.          while (1) {
3.              decide_next_branch(φ, μ);
4.              while (1) {
5.                  status = deduce(φ, μ);
6.                  if (status == sat)
7.                      return sat;
8.                  else if (status == conflict) {
9.                      blevel = analyze_conflict(φ, μ);
10.                     if (blevel == 0) return unsat;
11.                     else backtrack(blevel,φ, μ);
12.                 }
13.                 else break;
14.     }}}
```

Figure 1: Schema of a modern DPLL engine.

if there exist $\mathcal{M}$ and $u \in \mathcal{M}$ s.t. $\mathcal{M}, u \models \varphi$. (When this causes no ambiguity, we sometimes drop the prefix "K_m-".) We say that w is a *successor* of u through $\mathcal{R}_r$ iff $\mathcal{R}_r(u, w)$ holds in $\mathcal{M}$.

The problem of determining the K_m-satisfiability of a K_m-formula φ is decidable and PSPACE-complete (Ladner, 1977; Halpern & Moses, 1992), even restricting the language to a single Boolean atom (i.e., $\mathcal{A} = \{A_1\}$; Halpern, 1995); if we impose a bound on the modal depth of the K_m-formulas, the problem reduces to NP-complete (Halpern, 1995). For a more detailed description on K_m— including, e.g., axiomatic characterization, decidability and complexity results — we refer the reader to the Halpern and Moses's (1992), Halpern's (1995) works.

A K_m-formula is said to be in *Negative Normal Form (NNF)* if it is written in terms of the symbols $\Box_r$, $\Diamond_r$, $\wedge$, $\vee$ and propositional literals A_i, $\neg A_i$ (i.e., if all negations occur only before propositional atoms in $\mathcal{A}$). Every K_m-formula φ can be converted into an equivalent one $NNF(\varphi)$ by recursively applying the rewriting rules: $\neg\Box_r\varphi \Longrightarrow \Diamond_r\neg\varphi$, $\neg\Diamond_r\varphi \Longrightarrow \Box_r\neg\varphi$, $\neg(\varphi_1 \wedge \varphi_2) \Longrightarrow (\neg\varphi_1 \vee \neg\varphi_2)$, $\neg(\varphi_1 \vee \varphi_2) \Longrightarrow (\neg\varphi_1 \wedge \neg\varphi_2)$, $\neg\neg\varphi \Longrightarrow \varphi$.

A K_m-formula is said to be in *Box Normal Form (BNF)* (Pan et al., 2002; Pan & Vardi, 2003) if it is written in terms of the symbols $\Box_r$, $\neg\Box_r$, $\wedge$, $\vee$, and propositional literals A_i, $\neg A_i$ (i.e., if no diamonds are there, and all negations occurs only before boxes or before propositional atoms in $\mathcal{A}$). Every K_m-formula φ can be converted into an equivalent one $BNF(\varphi)$ by recursively applying the rewriting rules: $\Diamond_r\varphi \Longrightarrow \neg\Box_r\neg\varphi$, $\neg(\varphi_1 \wedge \varphi_2) \Longrightarrow (\neg\varphi_1 \vee \neg\varphi_2)$, $\neg(\varphi_1 \vee \varphi_2) \Longrightarrow (\neg\varphi_1 \wedge \neg\varphi_2)$, $\neg\neg\varphi \Longrightarrow \varphi$.

2.2 Propositional Satisfiability with the DPLL Algorithm

Most state-of-the-art SAT procedures are evolutions of the Davis-Putnam-Longeman-Loveland (DPLL) procedure (Davis & Putnam, 1960; Davis et al., 1962). A high-level schema of a modern DPLL engine, adapted from the one presented by Zhang and Malik (2002), is

reported in Figure 1. The Boolean formula φ is in CNF; the assignment μ is initially empty, and it is updated in a stack-based manner.

In the main loop, `decide_next_branch`(φ, μ) chooses an unassigned literal l from φ according to some heuristic criterion, and adds it to μ. (This operation is called *decision*, l is called *decision literal* end the number of decision literals in μ after this operation is called the *decision level* of l.) In the inner loop, `deduce`(φ, μ) iteratively deduces literals l deriving from the current assignment and updates φ and μ accordingly; this step is repeated until either μ satisfies φ, or μ falsifies φ, or no more literals can be deduced, returning sat, conflict and unknown respectively. (The iterative application of Boolean deduction steps in deduce is also called *Boolean Constraint Propagation, BCP.*) In the first case, DPLL returns sat. If the second case, `analyze_conflict`(φ, μ) detects the subset η of μ which caused the conflict (*conflict set*) and the decision level `blevel` to backtrack. If `blevel` $== 0$, then a conflict exists even without branching, so that DPLL returns unsat. Otherwise, `backtrack`$(\texttt{blevel}, \varphi, \mu)$ adds the clause $\neg\eta$ to φ (*learning*) and backtracks up to blevel (*backjumping*), updating φ and μ accordingly. In the third case, DPLL exits the inner loop, looking for the next decision.

Notably, modern DPLL implementations implement techniques, like the two-watched-literal scheme, which allow for extremely efficient handling of BCP (Zhang & Malik, 2002). Old versions of DPLL used to implement also the *Pure-Literal Rule (PLR)* (Davis et al., 1962): when one proposition occurs only positively (resp. negatively) in the formula, it can be safely assigned to *true* (resp. *false*). Modern DPLL implementations, however, no more implement it due to its computational cost. For a much deeper description of modern DPLL-based SAT solvers, we refer the reader to the literature (e.g., Zhang & Malik, 2002).

3. The Basic Encoding

We borrow some notation from the *Single Step Tableau (SST)* framework (Massacci, 2000; Donini & Massacci, 2000). We represent univocally states in $\mathcal{M}$ as labels σ, represented as non empty sequences of integers $1.n_1^{r_1}.n_2^{r_2}. \ldots .n_k^{r_k}$, s.t. the label 1 represents the root state, and $\sigma.n^r$ represents the n-th successor of σ through the relation $\mathcal{R}_r$. With a little abuse of notation, hereafter we may say "a state σ" meaning "a state labeled by σ". We call a *labeled formula* a pair $\langle \sigma, \psi \rangle$, s.t. σ is a state label and ψ is a K_m-formula.

Let $A_{[\ ,\ ]}$ be an *injective* function which maps a labeled formula $\langle \sigma, \psi \rangle$, s.t. ψ is not in the form $\neg\phi$, into a Boolean variable $A_{[\sigma,\ \psi]}$. We conventionally assume that $A_{[\sigma,\ \top]}$ is $\top$ and $A_{[\sigma,\ \bot]}$ is $\bot$. Let $L_{[\sigma,\ \psi]}$ denote $\neg A_{[\sigma,\ \phi]}$ if ψ is in the form $\neg\phi$, $A_{[\sigma,\ \psi]}$ otherwise. Given

a K_m-formula φ, the encoder K_m2SAT builds a Boolean CNF formula as follows: [3]

$$K_m2SAT(\varphi) \stackrel{\text{def}}{=} A_{[1,\ \varphi]} \wedge Def(1,\ \varphi) \tag{1}$$

$$Def(\sigma,\ \top) \stackrel{\text{def}}{=} \top \tag{2}$$

$$Def(\sigma,\ \bot) \stackrel{\text{def}}{=} \top \tag{3}$$

$$Def(\sigma,\ A_i) \stackrel{\text{def}}{=} \top \tag{4}$$

$$Def(\sigma,\ \neg A_i) \stackrel{\text{def}}{=} \top \tag{5}$$

$$Def(\sigma,\ \alpha) \stackrel{\text{def}}{=} (L_{[\sigma,\ \alpha]} \rightarrow (L_{[\sigma,\ \alpha_1]} \wedge L_{[\sigma,\ \alpha_2]})) \wedge Def(\sigma,\ \alpha_1) \wedge Def(\sigma,\ \alpha_2) \tag{6}$$

$$Def(\sigma,\ \beta) \stackrel{\text{def}}{=} (L_{[\sigma,\ \beta]} \rightarrow (L_{[\sigma,\ \beta_1]} \vee L_{[\sigma,\ \beta_2]})) \wedge Def(\sigma,\ \beta_1) \wedge Def(\sigma,\ \beta_2) \tag{7}$$

$$Def(\sigma,\ \pi^{r,j}) \stackrel{\text{def}}{=} (L_{[\sigma,\ \pi^{r,j}]} \rightarrow L_{[\sigma.j,\ \pi_0^{r,j}]}) \wedge Def(\sigma.j,\ \pi_0^{r,j}) \tag{8}$$

$$Def(\sigma,\ \nu^r) \stackrel{\text{def}}{=} \bigwedge_{\langle\sigma:\pi^{r,i}\rangle} \Big(((L_{[\sigma,\ \nu^r]} \wedge L_{[\sigma,\ \pi^{r,i}]}) \rightarrow L_{[\sigma.i,\ \nu_0^r]}) \ \wedge \ Def(\sigma.i,\ \nu_0^r)\Big). \tag{9}$$

Here by "$\langle\sigma : \pi^{r,j}\rangle$" we mean that $\pi^{r,j}$ is the j-th distinct π^r formula labeled by σ. Notice that (6) and (7) generalize to the case of n-ary $\wedge$ and $\vee$ in the obvious way: if ϕ is $\bigotimes_{i=1}^n \phi_i$ s.t. $\bigotimes \in \{\wedge, \vee\}$, then $Def(\sigma,\ \phi) \stackrel{\text{def}}{=} (L_{[\sigma,\ \phi]} \rightarrow \bigotimes_{i=1}^n L_{[\sigma,\ \phi_i]}) \wedge \bigwedge_{i=1}^n Def(\sigma,\ \phi_i)$. [4]

We assume that the K_m-formulas are represented as DAGs, so that to avoid the expansion of the same $Def(\sigma,\ \psi)$ more than once. Moreover, following what done by Massacci (2000), we assume that, for each σ, the $Def(\sigma,\ \psi)$'s are expanded in the order: $\alpha/\beta, \pi, \nu$. Thus, each $Def(\sigma,\ \nu^r)$ is expanded after the expansion of all $Def(\sigma,\ \pi^{r,i})$'s, so that $Def(\sigma,\ \nu^r)$ will generate one clause $((L_{[\sigma,\ \nu^r]} \wedge L_{[\sigma,\ \pi^{r,i}]}) \rightarrow L_{[\sigma.i,\ \nu_0^r]})$ and one novel definition $Def(\sigma.i,\ \nu_0^r)$ for each $Def(\sigma,\ \pi^{r,i})$ expanded.[5]

Intuitively, it is easy to see that $K_m2SAT(\varphi)$ mimics the construction of an SST tableau expansion (Massacci, 2000; Donini & Massacci, 2000). We have the following fact.

Theorem 1. *A K_m-formula φ is K_m-satisfiable if and only if the corresponding Boolean formula $K_m2SAT(\varphi)$ is satisfiable.*

The complete proof of Theorem 1 can be found in Appendix A.

Notice that, due to (9), the number of variables and clauses in $K_m2SAT(\varphi)$ may grow exponentially with $depth(\varphi)$. This is in accordance to what stated by Halpern (1995).

Example 3.1 (NNF)**.** Let φ_{nnf} be $(\Diamond A_1 \vee \Diamond(A_2 \vee A_3)) \ \wedge \ \Box\neg A_1 \ \wedge \ \Box\neg A_2 \ \wedge \ \Box\neg A_3$.[6] It is easy to see that φ_{nnf} is K_1-unsatisfiable: the $\Diamond$-atoms impose that at least one atom A_i

3. We say that the formula is in CNF because we represent clauses as implications, according to the notation described at the beginning of Section 2.
4. Although conceptually trivial, this fact has an important practical consequence: in order to encode $\bigotimes_{i=1}^n \phi_i$ one needs adding only one Boolean variable rather than up to $n-1$. See Section 4.2.
5. In practice, even if the definition of K_m2SAT is recursive, the Def expansions are performed grouped by states. More precisely, all the $Def(\sigma.n,\ \psi)$ expansions, for any formula ψ and every defined n, are done together (in the $\alpha/\beta, \pi, \nu$ order above exposed) and necessarily after that all the $Def(\sigma,\ \varphi)$ expansions have been completed.
6. For K_1-formulas we omit the box and diamond indexes, i.e., we write $\Box$, $\Diamond$ for $\Box_1$, $\Diamond_1$.

is true in at least one successor of the root state, whilst the $\Box$-atoms impose that all atoms A_i are false in all successor states of the root state. $K_m2SAT(\varphi_{nnf})$ is:

$$
\begin{array}{lll}
1. & & A_{[1,\ \varphi_{nnf}]} \\
2. & \wedge & (\ A_{[1,\ \varphi_{nnf}]} \rightarrow (A_{[1,\ \Diamond A_1 \vee \Diamond(A_2 \vee A_3)]} \wedge A_{[1,\ \Box\neg A_1]} \wedge A_{[1,\ \Box\neg A_2]} \wedge A_{[1,\ \Box\neg A_3]})\ ) \\
3. & \wedge & (\ A_{[1,\ \Diamond A_1 \vee \Diamond(A_2 \vee A_3)]} \rightarrow (A_{[1,\ \Diamond A_1]} \vee A_{[1,\ \Diamond(A_2 \vee A_3)]})\ ) \\
4. & \wedge & (\ A_{[1,\ \Diamond A_1]} \rightarrow A_{[1.1,\ A_1]}\ ) \\
5. & \wedge & (\ A_{[1,\ \Diamond(A_2 \vee A_3)]} \rightarrow A_{[1.2,\ A_2 \vee A_3]}\ ) \\
6. & \wedge & (\ (A_{[1,\ \Box\neg A_1]} \wedge A_{[1,\ \Diamond A_1]}) \rightarrow \neg A_{[1.1,\ A_1]}\ ) \\
7. & \wedge & (\ (A_{[1,\ \Box\neg A_2]} \wedge A_{[1,\ \Diamond A_1]}) \rightarrow \neg A_{[1.1,\ A_2]}\ ) \\
8. & \wedge & (\ (A_{[1,\ \Box\neg A_3]} \wedge A_{[1,\ \Diamond A_1]}) \rightarrow \neg A_{[1.1,\ A_3]}\ ) \\
9. & \wedge & (\ (A_{[1,\ \Box\neg A_1]} \wedge A_{[1,\ \Diamond(A_2 \vee A_3)]}) \rightarrow \neg A_{[1.2,\ A_1]}\ ) \\
10. & \wedge & (\ (A_{[1,\ \Box\neg A_2]} \wedge A_{[1,\ \Diamond(A_2 \vee A_3)]}) \rightarrow \neg A_{[1.2,\ A_2]}\ ) \\
11. & \wedge & (\ (A_{[1,\ \Box\neg A_3]} \wedge A_{[1,\ \Diamond(A_2 \vee A_3)]}) \rightarrow \neg A_{[1.2,\ A_3]}\ ) \\
12. & \wedge & (\ A_{[1.2,\ A_2 \vee A_3]} \rightarrow (A_{[1.2,\ A_2]} \vee A_{[1.2,\ A_3]})\ )
\end{array}
$$

After a run of Boolean constraint propagation (BCP), 3. reduces to the implicate disjunction. If the first element $A_{[1,\ \Diamond A_1]}$ is assigned to true, then by BCP we have a conflict on 4. and 6. If it is set to false, then the second element $A_{[1,\ \Diamond(A_2 \vee A_3)]}$ is assigned to true, and by BCP we have a conflict on 12. Thus $K_m2SAT(\varphi_{nnf})$ is unsatisfiable. $\Diamond$

4. Optimizations

The basic encoding of Section 3 is rather naive, and can be much improved to many extents, in order to reduce the size of the output propositional formula, or to make it easier to solve by DPLL, or both. We distinguish two main kinds of optimizations:

Preprocessing steps, which are applied on the input modal formula before the encoding. Among them, we have *Pre-conversion into BNF* (Section 4.1), *Atom Normalization* (Section 4.2), *Box Lifting* (Section 4.3), and *Controlled Box Lifting* (Section 4.4).

On-the-fly simplification steps, which are applied to the Boolean formula under construction. Among them, we have *On-the-fly Boolean Simplification* and *Truth Propagation Through Boolean Operators* (Section 4.5) and *Truth Propagation Through Modal Operators* (Section 4.6), *On-the-fly Pure-Literal Reduction* (Section 4.7), and *On-the-fly Boolean Constraint Propagation* (Section 4.8).

We analyze these techniques in details.

4.1 Pre-conversion into BNF

Many systems use to pre-convert the input K_m-formulas into NNF (e.g., Baader et al., 1994; Massacci, 2000). In our approach, instead, we pre-convert them into BNF (like, e.g., Giunchiglia & Sebastiani, 1996; Pan et al., 2002). For our approach, the advantage of the latter representation is that, when one $\Box_r\psi$ occurs both positively and negatively (like, e.g., in $(\Box_r\psi \vee ...) \wedge (\neg\Box_r\psi \vee ...) \wedge ...$), then both occurrences of $\Box_r\psi$ are labeled by the same Boolean atom $A_{[\sigma,\ \Box_r\psi]}$, and hence they are always assigned the same truth value by DPLL.

With NNF, instead, the negative occurrence $\neg\Box_r\psi$ is rewritten into $\Diamond_r(nnf(\neg\psi))$, so that two distinct Boolean atoms $A_{[\sigma,\ \Box_r(nnf(\psi))]}$ and $A_{[\sigma,\ \Diamond_r(nnf(\neg\psi))]}$ are generated; DPLL can assign them the same truth value, creating a hidden conflict which may require some extra Boolean search to reveal.[7]

Example 4.1 (BNF)**.** We consider the BNF variant of the φ_{nnf} formula of Example 3.1, $\varphi_{bnf} = (\neg\Box\neg A_1 \vee \neg\Box(\neg A_2 \wedge \neg A_3)) \wedge \Box\neg A_1 \wedge \Box\neg A_2 \wedge \Box\neg A_3$. As before, it is easy to see that φ_{bnf} is K_1-unsatisfiable. $K_m2SAT(\varphi_{bnf})$ is:[8]

$$
\begin{array}{lll}
1. & & A_{[1,\ \varphi_{bnf}]} \\
2. & \wedge & (\ A_{[1,\ \varphi_{bnf}]} \rightarrow (A_{[1,\ (\neg\Box\neg A_1 \vee \neg\Box(\neg A_2 \wedge \neg A_3))]} \wedge A_{[1,\ \Box\neg A_1]} \wedge A_{[1,\ \Box\neg A_2]} \wedge A_{[1,\ \Box\neg A_3]})\ ) \\
3. & \wedge & (\ A_{[1,\ (\neg\Box\neg A_1 \vee \neg\Box(\neg A_2 \wedge \neg A_3))]} \rightarrow (\neg A_{[1,\ \Box\neg A_1]} \vee \neg A_{[1,\ \Box(\neg A_2 \wedge \neg A_3)]})\ ) \\
4. & \wedge & (\ \neg A_{[1,\ \Box\neg A_1]} \rightarrow A_{[1.1,\ A_1]}\ ) \\
5. & \wedge & (\ \neg A_{[1,\ \Box(\neg A_2 \wedge \neg A_3)]} \rightarrow \neg A_{[1.2,\ (\neg A_2 \wedge \neg A_3)]}\ ) \\
6. & \wedge & (\ (A_{[1,\ \Box\neg A_1]} \wedge \neg A_{[1,\ \Box\neg A_1]}) \rightarrow \neg A_{[1.1,\ A_1]}\ ) \\
7. & \wedge & (\ (A_{[1,\ \Box\neg A_2]} \wedge \neg A_{[1,\ \Box\neg A_1]}) \rightarrow \neg A_{[1.1,\ A_2]}\ ) \\
8. & \wedge & (\ (A_{[1,\ \Box\neg A_3]} \wedge \neg A_{[1,\ \Box\neg A_1]}) \rightarrow \neg A_{[1.1,\ A_3]}\ ) \\
9. & \wedge & (\ (A_{[1,\ \Box\neg A_1]} \wedge \neg A_{[1,\ \Box(\neg A_2 \wedge \neg A_3)]}) \rightarrow \neg A_{[1.2,\ A_1]}\ ) \\
10. & \wedge & (\ (A_{[1,\ \Box\neg A_2]} \wedge \neg A_{[1,\ \Box(\neg A_2 \wedge \neg A_3)]}) \rightarrow \neg A_{[1.2,\ A_2]}\ ) \\
11. & \wedge & (\ (A_{[1,\ \Box\neg A_3]} \wedge \neg A_{[1,\ \Box(\neg A_2 \wedge \neg A_3)]}) \rightarrow \neg A_{[1.2,\ A_3]}\ ) \\
12. & \wedge & (\ \neg A_{[1.2,\ (\neg A_2 \wedge \neg A_3)]} \rightarrow (A_{[1.2,\ A_2]} \vee A_{[1.2,\ A_3]})\ )
\end{array}
$$

Unlike with the NNF formula φ_{nnf} in Example 3.1, $K_m2SAT(\varphi_{bnf})$ is found unsatisfiable directly by BCP. In fact, the unit-propagation of $A_{[1,\ \Box\neg A_1]}$ from 2. causes $\neg A_{[1,\ \Box\neg A_1]}$ in 3. to be false, so that one of the two (unsatisfiable) branches induced by the disjunction is cut a priori. With φ_{nnf}, K_m2SAT does not recognize $\Box\neg A_1$ and $\Diamond A_1$ to be one the negation of the other, so that two distinct atoms $A_{[1,\ \Box\neg A_1]}$ and $A_{[1,\ \Diamond A_1]}$ are generated. Hence $A_{[1,\ \Box\neg A_1]}$ and $A_{[1,\ \Diamond A_1]}$ cannot be recognized by DPLL to be one the negation of the other, s.t. DPLL may need exploring one Boolean branch more. $\Diamond$

In the following we will assume the formulas are in BNF (although most of the optimizations which follow work also for other representations).

