

- [HNSS90] B. Hollunder, W. Nutt, and M. Schmidt-Schauß. Subsumption Algorithms for Concept Description Languages. In *Proc. 8th European Conference on Artificial Intelligence*, pages 348–353, 1990.
- [Lad77] R. Ladner. The computational complexity of provability in systems of modal propositional logic. *SIAM J. Comp.*, 6(3):467–480, 1977.
- [Mas94] F. Massacci. Strongly analytic tableaux for normal modal logics. In *Proc. of the 12th Conference on Automated Deduction*, 1994.
- [MSL92] D. Mitchell, B. Selman, and H. Levesque. Hard and Easy Distributions of SAT Problems. In *Proc. of the 10th National Conference on Artificial Intelligence*, pages 459–465, 1992.
- [Sch91] K. D. Schild. A correspondence theory for terminological logics: preliminary report. In *Proc. of the 12th International Joint Conference on Artificial Intelligence*, pages 466–471, Sydney, Australia, 1991.
- [Seb94] R. Sebastiani. Applying GSAT to Non-Clausal Formulas. *Journal of Artificial Intelligence Research*, 1:309–314, 1994. Also DIST-Technical Report 94-0018, DIST, University of Genova, Italy.
- [SSS91] M. Schmidt-Schauß and G. Smolka. Attributive Concept Descriptions with Complements. *Artificial Intelligence*, 48:1–26, 1991.
- [US94] T. E. Uribe and M. E. Stickel. Ordered Binary Decision Diagrams and the Davis-Putnam Procedure. In *Proc. of the 1st International Conference on Constraints in Computational Logics*, 1994.
- [WH94] C. P. Williams and T. Hogg. Exploiting the deep structure of constraint problems. *Artificial Intelligence*, 70:73–117, 1994.
- [ZS94] H. Zhang and M. Stickel. Implementing the Davis-Putnam algorithm by tries. Technical report, University of Iowa, August 1994.

# References

- [AG93] A. Armando and E. Giunchiglia. Embedding Complex Decision Procedures inside an Interactive Theorem Prover. *Annals of Mathematics and Artificial Intelligence*, 8(3-4):475-502, 1993.
- [BB92] M. Buro and H. Buning. Report on a SAT competition. Technical Report 110, University of Paderborn, Germany, November 1992.
- [Bry92] R. E. Bryant. Symbolic Boolean manipulation with ordered binary-decision diagrams. *ACM Computing Surveys*, 24(3):293-318, September 1992.
- [CA93] J. Crawford and L. Auton. Experimental results on the crossover point in satisfiability problems. In *Proc. of the 11th National Conference on Artificial Intelligence*, pages 21-27, 1993.
- [CKT91] P. Cheeseman, B. Kanefski, and W. M. Taylor. Where the really hard problems are. In *Proc. of the 12th International Joint Conference on Artificial Intelligence*, pages 163-169, 1991.
- [D'A92] M. D'Agostino. Are Tableaux an Improvement on Truth-Tables? *Journal of Logic, Language and Information*, 1:235-252, 1992.
- [DLL62] M. Davis, G. Longemann, and D. Loveland. A machine program for theorem proving. *Journal of the ACM*, 5(7), 1962.
- [DM94] M. D'Agostino and M. Mondadori. The Taming of the Cut. *Journal of Logic and Computation*, 4(3):285-319, 1994.
- [DP60] M. Davis and H. Putnam. A computing procedure for quantification theory. *Journal of the ACM*, 7:201-215, 1960.
- [Fit83] M. Fitting. *Proof Methods for Modal and Intuitionistic Logics*. D. Reidel Publishg, 1983.
- [GS94] F. Giunchiglia and L. Serafini. Multilanguage hierarchical logics (or: how we can do without modal logics). *Artificial Intelligence*, 65:29-70, 1994. Also IRST-Technical Report 9110-07, IRST, Trento, Italy.
- [GS96] F. Giunchiglia and R. Sebastiani. Building decision procedures for modal logics from propositional decision procedures - the case study of modal K. In *Proc. of the 13th Conference on Automated Deduction*, August 1996.
- [GSGF93] F. Giunchiglia, L. Serafini, E. Giunchiglia, and M. Frixione. Non-Omniscient Belief as Context-Based Reasoning. In *Proc. of the 13th International Joint Conference on Artificial Intelligence*, pages 548-554, Chambéry, France, 1993. Also IRST-Technical Report 9206-03, IRST, Trento, Italy.
- [HM92] J.Y. Halpern and Y. Moses. A guide to the completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54(3):319-379, 1992.

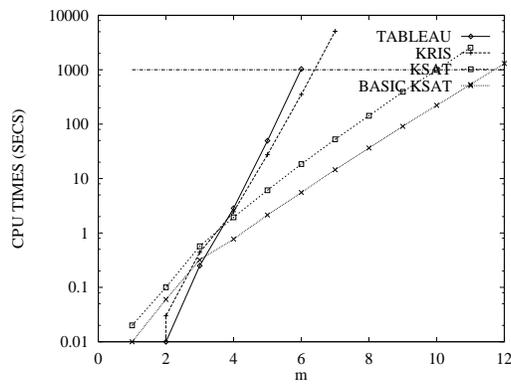


Figure 5: CPU times for the class of  $\varphi_d^K$  formulas.

(i.e., with no factorization of  $\bigwedge_i \alpha_{r_i}$ 's and no checking of intermediate assignments), called below BASIC KSAT, on these formulas, for increasing values of  $d$ . The results are plotted in Figure 5. The TABLEAU, KRIS, KSAT and BASIC KSAT curves grow exponentially, approximatively as  $(16.0)^d$ ,  $(12.7)^d$ ,  $(2.6)^d$  and  $(2.4)^d$  respectively, exceeding 1000s for  $d = 6$ ,  $d = 7$ ,  $d = 11$  and  $d = 12$  respectively. The slight difference between KSAT and BASIC KSAT is due to the overhead introduced by the  $\bigwedge_i \alpha_{r_i}$  factorization, which is useless with these formulas. It is worth observing that the result of tracing the global number of truth assignments  $\mu$ , recursively found by both KSAT and BASIC KSAT, gave exactly  $2^{d+1} - 1$  for every  $d$ , that is the minimum number of Kripke states. KSAT and BASIC KSAT found no redundant truth assignments.