4.2 Normalization of Modal Atoms

One potential source of inefficiency for DPLL-based procedures is the occurrence in the input formula of semantically-equivalent though syntactically-different atoms ψ' and ψ'' (e.g., $\Box_1(A_1 \vee A_2)$ and $\Box_1(A_2 \vee A_1)$), which are not recognized as such by K_m2SAT. This causes the introduction of duplicated Boolean atoms $A_{[\sigma,\ \psi']}$ and $A_{[\sigma,\ \psi'']}$ and —much worse— of duplicated subformulas $Def(\sigma,\ \psi')$ and $Def(\sigma,\ \psi'')$. This fact can have very negative consequences, in particular when ψ' and ψ'' occur with negative polarity, because this causes the creation of distinct versions of the same successor states, and the duplication of whole parts of the output formula.

7. Notice that this consideration holds for every representation involving both boxes and diamonds; we refer to NNF simply because it is the most popular of these representations.
8. Notice that the valid clause 6. can be dropped. See the explanation in Section 4.5.

Example 4.2. Consider the K_m-formula $(\phi_1 \vee \neg\Box_1(A_2 \vee A_1)) \wedge (\phi_2 \vee \neg\Box_1(A_1 \vee A_2)) \wedge \phi_3$, s.t. ϕ_1, ϕ_2, ϕ_3 are possibly-big K_m-formulas. Then K_m2SAT creates two distinct atoms $A_{[1,\ \neg\Box_1(A_2 \vee A_1)]}$ and $A_{[1,\ \neg\Box_1(A_1 \vee A_2)]}$ and two distinct formulas $Def(1,\ \neg\Box_1(A_2 \vee A_1))$ and $Def(1,\ \neg\Box_1(A_1 \vee A_2))$. The latter will cause the creation of two distinct states 1.1 and 1.2. Thus, the recursive expansion of all $\Box_1$-formulas occurring positively in ϕ_1, ϕ_2, ϕ_3 will be duplicated for these two states. $\diamond$

In order to cope to this problem, like done by Giunchiglia and Sebastiani (1996), we apply some normalization steps to modal atoms with the intent of rewriting as many as possible syntactically-different but semantically-equivalent modal atoms into syntactically-identical ones. This can be achieved by a recursive application of some simple validity-preserving rewriting rules.

Sorting: modal atoms are internally sorted according to some criterion, so that atoms which are identical modulo reordering are rewritten into the same atom (e.g., $\Box_i(\varphi_2 \vee \varphi_1)$ and $\Box_i(\varphi_1 \vee \varphi_2)$ are both rewritten into $\Box_i(\varphi_1 \vee \varphi_2)$).

Flattening: the associativity of $\wedge$ and $\vee$ is exploited and combinations of $\wedge$'s or $\vee$'s are "flattened" into n-ary $\wedge$'s or $\vee$'s respectively (e.g., $\Box_i(\varphi_1 \vee (\varphi_2 \vee \varphi_3))$ and $\Box_i((\varphi_1 \vee \varphi_2) \vee \varphi_3)$ are both rewritten into $\Box_i(\varphi_1 \vee \varphi_2 \vee \varphi_3)$).

Flattening has also the advantage of reducing the number of novel atoms introduced in the encoding, as a consequence of the fact noticed in Footnote 4 of Section 3. One possible drawback of this technique is that it can reduce the sharing of subformulas (e.g., with $\Box_i((\varphi_1 \vee \varphi_2) \vee \varphi_3)$ and $\Box_i((\varphi_1 \vee \varphi_2) \vee \varphi_4)$, the common part is no more shared). However, we have empirically experienced that this drawback is negligible wrt. the advantages of flattening.

4.3 Box Lifting

As second preprocessing the K_m-formula can also be rewritten by recursively applying the K_m-validity-preserving "box lifting rules":

$$(\Box_r\varphi_1 \wedge \Box_r\varphi_2) \implies \Box_r(\varphi_1 \wedge \varphi_2), \quad (\neg\Box_r\varphi_1 \vee \neg\Box_r\varphi_2) \implies \neg\Box_r(\varphi_1 \wedge \varphi_2). \tag{10}$$

This has the potential benefit of reducing the number of $\pi^{r,i}$ formulas, and hence the number of labels $\sigma.i$ to take into account in the expansion of the $Def(\sigma,\ \nu^r)$'s (9). We call *lifting* this preprocessing.

Example 4.3 (Box lifting)**.** If we apply the rules (10) to the formula of Example 4.1, then we have $\varphi_{bnflift} = \neg\Box(\neg A_1 \wedge \neg A_2 \wedge \neg A_3) \wedge \Box(\neg A_1 \wedge \neg A_2 \wedge \neg A_3)$. Consequently, $K_m2SAT(\varphi_{bnflift})$ is:

1. $A_{[1,\ \varphi_{bnflift}]}$
2. $\wedge$ $(A_{[1,\ \varphi_{bnflift}]} \rightarrow (\neg A_{[1,\ \Box(\neg A_1 \wedge \neg A_2 \wedge \neg A_3)]} \wedge A_{[1,\ \Box(\neg A_1 \wedge \neg A_2 \wedge \neg A_3)]}))$
3. $\wedge$ $(\neg A_{[1,\ \Box(\neg A_1 \wedge \neg A_2 \wedge \neg A_3)]} \rightarrow \neg A_{[1.1,\ (\neg A_1 \wedge \neg A_2 \wedge \neg A_3)]})$
4. $\wedge$ $((A_{[1,\ \Box(\neg A_1 \wedge \neg A_2 \wedge \neg A_3)]} \wedge \neg A_{[1,\ \Box(\neg A_1 \wedge \neg A_2 \wedge \neg A_3)]}) \rightarrow A_{[1.1,\ (\neg A_1 \wedge \neg A_2 \wedge \neg A_3)]})$
5. $\wedge$ $(\neg A_{[1.1,\ (\neg A_1 \wedge \neg A_2 \wedge \neg A_3)]} \rightarrow (A_{[1.1,\ A_1]} \vee A_{[1.1,\ A_2]} \vee A_{[1.1,\ A_3]}))$
6. $\wedge$ $(A_{[1.1,\ (\neg A_1 \wedge \neg A_2 \wedge \neg A_3)]} \rightarrow (\neg A_{[1.1,\ A_1]} \wedge \neg A_{[1.1,\ A_2]} \wedge \neg A_{[1.1,\ A_3]})).$

$K_m2SAT(\varphi_{bnflift})$ is found unsatisfiable directly by BCP on clauses 1. and 2.. Only one successor state (1.1) is considered. Notice that 3., 4., 5. and 6. are redundant, because 1. and 2. alone are unsatisfiable. $\diamond$

4.4 Controlled Box Lifting

One potential drawback of applying the lifting rules is that, by collapsing the formula $(\Box_r\varphi_1 \wedge \Box_r\varphi_2)$ into $\Box_r(\varphi_1 \wedge \varphi_2)$ and $(\neg\Box_r\varphi_1 \vee \neg\Box_r\varphi_2)$ into $\neg\Box_r(\varphi_1 \wedge \varphi_2)$, the possibility of sharing box subformulas in the DAG representation of the input K_m-formula is reduced.

In order to cope with this problem we provide an alternative policy for applying box lifting, that is, to apply the rules (10) only when neither box subformula occurring in the implicant in (10) has multiple occurrences. We call this policy *controlled box lifting*.

Example 4.4 (Controlled Box Lifting)**.** We apply *Controlled Box Lifting* to the formula of Example 4.1, then we have $\varphi_{bnfclift} = (\neg\Box\neg A_1 \vee \neg\Box(\neg A_2 \wedge \neg A_3)) \wedge \Box\neg A_1 \wedge \Box(\neg A_2 \wedge \neg A_3)$ since the rules (10) are applied among all the box subformulas except for $\Box\neg A_1$, which is shared. It follows that $K_m2SAT(\varphi_{bnfclift})$ is:

$$
\begin{array}{lll}
1. & & A_{[1,\ \varphi_{bnfclift}]} \\
2. & \wedge & (\ A_{[1,\ \varphi_{bnfclift}]} \rightarrow (A_{[1,\ (\neg\Box\neg A_1 \vee \neg\Box(\neg A_2 \wedge \neg A_3))]} \wedge A_{[1,\ \Box\neg A_1]} \wedge A_{[1,\ \Box(\neg A_2 \wedge \neg A_3)]}\ ) \\
3. & \wedge & (\ A_{[1,\ (\neg\Box\neg A_1 \vee \neg\Box(\neg A_2 \wedge \neg A_3))]} \rightarrow (\neg A_{[1,\ \Box\neg A_1]} \vee \neg A_{[1,\ \Box(\neg A_2 \wedge \neg A_3)]})\ ) \\
4. & \wedge & (\ \neg A_{[1,\ \Box\neg A_1]} \rightarrow A_{[1.1,\ A_1]}\ ) \\
5. & \wedge & (\ \neg A_{[1,\ \Box(\neg A_2 \wedge \neg A_3)]} \rightarrow \neg A_{[1.2,\ (\neg A_2 \wedge \neg A_3)]}\ ) \\
6. & \wedge & (\ (A_{[1,\ \Box\neg A_1]} \wedge \neg A_{[1,\ \Box\neg A_1]}) \rightarrow \neg A_{[1.1,\ A_1]}\ ) \\
7. & \wedge & (\ (A_{[1,\ \Box(\neg A_2 \wedge \neg A_3)]} \wedge \neg A_{[1,\ \Box\neg A_1]}) \rightarrow A_{[1.1,\ (\neg A_2 \wedge \neg A_3)]}\ ) \\
8. & \wedge & (\ (A_{[1,\ \Box\neg A_1]} \wedge \neg A_{[1,\ \Box(\neg A_2 \wedge \neg A_3)]}) \rightarrow \neg A_{[1.2,\ A_1]}\ ) \\
9. & \wedge & (\ (A_{[1,\ \Box(\neg A_2 \wedge \neg A_3)]} \wedge \neg A_{[1,\ \Box(\neg A_2 \wedge \neg A_3)]}) \rightarrow A_{[1.2,\ (\neg A_2 \wedge \neg A_3)]}\ ) \\
10. & \wedge & (\ A_{[1.1,\ (\neg A_2 \wedge \neg A_3)]} \rightarrow (\neg A_{[1.1,\ A_2]} \wedge \neg A_{[1.1,\ A_3]})\ ) \\
11. & \wedge & (\ \neg A_{[1.2,\ (\neg A_2 \wedge \neg A_3)]} \rightarrow (A_{[1.2,\ A_2]} \vee A_{[1.2,\ A_3]})\ ) \\
12. & \wedge & (\ A_{[1.2,\ (\neg A_2 \wedge \neg A_3)]} \rightarrow (\neg A_{[1.2,\ A_2]} \wedge \neg A_{[1.2,\ A_3]})\ )
\end{array}
$$

$K_m2SAT(\varphi_{bnfclift})$ is found unsatisfiable directly by BCP on clauses 1., 2. and 3.. Notice that the unit propagation of $A_{[1,\ \Box\neg A_1]}$ and $A_{[1,\ \Box(\neg A_2 \wedge \neg A_3)]}$ from 2. causes the implicate disjunction in 3. to be false. $\diamond$

4.5 On-the-fly Boolean Simplification and Truth Propagation

A first straightforward on-the-fly optimization is that of applying recursively the standard rewriting rules for the Boolean simplification of the formula like, e.g.,

$$
\begin{array}{llll}
\langle\sigma : \varphi\rangle \wedge \langle\sigma : \varphi\rangle & \Longrightarrow \langle\sigma : \varphi\rangle, & \langle\sigma : \varphi\rangle \vee \langle\sigma : \varphi\rangle & \Longrightarrow \langle\sigma : \varphi\rangle, \\
\langle\sigma : \varphi_1\rangle \wedge \langle\sigma : (\varphi_1 \vee \varphi_2)\rangle & \Longrightarrow \langle\sigma : \varphi_1\rangle, & \langle\sigma : \varphi_1\rangle \vee \langle\sigma : (\varphi_1 \wedge \varphi_2)\rangle & \Longrightarrow \langle\sigma : \varphi_1\rangle, \\
\langle\sigma : \varphi\rangle \wedge \neg\langle\sigma : \varphi\rangle & \Longrightarrow \langle\sigma : \bot\rangle, & \langle\sigma : \varphi\rangle \vee \neg\langle\sigma : \varphi\rangle & \Longrightarrow \langle\sigma : \top\rangle, \\
\ldots, & & &
\end{array}
$$

and for the propagation of truth/falsehood through Boolean operators like, e.g.,

$$\begin{array}{llll} \neg\langle\sigma : \bot\rangle & \Longrightarrow \langle\sigma : \top\rangle, & \neg\langle\sigma : \top\rangle & \Longrightarrow \langle\sigma : \bot\rangle, \\ \langle\sigma : \varphi\rangle \wedge \langle\sigma : \top\rangle & \Longrightarrow \langle\sigma : \varphi\rangle, & \langle\sigma : \varphi\rangle \wedge \langle\sigma : \bot\rangle & \Longrightarrow \langle\sigma : \bot\rangle, \\ \langle\sigma : \varphi\rangle \vee \langle\sigma : \top\rangle & \Longrightarrow \langle\sigma : \top\rangle, & \langle\sigma : \varphi\rangle \vee \langle\sigma : \bot\rangle & \Longrightarrow \langle\sigma : \varphi\rangle, \\ \ldots. & & & \end{array}$$

Example 4.5. If we consider the K_m-formula $\varphi_{bnflift} = \neg\Box(\neg A_1 \wedge \neg A_2 \wedge \neg A_3) \wedge \Box(\neg A_1 \wedge \neg A_2 \wedge \neg A_3)$ of Example 4.3 and we apply the Boolean simplification rule $\langle\sigma : \varphi\rangle \wedge \neg\langle\sigma : \varphi\rangle \Longrightarrow \langle\sigma : \bot\rangle$, then $\langle\sigma : \varphi_{bnflift}\rangle$ is simplified into $\langle\sigma : \bot\rangle$. ◇

One important subcase of on-the-fly Boolean simplification avoids the useless encoding of incompatible π^r and ν^r formulas. In BNF, in fact, the same subformula $\Box_r\psi$ may occur in the same state σ both positively and negatively (like $\pi^{r,j} = \neg\Box_r\psi$ and $\nu^r = \Box_r\psi$). If so, K_m2SAT labels both those occurrences of $\Box_r\psi$ with the same Boolean atom $A_{[\sigma,\ \Box_r\psi]}$, and produces recursively two distinct subsets of clauses in the encoding, by applying (8) to $\neg\Box_r\psi$ and (9) to $\Box_r\psi$ respectively. However, the latter step (9) generates a *valid* clause $(A_{[\sigma,\ \Box_r\psi]} \wedge \neg A_{[\sigma,\ \Box_r\psi]}) \rightarrow A_{[\sigma.i,\ \psi]}$, so that we can avoid generating it. Consequently, if $A_{[\sigma.i,\ \psi]}$ no more occurs in the formula, then $Def(\sigma.i,\ \psi)$ should not be generated, as there is no more need of defining $\langle\sigma.i : \psi\rangle$.[9]

Example 4.6. If we apply this observation in the constructions of the formulas of Examples 4.1 and 4.4, we have the following facts:

- In the formula $K_m2SAT(\varphi_{bnf})$ of Example 4.1, clause 6. is valid and thus it is dropped.
- In the formula $K_m2SAT(\varphi_{bnfclift})$ of Example 4.4, both valid clauses 6. and 9. are dropped, so that 12. is not generated. ◇

Hereafter we assume that on-the-fly Boolean simplification is applied also in combination with the techniques described in the next sections.

4.6 On-the-fly Truth Propagation Through Modal Operators

Truth and falsehood —which can derive by the application of the techniques in Section 4.5, Section 4.7 and Section 4.8— may be propagated on-the-fly also though modal operators. First, for every σ, both positive and negative instances of $\langle\sigma : \Box_r\top\rangle$ can be safely simplified by applying the rewriting rule $\langle\sigma : \Box_r\top\rangle \Longrightarrow \langle\sigma : \top\rangle$.

Second, we notice the following fact. When we have a positive occurrence of $\neg\Box_r\bot$ for some σ,[10] by (8) and (9) we have

$$Def(\sigma,\ \neg\Box_r\bot) \stackrel{\text{def}}{=} (L_{[\sigma,\ \neg\Box_r\bot]} \rightarrow A_{[\sigma.j,\ \top]}) \wedge Def(\sigma.j,\ \top), \tag{11}$$

$$Def(\sigma,\ \Box_r\psi) \stackrel{\text{def}}{=} ((L_{[\sigma,\ \Box_r\psi]} \wedge L_{[\sigma,\ \neg\Box_r\bot]}) \rightarrow L_{[\sigma.j,\ \psi]}) \wedge Def(\sigma.j,\ \psi) \tag{12}$$

9. Here the "if" is due to the fact that it may be the case that $A_{[\sigma.i,\ \psi]}$ is generated anyway from the expansion of some other subformula, like, e.g., $\Box_r(\psi \vee \phi)$. If this is the case, $Def(\sigma.i,\ \psi)$ must be generated anyway.
10. E.g., $\neg\Box_r\bot$ may result from applying the steps of Section 4.1 and of Section 4.5 to $\neg\Box_r(\Box_r A_1 \wedge \Diamond_r\neg A_1)$.

for some new label $\sigma.j$ and for every $\Box_r\psi$ occurring positively in σ. $Def(\sigma,\ \neg\Box_r\bot)$ reduces to $\top$ because both $A_{[\sigma.j,\ \top]}$ and $Def(\sigma.j,\ \top)$ reduce to $\top$. If at least another distinct π-formula $\neg\Box_r\varphi$ occurs positively in σ, however, there is no need for the $\sigma.j$ label in (11) and (12) to be a *new* label, and we can re-use instead the label $\sigma.i$ introduced in the expansion of $Def(\sigma,\ \neg\Box_r\varphi)$:

$$Def(\sigma,\ \neg\Box_r\varphi) \stackrel{\text{def}}{=} (L_{[\sigma,\ \neg\Box_r\varphi]} \rightarrow L_{[\sigma.i,\ \neg\varphi]}) \wedge Def(\sigma.i,\ \neg\varphi). \tag{13}$$

Thus (11) is dropped and, for every $\Box_r\psi$ occurring positively in σ, (12) is instead written as

$$Def(\sigma,\ \Box_r\psi) \stackrel{\text{def}}{=} ((L_{[\sigma,\ \Box_r\psi]} \wedge L_{[\sigma,\ \neg\Box_r\bot]}) \rightarrow L_{[\sigma.i,\ \psi]}) \wedge Def(\sigma.i,\ \psi) \tag{14}$$

(Notice the label $\sigma.i$ introduced in (13) rather than the label $\sigma.j$ of (11).)

This is motivated by the fact that $Def(\sigma,\ \neg\Box_r\bot)$ forces the existence of at least one successor of σ but imposes no constraints on which formulas should hold there, so that we can use some other already-defined successor state, if any. This fact has the important benefit of eliminating useless successor states from the encoding.

Example 4.7. Let φ be the BNF K-formula:

$$(\neg A_1 \vee \neg\Box A_2) \wedge (A_1 \vee \neg\Box\bot) \wedge (\neg A_1 \vee A_3) \wedge (\neg A_1 \vee \neg A_3) \wedge (A_1 \vee \Box\neg A_4) \wedge \Box A_4.$$

φ is K-inconsistent, because the only possible assignment is $\{\neg A_1, \neg\Box\bot, \Box\neg A_4, \Box A_4\}$, which is K-inconsistent. $K_m2SAT(\varphi)$ is encoded as follows:

1. $A_{[1,\ \varphi]}$
2. $\wedge$ $(A_{[1,\ \varphi]} \rightarrow (A_{[1,\ (\neg A_1\vee\neg\Box A_2)]} \wedge A_{[1,\ (A_1\vee\neg\Box\bot)]} \wedge A_{[1,\ (\neg A_1\vee A_3)]} \wedge A_{[1,\ (A_1\vee\Box\neg A_4)]} \wedge A_{[1,\ \Box A_4]}))$
3. $\wedge$ $(A_{[1,\ (\neg A_1\vee\neg\Box A_2)]} \rightarrow (\neg A_{[1,\ A_1]} \vee \neg A_{[1,\ \Box A_2]}))$
4. $\wedge$ $(A_{[1,\ (A_1\vee\neg\Box\bot)]} \rightarrow (A_{[1,\ A_1]} \vee \neg A_{[1,\ \Box\bot]}))$
5. $\wedge$ $(A_{[1,\ (\neg A_1\vee A_3)]} \rightarrow (\neg A_{[1,\ A_1]} \vee A_{[1,\ A_3]}))$
6. $\wedge$ $(A_{[1,\ (\neg A_1\vee\neg A_3)]} \rightarrow (\neg A_{[1,\ A_1]} \vee \neg A_{[1,\ A_3]}))$
7. $\wedge$ $(A_{[1,\ (A_1\vee\Box\neg A_4)]} \rightarrow (A_{[1,\ A_1]} \vee A_{[1,\ \Box\neg A_4]}))$
8. $\wedge$ $(\neg A_{[1,\ \Box A_2]} \rightarrow \neg A_{[1.1,\ A_2]})$
9. $\wedge$ $((A_{[1,\ \Box\neg A_4]} \wedge \neg A_{[1,\ \Box A_2]}) \rightarrow \neg A_{[1.1,\ A_4]})$
10. $\wedge$ $((A_{[1,\ \Box A_4]} \wedge \neg A_{[1,\ \Box A_2]}) \rightarrow A_{[1.1,\ A_4]})$
11. $\wedge$ $(\neg A_{[1,\ \Box\bot]} \rightarrow \neg A_{[1.1,\ \bot]})$
12. $\wedge$ $((A_{[1,\ \Box\neg A_4]} \wedge \neg A_{[1,\ \Box\bot]}) \rightarrow \neg A_{[1.1,\ A_4]})$
13. $\wedge$ $((A_{[1,\ \Box A_4]} \wedge \neg A_{[1,\ \Box\bot]}) \rightarrow A_{[1.1,\ A_4]})$

Clause 11. is then simplified into $\top$. (In a practical implementation it is not even generated.) Notice that in clauses 11., 12. and 13. it is used the label 1.1 of clauses 8., 9. and 10. rather than a new label 1.2. Thus, only one successor label is generated.

When DPLL is run on $K_m2SAT(\varphi)$, by BCP 1. and 2. are immediately satisfied and the implicants are removed from 3., 4., 5., 6.. Thanks to 5. and 6., $A_{[1,\ A_1]}$ can be assigned only to false, which causes 3. to be satisfied and forces the assignment of the literals $\neg A_{[1,\ \Box\bot]}$, $A_{[1,\ \Box\neg A_4]}$ by BCP on 3. and 7. and hence of $\neg A_{[1.1,\ \bot]}$, $\neg A_{[1.1,\ A_4]}$ and $A_{[1.1,\ A_4]}$ by BCP on 12. and 13., causing a contradiction. $\diamond$

It is worth noticing that (14) is strictly necessary for the correctness of the encoding even when another π-formula occurs in σ. (E.g., in Example 4.7, without 12. and 13. the formula $K_m2SAT(\varphi)$ would become satisfiable because $A_{[1,\ \Box A_2]}$ could be safely be assigned to true by DPLL, which would satisfy 8., 9. and 10..)

Hereafter we assume that this technique is applied also in combination with the techniques described in Section 4.5 and in the next sections.

4.7 On-the-fly Pure-Literal Reduction

Another technique, evolved from that proposed by Pan and Vardi (2003), applies Pure-Literal Reduction (PLR) on-the-fly during the construction of $K_m2SAT(\varphi)$. When for a label σ all the clauses containing atoms in the form $A_{[\sigma,\ \psi]}$ have been generated, if some of them occurs only positively [resp. negatively], then it can be safely assigned to true [resp. to false], and hence the clauses containing $A_{[\sigma,\ \psi]}$ can be dropped.[11] As a consequence, some other atom $A_{[\sigma,\ \psi']}$ can become pure, so that the process is repeated until a fixpoint is reached.

Example 4.8. Consider the formula φ_{bnf} of Example 4.1. During the construction of $K_m2SAT(\varphi_{bnf})$, after 1.-8. are generated, no more clause containing atoms in the form $A_{[1.1,\ \psi]}$ is to be generated. Then we notice that $A_{[1.1,\ A_2]}$ and $A_{[1.1,\ A_3]}$ occur only negatively, so that they can be safely assigned to false. Therefore, 7. and 8. can be safely dropped. Same discourse applies lately to $A_{[1.2,\ A_1]}$ and 9.. The resulting formula is found inconsistent by BCP. (In fact, notice from Example 4.1 that the atoms $A_{[1.1,\ A_2]}$, $A_{[1.1,\ A_3]}$, and $A_{[1.2,\ A_1]}$ play no role in the unsatisfiability of $K_m2SAT(\varphi_{bnf})$.) $\diamond$

We remark the differences between PLR and the Pure-Literal Reduction technique proposed by Pan and Vardi (2003). In KBDD (Pan et al., 2002; Pan & Vardi, 2003), the Pure-Literal Reduction is a preprocessing step which is applied to the input modal formula, either at global level (i.e. looking for pure-polarity primitive propositions for the whole formula) or, more effectively, at different modal depths (i.e. looking for pure-polarity primitive propositions for the subformulas at the same nesting level of modal operators).

Our technique is much more fine-grained, as PLR is applied on-the-fly with a single-state granularity, obtaining a much stronger reduction effect.

Example 4.9. Consider again the BNF K_m-formula φ_{bnf} discussed in Examples 4.1 and 4.8: $\varphi_{bnf} = (\neg\Box\neg A_1 \vee \neg\Box(\neg A_2 \wedge \neg A_3)) \wedge \Box\neg A_1 \wedge \Box\neg A_2 \wedge \Box\neg A_3$. It is immediate to see that all primitive propositions A_1, A_2, A_3 occur at every modal depth with both polarities, so that the technique of Pan and Vardi (2003) produces no effect on this formula. $\diamond$

4.8 On-the-fly Boolean Constraint Propagation

One major problem of the basic encoding of Section 3 is that it is "purely-syntactic", that is, it does not consider the possible truth values of the subformulas, and the effect of their propagation through the Boolean and modal connectives. In particular, K_m2SAT applies

11. In our actual implementation this reduction is performed directly within an intermediate data structure, so that these clauses are never generated.

(8) [resp. (9)] to *every* π-subformula [resp. ν-subformula], regardless the fact that the truth values which can be deterministically assigned to the labeled subformulas of $\langle 1 : \varphi \rangle$ may allow for dropping some labeled π-/ν-subformulas, and thus prevent the need of encoding them.

One solution to this problem is that of applying Boolean Constraint Propagation (BCP) on-the-fly during the construction of $K_m2SAT(\varphi)$, starting from the fact that $A_{[1,\ \varphi]}$ must be true. If a contradiction is found, then $K_m2SAT(\varphi)$ is unsatisfiable, so that the formula is not expanded any further, and the encoder returns the formula "$\bot$".[12] When BCP allows for dropping one implication in (6)-(9) without assigning some of its implicate literals, namely $L_{[\sigma,\ \psi_i]}$, then $\langle \sigma : \psi_i \rangle$ needs not to be defined, so that $Def(\sigma,\ \psi_i)$ must not be expanded.[13] Importantly, dropping $Def(\sigma,\ \pi^{r,j})$ for some π-formula $\langle \sigma : \pi^{r,j} \rangle$ prevents generating the label $\sigma.j$ (8) and all its successor labels $\sigma.j.\sigma'$ (corresponding to the subtree of states rooted in $\sigma.j$), so that all the corresponding labeled subformulas are not encoded.

Example 4.10. Consider Example 4.1, and suppose we apply on-the-fly BCP. During the construction of 1., 2. and 3. in $K_m2SAT(\varphi_{bnf})$, the atoms $A_{[1,\ \varphi_{bnf}]}$, $A_{[1,\ (\neg\Box\neg A_1 \vee \neg\Box(\neg A_2 \wedge \neg A_3))]}$, $A_{[1,\ \Box\neg A_1]}$, $A_{[1,\ \Box\neg A_2]}$ and $A_{[1,\ \Box\neg A_3]}$ are deterministically assigned to true by BCP. This causes the removal from 3. of the first-implied disjunct $\neg A_{[1,\ \Box\neg A_1]}$, so that there is no need to generate $Def(1,\ \neg\Box\neg A_1)$, and hence label 1.1. is not defined and 4. is not generated. While building 5., $A_{[1.2,\ (\neg A_2 \wedge \neg A_3)]}$, is unit-propagated. As label 1.1. is not defined, 6., 7. and 8. are not generated. Then during the construction of 5., 9., 10., 11. and 12., by applying BCP a contradiction is found, so that $K_m2SAT(\varphi)$ is $\bot$.

An analogous situation happens with $\varphi_{bnflift}$ in Example 4.3: while building 1. and 2. a contradiction is found by BCP, s.t. K_m2SAT returns $\bot$ without expanding the formula any further. Same discourse holds for $\varphi_{bnfclift}$ in Example 4.4: while building 1., 2. and 3. a contradiction is found by BCP, s.t. K_m2SAT returns $\bot$ without expanding the formula any further. $\diamond$

4.9 A Paradigmatic Example: Halpern & Moses Branching Formulas.

Among all optimizations described in this section, on-the-fly BCP is by far the most effective. In order to better understand this fact, we consider as a paradigmatic example the branching formulas φ_h^K by Halpern and Moses (1992, 1995) (also called "`k_branch_n`" in the set of benchmark formulas proposed by Heuerding and Schwendimann, 1996) and their unsatisfiable version (called "`k_branch_p`" in the above-mentioned benchmark suite).

12. For the sake of compatibility with standard SAT solvers, our actual implementation returns the formula $A_1 \wedge \neg A_1$.
13. Here we make the same consideration as in Footnote 9: if $L_{[\sigma.j,\ \psi]}$ is generated also from the expansion of some other subformula, (e.g., $\Box_r(\psi \vee \phi)$), then (another instance of) $Def(\sigma.i,\ \psi)$ must be generated anyway.

Given a single modality $\Box$, an integer parameter h, and the primitive propositions $D_0, \ldots, D_{h+1}, P_1, \ldots, P_h$, the formulas φ_h^K are defined as follows:[14]

$$\varphi_h^K \stackrel{\text{def}}{=} D_0 \ \wedge \ \neg D_1 \ \wedge \ \bigwedge_{i=0}^{h} \Box^i(\mathit{depth} \wedge \mathit{determined} \wedge \mathit{branching}), \tag{15}$$

$$\mathit{depth} \stackrel{\text{def}}{=} \bigwedge_{i=1}^{h+1} (D_i \rightarrow D_{i-1}), \tag{16}$$

$$\mathit{determined} \stackrel{\text{def}}{=} \bigwedge_{i=1}^{h} \left(D_i \ \rightarrow \ \left(\begin{array}{ll} (& P_i \rightarrow \Box(D_i \rightarrow \ \ P_i)) \ \ \wedge \\ (\neg & P_i \rightarrow \Box(D_i \rightarrow \neg P_i)) \end{array} \right) \right), \tag{17}$$

$$\mathit{branching} \stackrel{\text{def}}{=} \bigwedge_{i=0}^{h-1} \left((D_i \wedge \neg D_{i+1}) \ \rightarrow \ \left(\begin{array}{ll} \Diamond(D_{i+1} \wedge \neg D_{i+2} \wedge & P_{i+1}) \ \ \wedge \\ \Diamond(D_{i+1} \wedge \neg D_{i+2} \wedge & \neg P_{i+1}) \end{array} \right) \right). \tag{18}$$

A conjunction of the formulas *depth*, *determined* and *branching* is repeated at every nesting level of modal operators (i.e. at every depth): *depth* captures the relation between the D_i's at every level; *determined* states that, if P_i is true [false] in a state at depth $\geq i$, then it is true [false] in all the successor states of depth $\geq i$; *branching* states that, for every node at depth i, it is possible to find two successor states at depth $i+1$ such that P_{i+1} is true in one and false in the other. For each value of the parameter h, φ_h^K is K-satisfiable, and every Kripke model M that satisfies it has at least $2^{h+1}-1$ states. In fact, φ_h^K is build in such a way to force the construction of a binary-tree Kripke model of depth $h+1$, each of whose leaves encodes a distinct truth assignment to the primitive propositions $P_1, \ldots, P_h$, whilst each D_i is true in all and only the states occurring at a depth $\geq i$ in the tree (and thus denotes the level of nesting).

The unsatisfiable counterpart formulas `k_branch_p` proposed by Heuerding and Schwendimann (1996) are obtained by conjoining to (15) the formula:

$$\Box^h P_{(h \ div \ 3+1)} \tag{19}$$

which forces the atom $P_{(h \ div \ 3+1)}$ to be true in all depth-h states of the candidate Kripke model, which is incompatible with the fact that the remaining specifications say that it has to be false in half depth-h states.[15]

These formulas are very pathological for many approaches (Giunchiglia & Sebastiani, 2000; Giunchiglia, Giunchiglia, Sebastiani, & Tacchella, 2000; Horrocks et al., 2000). In particular, before introducing on-the-fly BCP, they used to be the pet hate of our $K_m 2SAT$ approach, as they caused the generation of huge Boolean formulas. In fact, due to *branching* (18), φ_h^K contains $2h$ $\Diamond$-formulas (i.e., π-formulas) at every depth. Therefore, the $K_m 2SAT$ encoder of Section 3 has to consider $1 + 2h + (2h)^2 + ... + (2h)^{h+1} = ((2h)^{h+2} - 1)/(2h-1)$ distinct labels, which is about h^{h+1} times the number of those labeling the states which are actually needed. (None of the optimizations of Sections 4.1-4.7 is of any help with

14. For the sake of better readability, here we adopt the description given by Halpern and Moses (1992) without converting the formulas into BNF. This fact does not affect the discussion.

15. Heuerding and Schwendimann (1996) do not explain the choice of the index "$(h \ div \ 3 + 1)$". We understand that also other choices would have done the job.

these formulas, because neither BNF encoding nor atom normalization causes any sharing of subformulas, the formulas are already in lifted form, and no literal occurs pure.[16])

This pathological behavior can be mostly overcome by applying on-the-fly-BCP, because some truth values can be deterministically assigned to some subformulas of φ_h^K by on-the-fly-BCP, which prevent encoding some or even most $\Box/\Diamond$-subformulas.

In fact, consider the *branching* and *determined* formulas occurring in φ_h^K at a generic depth $d \in \{0...h\}$, which determine the states at level d in the tree. As in these states $D_0, ..., D_d$ are forced to be true and $D_{d+1}, ..., D_{h+1}$ are forced to be false, then all but the d-th conjunct in *branching* (all conjuncts if $d = h$) are forced to be true and thus they could be dropped. Therefore, only 2 $\Diamond$-formulas per non-leaf level could be considered instead, causing the generation of $2^{h+1} - 1$ labels overall. Similarly, in all states at level d the last $h-d$ conjuncts in *determined* are forced to be true and could be dropped, reducing significantly the number of $\Box$-formulas to be considered.

It is easy to see that this is exactly what happens by applying on-the-fly-BCP. In fact, suppose that the construction of $K_m2SAT(\varphi_h^K)$ has reached depth d (that is, the point where for every state σ at level d, the $Def(\sigma,\ \alpha)$'s and $Def(\sigma,\ \beta)$'s are expanded but no $Def(\sigma,\ \pi)$ and $Def(\sigma,\ \nu)$ is expanded yet). Then, BCP deterministically assigns true to the literals $L_{[\sigma,\ D_0]}, ..., L_{[\sigma,\ D_d]}$ and false to $L_{[\sigma,\ D_{d+1}]}, ..., L_{[\sigma,\ D_{h+1}]}$, which removes all but one conjuncts in *branching*, so that only two $Def(\sigma,\ \pi)$'s out of $2h$ ones are actually expanded; similarly, the last $h - d$ conjuncts in *determined* are removed, so that the corresponding $Def(\sigma,\ \nu)$'s are not expanded.

As far as the unsatisfiable version $K_m2SAT(\varphi_h^K \wedge \Box^h P_{(h\ div\ 3+1)})$ is concerned, when the expansion reaches depth h, thanks to (19), $L_{[\sigma,\ P_{(h\ div\ 3+1)}]}$ is generated and deterministically assigned to true by BCP for every depth-h label σ; thanks to *determined* and *branching*, BCP assigns all literals $L_{[\sigma,\ P_1]}, ..., L_{[\sigma,\ P_h]}$ deterministically, so that $L_{[\sigma,\ P_{(h\ div\ 3+1)}]}$ is assigned to false for 50% of the depth-h labels σ. This causes a contradiction, so that the encoder stops the expansion and returns $\bot$.

Figure 2 shows the growth in size and the CPU time required to encode and solve $K_m2SAT(\varphi_h^K)$ (1st row) and $K_m2SAT(\varphi_h^K \wedge \Box^h P_{(h\ div\ 3+1)})$ (2nd row) wrt. the parameter h, for eight combinations of the following options of the encoder: with and without box-lifting, with and without on-the-fly PLR, with and without on-the-fly BCP. (Notice the log scale of the y axis.) In Figure 2(d) the plots of the four versions "`-xxx-bcp`" (with on-the-fly BCP) coincide with the x axis (i.e, one variable) and in Figure 2(e) they coincide with an horizontal line of value 2 (i.e, two clauses), corresponding to the fact that the 1-variable/2-clause formula $A_1 \wedge \neg A_1$ is returned (see Footnote 12).

We notice a few facts. First, for both formulas, the eight plots always collapse into two groups of overlapping plots, representing the four variants with and without on-the-fly BCP respectively. This shows that box-lifting and on-the-fly PLR are irrelevant in the encoding of these formulas, causing just little variations in the time required by the encoder (Figures 2(c) and 2(f)). Second, the four versions with on-the-fly-BCP always outperform of several orders magnitude these without this option, in terms of both size of encoded formulas and of CPU time required to encode and solve them. In particular, in the case of the unsatisfiable variant (Figure 2, second row) the encoder returns the $\bot$ formula, so that

16. More precisely, only one literal, $\neg D_{h+1}$, occurs pure in *branching*, but assigning it plays no role in simplifying the formula.

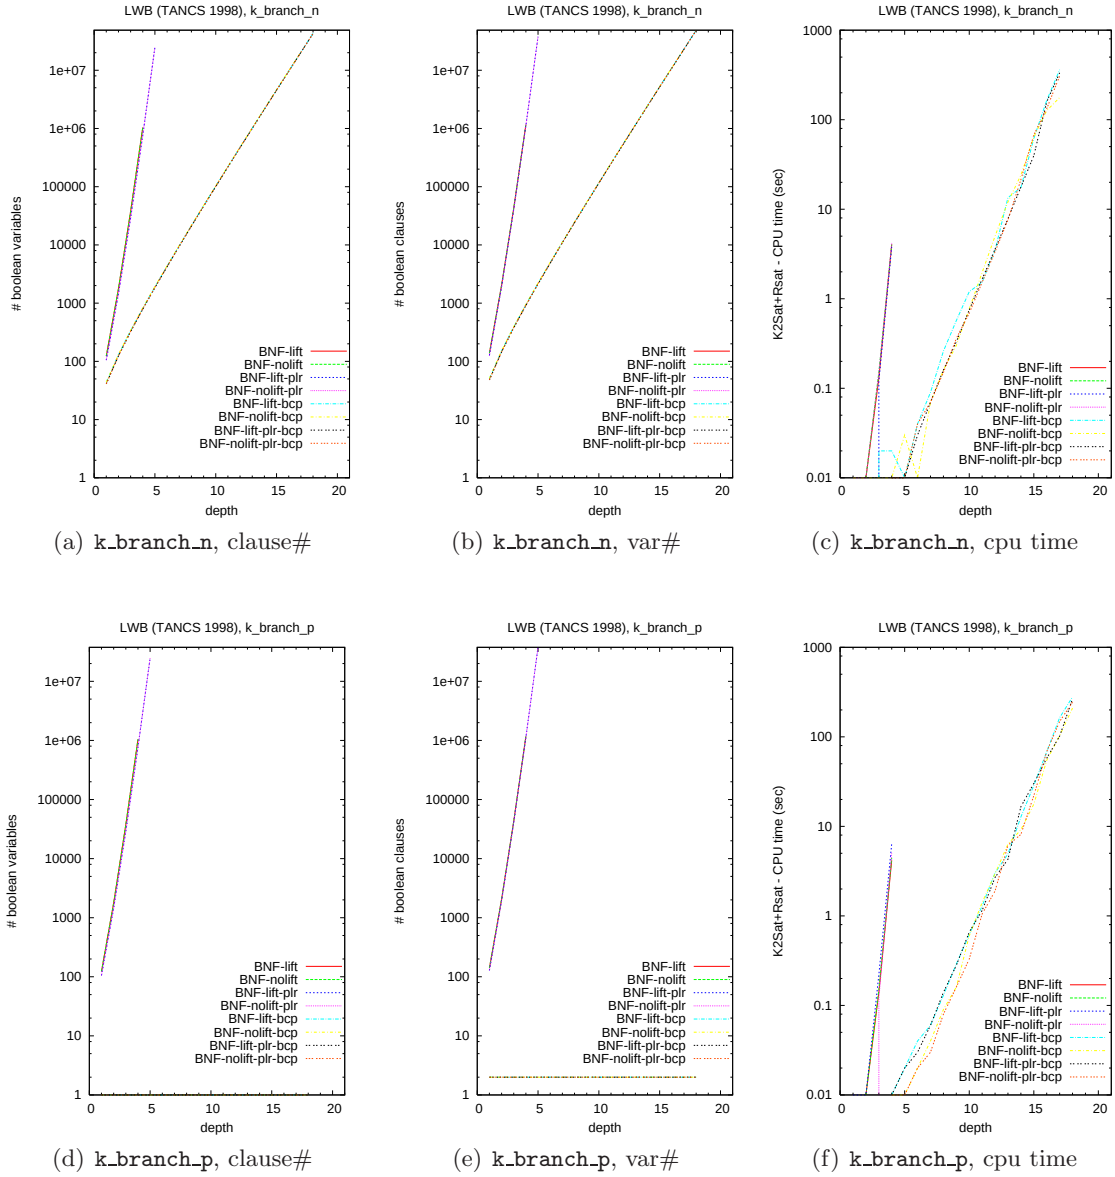

(a) `k_branch_n`, clause# (b) `k_branch_n`, var# (c) `k_branch_n`, cpu time

(d) `k_branch_p`, clause# (e) `k_branch_p`, var# (f) `k_branch_p`, cpu time

Figure 2: Empirical analysis of $K_m 2SAT$ on Halpern & Moses formulas wrt. the depth parameter h, for different options of the encoder. 1st row: `k_branch_n` formulas (satisfiable); 2nd row: `k_branch_p` formulas (unsatisfiable). Left: number of Boolean variables; center: number of clauses; right: total CPU time requested to encoding+solving. Note: in Figure 2(d) the plots of the four versions "`-xxx-bcp`" coincide with the x axis, whilst in Figure 2(e) they coincide with the horizontal line of value 2.

no actual work is required to the SAT solver (the plot of Figure 2(f) refers only to encoding time).

5. Empirical Evaluation

In order to verify empirically the effectiveness of this approach, we have performed a very-extensive empirical test session on about 14,000 $K_m/\mathcal{ALC}$ formulas. We have implemented the encoder $K_m 2SAT$ in C++, with some flags corresponding to the optimizations exposed in the previous section: (i) `NNF/BNF`, performing a pre-conversion into NNF/BNF before the encoding; (ii) `lift/ctrl.lift/nolift`, performing respectively Box Lifting, Controlled Box Lifting or no Box Lifting before the encoding; (iii) `plr` if on-the-fly Pure Literal Reduction is performed and (iv) `bcp` if on-the-fly Boolean Constraint Propagation is performed. The techniques introduced in Section 4.2, Section 4.5 and Section 4.6 are hardwired in the encoder. Moreover, as pre-conversion into `BNF` almost always produces smaller formulas than `NNF`, we have set the `BNF` flag as a default.