## 6 Conclusions

This paper presents what we think are three very important results:

1. it provides a new implemented algorithm, KSAT, for deciding satisfiability in  $\mathcal{ALC}$  ( $K(m)$ ) which outperforms of orders of magnitude the previous state-of-the-art decision procedures;
2. it shows that the results provided are not by chance, and that all SAT-based modal decision procedures (that is, all the modal decision procedures based on SAT decision procedures) are intrinsically bound to be more efficient than tableau-based decision procedures; and
3. it provides evidence, though very partial, of an easy-hard-easy pattern independent of all the parameters of evaluation considered. If the current partial evidence is confirmed, this is the first time that this phenomenon, well known for SAT and other NP-hard problems, is found in modal logics.

non-interfering search trees. Therefore, the bigger is  $m$ , the more partitioned is the search space, and the easier is the problem to solve. Second, a careful look reveals that the satisfiability percentage increases with  $m$ . Again, there is no mutual dependency between the satisfiability of the distinct  $\mu_r$ 's. Therefore the bigger is  $m$ , the less constrained is  $\mu$ , and the more likely satisfiable is  $\varphi$ .

The results of the third experiment (right column) provide evidence of the fact that the complexity increases with the modal depth  $d$ , for both KSAT and KRIS. This is rather intuitive: the higher is  $d$ , the deeper are the Kripke models to be searched, and the higher is the complexity of the search.

## 5.2 A global analysis

The KSAT curves (top and middle rows) highlight the existence a linear and an easy-hard-easy component. In fact, if we increase  $N$  from 3 to 5 and  $L$  accordingly (left column), the size of the search space has a relevant increase. Therefore, while for  $N = 3$  the linear component prevails, for  $N = 5$  the easy-hard-easy component dominates. Moreover, when varying the number of modalities (center column), the wff sizes are kept the same for all curves. Therefore, when the effect of the easy-hard-easy component vanishes ( $L/N > 20$ ), the curves collapse together, as the time for preprocessing and *assign* does not depend on the number of modalities  $m$ . Notice that the locations of the easy-hard-easy zones do not seem to vary significantly, neither with the number of variables  $N$  (left column), nor with the number of modalities  $m$  (center column), nor with the depth  $d$  (right column).

Let us now consider the satisfiability plots in Figure 4 (bottom row). Despite the noise and the approximations due to timeouts, it is easy to notice that the 50% satisfiability point is centered around  $L = 15N \sim 20N$  in all the experiments. Moreover, in the first experiment a careful look reveals that the satisfiability transition becomes steeper when increasing  $N$  (e.g., compare the  $N = 3$  and  $N = 5$  plots). Finally, in all experiments, the curves representing the median number of DPLL calls (middle row) generally locate the peaks around the satisfiability transition, although they seem to anticipate a little the 50% crossover point. From these facts we may conjecture (to be verified!) the existence for  $K(m)/\mathcal{ALC}$  of a phase transition phenomenon, similar to that already known for SAT and other NP-hard problems (see, e.g, [CKT91, MSL92, WH94]).

The final observation comes from the three sets of median CPU times curves (top row): KSAT outperforms KRIS in all the testbeds, independently on the number of variables  $N$ , the number of modalities  $m$  or the depth  $d$  considered. This confirms the analysis done in Section 4. Again, this is not only a quantitative performance gap (up to 3-4 orders of magnitude) but also a qualitative one, as all KRIS curves grow exponentially with  $L$ , while all KSAT curves grow polinomially. To provide further evidence, we have performed another, quite different, test, based on the class of wffs  $\{\varphi_d^K\}_{d=1,2,\dots}$  presented in [HM92]. This is a class of  $K(1)$ -satisfiable wffs, with depth  $d$  and  $2d + 1$  propositional variables. These wffs are paradigmatic for modal  $K$ , as every Kripke structure satisfying  $\varphi_d^K$  has at least  $2^{d+1} - 1$  distinct states, while  $|\varphi_d^K|$  is  $O(d^2)$ . From the results in [HM92] we can reasonably assume a minimum exponential growth factor of  $2^d$  for any ordinary algorithm based on Kripke semantics. We run TABLEAU, KRIS, KSAT and a “basic” version of KSAT

we investigate the effects of varying one parameter while fixing the others. In Experiment 1 (left column) we fix  $d = 2$ ,  $m = 1$ ,  $p = 0.5$  and plot different curves for increasing numbers of variables  $N = 3, 4, 5$ <sup>13</sup>. In Experiment 2 (center column) we fix  $d = 2$ ,  $N = 4$ ,  $p = 0.5$  and plot different curves for increasing number of distinct modalities  $m = 1, 2, 5, 10, 20$ . In Experiment 3 (right column) we fix  $m = 1$ ,  $N = 3$ ,  $p = 0.5$  and plot different curves for increasing modal depths  $d = 2, 3, 4, 5$ . For each experiment, we present three distinct sets of curves, each corresponding to a distinct row. In the first (top row) we plot the median CPU time obtained by running both KSAT and KRIS. This gives an overall picture of the qualitative behaviour of KSAT and KRIS and allows for a direct comparison between them. In the second (middle row) we plot the KSAT median number of recursive DPLL calls, that is, the size of the space effectively searched. This allows us to drop the linear component due to the preprocessing and the function calls to *assign*. In the third (bottom row) we plot the percentage of satisfiable wffs evaluated by KSAT. This gives a coarse indication of the average level of constraintness of the test wffs<sup>14</sup>.

Despite the big noise, due to the small samples/point rate (100), the results indicated in Figure 4 provide interesting indications. We report below (Subsection 5.1) a first pass, experiment by experiment, analysis of the results. This gives us an idea of how efficiency and satisfiability are affected by each single parameter. In Subsection 5.2 we report a global and, in some respects, more interesting analysis of the results we have.

## 5.1 A testwise analysis

The results of the first experiment (left column) show that increasing  $N$  (and  $L$  accordingly) causes a relevant increase in complexity — up to one order of magnitude per variable in the “hard” zone for KSAT, up to two orders of magnitude per variable, as far as we can see, for KRIS. This should not be a surprise, as in  $K/K(m)$ , adding few variables may cause an exponential increase of the search space. Each variable may in fact assume distinct truth values inside distinct states/possible worlds, that is, each variable must be considered with an “implicit multiplicity” equal to the number of states of a potential Kripke model.

The results of the second experiment (center column) present two interesting aspects. First, the complexity of the search monotonically decreases with the increase of the number  $m$  of modalities, for both KSAT and KRIS (top and middle box). At a first sight it may sound like a surprise, but it should not be so. In fact, each truth assignment  $\mu$  is partitioned into  $m$  independent sub-assignments  $\mu_r$ ’s, each restricted to a single  $\square_r$  (see Equations (1) and (2)). This means “dividing and conquering” the search tree into  $m$

---

<sup>13</sup>If we compare the KSAT and KRIS plots in Figure 2 (left) with the  $L = 3$  KSAT and KRIS plots in Figure 4 (top left), we notice that the plots are different, although they are computed on sample wffs with the same parameter values. This is due to the fact that (i) the former ones are run on a much faster machine; (ii) the starting seeds are different, causing thus the generation of different sample sets.

<sup>14</sup>This percentage is evaluated on the number of samples which *effectively* ended computation within the timeout. Therefore this datum should be considered only as a coarse indication. To obtain an accurate evaluation, we should drop the timeout mechanism and evaluate the satisfiability percentage on at least 1000 samples/point, like in Figure 2 (right).

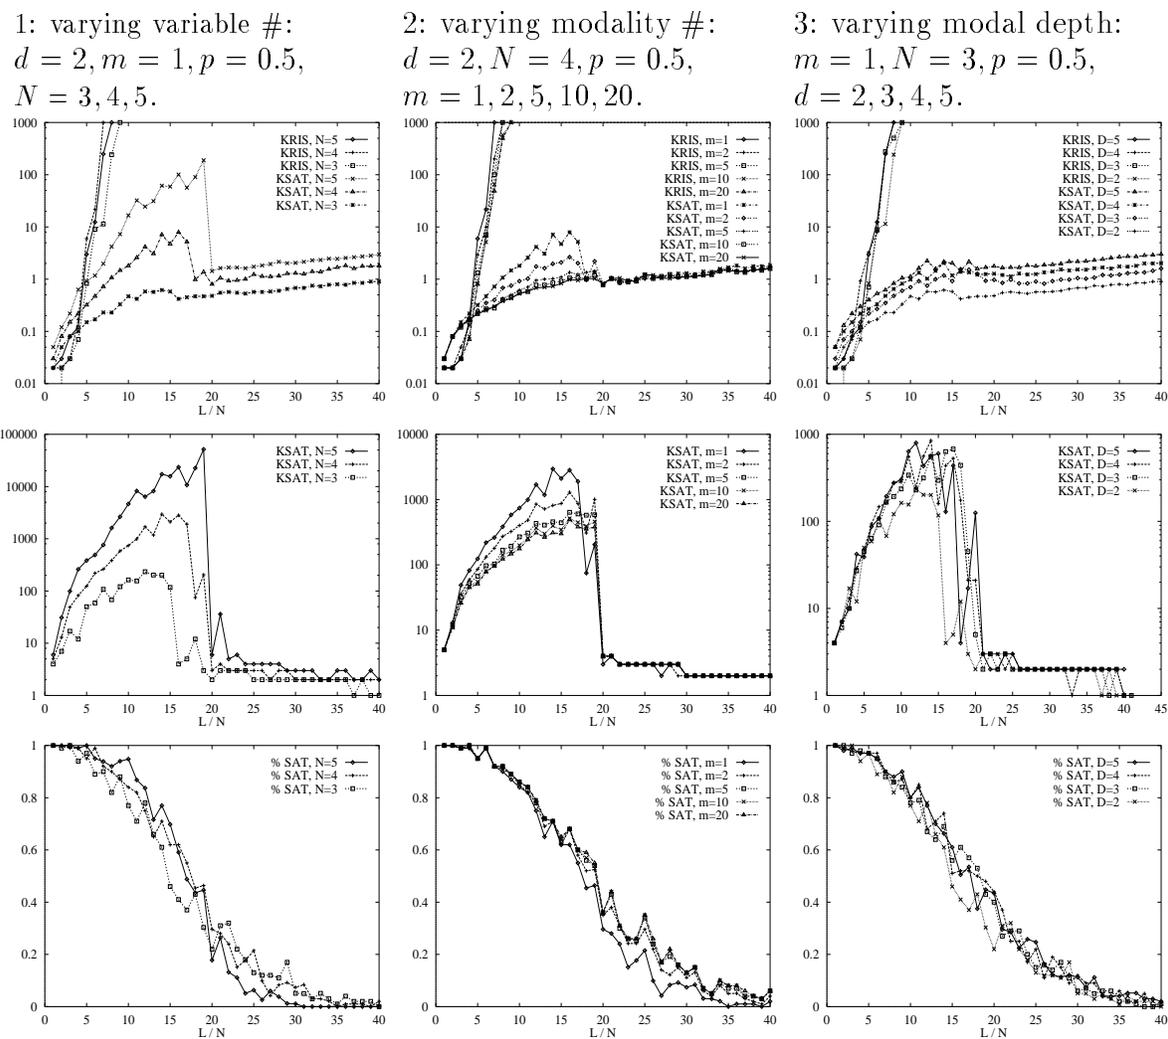


Figure 4: The results of the three experiments.

the bigger is  $L$  in a CNF wff, the more likely an assignment generates an empty clause.) Therefore, as  $\varphi$  becomes highly constrained (e.g., when  $L$  is big enough) the search tree is very heavily pruned. As a consequence, for  $L$  bigger than a certain value, the size of the search tree *decreases* with  $L$ , as it can be easily noticed in Figure 2 (right).