In combination with $K_m 2SAT$ we have tried several SAT solvers on our encoded formulas (including ZChaff `2004.11.15`, Siege `v4`, BerkMin `5.6.1`, MiniSat `v1.13`, SAT-Elite `v1.0`, SAT-Elite GTI `2005 submission`, MiniSat 2.0 `061208` and Rsat `1.03`). After a preliminary evaluation and further intensive experiments we have selected Rsat `1.03` (Pipatsrisawat & Darwiche, 2006), because it produced the best overall performances on our benchmark suites (although the performance gaps wrt. other SAT tools, e.g. MiniSat 2.0, were not dramatic).

Among the state-of-the-art tools for K_m-satisfiability, we have selected Racer (V. Haarslev, 2001) and *SAT (Tacchella, 1999) as the best representants of the tableaux/DPLL-based tools, Mspass (Hustadt & Schmidt, 1999; Hustadt et al., 1999)[17] as the best representant of the FOL-encoding approach, KBDD (Pan et al., 2002; Pan & Vardi, 2003)[18] as the representant of the automata-theoretic approach. No representative of the CSP-based and of the inverse method approaches could be used.[19] Finally, as representant of the QBF-encoding approach, we have selected the K-QBF translator (Pan & Vardi, 2003) combined with the sKizzo version `0.8.2` QBF solver (Benedetti, 2005), which turned out to be by far[20] the best QBF solver on our benchmarks among the freely-available QBF solvers from the QBF2006 competition (Narizzano, Pulina, & Tacchella, 2006). (In our evaluation we

17. We have run Mspass with the options `-EMLTranslation=2 -EMLFuncNary=1 -Sorts=0 -CNFOptSkolem=0 -CNFStrSkolem=0 -Select=2 -Split=-1 -DocProof=0 -PProblem=0 -PKept=0 -PGiven=0`, which are suggested for K_m-formulas in the Mspass README file. We have also tried other options, but the former gave the best performances.
18. KBDD has been recompiled to be run with an increased internal memory bound of 1 GB.
19. At the moment KK is not freely available, and we failed in the attempt of obtaining it from the authors. KCSP is a prolog piece of software, which is difficult to compare in performances wrt. other optimized tools on a common platform; moreover, KCSP is no more maintained since 2005, and it is not competitive wrt. state-of-the-art tools (Brand, 2008). Other tools like leanK, □KE, LWB, Kris are not competitive with the ones listed above (Horrocks et al., 2000). Tools like DLP and FaCT are similar in spirit and performances to Racer and *SAT. KSAT (Giunchiglia & Sebastiani, 1996, 2000; Giunchiglia et al., 2000) has been reimplemented into *SAT.
20. Unlike with the choice of SAT solver, the performance gaps from the best choice and the others were very significant: e.g., in the LWB benchmark (see later), sKizzo was able to solve nearly 90 problems more than its best QBF competitor.

have considered the tools : 2clsQ, SQBF, preQuantor—i.e. preQuel +Quantor—Quantor 2.11, and Semprop 010604.)

All tests presented in this section have been performed on a two-processor Intel Xeon 3.0GHz computer, with 1 MByte Cache each processor, 4 GByte RAM, with Red Hat Linux 3.0 Enterprise Server, where four processes can run in parallel. When reporting the results for one K_m2SAT +Rsat version, the CPU times reported are the sums of both the encoding and Rsat solving times. When reporting the results for K-QBF +sKizzo, the CPU times reported are only due to sKizzo because the time spent by the K-QBF converter is negligible.

We anticipate that, for all formulas of all benchmark suites, all tools under test —i.e. all the variants of K_m2SAT +Rsat and all the state-of-the-art K_m-satisfiability solvers— agreed on the satisfiability/unsatisfiability result when terminating within the timeout.

Remark 1. Due to the big number of empirical tests performed and to the huge amount of data plotted, and due to limitations in size, the plots are not always clearly distinguishable in the figures. Thus, for the sake of the reader's convenience, full-size versions of all plots are available in the online appendix, together with the files with all data.

5.1 Test Description

We have performed our empirical evaluation on three different well-known benchmarks suites of $K_m/\mathcal{ALC}$ problems: the LWB (Heuerding & Schwendimann, 1996), the random $\Box_m$-CNF (Horrocks et al., 2000; Patel-Schneider & Sebastiani, 2003) and the TANCS 2000 (Massacci & Donini, 2000) benchmark suites. We are not aware of any other publicly-available benchmark suite on $K_m/\mathcal{ALC}$-satisfiability from the literature. These three groups of benchmark formulas allow us for test the effectiveness of our approach on a large number of problems of various sizes, hardness and characteristics, for a total amount of about 14,000 formulas.

In the following we describe these three benchmark suites.

The LWB Benchmark Suite

As a first group of benchmark formulas we used the LWB benchmark suite used in a comparison at Tableaux'98 (Heuerding & Schwendimann, 1996). It consists on 9 classes of parametrized formulas (each in two versions, valid "_p" or invalid "_n" [21]), for a total amount of 378 formulas. The parameter allows for creating formulas of increasing size and difficulty.

The benchmark methodology is to test formulas from each class, in increasing difficulty, until one formula cannot be solved within a given timeout, 1000 seconds in our tests. [22] The result from this class is the parameter's value of the largest (and hardest) formula that can be solved within the time limit. The parameter ranges only from 1 to 21 so that, if a system can solve all 21 instances of a class, the result is given as 21. For a discussion on this benchmark suite, we refer the reader to the work of Heuerding and Schwendimann (1996) and of Horrocks et al. (2000).

21. As all tools check K_m-(un)satisfiability, all formulas are negated.
22. We also set a 1 GB file-size limit for the encoding produced by K_m2SAT.

The Random $\Box_m$-CNF Benchmark Suite

As a second group of benchmark formulas, we have selected the random $\Box_m$-CNF testbed described by Horrocks et al. (2000), Patel-Schneider and Sebastiani (2003). This is a generalization of the well-known random k-SAT test methods, and is the final result of a long discussion in the communities of modal and description logics on how to to obtain significant and flawless random benchmarks for modal/description logics (Giunchiglia & Sebastiani, 1996; Hustadt & Schmidt, 1999; Giunchiglia et al., 2000; Horrocks et al., 2000; Patel-Schneider & Sebastiani, 2003).

In the $\Box_m$-CNF test methodology, a $\Box_m$-CNF formula is randomly generated according to the following parameters:

- the (maximum) modal depth d;
- the number of top-level clauses L;
- the number of literal per clause clauses k;
- the number of distinct propositional variables N;
- the number of distinct box symbols m;
- the percentage p of purely-propositional literals in clauses occurring at depth $< d$, s.t. each clause of length k contains on average $p \cdot k$ randomly-picked Boolean literals and $k - p \cdot k$ randomly-generated modal literals $\Box_r \psi$, $\neg \Box_r \psi$. [23]

(We refer the reader to the works of Horrocks et al., 2000 and Patel-Schneider & Sebastiani, 2003 for a more detailed description.)

A typical problem set is characterized by fixed values of d, k, N, m, and p: L is varied in such a way as to empirically cover the "100% satisfiable / 100% unsatisfiable" transition. Then, for each tuple of the five values in a problem set, a certain number of $\Box_m$-CNF formulas are randomly generated, and the resulting formulas are given in input to the procedure under test, with a maximum time bound. The fraction of formulas which were solved within a given timeout, and the median/percentile values of CPU times are plotted against the ratio L/N. Also, the fraction of satisfiable/unsatisfiable formulas is plotted for a better understanding.

Following the methodology proposed by Horrocks et al. (2000) and by Patel-Schneider and Sebastiani (2003), we have fixed $m = 1$, $k = 3$ and 100 samples per point in all tests, and we have selected two groups: an "easier" one, with $d = 1$, $p = 0.5$, $N = 6, 7, 8, 9$, $L/N = 10..60$, and a "harder" one, with $d = 2$, $p = 0.6, 0.5$, $N = 3, 4$, $L/N = 30..150$ with $p = 0.6$ and $L/N = 50..140$ with $p = 0.5$, varying the L/N ratio in steps of 5, for a total amount of 13,200 formulas.

In each test, we imposed a timeout of 500 seconds per sample[24] and we plot the number of samples which were solved within the timeout, and the 50%th and 90%th percentiles of

23. More precisely, the number of Boolean literals in a clause is $\lfloor p \cdot k \rfloor$ (resp. $\lceil p \cdot k \rceil$) with probability $\lceil p \cdot k \rceil - p \cdot k$ (resp. $1 - (\lceil p \cdot k \rceil - p \cdot k)$). Notice that typically the smaller is p, the harder is the problem (Horrocks et al., 2000; Patel-Schneider & Sebastiani, 2003).
24. With also a 512 MB file-size limit for the encoding produced by $K_m 2SAT$.

CPU time. In order to correlate the performances with the (un)satisfiability of the sample formulas, in the background of each plot we also plot the satisfiable/unsatisfiable ratio.

The TANCS 2000 Benchmark Suite

Finally, as a third group of benchmark formulas, we used the MODAL PSPACE division benchmark suite used in the comparison at TANCS 2000 (Massacci & Donini, 2000). It contains both satisfiable and unsatisfiable formulas, with scalable hardness. In this benchmark suite, that we call TANCS 2000, the formulas are constructed by translating QBF formulas into K using three translation schemes, namely the Schmidt-Schauss-Smolka translation, which gives easy formulas (240 problems), the Ladner translation, which gives medium-difficulty formulas (240 problems), and the Halpern translation, which gives hard formulas (56 problems) (Massacci & Donini, 2000). As done by Massacci and Donini (2000), we call these classes *easy*, *medium* and *hard* respectively.

All formulas from each class are tested within a timeout of 1000 seconds.[25] For each class, we report the number of solved formulas (X axis) and the total (cumulative) CPU time spent for solving these formulas (Y axes). For each class the results are plotted sorting the solved problems from the easiest one to the hardest one.

5.2 An Empirical Comparison of the Different Variants of K_m2SAT

We have first evaluated the various variants of the encoding in combination with Rsat. In order to avoid considering too many combinations of the flags, we have considered the `BNF` format, and we have grouped `plr` and `bcp` into one parameter `plr-bcp`, restricting thus our investigation to 6 combinations: `BNF`, `lift/ctrl.lift/nolift`, and `plr-bcp` on/off. (We recall that the techniques introduced in Section 4.2, Section 4.5 and Section 4.6 are hardwired in the encoder.) Here we expose and analyze the results w.r.t. the three different suites of benchmark problems.

Results on the LWB Benchmark Suite

The results on the LWB benchmark suite are summarized in Table 1 and Figure 3.

Table 1(a) reports in the left block the indexes of the hardest formulas encoded within the file-size limit and, in the right block, those of the hardest formulas solved within the timeout by Rsat; Table 1(b) reports the numbers of variables and clauses of $K_m2SAT(\varphi)$, referring to the hardest formulas solved within the timeout by Rsat (i.e., those reported in the right block of Table 1(a)). For instance, the `BNF-ctrl.lift-plr-bcp` encoding of `k_dum_n(21)` contains $11{\cdot}10^6$ variables and $14{\cdot}10^6$ clauses; it is the hardest `k_dum_n` problem solved by Rsat with `BNF-ctrl.lift-plr-bcp` and it is the first which is not solved with `BNF-ctrl.lift`.

Looking at the numbers of cases solved in Table 1(a), we notice that the introduction of the on-the-fly Pure Literal Reduction and Boolean Constraint Propagation optimizations is really effective and produces a consistent performance enhancement (the effect of these optimizations is eye-catching in the branching formulas `k_branch_*` – see Section 4.9 – and

25. We also set a 1 GB file-size limit for the encoding produced by K_m2SAT.

	K_m2SAT, encoded						K_m2SAT + Rsat, solved					
				plr-bcp						plr-bcp		
lifting	no	yes	ctrl	no	yes	ctrl	no	yes	ctrl	no	yes	ctrl
k_branch_n	4	4	4	**18**	**18**	**18**	4	4	4	**17**	**17**	**17**
k_branch_p	4	4	4	**18**	**18**	**18**	4	4	4	**18**	**18**	**18**
k_d4_n	8	8	8	8	**9**	8	**8**	**8**	**8**	**8**	**8**	**8**
k_d4_p	**14**	**14**	**14**	**14**	**14**	**14**	**14**	**14**	**14**	**14**	**14**	**14**
k_dum_n	20	20	20	**21**	**21**	**21**	20	20	20	**21**	**21**	**21**
k_dum_p	19	19	19	**21**	**21**	**21**	18	18	18	**21**	**21**	**21**
k_grz_n	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**
k_grz_p	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**
k_lin_n	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**
k_lin_p	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**
k_path_n	7	7	7	14	**15**	14	7	7	7	13	**14**	13
k_path_p	8	8	8	15	**16**	15	8	8	8	15	**16**	15
k_ph_n	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**
k_ph_p	**21**	**21**	**21**	**21**	**21**	**21**	10	**11**	10	10	10	**11**
k_poly_n	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**
k_poly_p	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**	**21**
k_t4p_n	**6**	**6**	**6**	**6**	**6**	**6**	5	**6**	5	**6**	**6**	**6**
k_t4p_p	**11**	**11**	**11**	**11**	**11**	**11**	10	10	10	**11**	**11**	**11**

(a) Indexes of the hardest problems encoded (left) and of the hardest problems solved (right).

	number of variables (10^3)						number of clauses (10^3)					
				plr-bcp						plr-bcp		
lifting	no	yes	ctrl	no	yes	ctrl	no	yes	ctrl	no	yes	ctrl
k_branch_n	1000	1000	1000	20000	20000	20000	1000	1000	1000	23000	23000	23000
k_branch_p	1000	1000	1000	0	0	0	1000	1000	1000	0	0	0
k_d4_n	12000	6000	12000	10000	26000	10000	17000	9000	17000	16000	43000	16000
k_d4_p	19000	18000	19000	0	0	0	28000	25000	28000	0	0	0
k_dum_n	19000	19000	19000	11000	11000	11000	23000	23000	23000	14000	14000	14000
k_dum_p	11000	11000	11000	20000	19000	20000	14000	13000	14000	26000	25000	26000
k_grz_n	10	10	10	5	5	5	10	10	10	6	6	6
k_grz_p	8	8	8	0.2	0.1	0.2	8	8	8	0.3	0.2	0.2
k_lin_n	30	30	20	20	10	20	50	50	20	30	30	30
k_lin_p	0	0	0	0	0	0	0	0	0	0	0	0
k_path_n	11000	12000	11000	10000	7000	10000	13000	14000	13000	14000	9000	13000
k_path_p	11000	12000	11000	26000	16000	26000	13000	14000	13000	35000	20000	35000
k_ph_n	50	300	50	50	300	50	50	300	50	50	600	50
k_ph_p	3	13	3	3	8	4	3	14	3	3	14	5
k_poly_n	200	20	20	200	20	20	200	20	20	200	20	20
k_poly_p	200	20	20	200	20	20	200	20	20	200	20	20
k_t4p_n	4000	21000	4000	17000	14000	17000	4000	22000	4000	20000	17000	20000
k_t4p_p	12000	10000	12000	20000	18000	20000	12000	11000	12000	24000	21000	24000

(b) # of variables and # of clauses of the hardest problems solved.
Note: Here "0" means that the formula is simplified into $\bot$ by K_m2SAT.

Table 1: Comparison of the variants of K_m2SAT +Rsat on the LWB benchmarks.

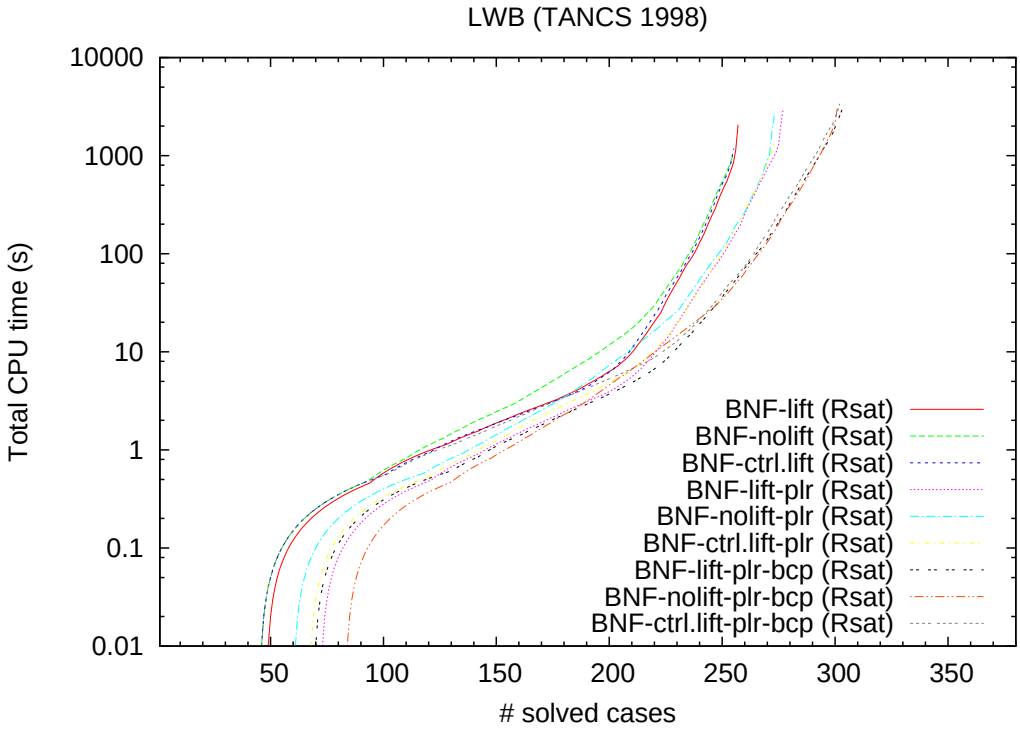

Figure 3: Comparison of different variants of $K_m 2SAT$ +Rsat on the LWB problems. X axis: number of solved problems; Y axis: total CPU time spent (sorting problems from the easiest to the hardest).

in the `k_path_*` formulas). We also notice that `lift` sometimes introduces some slight further improvement.

The view of Tables 1(a) and 1(b) hides the actual CPU times required to encode and solve the problems. Small gaps in the numbers of Table 1(a) may correspond to big gaps in CPU time. In order to analyze also this aspect, in Figure 3 we plotted the total cumulative amount of CPU time spent by all the variants of $K_m 2SAT$ +Rsat to solve all the problems of the LWB benchmark, sorted by hardness. For this plot, we also considered three more options —`BNF`, `lift/ctrl.lift/nolift`, with `plr` on and `bcp` off— so that to evaluate also the effect of `plr` and `bcp` separately. We notice that the plots are clearly clustered into three groups of increasing performance: `BNF-*`, `BNF-*-plr`, and `BNF-*-plr-bcp`., "`*`" representing the three options `lift/ctrl.lift/nolift`. This highlights the fact that on this suite on-the-fly Pure Literal Reduction improves the performances significantly, that on-the-fly Boolean Constraint Propagation introduce drastic improvements, and that the variations due to Box Lifting are minor wrt. the other two optimizations.

Overall, the configuration `BNF-lift-plr-bcp` turns out to be the best performer on this suite, with a tiny advantage wrt. `BNF-ctrl.lift-plr-bcp`.

Results on the Random $\Box_m$-CNF Benchmark Suite

The results on the random $\Box_m$-CNF benchmark suite are reported in Figures 4, 5, 6 and 7.

In Figures 4 and 5 we report the percentage of problems solved within the timeout (1st column) and the 50%- and 90%-percentile CPU times required to encode and solve the formulas by the different K_m2SAT +Rsat variants for all the benchmarks problems (respectively 2nd and 3rd columns).

The tests of Figure 4 are simply too easy for K_m2SAT +Rsat (but not for its competitors, see Section 5.3) which solved every sample formula in less than 1 second. Although the tests of Figure 5 are more challenging, they are all solved within the timeout as well. Notice also that the results are rather regular, as there are no big gaps between 50%- and 90%-percentile values.

In general, we do not have relevant performance gaps between the various configurations on this benchmark suite; it seems that in the majority of cases `ctrl.lift` slightly beats `nolift` and `nolift` slightly beats `lift`. These gaps are even more relevant if we consider the size of the formulas generated (Figures 6 and 7). We believe that this may be due to the fact that random $\Box_m$-CNF formulas may contain lots of shared subformulas $\Box_r\psi$, so that lifting may cause a reduction of such sharing (see Section 3). Further, `plr-bcp` does not seem to introduce relevant improvements here. We believe that this is due to the fact that these random formulas produce pure literals and unit literals with very low or even zero probability.

Overall, the configuration `BNF-nolift` turns out to be the best performer on this suite, with a slight advantage wrt. `BNF-ctrl.lift-plr-bcp`.

Finally, from Figure 5 we notice that for K_m2SAT +Rsat the problems tend to be harder within the satisfiability/unsatisfiability transition area. (This fact holds especially for Racer and *SAT, see Section 5.3.) This seems to confirm the fact that the easy-hard-easy pattern of random k-SAT extends also to $\Box_m$-CNF, as already observed in literature (Giunchiglia & Sebastiani, 1996, 2000; Giunchiglia et al., 2000; Horrocks et al., 2000; Patel-Schneider & Sebastiani, 2003).

Results on the TANCS 2000 Benchmark Suite

The comparison among the K_m2SAT variants on the TANCS 2000 benchmark is presented in Figures 8 and 9, where different `BNF` variants of K_m2SAT are compared both enabling or disabling `lift/ctrl.lif` and `plr-bcp`.

On the left side of Figure 8, from top to bottom, we present the cumulative CPU times spent by K_m2SAT +Rsat on the easy, medium and hard categories respectively, whereas on the right side the corresponding non-cumulative CPU times are reported. In Figure 9 we present the non-cumulative plots of the number of variables and clauses of the formulas solved.

We notice that there are only slight differences among the different variants of K_m2SAT; `BNF` with `lift` is the best option which allows for solving more problems in the hard class and requiring less CPU time.