## 5 An exhaustive empirical analysis

In the tests described in this section we have tested and compared KSAT and KRIS, that is, the fastest SAT-based and the fastest tableau-based decision procedure at our disposal. We have performed three experiments on 48,000 randomly generated wffs, run according to our test methodology, whose results are all described in Figure 4. All curves represent 100 samples/point. As above, the range  $N \dots 40N$  for  $L$  has been chosen empirically to cover coarsely the “100% satisfiable – 100% unsatisfiable” transition. In each experiment



KSAT and KSAT<sub>0</sub> qualitatively outperform TABLEAU and KRIS. In fact, while TABLEAU and KRIS present an exponential growth against the number of clauses, both KSAT<sub>0</sub> and KSAT curves present a polynomial growth. In particular, the KSAT CPU time curve (like that of KSAT<sub>0</sub>) results from a combination of (i) a linear component and (ii) an easy-hard-easy component, centered in the satisfiability transition zone. Both components above are straightforward to notice in Figure 2 (right). The former is due to the preprocessing and to the linear-time function *assign*, which is invoked at every DPLL recursive call. The latter represents the number of recursive DPLL calls, i.e., the size of the tree effectively searched.

The quantitative and qualitative performance gaps pointed out by the third and the fourth observation above are very important and deserve some explanation. Let us consider for instance KSAT and KRIS. Both procedures work (i) by enumerating truth assignments which propositionally satisfy the input wff  $\varphi$  and (ii) by recursively checking the K(m)-satisfiability of the assignments found. Both algorithms perform the latter step in the same way. The key difference is in the first step, that is, in the way KSAT and KRIS handle propositional inference.

In KRIS truth assignments are (implicitly) generated as branches of an analytic propositional tableau, that is, by the recursive application of the rules:

$$\frac{\varphi_1 \wedge \varphi_2}{\varphi_1, \varphi_2} (\wedge\text{-rule}) \qquad \frac{\varphi_1 \vee \varphi_2}{\varphi_1 \quad \varphi_2} (\vee\text{-rule}) \qquad (4)$$

and the equivalent rules for  $\{\neg\vee, \neg\wedge, \supset\}$ . Analytic propositional tableaux perform what we call *syntactic branching*, that is, a branching on the syntactic structure of  $\varphi$ . As widely discussed in [D'A92, DM94], any application of the  $\vee$ -rule generates two subtrees which are *not mutually inconsistent*<sup>11</sup>, that is, two subtrees which may share propositional models. The set of truth assignments enumerated by propositional tableau procedures is intrinsically redundant, and may contain many duplicate and/or subsumed assignments. As a consequence, the number of truth assignments generated grows exponentially with the number of disjunctions occurring positively in  $\varphi$  (in our tests, the number of clauses  $L$ ), although the actual number of non-redundant assignments propositionally satisfying  $\varphi$  is much smaller. This redundancy is a source of a high degree of inefficiency when using analytic tableaux for propositional satisfiability.

Things get much worse in the modal case. Unlike the propositional case — where tableaux look for *one* assignment satisfying the input formula — in K(m) propositional tableaux enumerate all the truth assignments, which must be recursively checked for K(m)-consistency. (The number of assignments may be huge: up to ten thousands in our tests.) This requires checking recursively (possibly many) subwffs of the form  $\bigwedge_i \alpha_{ri} \wedge \beta_j$  of depth  $d - 1$ , for which a propositional tableau will enumerate truth assignments, and so forth. Any redundant truth assignment enumerated at depth  $d$  introduces a redundant modal search tree of depth  $d$ . Even worse, this propositional redundancy propagates exponentially with the depth  $d$ , following the analysis of the subwffs of decreasing depth.

---

<sup>11</sup>As pointed out in [D'A92, DM94], in Analytic tableaux rules are unable to represent *bivalence*: “every proposition is either true or false, *tertium non datur*”. This is a consequence of the elimination of the cut rule in cut-free sequent calculi, from which analytic tableaux are derived.

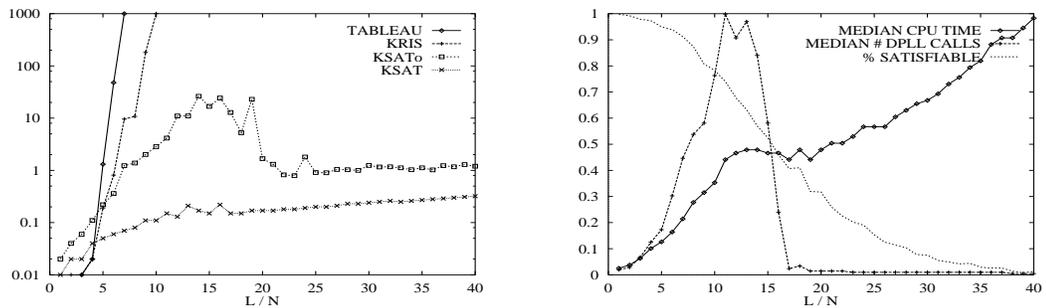


Figure 2:  $d = 2$ ,  $m = 1$ ,  $N = 3$ ,  $p = 0.5$ ,  $L = N \dots 40N$ . Left: TABLEAU, KRIS,  $\text{KSAT}_0$  and KSAT. Median CPU time, 100 samples/point. Right: KSAT. Normalized plots of Median CPU time, Median # of DPLL calls, satisfiability ratio, 1000 samples/point.

time. Figure 2 (left) presents the median CPU time plots for all four systems <sup>9</sup>. Notice the logarithmic scale on the vertical axis. In Figure 2 (right) we plot respectively the median CPU time, the median number of DPLL calls and the satisfiability percentage curves obtained by running KSAT on the same problem above, with 1000 sample wffs/point. In Figure 2 (right) the curves are all normalized to  $1^{10}$ .

Four observations can be made, given below in increasing order of importance.

First, improving the quality of the implementation, e.g., from TABLEAU to KRIS or from  $\text{KSAT}_0$  to KSAT, may introduce good quantitative performance improvements. In fact, KRIS reaches the time bound at the 10th step, while TABLEAU reaches the time bound at the 7th step, about two orders of magnitude above the corresponding KRIS value. Similarly,  $\text{KSAT}_0$  has a maximum at the 14th step, more than 2 orders of magnitude above the corresponding KSAT value. However, and this is the second observation, improving the quality of the implementation does not seem to affect the qualitative behaviour of the procedures. In fact, both the TABLEAU and KRIS curves present an exponential growth with the number of clauses, while both  $\text{KSAT}_0$  and KSAT curves flatten when the number of clauses exceeds a certain value.

Third, independently from the quality of implementation, KSAT and  $\text{KSAT}_0$  quantitatively outperform TABLEAU and KRIS. For instance, the performance gap between KSAT and KRIS at the 10th step is about 4 orders of magnitude. Moreover, the extrapolation of the KRIS curve suggests that its value — and the performance gap with KSAT — would reach several orders of magnitude for problems at the right end side of the plots. To support this consideration, we ran KRIS on 100 samples of the same problem, for  $L = 40N$ . No sample wff was solved within the timeout. When releasing the timeout mechanism, KRIS was not able to end successfully the computation of the first sample wff after a run of one week. Fourth, and most important, independently of the quality of implementation,

<sup>9</sup>We compare median values rather than mean values, as the former are much less sensitive to the noise introduced by outliers.

<sup>10</sup>The tests in Figures 2 (left) and 5 have been compiled and run under Allegro CL 4.2 on a SUN SPARC10 32M workstation. The test in Figure 2 (right) has been compiled and run under AKCL 1.623 on another SUN SPARC10 32M workstation. The tests in Figure 4 have been compiled and run under Allegro CL 4.1 on two identical SUN SPARC2 32M workstations.

### 3 The test method

The methodology we use generalizes the *fixed-clause-length* model commonly used in propositional SAT testing (see, e.g., [MSL92, BB92]).

Let a  $3\text{CNF}_{K(m)}$  wff be a conjunction of  $3\text{CNF}_{K(m)}$  clauses. Let a  $3\text{CNF}_{K(m)}$  clause be a disjunction of three  $3\text{CNF}_{K(m)}$  literals, i.e.,  $3\text{CNF}_{K(m)}$  atoms or their negations. Let a  $3\text{CNF}_{K(m)}$  atom be either a propositional atom or a wff in the form  $\Box_r C'$ ,  $C'$  being a  $3\text{CNF}_{K(m)}$  clause. Then  $3\text{CNF}_{K(m)}$  wffs are randomly generated according to the following parameters:

- (i) the modal depth  $d$ ;
- (ii) the number of distinct boxes  $m$ ;
- (iii) the number of clauses  $L$ ;
- (iv) the number of propositional variables  $N$ ;
- (v) a probability value  $p$ , stating the mean ratio of the propositional atoms at every level of the wff tree.

(Notice that, if we set  $d = 0$ , we have the standard 3SAT test method [MSL92].)

For fixed  $N$ ,  $d$ ,  $m$  and  $p$ , for increasing values of  $L$ , a certain number (100, 500, 1000...) of random  $3\text{CNF}_{K(m)}$  wffs are generated, internally sorted, and then given in input to the procedure under test. Satisfiability percentages and mean/median CPU times are plotted against the  $L/N$  ratio.

Similarly to the propositional 3CNF case, the methodology proposed above presents three main features. First, the method is very general:  $3\text{CNF}_{K(m)}$  wffs represent all  $K(m)$  wffs, as there is a  $K(m)$ -satisfiability-preserving way of converting any  $K(m)$  wff into  $3\text{CNF}_{K(m)}$ . Second, the usage of  $3\text{CNF}_{K(m)}$  form minimizes the number of parameters to handle. Finally, the parameters  $L$  and  $N$  allow for a coarse “tuning” of both the satisfiability probability and the hardness of random 3CNF modal wffs, so that it is possible to generate very hard problems with near 0.5 satisfiability probability.

### 4 Tableau-based vs. SAT-based procedures

In the tests described in this section we have tested and compared TABLEAU, KRIS, KSAT<sub>0</sub> and KSAT on the same group of 4,000 random formulas, with  $d = 2$ ,  $m = 1$ ,  $N = 3$ ,  $p = 0.5$ ,  $L \in \{N \dots 40N\}$ , 100 samples/point. These values have been chosen as in the analysis described in [GS96] they gave the highest execution times with both TABLEAU and KSAT<sub>0</sub>. The range  $N \dots 40N$  for  $L$  has been chosen empirically to cover coarsely the “100% satisfiable – 100% unsatisfiable” transition. As a general test rule we have introduced a timeout of 1000s on each sample wff. If the decision procedure under test exceeds the timeout for a given wff, a failure value is returned and the CPU time value is conventionally set to 1000s. Furthermore, we have stopped running the whole test whenever more than 50% samples (e.g., 50 out of 100 samples) have taken more than 1000s each to execute. These two choices have caused a relevant reduction of the testing