We remark that, despite the expected exponential growth of the encoded formulas wrt. the modal depth, in this benchmark K_m2SAT +Rsat allows for encoding and solving

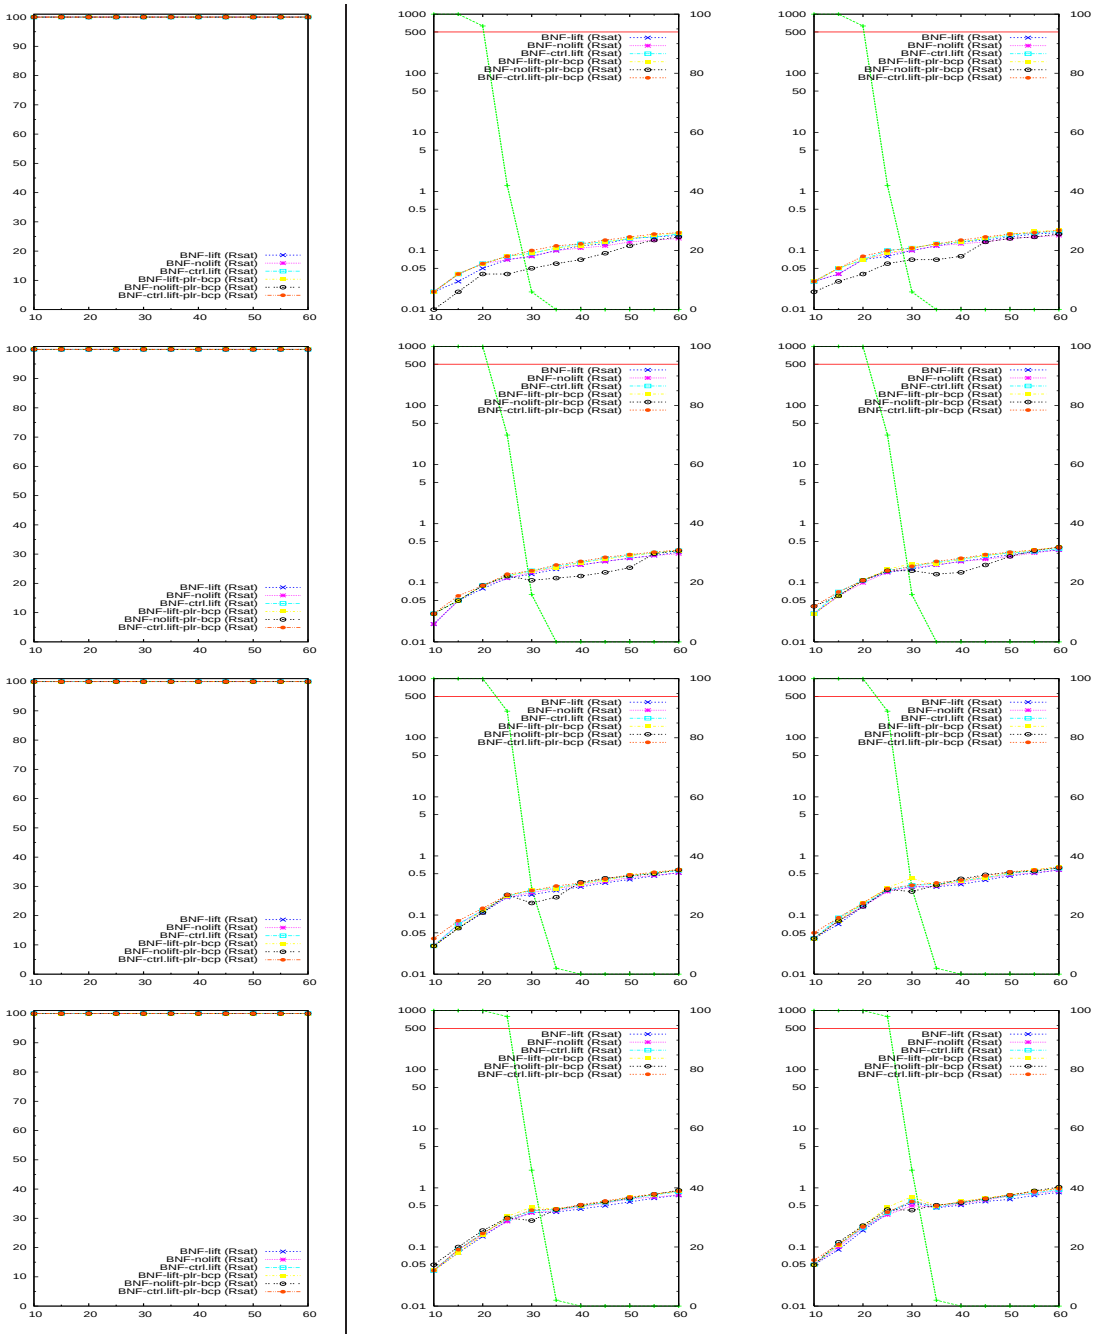

Figure 4: Comparison among different variants of $K_m 2SAT$ +Rsat on random problems, $d = 1$, $p = 0.5$, 100 samples/point. X axis: $\#clauses/N$. Y axis: 1st column: % of problems solved within the timeout; 2nd and 3rd columns: CPU time, 50th and 90th percentiles. 1st to 4th row: $N = 6, 7, 8, 9$. Background: % of satisfiable instances.

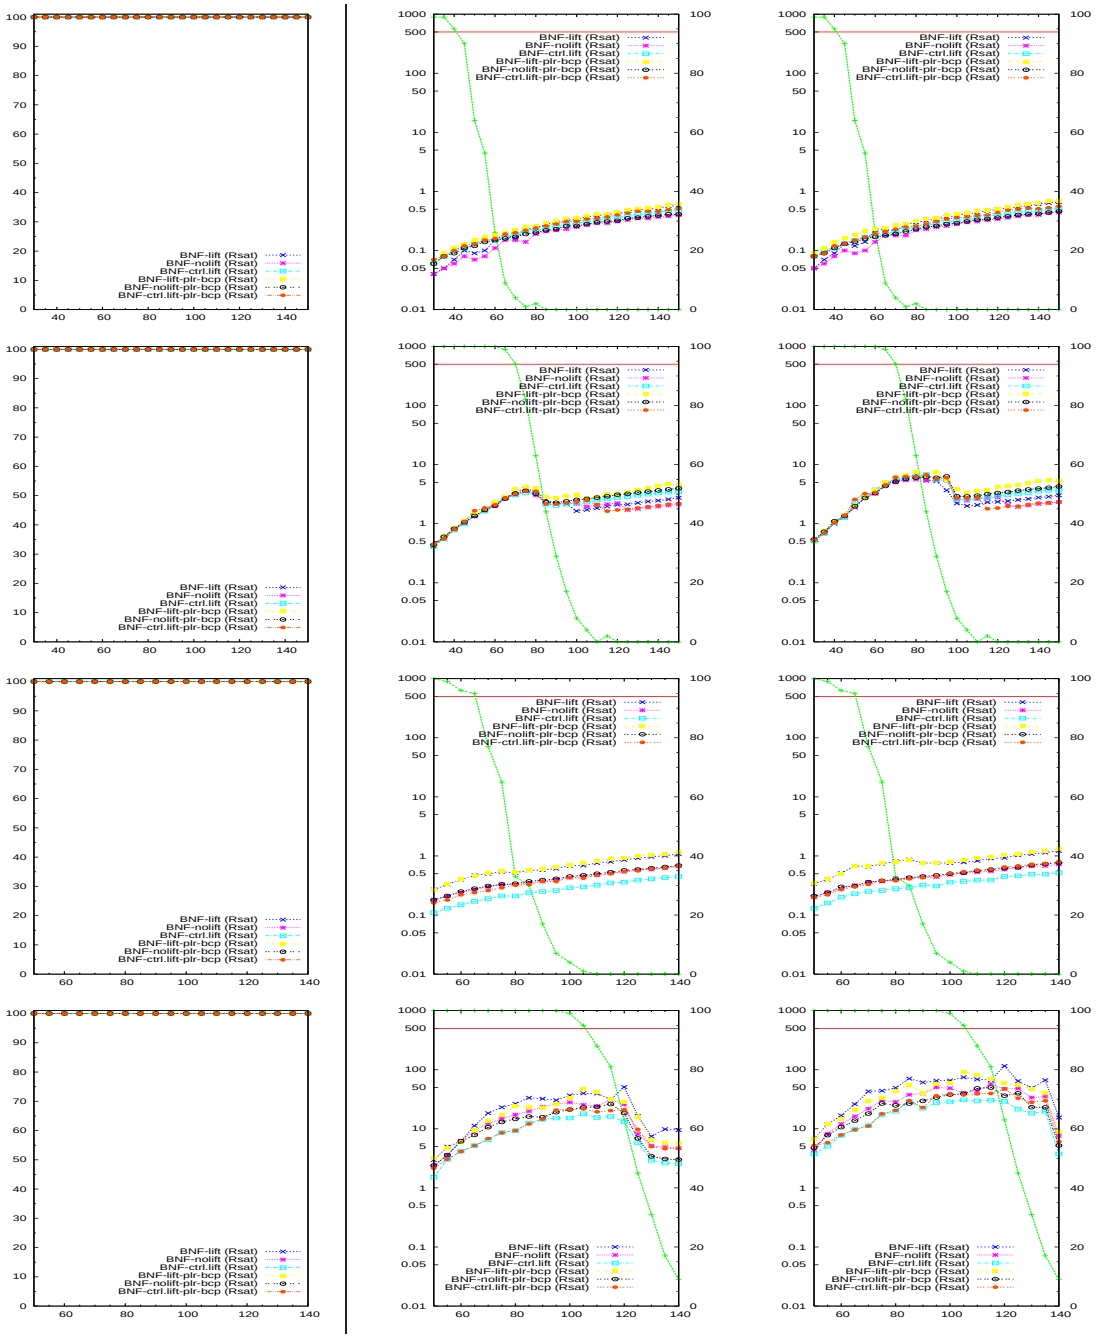

Figure 5: Comparison among different variants of $K_m 2SAT$ +Rsat on random problems, $d = 2$, 100 samples/point. X axis: $\#clauses/N$. Y axis: 1st column: % of problems solved within the timeout; 2nd and 3rd columns: CPU time, 50th and 90th percentiles. 1st and 2nd row: $p = 0.6$, $N = 3, 4$; 3rd and 4th row: $p = 0.5$, $N = 3, 4$. Background: % of satisfiable instances.

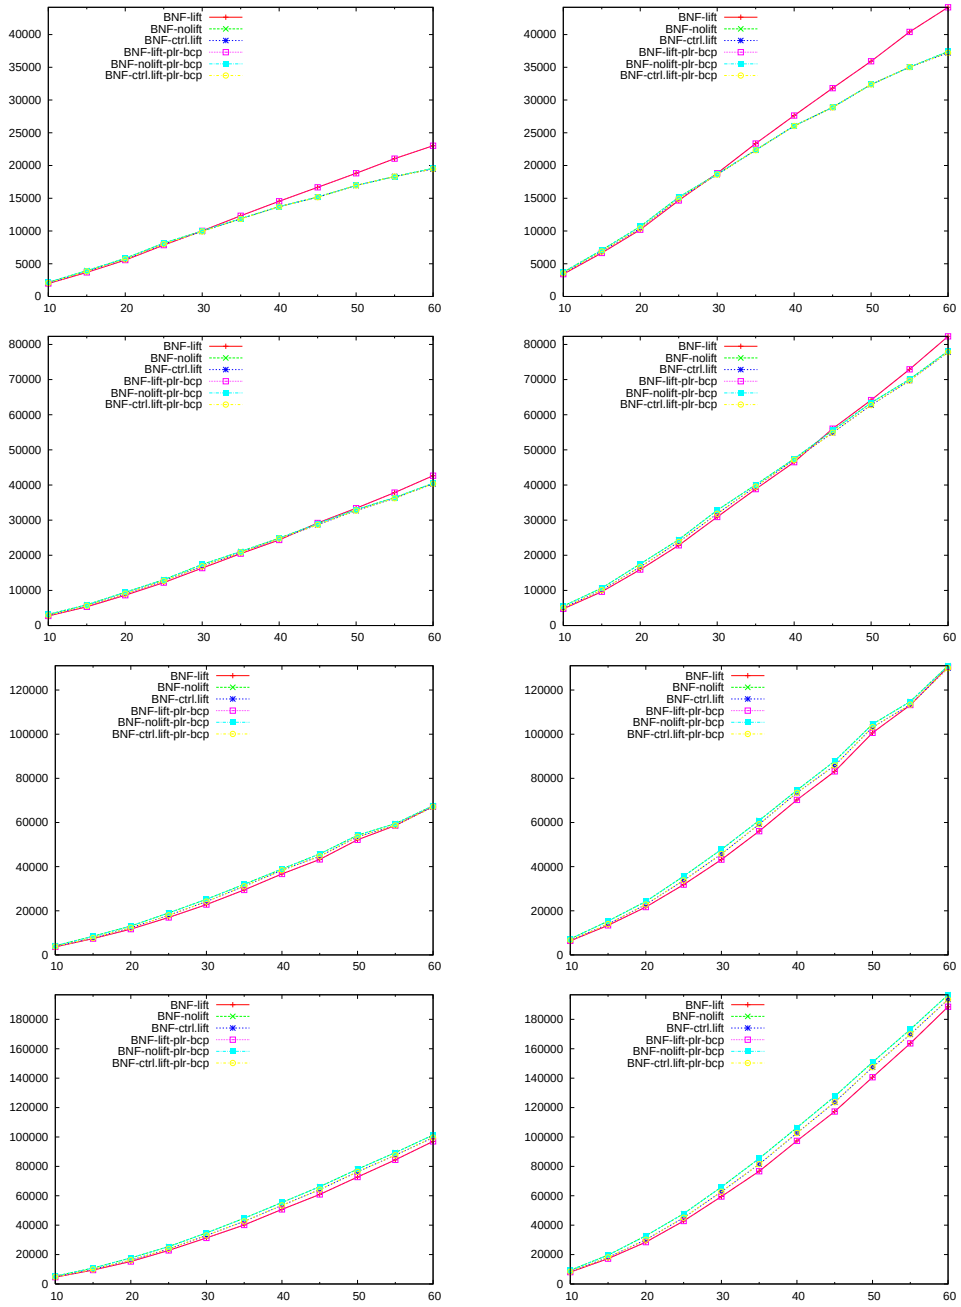

Figure 6: Comparison among different variants of $K_m 2SAT$ on random problems, $d = 1$, $p = 0.5$, 100 samples/point. X axis: $\#clauses/N$. Y axis: 1st column: 90th percentiles of the $\#variables$ in the SAT encoding; 2nd column: 90th percentiles of the $\#clauses$ in the SAT encoding. 1st to 4th row: $N = 6, 7, 8, 9$.

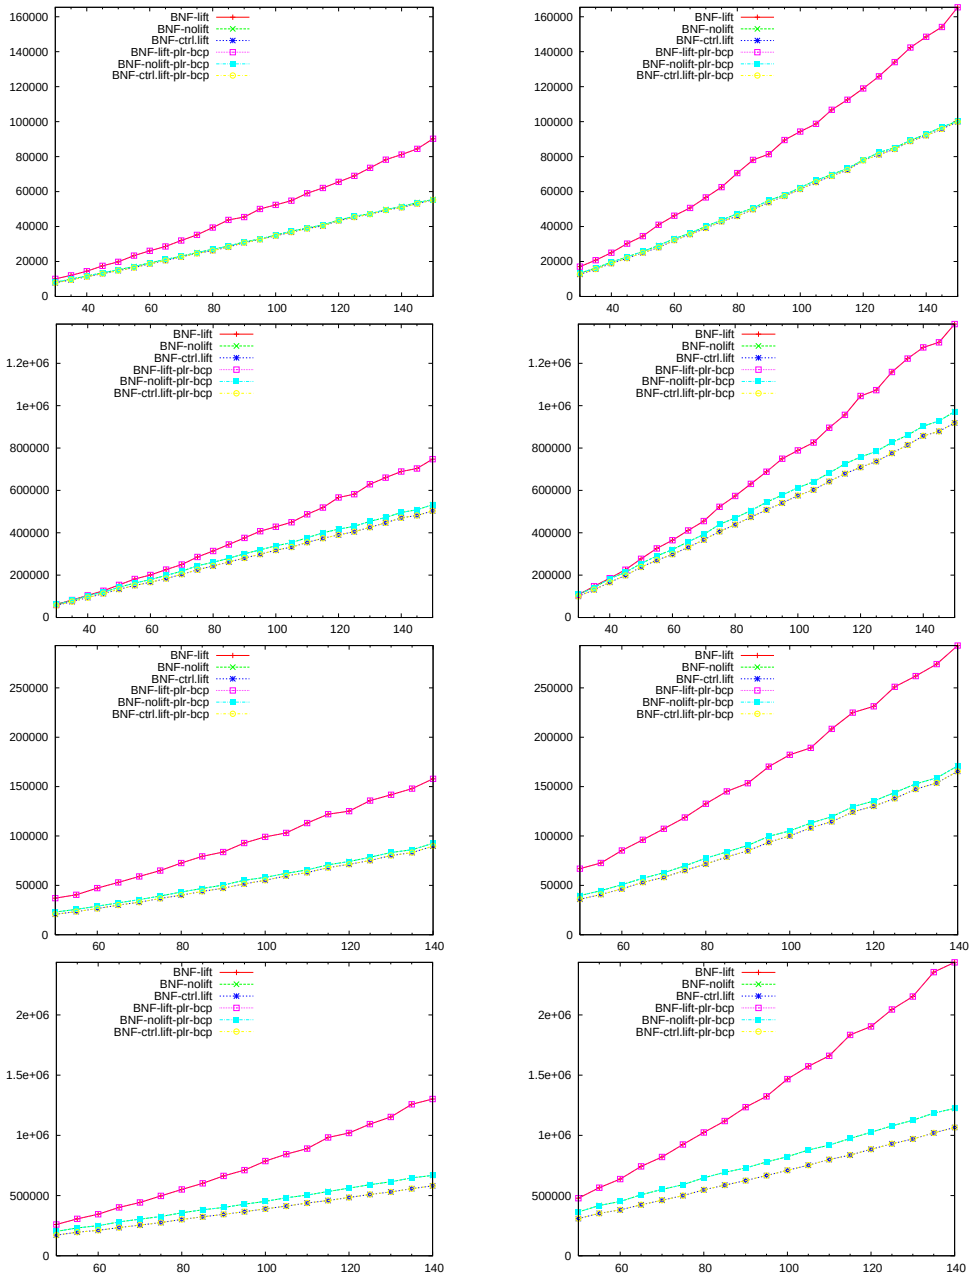

Figure 7: Comparison among different variants of $K_m 2SAT$ on random problems, $d = 2$, 100 samples/point. X axis: $\#clauses/N$. Y axis: 1st column: 90th percentiles of the $\#variables$ in the SAT encoding; 2nd column: 90th percentiles of the $\#clauses$ in the SAT encoding. 1st and 2nd row: $p = 0.6$, $N = 3, 4$; 3rd and 4th row: $p = 0.5$, $N = 3, 4$.

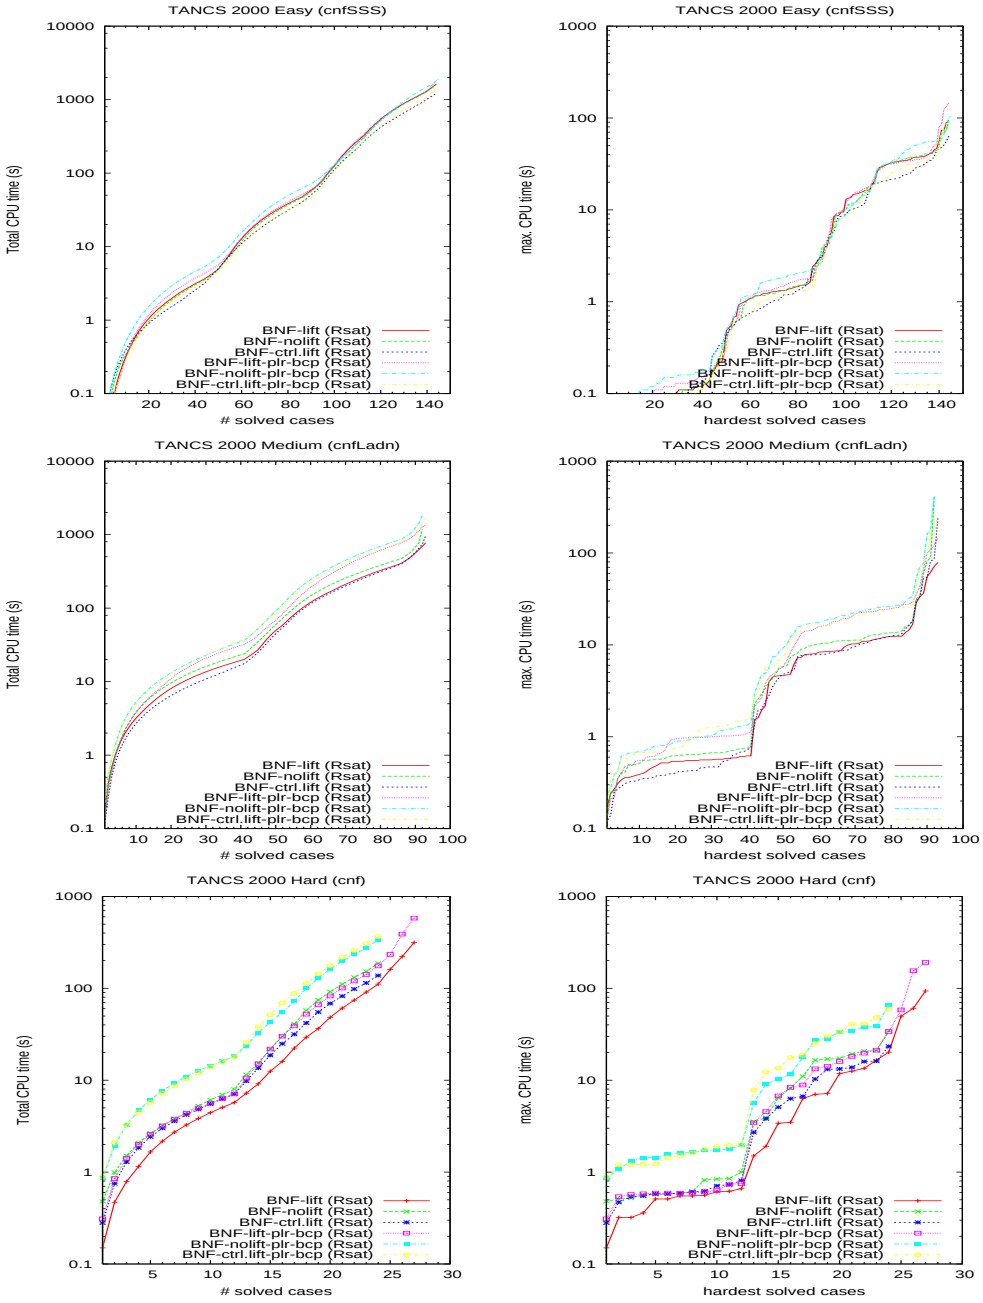

Figure 8: Comparison among different variants of $K_m 2SAT$ +Rsat on TANCS 2000 problems. 1st column: X axis: number of solved problems, Y axis: total CPU time spent; 2nd column: X axis: number of the solved problem, Y axis: CPU time spent to solve the problem; 1st to 3th row: easy, medium, hard problems. (Problems are sorted from the easiest to the hardest).