- (base) If  $\varphi = T$ , then  $\mu$  satisfies  $\varphi$ . Thus, if  $\mu$  is  $K(m)$ -satisfiable, then  $\varphi$  is  $K(m)$ -satisfiable. Therefore  $KSAT_W$  invokes  $KSAT_A(\mu)$  (where “ $A$ ” stands for (truth) Assignment).  $KSAT_A$  returns a truth value asserting whether  $\mu$  is  $K(m)$ -satisfiable or not.
- (backtrack) If  $\varphi = F$ , then  $\mu$  can not be a truth assignment for  $\varphi$ . Therefore  $KSAT_W$  returns *False*.
- (unit) If a literal  $l$  occurs in  $\varphi$  as a unit clause, then  $l$  must be assigned  $T$ <sup>8</sup>. To obtain this,  $KSAT_W$  is invoked recursively with arguments the wff returned by  $assign(l, \varphi)$  and the assignment obtained by adding  $l$  to  $\mu$ .  $assign(l, \varphi)$  substitutes every occurrence of  $l$  in  $\varphi$  with  $T$  and evaluates the result.
- (split) If none of the above situations occurs, then  $choose\_literal(\varphi)$  returns an unsigned literal  $l$  according to some heuristic criterion. Then  $KSAT_W$  is first invoked recursively with arguments  $assign(l, \varphi)$  and  $\mu \wedge l$ . If the result is negative, then  $KSAT_W$  is invoked with arguments  $assign(\neg l, \varphi)$  and  $\mu \wedge \neg l$ .

$KSAT_A(\mu)$  invokes  $KSAT_{RA}(\mu_r)$  (where “ $RA$ ” stands for Assignment Restricted to one modality) for any index  $r$  such that  $\Box_r$  occurs in  $\mu$ .  $KSAT_{RA}$  returns a truth value asserting whether  $\mu_r$  is  $K(m)$ -satisfiable or not.

The correctness and completeness of  $KSAT$  can be easily seen, for instance by noticing the close parallel with Fitting’s tableau described in [Fit83]. It is important to notice that  $KSAT_W$  is a variant of the non-CNF version of the Davis-Putnam-Longemann-Loveland SAT procedure [DP60, DLL62] (DPLL from now on), as described in [AG93]. Unlike DPLL, whenever an assignment  $\mu$  has been found,  $KSAT_W$ , instead of returning *True*, invokes  $KSAT_A(\mu)$ . Essentially, DPLL is used to generate truth assignments, whose  $K(m)$ -satisfiability is recursively checked by  $KSAT_A$ . We have implemented the algorithm described in Figure 1 as a procedure, also called  $KSAT$ , implemented in Common Lisp on top of the non-CNF DPLL decision procedure described in [AG93]. DPLL is well known to be one of the fastest decision procedures for SAT (see, e.g., [BB92, US94]). However the implementation we use, though relatively fast, is much slower than the state-of-the-art SAT decision procedures (see, e.g., [BB92, ZS94]). The basic version of the algorithm described in Figure 1 is improved in the following way. First, all modal atoms are internally ordered. This avoids assigning different truth values to permutations of the same sub-wffs. Secondly,  $KSAT_{RA}$  is implemented in such a way to “factorize” the common component  $\bigwedge_i \alpha_{ri}$  in searching truth assignments for  $\bigwedge_i \alpha_{ri} \wedge \neg \beta_{r1}, \bigwedge_i \alpha_{ri} \wedge \neg \beta_{r2}, \dots$ . Finally,  $KSAT_W$  is modified in such a way that  $KSAT_A$  is (optionally) invoked on intermediate assignments before every split. This drastically prunes search whenever inconsistent intermediate assignments are detected. These topics are described in detail in [GS96]. More recently we have also introduced a form of preprocessing — essentially, a recursive removal of duplicate and contradictory subwffs — of the input formulas.

---

<sup>8</sup>A notion of unit clause for non-CNF propositional wffs is given in [AG93]. More generally, [AG93] and [Seb94] show how decision procedures for CNF formulas can be modified to work for non-CNF formulas

```

function KSAT( $\varphi$ )
  return KSATW( $\varphi, T$ );

function KSATW( $\varphi, \mu$ )
  if  $\varphi = T$                                      /* base */
    then return KSATA( $\mu$ );
  if  $\varphi = F$                                      /* backtrack */
    then return False;
  if {a unit clause ( $l$ ) occurs in  $\varphi$ }          /* unit */
    then return KSATW(assign( $l, \varphi$ ),  $\mu \wedge l$ );
   $l :=$  choose-literal( $\varphi$ );                       /* split */
  return KSATW(assign( $l, \varphi$ ),  $\mu \wedge l$ ) or
         KSATW(assign( $\neg l, \varphi$ ),  $\mu \wedge \neg l$ );

function KSATA( $\bigwedge_i \Box_1 \alpha_{1i} \wedge \bigwedge_j \neg \Box_1 \beta_{1j} \wedge \dots \wedge \bigwedge_i \Box_m \alpha_{mi} \wedge \bigwedge_j \neg \Box_m \beta_{mj} \wedge \gamma$ )
  for any box index  $r$  do
    if not KSATRA( $\bigwedge_i \Box_r \alpha_{ri} \wedge \bigwedge_j \neg \Box_r \beta_{rj}$ )
      then return False;
  return True;

function KSATRA( $\bigwedge_i \Box_r \alpha_{ri} \wedge \bigwedge_j \neg \Box_r \beta_{rj}$ )
  for any conjunct “ $\neg \Box_r \beta_{rj}$ ” do
    if not KSAT( $\bigwedge_i \alpha_{ri} \wedge \neg \beta_{rj}$ )
      then return False;
  return True;

```

Figure 1: The basic version of KSAT algorithm.

where  $\gamma = \bigwedge_{k=1}^R A_k \wedge \bigwedge_{h=R+1}^S \neg A_h$  is a conjunction of propositional literals. Furthermore we use the greek letters  $\mu, \eta$  to represent truth assignments.

2. Prove that the input wff  $\varphi$  is K(m)-satisfiable by finding (among all the possible truth assignments) a K(m)-satisfiable truth assignment  $\mu$  of form as in (1).  $\mu$  is K(m)-satisfiable iff the restricted assignment

$$\mu^r = \bigwedge_i \Box_r \alpha_{ri} \wedge \bigwedge_j \neg \Box_r \beta_{rj} \quad (2)$$

is K(m)-satisfiable, for every  $r$ .  $\mu^r$  is K(m)-satisfiable iff the wff

$$\varphi^{rj} = \bigwedge_i \alpha_{ri} \wedge \neg \beta_{rj} \quad (3)$$

is K(m)-satisfiable, for every  $j$ . If no truth assignment is found which is K(m)-satisfiable, then  $\varphi$  is not K(m)-satisfiable (K(m)-unsatisfiable).

The two steps recurse until we get to a truth assignment with no modal atoms.

The algorithm is implemented by the function KSAT in Figure 1. KSAT takes in input a modal propositional wff  $\varphi$  and returns a truth value asserting whether  $\varphi$  is K(m)-satisfiable or not. KSAT invokes directly KSAT<sub>W</sub>(where “W” stands for “Wff”), passing as arguments  $\varphi$  and the truth value  $T$  (i.e., by (1), the empty truth assignment). KSAT<sub>W</sub> tries to build a K(m)-satisfiable truth assignment  $\mu$  satisfying  $\varphi$ . This is done recursively, according to the following steps:

like a phase transition [MSL92, WH94]. To our knowledge this is the first time that this phenomenon has been found in a modal logic.

The analysis presented in this paper builds on and takes to its conclusion the work preliminarily described in [GS96]. It improves on the previous material in three important aspects. Let us call  $\text{KSAT}_0$  the decision procedure presented in [GS96] (called  $\text{KSAT}$  in [GS96])<sup>6</sup>. First, the algorithm and its heuristics, are extended from dealing with a single modality to dealing with multiple modalities. Second, the implementation is improved.  $\text{KSAT}$  is much faster than  $\text{KSAT}_0$  (in our tests, up to two orders of magnitude, see Sections 4 and 5). This has been obtained essentially by adding an initial phase of wff preprocessing. Other — relatively minor — implementational variations can be understood by comparing the code of the two systems. Third, and more important, the testing in [GS96] was not exhaustive and only compared  $\text{KSAT}_0$  with  $\text{TABLEAU}$ . This made us miss some important points, and the phenomena described in Sections 4 and 5 went unnoticed. Furthermore, the increased efficiency of  $\text{KSAT}_0$  with respect to  $\text{TABLEAU}$  in [GS96] was wrongly motivated by the efficiency of the propositional decision procedure.  $\text{KSAT}$  and  $\text{KSAT}_0$  (when applied to a single modality) implement essentially the same algorithm.  $\text{KSAT}$  is only a more efficient implementation. The same applies to  $\text{KRIS}$  and  $\text{TABLEAU}$ . As the results in Section 5 show, the move from  $\text{KSAT}_0$  to  $\text{KSAT}$ , or from  $\text{TABLEAU}$  to  $\text{KRIS}$  causes an increase in efficiency, but it does not change the shape of the efficiency curves, as it happens in the move from  $\text{TABLEAU}$  to  $\text{KSAT}_0$  (or from  $\text{KRIS}$  to  $\text{KSAT}$ ).

## 2 The algorithm

Let us write  $\Box_r$  to mean the  $r$ -th modality. Let us call *atom* any wff which can not be decomposed propositionally, and *modal atom* any atom of the form  $\Box_r\psi$ . Let  $\varphi$  be the modal wff to be proved satisfiable. The algorithm for testing  $\text{K}(m)$ -satisfiability follows two basic steps, the first implementing the propositional reasoning, the second implementing the modal reasoning:

1. Using a decision procedure for propositional satisfiability, assign a truth value to (a subset of) the atoms occurring in  $\varphi$  in a way to make  $\varphi$  evaluate to  $T$ . Let us call *truth assignment (for  $\varphi$ )* the resulting set  $\mu$  of truth value assignments. We say that  $\mu$  *propositionally satisfies  $\varphi$* <sup>7</sup>. Then  $\mu$  is of the form

$$\begin{aligned} \mu = \{ & \Box_1\alpha_{11} = T, \Box_1\alpha_{12} = T, \dots, \Box_1\beta_{11} = F, \Box_1\beta_{12} = F, \dots, \\ & \dots \\ & \Box_m\alpha_{m1} = T, \Box_m\alpha_{m2} = T, \dots, \Box_m\beta_{m1} = F, \Box_m\beta_{m2} = F, \dots, \\ & A_1 = T, A_2 = T, \dots, A_{R+1} = F, A_{R+2} = F, \dots \} \end{aligned}$$

Notationally, from now on we write  $\mu$  as

$$\mu = \bigwedge_i \Box_1\alpha_{1i} \wedge \bigwedge_j \neg\Box_1\beta_{1j} \wedge \dots \wedge \bigwedge_i \Box_m\alpha_{mi} \wedge \bigwedge_j \neg\Box_m\beta_{mj} \wedge \gamma \quad (1)$$

---

<sup>6</sup> $\text{KSAT}_0$ , the test code and all the results presented in [GS96] are available via anonymous FTP at <ftp.mrg.dist.unige.it> in `pub/mrg-systems/ksat/ksat0/`.