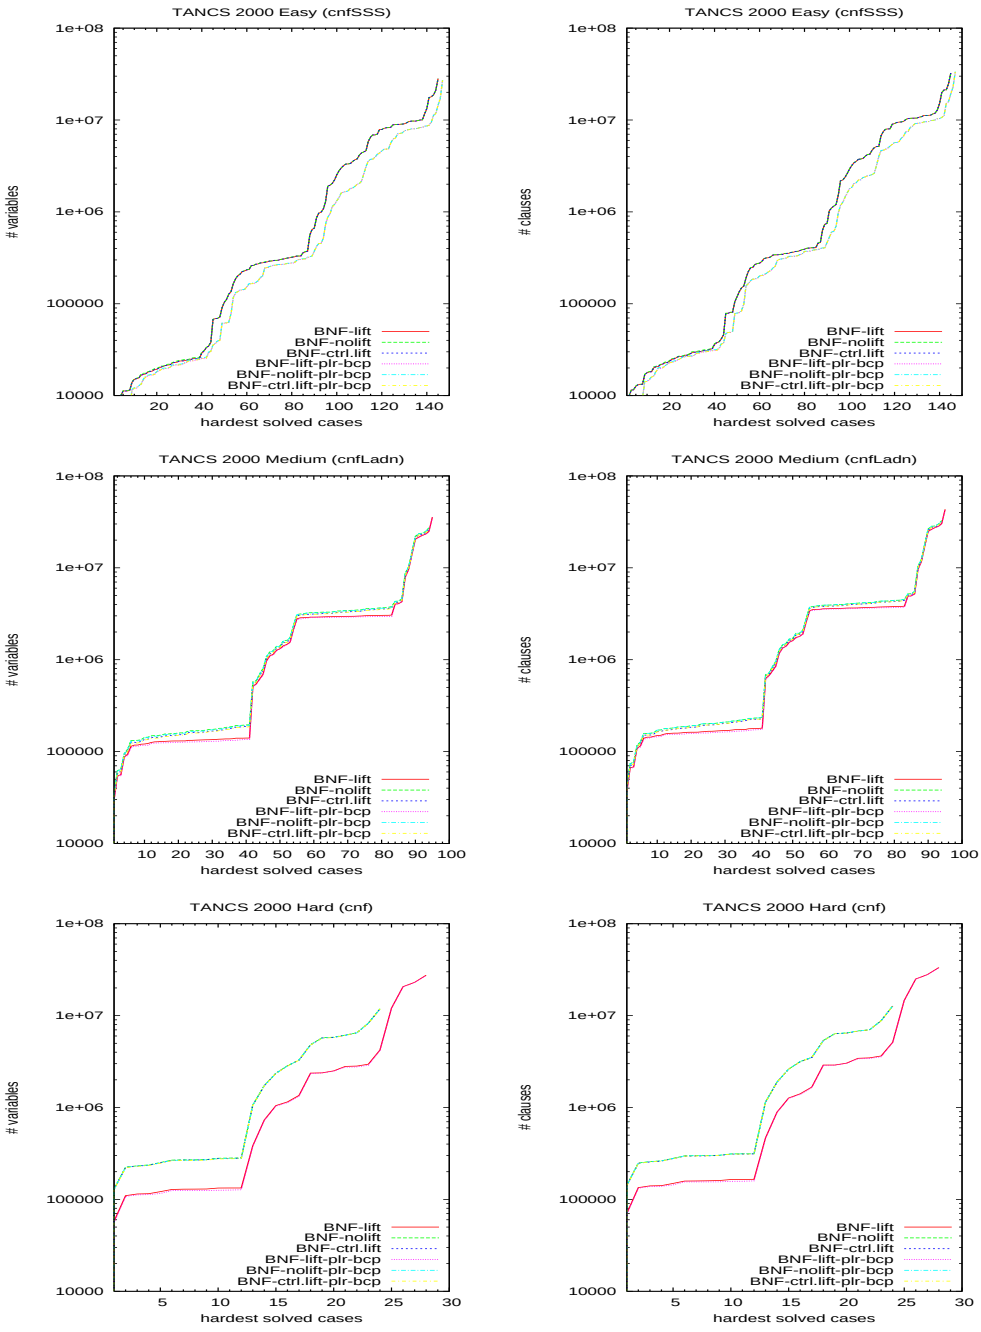

Figure 9: Comparison among different variants of $K_m 2SAT$ on TANCS 2000 problems. X axis: number of the solved problem; Y axis: 1st column: $\#variables$ in the SAT encoding of the problem; 2nd column: $\#clauses$ in the SAT encoding of the problem. 1st to 3th row: easy, medium, hard problems.

problems of modal depth greater than 100 for the easy class and greater than 50 for the medium and hard classes, producing and solving SAT-encoded formulas with more than $10 \cdot 10^6$ variables and $14 \cdot 10^6$ clauses.

5.3 An Empirical Comparison wrt. the Other Approaches

Now we proceed with the comparison of our approach w.r.t. the current state-of-the-art evaluating K_m2SAT +Rsat against the other K_m-satisfiability solvers listed above. In more details, we chose to compare the performance of the other solvers against both the best "local" K_m2SAT +Rsat variant for the single benchmark suite and the best "global" K_m2SAT +Rsat variant among all the benchmark suites, which we have identified in `BNF-ctrl.lift-plr-bcp`.

Comparison on the LWB Benchmark Suite

The results on the LWB benchmark suite are summarized numerically and graphically in Table 2.

From Table 2(a) we notice a few facts: Racer and *SAT are the best performers (confirming the analysis done by Horrocks et al., 2000) with a significant gap wrt. the others; then, K-QBF +sKizzo solves a few more problems than K_m2SAT +Rsat; then follows KBDD which outperforms Mspass, which is the worst performer. In detail, K_m2SAT +Rsat is (one of) the worst performer(s) on `k_d4_*` and `k_t4_*`, the fourth best performer on `k_path_n`, the third best performer on `k_path_p` and `k_branch_p`, and it is (one of) the best performer(s) on `k_branch_n`, `k_dum_*`, `k_grz_*`, `k_lin_*`, `k_ph_*` and `k_poly_*`; it is the absolute best performer on `k_branch_n` and `k_ph_p`.

In Table 2(b) we give a graphical representation of this comparison, plotting the number of solved problems by each approach against the total cumulative amount of CPU time spent. We notice that, even if K_m2SAT +Rsat solves a few problems less than K-QBF +sKizzo, K_m2SAT +Rsat is mostly faster than K-QBF +sKizzo.

Comparison on the Random $\Box_m$-CNF Benchmark suite

In the random $\Box_m$-CNF benchmark suite the results are dominated by K_m2SAT +Rsat.

Looking at the left columns of Figures 10 and 11 (# of problems solved), we notice that K_m2SAT +Rsat (in both versions) is the only tool which always terminates within the timeout, whilst *SAT and Racer sometimes do not terminate in the hardest problems, K-QBF +sKizzo very often does not terminate, and Mspass and KBDD do not terminate for most values.

In Figures 10 and 11 (center and right) we notice that K_m2SAT +Rsat is most often the best performer (in particular with the hardest problems), followed by *SAT and Racer. In all these tests, K-QBF +sKizzo, Mspass and KBDD are drastically outperformed by the others.

Comparison on the TANCS 2000 Benchmark Suite

The results of the TANCS 2000 benchmark are summarized in Figure 12, from the easy category (up) to the hard category (bottom).

test	other tools					K_m2SAT + Rsat	
	K-QBF + sKizzo	KBDD	Mspass	Racer	*SAT	BNF-plr-bcp -ctrl.lift	BNF-plr-bcp -lift
k_branch_n	4	8	10	15	14	**17**	**17**
k_branch_p	16	8	10	**21**	**21**	18	18
k_d4_n	8	**21**	**21**	**21**	**21**	8	8
k_d4_p	**21**	**21**	**21**	**21**	**21**	14	14
k_dum_n	**21**	**21**	**21**	**21**	**21**	**21**	**21**
k_dum_p	**21**	**21**	**21**	**21**	**21**	**21**	**21**
k_grz_n	19	**21**	**21**	**21**	**21**	**21**	**21**
k_grz_p	**21**	**21**	**21**	**21**	**21**	**21**	**21**
k_lin_n	20	**21**	**21**	**21**	**21**	**21**	**21**
k_lin_p	**21**	**21**	3	**21**	**21**	**21**	**21**
k_path_n	9	**21**	4	**21**	**21**	13	14
k_path_p	13	17	5	**21**	**21**	15	16
k_ph_n	**21**	4	12	**21**	13	**21**	**21**
k_ph_p	10	4	8	9	9	**11**	10
k_poly_n	**21**	8	7	**21**	**21**	**21**	**21**
k_poly_p	**21**	8	7	**21**	**21**	**21**	**21**
k_t4p_n	**21**	**21**	12	**21**	**21**	6	6
k_t4p_p	**21**	**21**	**21**	**21**	**21**	11	11

(a) Indexes of the hardest problems solved.

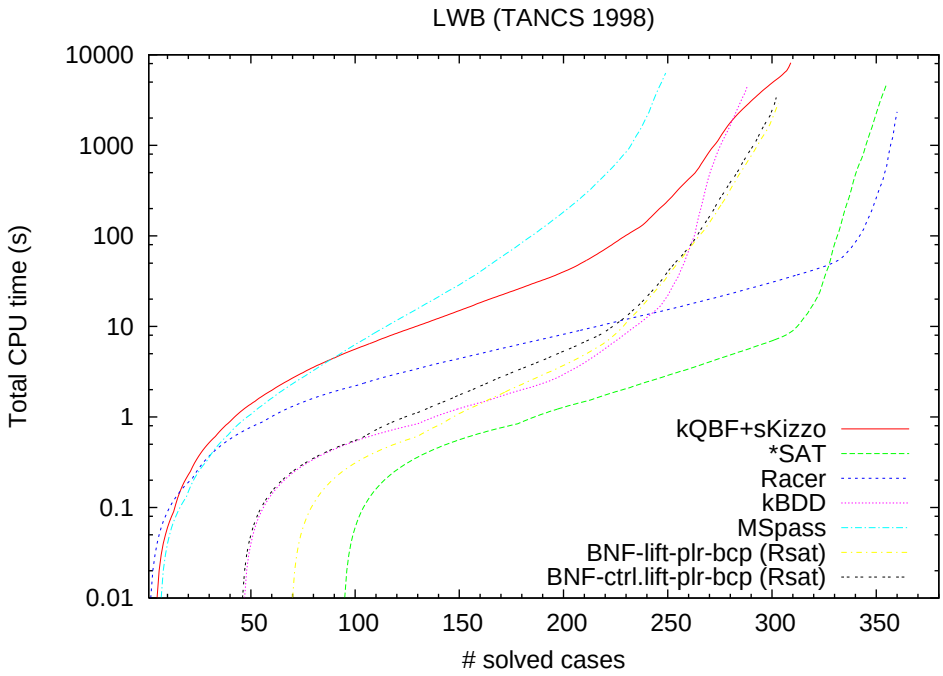

(b) X axis: # of problems solved; Y axis: total (cumulative) CPU time spent.

Table 2: Comparison of K_m2SAT +Rsat against the other tools on the LWB benchmarks.

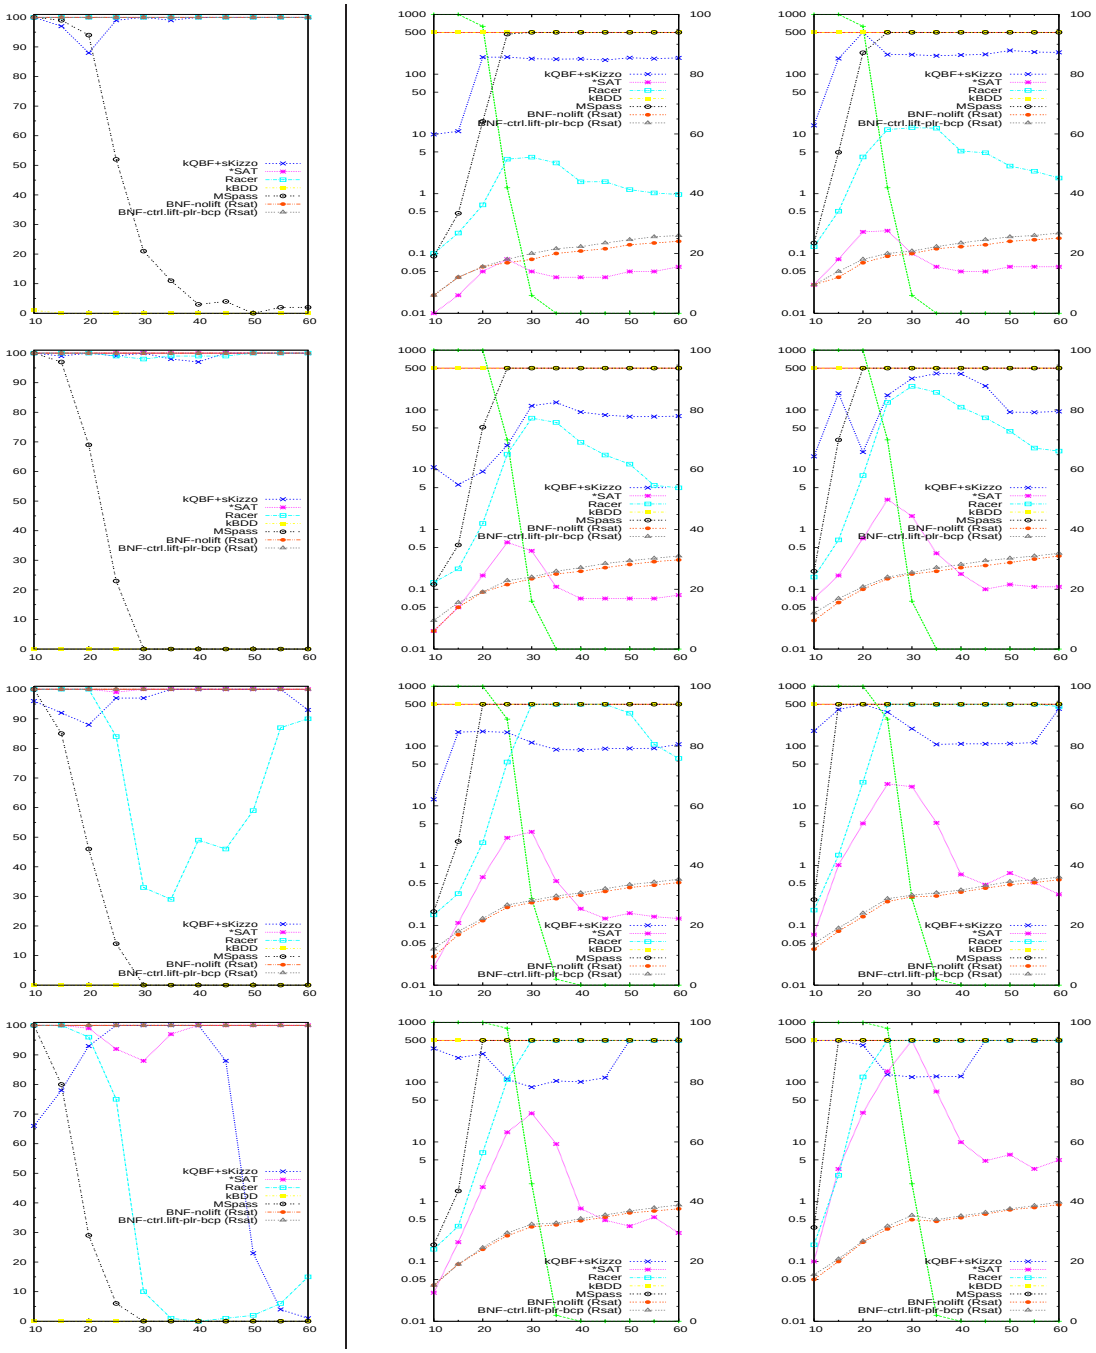

Figure 10: Comparison against other approaches on random problems, $d = 1$, $p = 0.5$, 100 samples/point. X axis: $\#clauses/N$. Y axis: 1st column: % of problems solved within the timeout; 2nd and 3rd columns: CPU time, 50th and 90th percentiles. 1st to 4th row: $N = 6, 7, 8, 9$. Background: % of satisfiable instances.

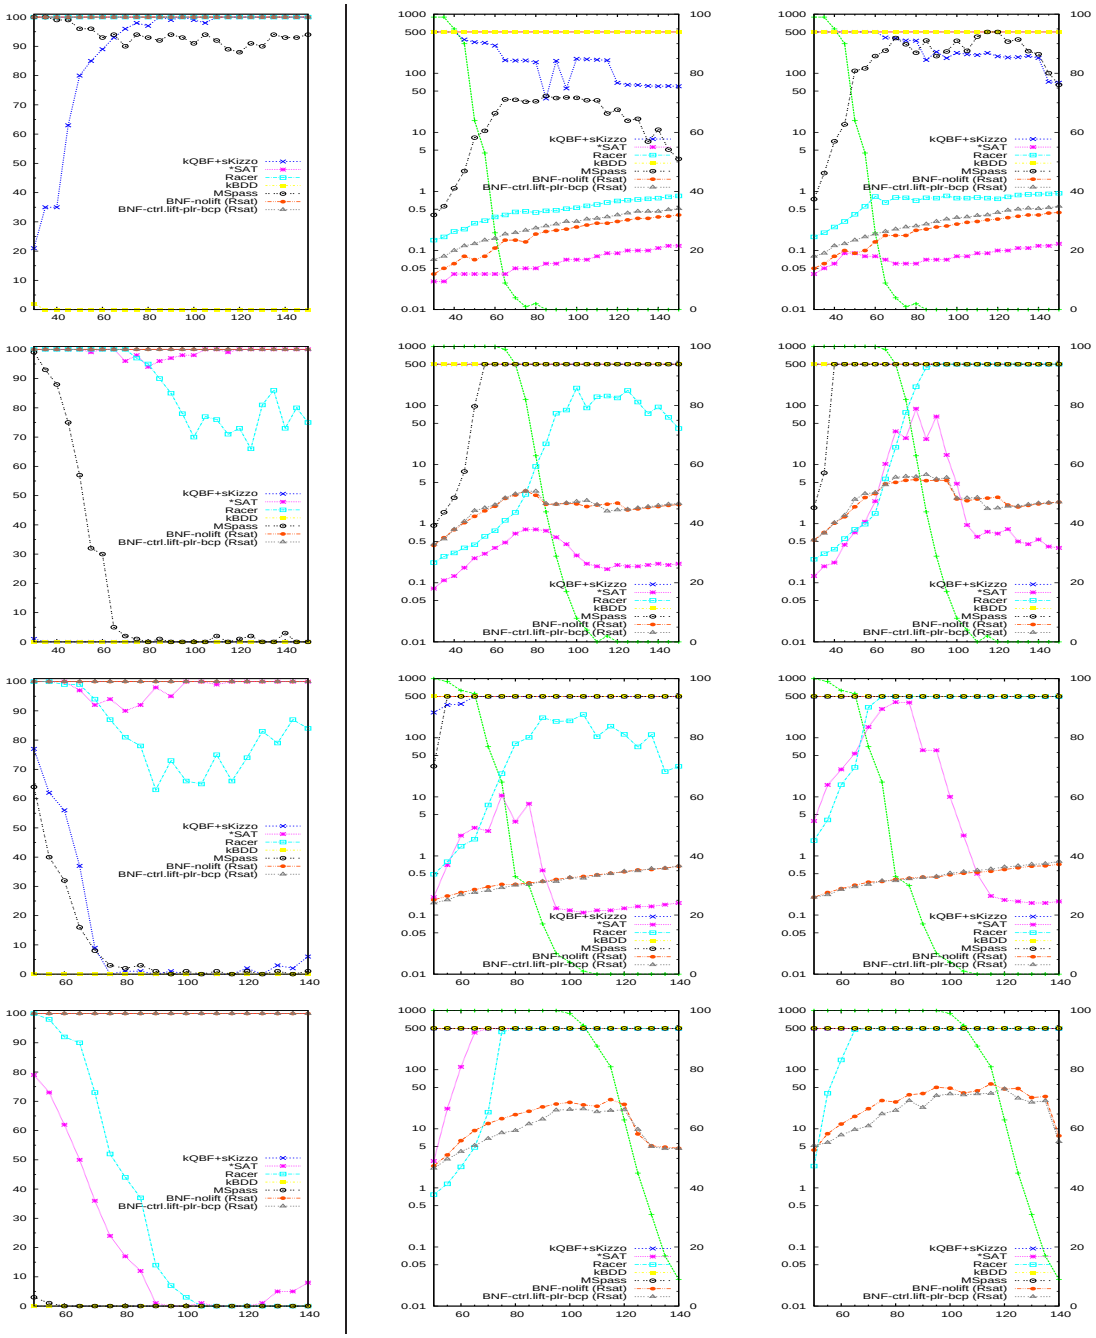

Figure 11: Comparison against other approaches on random problems, $d = 2$, 100 samples/point. X axis: $\#clauses/N$. Y axis: 1st column: % of problems solved within the timeout; 2nd and 3rd columns: CPU time, 50th and 90th percentiles. 1st and 2nd row: $p = 0.6$, $N = 3, 4$. 3rd and 4th row: $p = 0.5$, $N = 3, 4$; Background: % of satisfiable instances.

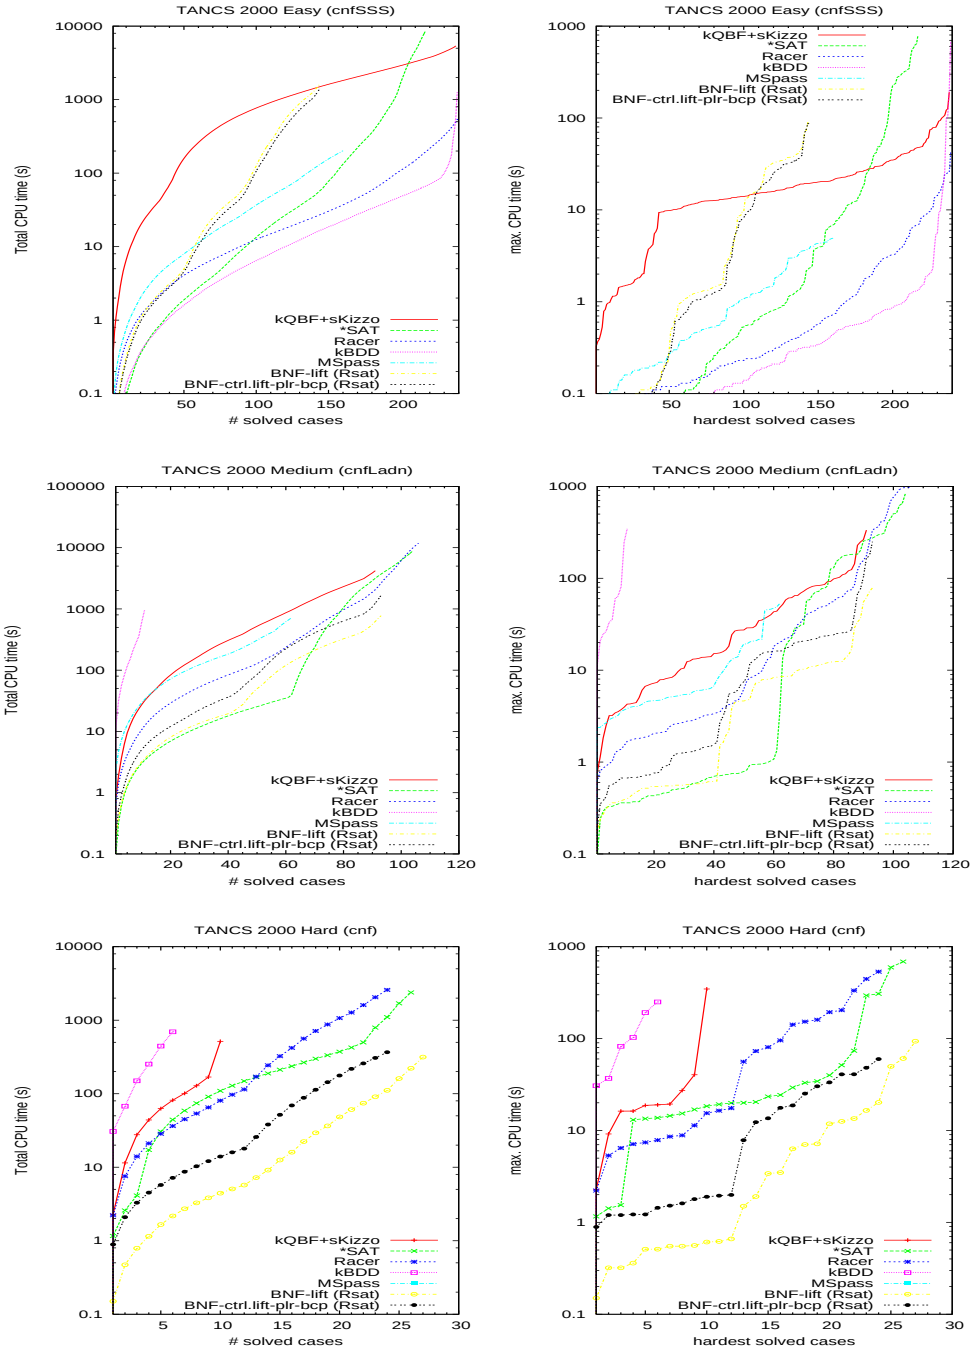

Figure 12: Comparison against other approaches on TANCS 2000 problems. X axis: number of problems solved; Y axis: cumulative (left) and non-cumulative (right) CPU times.

From the plots we notice that the relative performances of the tools under test vary significantly with the category: Racer and *SAT are among the best performers in all categories; K-QBF +sKizzo behaves well on the easy and medium categories but solvers very few problems on the hard one; KBDD behaves very well on the easy category, but solves very few problems in the medium and hard ones. Mspass is among the worst performers in all categories: in particular, it does not solve any problem in the hard category; finally, K_m2SAT +Rsat is the worst performer on the easy category, it outperforms all competitors but *SAT and Racer on the medium category, and competes head-to-head with both Racer and *SAT for the first position on the hard category (the "local-best" configuration `BNF-lift` beats both competitors; the "global-best" configuration `BNF-ctrl.lift-prl-bcp` solves as many problems as Racer, with one-order-magnitude CPU-time performance gap, and two problems less than *SAT).

5.4 Discussion

As highlighted by Giunchiglia et al. (2000) and Horrocks et al. (2000), the satisfiability problem in modal logics like K_m is characterized by the alternation of two orthogonal components of reasoning: a *Boolean* component, performing Boolean reasoning within each state, and a *modal* component, generating the successor states of each state. The latter must cope with the fact that the candidate models may be up to exponentially big wrt. $depth(\varphi)$, whilst the former must cope with the fact that there may be up to exponentially many candidate (sub)models to explore. In the K_m2SAT +DPLL approach the encoder has to handle the whole modal component (rules (8) and (9)), whilst the handling of the whole Boolean component is delegated to an external SAT solver.

From the results displayed in Section 5.3 we notice that the relative performances of the K_m2SAT +DPLL approach wrt. other state-of-the-art tools range from cases where K_m2SAT +Rsat is much less efficient than other state-of-the-art approaches (e.g., the `k_d4` and `k_t4p` formulas) up to formulas where it is much more efficient (e.g., the `k_ph_p` or the $\Box_m$-CNF formulas with $d = 1$). In the middle stands a large majority of formulas in which the approach competes well against the other state-of-the art tools; in particular, K_m2SAT +Rsat competes very well or even outperforms the other approaches based on translations into different formalisms (the translational approach, the automata-theoretic approach and the QBF-encoding approach).

A simple explanation of the former fact could be that the K_m2SAT +DPLL approach loses on problems with high modal depth, or where the modal component of reasoning dominates (like, e.g., the `k_d4` and `k_t4p` formulas), and wins on problems where the Boolean component of reasoning dominates (like, e.g., the `k_ph_n` or the $\Box_m$-CNF formulas with $d = 1$), and it is competitive for formulas in which both components are relevant.

We notice, however, that K_m2SAT +Rsat wins in the hard category of TANCS 2000 benchmarks, with modal depths greater than 50, and on `k_branch_n`, where the search is dominated by the modal component.[26] In fact, we recall that reducing the Boolean component of reasoning may produce a reduction also of the modal reasoning effort, because

26. The `k_branch_n` formulas are very hard from the perspective of modal reasoning, because they require finding one model $\mathcal{M}$ with $2^{d+1}-1$ states (Halpern & Moses, 1992), but no Boolean reasoning within each state is really required (Giunchiglia et al., 2000; Horrocks et al., 2000): e.g., *SAT solves `k_branch_n(d)` with $2^{d+1}-1$ calls to its embedded DPLL engine, one for each state of $\mathcal{M}$, each call solved by BCP only.

it may reduce the number of successor states to analyze (e.g. Sebastiani, 2007a, 2007b). Thus, e.g., techniques like on-the-fly BCP, although exploiting only purely-Boolean properties, may produce not only a drastic pruning of the Boolean search, but also a drastic reduction in the size of the model investigated, because they cut a priori the amount of successor states to expand.

6. Related Work and Research Trends

In the last fifteen years one main research line in description logic has focused on investigating increasingly expressive logics, with the goal of establishing the theoretical boundaries of decidability and of allowing for more expressive power in the languages defined (Baader, Calvanese, McGuinness, Nardi, & Patel-Schneider, 2003). Consequently, very expressive — though very hard— description logics have today notable application in the field of Semantic Web. For example, the $\mathcal{SHOIN}^{(\mathcal{D})}$ logic (which has NExpTime complexity) captures the sub-language OWL DL of the full OWL language (Bechhofer, van Harmelen, Hendler, Horrocks, McGuinness, Patel-Schneider, & Stein, 2004).

In contrast, the recent quest for tractable description logic-based languages arising from the field of bio-medical ontologies (e.g., Spackman, Campbell, & Cote, 1997; Sioutos, de Coronado, Haber, Hartel, Shaiu, & Wright, 2007; The Gene Ontology Consortium, 2000; Rector & Horrocks, 1997) has stimulated the opening of another research line on tractable description logics (also called *lightweight description logics*), which are suitable for reasoning on these very big bio-medical ontologies. In particular, Baader et al. (2005, 2006, 2007, 2008) have spent a considerable effort in the attempt of defining a small but maximal subset of logical constructors, expressive enough to cover the needs of these practical applications, but whose inference problems must remain tractable. For example, simple and tractable description logics like $\mathcal{EL}$, $\mathcal{EL}^+$ and $\mathcal{EL}^{++}$ (Baader et al., 2005) are expressive enough to describe several important bio-medical ontologies such as SNoMed (Spackman et al., 1997), NCI (Sioutos et al., 2007), the Gene Ontology (The Gene Ontology Consortium, 2000) and the majority of Galen (Rector & Horrocks, 1997).

Reasoning on these ontologies represents not only an important application of lightweight description logics, but also a challenge due to the required efficiency and the huge dimensions of the ontologies. In this perspective, it is important to face efficiently not only the standard reasoning services (e.g., satisfiability, subsumption, queries) on logics like $\mathcal{EL}$, $\mathcal{EL}^+$ and $\mathcal{EL}^{++}$, but also non-standard reasoning problems like e.g., *axiom pinpointing* (Baader et al., 2007; Baader & Penaloza, 2008) and *logical difference* between terminologies (Konev, Walther, & Wolter, 2008).

7. Conclusions and Future Work

In this paper we have explored the idea of encoding $K_m/\mathcal{ALC}$-satisfiability into SAT, so that to be handled by state-of-the-art SAT tools. We have showed that, despite the intrinsic risk of blowup in the size of the encoded formulas, the performances of this approach are comparable with those of current state-of-the-art tools on a rather extensive variety of empirical tests.

We see many possible directions to explore in order to enhance and extend our approach. An important open research line is to explore the feasibility of SAT encodings for other and more expressive modal and description logics (e.g., whilst for logics like T_m the extension should be straightforward, logics like $S4_m$, or more elaborated description logics than $\mathcal{ALC}$, should be challenging) and for more complex form of reasoning (e.g., including TBoxes and ABoxes).

Our current investigation, however, focuses on the lightweight logics of Baader et al. (2005). We are currently investigating an encoding of the main reasoning services in $\mathcal{EL}$, $\mathcal{EL}^+$ and $\mathcal{EL}^{++}$ into (Horn) SAT, which should allow for reasoning efficiently on the (often huge) bio-medical ontologies mentioned in Section 6, and also for handling the non-standard inference problems mentioned there (e.g., axiom pinpointing), by exploiting some of the advanced functionalities which can be implemented on top of a modern SAT solver.

References

Areces, C., Gennari, R., Heguiabehere, J., & de Rijke, M. (2000). Tree-Based Heuristics in Modal Theorem Proving. In *Proc. ECAI'00*.

Baader, F., Brandt, S., & Lutz, C. (2005). Pushing the $\mathcal{EL}$ envelope. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence IJCAI-05*, Edinburgh, UK. Morgan-Kaufmann Publishers.

Baader, F., Franconi, E., Hollunder, B., Nebel, B., & Profitlich, H. (1994). An Empirical Analysis of Optimization Techniques for Terminological Representation Systems or: Making KRIS get a move on. *Applied Artificial Intelligence. Special Issue on Knowledge Base Management*, *4*, 109–132.

Baader, F., & Hollunder, B. (1991). A Terminological Knowledge Representation System with Complete Inference Algorithms. In *Proceedings of the First International Workshop on Processing Declarative Knowledge*, Vol. 572 of *Lecture Notes in Computer Science*, pp. 67–85, Kaiserslautern (Germany). Springer–Verlag.

Baader, F., Lutz, C., & Suntisrivaraporn, B. (2006). CEL—a polynomial-time reasoner for life science ontologies. In Furbach, U., & Shankar, N. (Eds.), *Proceedings of the 3rd International Joint Conference on Automated Reasoning (IJCAR'06)*, Vol. 4130 of *LNAI*, pp. 287–291. Springer-Verlag.

Baader, F., Brandt, S., & Lutz, C. (2008). Pushing the $\mathcal{EL}$ envelope further. In Clark, K., & Patel-Schneider, P. F. (Eds.), *In Proceedings of the OWLED 2008 DC Workshop on OWL: Experiences and Directions*.

Baader, F., Calvanese, D., McGuinness, D., Nardi, D., & Patel-Schneider, P. F. (Eds.). (2003). *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press.

Baader, F., & Penaloza, R. (2008). Automata-Based Axiom Pinpointing. In *Proc. IJCAR*, Vol. 5195 of *LNAI*. Springer.

Baader, F., Peñaloza, R., & Suntisrivaraporn, B. (2007). Pinpointing in the description logic $\mathcal{EL}^+$. In *Proceedings of the 30th German Conference on Artificial Intelligence (KI2007)*, LNAI, Osnabrück, Germany. Springer.

Balsiger, P., Heuerding, A., & Schwendimann, S. (1998). Logics workbench 1.0. In *Proc. TABLEAUX*, Vol. 1397 of *LNCS*, pp. 35–37. Springer.

Bechhofer, S., van Harmelen, F., Hendler, J., Horrocks, I., McGuinness, D. L., Patel-Schneider, P. F., & Stein, L. A. (2004). OWL Web Ontology Language reference. W3C Recommendation. Available at `http://www.w3.org/TR/owl-ref/`.

Benedetti, M. (2005). sKizzo: A Suite to Evaluate and Certify QBFs. In *CADE*, Vol. 3632 of *LNCS*, pp. 369–376. Springer.

Biere, A., Cimatti, A., Clarke, E. M., & Zhu, Y. (1999). Symbolic Model Checking without BDDs. In *Proc. TACAS'99*, pp. 193–207.

Brand, S. (2008). Personal communication..

Brand, S., Gennari, R., & de Rijke, M. (2003). Constraint Programming for Modelling and Solving Modal Satisfability. In *Proc. CP 2003*, Vol. 3010 of *LNAI*. Springer.

Davis, M., Longemann, G., & Loveland, D. (1962). A machine program for theorem-proving. *Journal of the ACM*, *5*(7), 394–397.

Davis, M., & Putnam, H. (1960). A computing procedure for quantification theory. *Journal of the ACM*, *7*, 201–215.

Donini, F., & Massacci, F. (2000). EXPTIME tableaux for ALC. *Artificial Intelligence*, *124*(1), 87–138.

Eén, N., & Sörensson, N. (2004). An Extensible SAT-solver. In *Theory and Applications of Satisfiability Testing (SAT 2003)*, Vol. 2919 of *LNCS*, pp. 502–518. Springer.

Fitting, M. (1983). *Proof Methods for Modal and Intuitionistic Logics*. D. Reidel Publishing.

Franconi, E. (1998). CRACK. In *Description Logics*, Vol. 11 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Giunchiglia, E., Giunchiglia, F., Sebastiani, R., & Tacchella, A. (2000). SAT vs. Translation based decision procedures for modal logics: a comparative evaluation. *Journal of Applied Non-Classical Logics*, *10*(2).

Giunchiglia, E., Giunchiglia, F., & Tacchella, A. (2001). SAT Based Decision Procedures for Classical Modal Logics.. Journal of Automated Reasoning. Special Issue: Satisfiability at the start of the year 2000.

Giunchiglia, F., Roveri, M., & Sebastiani, R. (1996). A new method for testing decision procedures in modal and terminological logics. In *Proc. of 1996 International Workshop on Description Logics - DL'96*, Cambridge, MA, USA.

Giunchiglia, F., & Sebastiani, R. (1996). Building decision procedures for modal logics from propositional decision procedures - the case study of modal K. In *Proc. CADE'13*. Springer.

Giunchiglia, F., & Sebastiani, R. (2000). Building decision procedures for modal logics from propositional decision procedures - the case study of modal K(m). *Information and Computation*, *162*(1/2).

Halpern, J. Y. (1995). The effect of bounding the number of primitive propositions and the depth of nesting on the complexity of modal logic. *Artificial Intelligence*, *75*(3), 361–372.

Halpern, J., & Moses, Y. (1992). A guide to the completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, *54*(3), 319–379.

Heuerding, A., & Schwendimann, S. (1996). A benchmark method for the propositional modal logics K, KT, S4. Tech. rep. IAM-96-015, University of Bern, Switzerland.

Horrocks, I. (1998). Using an expressive description logic: FaCT or fiction?. In *Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*, pp. 636–647.

Horrocks, I., & Patel-Schneider, P. F. (1999). Optimizing Description Logic Subsumption. *Journal of Logic and Computation*, *9*(3), 267–293.

Horrocks, I., Patel-Schneider, P. F., & Sebastiani, R. (2000). An Analysis of Empirical Testing for Modal Decision Procedures. *Logic Journal of the IGPL*, *8*(3).

Hustadt, U., Schmidt, R. A., & Weidenbach, C. (1999). MSPASS: Subsumption Testing with SPASS. In *Proc. DL'99*, pp. 136–137.

Hustadt, U., & Schmidt, R. (1999). An empirical analysis of modal theorem provers. *Journal of Applied Non-Classical Logics*, *9*(4).

Kautz, H., McAllester, D., & Selman, B. (1996). Encoding Plans in Propositional Logic. In *Proceedings KR'96*. AAAI Press.

Konev, B., Walther, D., & Wolter, F. (2008). The Logical Difference Problem for Description Logic Terminologies. In *Proc. IJCAR*, Vol. 5195 of *LNAI*. Springer.

Ladner, R. (1977). The computational complexity of provability in systems of modal propositional logic. *SIAM J. Comp.*, *6*(3), 467–480.

Massacci, F. (1999). Design and Results of Tableaux-99 Non-Classical (Modal) System Competition. In *Automated Reasoning with Analytic Tableaux and Related Methods: International Conference (Tableaux'99)*.

Massacci, F. (2000). Single Step Tableaux for modal logics: methodology, computations, algorithms. *Journal of Automated Reasoning, Vol. 24*(3).

Massacci, F., & Donini, F. (2000). Design and results of TANCS-2000, Automated Reasoning with Analytic Tableaux and Related Methods. In *Int. Conf. TABLEAUX 2000*, pp. 52–56.

Moskewicz, M. W., Madigan, C. F., Zhao, Y., Zhang, L., & Malik, S. (2001). Chaff: Engineering an efficient SAT solver. In *Design Automation Conference*.

Narizzano, M., Pulina, L., & Tacchella, A. (2006). The QBFEVAL Web Portal. In *Proc. Logics in Artificial Intelligence, 10th European Conference (JELIA'06)*, Vol. 4160 of *LNCS*, pp. 494–497. Springer. See: http://www.qbflib.org/index_eval.php.

Pan, G., Sattler, U., & Vardi, M. Y. (2002). BDD-Based Decision Procedures for K. In *Proc. CADE*, LNAI. Springer.

Pan, G., & Vardi, M. Y. (2003). Optimizing a BDD-based modal solver. In *Proc. CADE*, LNAI. Springer.

Patel-Schneider, P. F. (1998). DLP system description. In *Proc. Tableaux'98*, pp. 87–89.

Patel-Schneider, P. F., & Sebastiani, R. (2001). A system and methodology for generating random modal formulae. In *Proc. of International Joint Conference on Automated reasoning - IJCAR-2001*, Vol. 2083 of *LNAI*, Siena, Italy. Springer-verlag.

Patel-Schneider, P. F., & Sebastiani, R. (2003). A New General Method to Generate Random Modal Formulae for Testing Decision Procedures. *Journal of Artificial Intelligence Research, (JAIR), 18*. Morgan Kaufmann.

Pipatsrisawat, T., & Darwiche, A. (2006). SAT Solver Description: Rsat.. http://fmv.jku.at/sat-race-2006/descriptions/9-rsat.pdf.

Rector, A., & Horrocks, I. (1997). Experience building a large, re-usable medical ontology using a description logic with transitivity and concept inclusions. In *Proceedings of the Workshop on Ontological Engineering, AAAI Spring Symposium (AAAI'97)*. AAAI Press, Menlo Park, California.

Schild, K. D. (1991). A correspondence theory for terminological logics: preliminary report. In *Proc. of the 12th International Joint Conference on Artificial Intelligence*, pp. 466–471, Sydney, Australia.

Sebastiani, R. (2007a). From KSAT to Delayed Theory Combination: Exploiting DPLL Outside the SAT Domain. In *Proc. Frontiers of Combining Systems, FroCoS'07*, Vol. 4720 of *LNCS*. Springer. Invited talk.

Sebastiani, R. (2007b). Lazy Satisfiability Modulo Theories. *Journal on Satisfiability, Boolean Modeling and Computation, JSAT.*, *3*, 141–224.

Sebastiani, R., & Vescovi, M. (2006). Encoding the Satisfiability of Modal and Description Logics into SAT: The Case Study of K(m)/ALC.. In *Theory and Applications of Satisfiability Testing - SAT 2006, 9th International Conference, 2006, Proceedings*, Vol. 4121 of *LNCS*, pp. 130–135. Springer.

Seshia, S. A., Lahiri, S. K., & Bryant, R. E. (2003). A Hybrid SAT-Based Decision Procedure for Separation Logic with Uninterpreted Functions. In *Proc. DAC'03*.