<sup>7</sup>Notice that it is not necessary for a truth assignment to assign *all* the atoms of  $\varphi$ . For instance,  $\{\Box_1\psi_1 = T\}$  propositionally satisfies  $\Box_1\psi_1 \vee \Box_2\psi_2$ .

# 1 Introduction

The goal of this paper is to describe and thoroughly test a new decision procedure, called KSAT, checking satisfiability in the terminological logic  $\mathcal{ALC}$ , as defined in [SSS91], comprising Boolean operations on concepts and value restrictions, and not restricted to CNF form <sup>1</sup>.

As it is well known,  $\mathcal{ALC}$  is a notational variant of  $K(m)$ , that is,  $K$  with  $m$  modalities [Sch91] <sup>2</sup>. The main idea underlying the definition of KSAT is that a decision procedure for satisfiability in  $K(m)$  ( $K(m)$ -satisfiability) can be defined in terms of a decision procedure for propositional satisfiability (SAT). As a matter of terminology, we call *SAT-based* all the decision procedures whose definition is based on this idea <sup>3</sup>.

KSAT outperforms all the decision procedures and systems for terminological and modal logics we have been able to acquire. In this paper we compare KSAT with two of them. The first is a tableau-based procedure — due to B. Nebel and E. Franconi — which is essentially a straightforward implementation of the algorithm described in [HNSS90]. This procedure is called TABLEAU from now on <sup>4</sup>. The second is the state-of-the-art system KRIS described in [HNSS90] <sup>5</sup>. There are many reasons why a system can be more efficient than another. A crucial one is the “smartness” of the implementation. The implementation of KSAT we use is naive in many respects, e.g., it is in Lisp and it does not use fancy optimized data structures. We still have to push our work in this direction. KSAT is smarter than its competitors for a much more interesting reason. Both TABLEAU and KRIS are tableau-based. As Section 4 describes in detail, tableau-based decision procedures have an intrinsic weakness which makes it very hard if not impossible to be as efficient as SAT-based decision procedures. In our opinion, this is the most interesting theoretical result of this paper.

The paper is structured as follows. In Section 2 we present the algorithm implemented by KSAT. In Section 3 we briefly survey our test methodology, originally defined in [GS96]. This material is needed for a correct understanding of the experimental results reported later. In Section 4 we perform a comparative analysis of a first set of experimental results. This analysis allows us to show why SAT-based decision procedures are intrinsically more efficient than tableau-based decision procedures. Finally, in Section 5 we perform an exhaustive empirical analysis of KSAT and KRIS, that is, the fastest SAT-based and the fastest tableau-based decision procedure at our disposal. This allows us to confirm the analysis done in Section 4 and, looking at the KSAT results, to get a better understanding of where the hardest cases are. Among other things, this allows us to reveal what looks

---

<sup>1</sup>KSAT, the test code and all the results presented in this paper are available via anonymous FTP at `ftp.mrg.dist.unige.it` in `pub/mrg-systems/ksat/ksat1/`.

<sup>2</sup>In this paper we always refer to  $K(m)$  rather than to  $\mathcal{ALC}$ . In particular, we speak of wffs rather than concepts, modalities rather than roles, and so on.  $K(m)$ 's syntax is simpler than  $\mathcal{ALC}$ 's. Notice however that the current implementation of KSAT works with  $\mathcal{ALC}$ 's syntax.

<sup>3</sup>Although this is beyond the goals of this paper, it is worth noticing that this methodology is general, and can be extended to the other normal and (we think) non normal logics, following the methodology and results presented in [GS94, GSGF93] (but see also, e.g., [Fit83, Mas94]).

<sup>4</sup>TABLEAU is available via anonymous FTP at `ftp.mrg.dist.unige.it` in `pub/mrg-systems/tableau/`.

<sup>5</sup>KRIS is available via anonymous FTP at `ftp.dfki.uni-sb.de` in `/pub/tacos/KRIS/`.

# A *SAT*-based decision procedure for $\mathcal{ALC}$ \*

Fausto Giunchiglia<sup>1,2</sup> Roberto Sebastiani<sup>3</sup>

<sup>1</sup> IRST, 38050 Povo (TN), Italy. ph: ++39.461.314436

<sup>2</sup> DISA, via Inama 5, 38100 Trento

<sup>3</sup> DIST, v. Causa 13, 16146 Genoa, Italy. ph: ++39.10.3532811

fausto@irst.itc.it rseba@mrg.dist.unige.it

May 13, 1996

## Abstract

The goal of this paper is to describe and thoroughly test a decision procedure, called KSAT, checking satisfiability in the terminological logic  $\mathcal{ALC}$ . KSAT is said to be *SAT-based* as it is defined in terms of a decision procedure for propositional satisfiability (SAT). The tests are performed comparing KSAT with, among other procedures, KRIS, a state-of-the-art tableau-based implementation of a decision procedure for  $\mathcal{ALC}$ . KSAT outperforms KRIS of orders of magnitude. Furthermore, the empirical results highlight an intrinsic weakness that tableau-based decision procedures have with respect to SAT-based decision procedures.

---

\*Franz Baader, Marco Cadoli, Enrico Franconi, Enrico Giunchiglia, Fabio Massacci and Bernhard Nebel have given very useful feedback. Fabio Massacci has suggested testing the Halpern & Moses formulas. Marco Roveri has given technical assistance in the testing phase. All the members of the Mechanized Reasoning Group in Genoa have put up with many weeks of CPU-time background processes.



ISTITUTO PER LA RICERCA SCIENTIFICA E TECNOLOGICA

I 38100 TRENTO – LOC. PANTÉ DI POVO – TEL. 0461–814444

TELEX 400874 ITCRST – TELEFAX 0461–810851

A SAT-BASED DECISION PROCEDURE FOR ALC

Fausto Giunchiglia

Roberto Sebastiani

September 1996

Technical Report # 9607-08



ISTITUTO TARENTINO DI CULTURA