Silva, J. P. M., & Sakallah, K. A. (1996). GRASP - A new Search Algorithm for Satisfiability. In *Proc. ICCAD'96*.

Sioutos, N., de Coronado, S., Haber, M. W., Hartel, F. W., Shaiu, W., & Wright, L. W. (2007). Nci thesaurus: A semantic model integrating cancer-related clinical and molecular information. *Journal of Biomedical Informatics*, *40*(1), 30–43.

Spackman, K. A., Campbell, K., & Cote, R. (1997). SNOMED RT: A reference terminology for healt care. *Journal of American Medical Informatics Association (Fall Symposium Supplement)*, 640–644.

Strichman, O. (2002). On Solving Presburger and Linear Arithmetic with SAT. In *Proc. of Formal Methods in Computer-Aided Design (FMCAD 2002)*, LNCS. Springer.

Strichman, O., Seshia, S., & Bryant, R. (2002). Deciding separation formulas with SAT. In *Proc. CAV'02*, LNCS. Springer.

Tacchella, A. (1999). *SAT system description. In *Proc. DL'99*, pp. 142–144.

The Gene Ontology Consortium (2000). Gene ontology: Tool for the unification of biology. *Nature Genetics*, *25*, 25–29.

V. Haarslev, R. M. (2001). RACER System Description. In *Proc. IJCAR-2001*, Vol. 2083 of *LNAI*. Springer.

Voronkov, A. (1999). KЯ: a theorem prover for K. In *CADE-16: Proceedings of the 16th International Conference on Automated Deduction*, No. 1632 in LNAI, pp. 383–387. Springer.

Voronkov, A. (2001). How to optimize proof-search in modal logics: new methods of proving redundancy criteria for sequent calculi. *ACM Transacrtions on Computational Logic*, *2*(2), 182–215.

Zhang, L., & Malik, S. (2002). The quest for efficient boolean satisfiability solvers. In *Proc. CAV'02*, No. 2404 in LNCS, pp. 17–36. Springer.

Appendix A. The Proof of Correctness & Completeness

Some Further Notation

Let ψ be a K_m-formula. We denote by $\overline{\psi}$ the representation of $\neg\psi$ in the current formalism: in NNF, $\overline{\Diamond_r\psi} \stackrel{\text{def}}{=} \Box_r\overline{\psi}$, $\overline{\Box_r\psi} \stackrel{\text{def}}{=} \Diamond_r\overline{\psi}$, $\overline{\psi_1 \wedge \psi_2} \stackrel{\text{def}}{=} \overline{\psi_1} \vee \overline{\psi_2}$, $\overline{\psi_1 \vee \psi_2} \stackrel{\text{def}}{=} \overline{\psi_1} \wedge \overline{\psi_2}$, $\overline{A_i} \stackrel{\text{def}}{=} \neg A_i$, $\overline{\neg A_i} \stackrel{\text{def}}{=} A_i$; in BNF, $\overline{\neg\Box_r\psi} \stackrel{\text{def}}{=} \Box_r\psi$, $\overline{\Box_r\psi} \stackrel{\text{def}}{=} \neg\Box_r\psi$, $\overline{\psi_1 \wedge \psi_2} \stackrel{\text{def}}{=} \overline{\psi_1} \vee \overline{\psi_2}$, $\overline{\psi_1 \vee \psi_2} \stackrel{\text{def}}{=} \overline{\psi_1} \wedge \overline{\psi_2}$, $\overline{A_i} \stackrel{\text{def}}{=} \neg A_i$, $\overline{\neg A_i} \stackrel{\text{def}}{=} A_i$.

We represent a truth assignment μ as a set of literals, with the intended meaning that a positive literal A_i (resp. negative literal $\neg A_i$) in μ means that A_i is assigned to true (resp. false). We say that μ *assigns* a literal l if it assigns a truth value to the atom of l. We say that a literal l occurs in a Boolean formula ϕ iff its atom occurs in ϕ.

Let $\mathcal{M}$ denote a Kripke model, and let σ (the label of) a generic state in $\mathcal{M}$. We denote by "1" is the root state of $\mathcal{M}$. By "$\langle\sigma : \psi\rangle \in \mathcal{M}$" we mean that $\sigma \in \mathcal{M}$ and $\mathcal{M}, \sigma \models \psi$. Thus, for every $\sigma \in \mathcal{M}$, either $\langle\sigma : \psi\rangle \in \mathcal{M}$ or $\langle\sigma : \overline{\psi}\rangle \in \mathcal{M}$.

For convenience, instead of (9) sometimes we use the equivalent definition:

$$Def(\sigma,\ \nu^r) \stackrel{\text{def}}{=} (L_{[\sigma,\ \nu^r]} \rightarrow \bigwedge_{\langle\sigma:\pi^{r,i}\rangle} (L_{[\sigma,\ \pi^{r,i}]} \rightarrow L_{[\sigma.i,\ \nu_0^r]})) \wedge \bigwedge_{\langle\sigma:\pi^{r,i}\rangle} Def(\sigma.i,\ \nu_0^r). \quad (20)$$

Notice that each $Def(\sigma,\ \psi)$ in (6), (7), (8), (20) is written in the general form

$$(L_{[\sigma,\ \psi]} \rightarrow \Phi_{\langle\sigma:\psi\rangle}) \wedge \bigwedge_{\langle\sigma':\psi'\rangle} Def(\sigma',\ \psi'). \quad (21)$$

We call *definition implication* for $Def(\sigma,\ \psi)$ the expressions "$(L_{[\sigma,\ \psi]} \rightarrow \Phi_{\langle\sigma:\psi\rangle})$".

The Proof

Let φ be a K_m-formula. We prove the following theorem, which states the soundness and completeness of K_m2SAT.

Theorem 1. *A K_m-formula φ is K_m-satisfiable if and only if the corrisponding $K_m2SAT(\varphi)$ is satisfiable.*

Proof. It is a direct consequence of the following Lemmas 2 and 3. □

Lemma 2. *Given a K_m-formula φ, if $K_m2SAT(\varphi)$ is satisfiable, then there exists a Kripke model $\mathcal{M}$ s.t. $\mathcal{M}, 1 \models \varphi$.*

Proof. Let μ be a total truth assignment satisfying $K_m2SAT(\varphi)$. We build from μ a Kripke model $\mathcal{M} = \langle\mathcal{U}, \mathcal{L}, \mathcal{R}_1, \ldots, \mathcal{R}_m\rangle$ as follows:

$$\mathcal{U} \stackrel{\text{def}}{=} \{\sigma\ :\ A_{[\sigma,\ \psi]}\ occurs\ in\ K_m2SAT(\varphi)\ for\ some\ \psi\} \quad (22)$$

$$\mathcal{L}(\sigma, A_i) \stackrel{\text{def}}{=} \begin{cases} True & if\ L_{[\sigma,\ A_i]} \in \mu \\ False & if\ \neg L_{[\sigma,\ A_i]} \in \mu \end{cases} \quad (23)$$

$$\mathcal{R}_r \stackrel{\text{def}}{=} \{\langle\sigma, \sigma.i\rangle\ :\ L_{[\sigma,\ \pi^{r,i}]} \in \mu\}. \quad (24)$$

We show by induction on the structure of φ that, for every $\langle \sigma : \psi \rangle$ s.t. $L_{[\sigma,\ \psi]}$ occurs on $K_m 2SAT(\varphi)$,

$$\langle \sigma : \psi \rangle \in \mathcal{M} \ \text{ if } \ L_{[\sigma,\ \psi]} \in \mu. \tag{25}$$

Base

$\psi = A_i$ or $\psi = \neg A_i$. Then (25) follows trivially from (23).

Step

$\psi = \alpha$. Let $L_{[\sigma,\ \alpha]} \in \mu$.

As μ satisfies (6), $L_{[\sigma,\ \alpha_i]} \in \mu$ for every $i \in \{1, 2\}$.

By inductive hypothesis, $\langle \sigma : \alpha_i \rangle \in \mathcal{M}$ for every $i \in \{1, 2\}$.

Then, by definition, $\langle \sigma : \alpha \rangle \in \mathcal{M}$.

Thus, $\langle \sigma : \alpha \rangle \in \mathcal{M}$ if $L_{[\sigma,\ \alpha]} \in \mu$.

$\psi = \beta$. Let $L_{[\sigma,\ \beta]} \in \mu$.

As μ satisfies (7), $L_{[\sigma,\ \beta_i]} \in \mu$ for some $i \in \{1, 2\}$.

By inductive hypothesis, $\langle \sigma : \beta_i \rangle \in \mathcal{M}$ for some $i \in \{1, 2\}$.

Then, by definition, $\langle \sigma : \beta \rangle \in \mathcal{M}$.

Thus, $\langle \sigma : \beta \rangle \in \mathcal{M}$ if $L_{[\sigma,\ \beta]} \in \mu$.

$\psi = \pi^{r,j}$. Let $L_{[\sigma,\ \pi^{r,j}]} \in \mu$.

As μ satisfies (8), $L_{[\sigma.j,\ \pi_0^{r,j}]} \in \mu$.

By inductive hypothesis, $\langle \sigma.j : \pi_0^{r,j} \rangle \in \mathcal{M}$.

Then, by definition and by (24), $\langle \sigma : \pi^{r,j} \rangle \in \mathcal{M}$.

Thus, $\langle \sigma : \pi^{r,j} \rangle \in \mathcal{M}$ if $L_{[\sigma,\ \pi^{r,j}]} \in \mu$.

$\psi = \nu^r$. Let $L_{[\sigma,\ \nu^r]} \in \mu$.

As μ satisfies (9), for every $\langle \sigma : \pi^{r,i} \rangle$ s.t. $L_{[\sigma,\ \pi^{r,i}]} \in \mu$, we have that $L_{[\sigma.i,\ \nu_0^r]} \in \mu$.

By inductive hypothesis, we have that $\langle \sigma : \pi^{r,i} \rangle \in \mathcal{M}$ and $\langle \sigma.i : \nu_0^r \rangle \in \mathcal{M}$.

Then, by definition and by (24), $\langle \sigma : \nu^r \rangle \in \mathcal{M}$.

Thus, $\langle \sigma : \nu^r \rangle \in \mathcal{M}$ if $L_{[\sigma,\ \nu^r]} \in \mu$.

If $\mu \models K_m 2SAT(\varphi)$, then $A_{[1,\ \varphi]} \in \mu$. Thus, by (25), $\langle 1 : \varphi \rangle \in \mathcal{M}$, i.e., $\mathcal{M}, 1 \models \varphi$. □ □

Lemma 3. *Given a K_m-formula φ, if there exists a Kripke model $\mathcal{M}$ s.t. $\mathcal{M}, 1 \models \varphi$, then $K_m 2SAT(\varphi)$ is satisfiable.*

Proof. Let $\mathcal{M}$ be a Kripke model s.t. $\mathcal{M}, 1 \models \varphi$. We build from $\mathcal{M}$ a truth assignment μ for $K_m2SAT(\varphi)$ recursively as follows: [27]

$$\mu \stackrel{\text{def}}{=} \mu_{\mathcal{M}} \cup \mu_{\overline{\mathcal{M}}} \tag{26}$$

$$\mu_{\mathcal{M}} \stackrel{\text{def}}{=} \{L_{[\sigma,\ \psi]} \in K_m2SAT(\varphi)\ :\ \langle \sigma : \psi \rangle \in \mathcal{M}\} \tag{27}$$
$$\cup\ \{\neg L_{[\sigma,\ \psi]} \in K_m2SAT(\varphi)\ :\ \langle \sigma : \overline{\psi} \rangle \in \mathcal{M}\}$$

$$\mu_{\overline{\mathcal{M}}} \stackrel{\text{def}}{=} \mu_{\pi\nu} \cup \mu_{\alpha\beta} \cup \mu_{\mathcal{A}} \tag{28}$$

$$\mu_{\pi\nu} \stackrel{\text{def}}{=} \{\neg L_{[\sigma,\ \pi^{r,i}]} \in K_m2SAT(\varphi)\ :\ \sigma \notin \mathcal{M}\} \tag{29}$$
$$\cup\ \{L_{[\sigma,\ \nu^r]} \in K_m2SAT(\varphi)\ :\ \sigma \notin \mathcal{M}\}$$

$$\mu_{\alpha\beta} \stackrel{\text{def}}{=} \{\neg L_{[\sigma,\ \alpha]} \in K_m2SAT(\varphi)\ :\ \sigma \notin \mathcal{M} \textit{ and } \neg L_{[\sigma,\ \alpha_i]} \in \mu_{\overline{\mathcal{M}}} \textit{ for some } i \in \{1,2\}\} \tag{30}$$
$$\cup\ \{\neg L_{[\sigma,\ \beta]} \in K_m2SAT(\varphi)\ :\ \sigma \notin \mathcal{M} \textit{ and } \neg L_{[\sigma,\ \beta_i]} \in \mu_{\overline{\mathcal{M}}} \textit{ for every } i \in \{1,2\}\}.$$

where $\mu_{\mathcal{A}}$ is a consistent truth assignment for the literals $L_{[\sigma,\ A_i]}$ s.t. $A_i \in \mathcal{A}$ and $\sigma \notin \mathcal{M}$.

By construction, for every $L_{[\sigma,\ \psi]}$ in $K_m2SAT(\varphi)$, μ assigns $L_{[\sigma,\ \psi]}$ to true iff it assigns $L_{[\sigma,\ \overline{\psi}]}$ to false and vice versa, so that μ is a consistent truth assignment.

First, we show that $\mu_{\mathcal{M}}$ satisfies the definition implications of all $Def(\sigma,\ \psi)$'s and $Def(\sigma,\ \overline{\psi})$' s.t. $\sigma \in \mathcal{M}$. Let $\sigma \in \mathcal{M}$. We distinguish four cases.

$\psi = \alpha$. Thus $\overline{\psi} = \beta$ s.t. $\beta_1 = \overline{\alpha_1}$ and $\beta_2 = \overline{\alpha_2}$.

- If $\langle \sigma : \alpha \rangle \in \mathcal{M}$ (and hence $\langle \sigma : \beta \rangle \notin \mathcal{M}$), then for both i's $\langle \sigma : \alpha_i \rangle \in \mathcal{M}$ and $\langle \sigma : \beta_i \rangle \notin \mathcal{M}$. Thus, by (27), $\{L_{[\sigma,\ \alpha_1]}, L_{[\sigma,\ \alpha_2]}, \neg L_{[\sigma,\ \beta]}\} \subseteq \mu_{\mathcal{M}}$, so that $\mu_{\mathcal{M}}$ satisfies the definition implications of both $Def(\sigma,\ \alpha)$ and $Def(\sigma,\ \beta)$.
- If $\langle \sigma : \alpha \rangle \notin \mathcal{M}$ (and hence $\langle \sigma : \beta \rangle \in \mathcal{M}$), then for some i $\langle \sigma : \alpha_i \rangle \notin \mathcal{M}$ and $\langle \sigma : \beta_i \rangle \in \mathcal{M}$. Thus, by (27), $\{\neg L_{[\sigma,\ \alpha]}, L_{[\sigma,\ \beta_i]}\} \subseteq \mu_{\mathcal{M}}$, so that $\mu_{\mathcal{M}}$ satisfies the definition implications of both $Def(\sigma,\ \alpha)$ and $Def(\sigma,\ \beta)$.

$\psi = \beta$. Like in the previous case, inverting ψ and $\overline{\psi}$.

$\psi = \pi^{r,j}$. Thus $\overline{\psi} = \nu^r$ s.t. $\nu^r_0 = \overline{\pi^{r,j}_0}$.

- If $\langle \sigma : \pi^{r,j} \rangle \in \mathcal{M}$ (and hence $\langle \sigma : \nu^r \rangle \notin \mathcal{M}$), then $\langle \sigma.j : \pi^{r,j}_0 \rangle \in \mathcal{M}$. Thus, by (27), $\{L_{[\sigma.j,\ \pi^{r,j}_0]}, \neg L_{[\sigma,\ \nu^r]}\} \subseteq \mu_{\mathcal{M}}$, so that $\mu_{\mathcal{M}}$ satisfies the definition implications of both $Def(\sigma,\ \pi^{r,j})$ and $Def(\sigma,\ \nu^r)$.
- If $\langle \sigma : \pi^{r,j} \rangle \notin \mathcal{M}$ (and hence $\langle \sigma : \nu^r \rangle \in \mathcal{M}$), then by (27) $\neg L_{[\sigma,\ \pi^{r,j}]} \in \mu_{\mathcal{M}}$, so that $\mu_{\mathcal{M}}$ satisfies the definition implications of $Def(\sigma,\ \pi^{r,j})$.

27. We assume that $\mu_{\mathcal{M}}$, $\mu_{\pi\nu}$ and $\mu_{\alpha\beta}$ are generated *in order*, so that $\mu_{\alpha\beta}$ is generated recursively starting from $\mu_{\pi\nu}$. Intuitively, $\mu_{\mathcal{M}}$ assigns the literals $L_{[\sigma,\ \psi]}$ s.t. $\sigma \in \mathcal{M}$ in such a way to mimic $\mathcal{M}$; $\mu_{\overline{\mathcal{M}}}$ assigns the other literals in such a way to mimic the fact that no state outside those in $\mathcal{M}$ is generated (i.e., all $L_{[\sigma,\ \pi]}$'s are assigned false and the $L_{[\sigma,\ \nu]}$'s, $L_{[\sigma,\ \alpha]}$'s, $L_{[\sigma,\ \beta]}$'s are assigned consequently).

As far as $Def(\sigma,\ \nu^r)$ is concerned, we partition the clauses in (9):

$$((L_{[\sigma,\ \nu^r]} \wedge L_{[\sigma,\ \pi^{r,i}]}) \rightarrow L_{[\sigma.i,\ \nu_0^r]}) \tag{31}$$

into two subsets. The first is the set of clauses (31) for which $\langle \sigma : \pi^{r,i} \rangle \in \mathcal{M}$. As $\langle \sigma : \nu^r \rangle \in \mathcal{M}$, we have that $\langle \sigma.i : \nu_0^r \rangle \in \mathcal{M}$. Thus, by (27), $L_{[\sigma.i,\ \nu_0^r]} \in \mu_{\mathcal{M}}$, so that $\mu_{\mathcal{M}}$ satisfies (31). The second is the set of clauses (31) for which $\langle \sigma : \pi^{r,i} \rangle \notin \mathcal{M}$. By (27) we have that $\neg L_{[\sigma,\ \pi^{r,i}]} \in \mu_{\mathcal{M}}$, so that $\mu_{\mathcal{M}}$ satisfies (31). Thus, $\mu_{\mathcal{M}}$ satisfies the definition implications also of $Def(\sigma,\ \nu^r)$.

$\psi = \nu^r$. Like in the previous case, inverting ψ and $\overline{\psi}$.

Notice that, if $\sigma \notin \mathcal{M}$, then $\sigma.i \notin \mathcal{M}$ for every i. Thus, for every $Def(\sigma,\ \psi)$ s.t. $\sigma \notin \mathcal{M}$, all atoms in the implication definition of $Def(\sigma,\ \psi)$ are not assigned by $\mu_{\mathcal{M}}$.

Second, we show by induction on the recursive structure of $\mu_{\overline{\mathcal{M}}}$ that $\mu_{\overline{\mathcal{M}}}$ satisfies the definition implications of all $Def(\sigma,\ \psi)$'s and $Def(\sigma,\ \overline{\psi})$'s s.t. $\sigma \notin \mathcal{M}$. Let $\sigma \notin \mathcal{M}$.

As a base step, by (29), $\mu_{\pi\nu}$ satisfies the definition implications of all $Def(\sigma,\ \pi^{r,i})$'s and $Def(\sigma,\ \nu^r)$'s because it assigns false to all $L_{[\sigma,\ \pi^{r,i}]}$'s. Indeed, $\mu_{\mathcal{A}}$ assigns every literal of the type $L_{[\sigma,\ A_i]}$ s.t $A_i \in \mathcal{A}$ and $\sigma \notin \mathcal{M}$ (notice that all the $Def(\sigma,\ A_i)$'s definitions are trivially satisfied and don't contain any definition implications).

As inductive step, we show on the inductive structure of $\mu_{\alpha\beta}$ that $\mu_{\alpha\beta}$ satisfies the definition implications of all $Def(\sigma,\ \alpha)$'s and $Def(\sigma,\ \beta)$'s

Let $\psi \stackrel{\text{def}}{=} \alpha$ and $\overline{\psi} = \beta$ s.t. $\beta_i = \overline{\alpha_i}$ (or vice versa). Then we have that:

- if both $L_{[\sigma,\ \alpha_i]}$'s (respectively at least one $L_{[\sigma,\ \beta_i]}$) are assigned true by $\mu_{\overline{\mathcal{M}}}$, then the definition implications of $Def(\sigma,\ \alpha)$ (respectively $Def(\sigma,\ \beta)$) is already trivially satisfied;

- if at least one $L_{[\sigma,\ \alpha_i]}$ (respectively both $L_{[\sigma,\ \beta_i]}$'s) is assigned false by $\mu_{\overline{\mathcal{M}}}$, then by (30) $L_{[\sigma,\ \alpha]}$ (respectively $L_{[\sigma,\ \beta]}$) is assigned false by $\mu_{\alpha\beta}$, which satisfies the definition implication of $Def(\sigma,\ \alpha)$ (respectively $Def(\sigma,\ \beta)$).

Thus $\mu_{\overline{\mathcal{M}}}$ satisfies the definition implications of all the $Def(\sigma,\ \psi)$'s and $Def(\sigma,\ \overline{\psi})$'s s.t. $\sigma \notin \mathcal{M}$.

On the whole, $\mu \models Def(\sigma,\ \psi)$ for every $Def(\sigma,\ \psi)$. By construction, $\mu_{\mathcal{M}} \models A_{[1,\ \varphi]}$ since $\langle 1 : \varphi \rangle \in \mathcal{M}$. Therefore $\mu \models K_m 2SAT(\varphi)$. □