

PARTITIONING SEARCH SPACES OF A RANDOMIZED SEARCH

Antti E. J. Hyvärinen, Tommi Juntila and Ilkka Niemelä



TEKNILLINEN KORKEAKOULU
TEKNISKA HÖGSKOLAN
HELSINKI UNIVERSITY OF TECHNOLOGY
TECHNISCHE UNIVERSITÄT HELSINKI
UNIVERSITE DE TECHNOLOGIE D'HELSINKI

PARTITIONING SEARCH SPACES OF A RANDOMIZED SEARCH

Antti E. J. Hyvärinen, Tommi Juntila and Ilkka Niemelä

Helsinki University of Technology
Faculty of Information and Natural Sciences
Department of Information and Computer Science

Teknillinen korkeakoulu
Informaatio- ja luonnontieteiden tiedekunta
Tietojenkäsittelytieteen laitos

Distribution:

Helsinki University of Technology
Faculty of Information and Natural Sciences
Department of Information and Computer Science
P.O.Box 5400
FI-02015 TKK
FINLAND
URL: <http://ics.tkk.fi>
Tel. +358 9 470 01
Fax +358 9 470 23369
E-mail: series@ics.tkk.fi

© Antti E. J. Hyvärinen, Tommi Junttila and Ilkka Niemelä

ISBN 978-952-248-230-3 (Print)
ISBN 978-952-248-231-0 (Online)
ISSN 1797-5034 (Print)
ISSN 1797-5042 (Online)
URL: <http://lib.tkk.fi/Reports/2009/isbn9789522482310.pdf>

TKK ICS
Espoo 2009

ABSTRACT: This work studies the following question: given an instance of the propositional satisfiability problem, a randomized satisfiability solver, and a cluster of n computers, what is the best way to use the computers to solve the instance? Two approaches, simple distribution and search space partitioning as well as their combinations are investigated both analytically and empirically. It is shown that the results depend heavily on the type of the problem (unsatisfiable, satisfiable with few solutions, and satisfiable with many solutions) as well as on how good the search space partitioning function is. In addition, the behavior of a real search space partitioning function is evaluated in the same framework. The results suggest that in practice one should combine the simple distribution and search space partitioning approaches.

KEYWORDS: Constraint-based Search, Distributed Search, Randomized Search, Search-Space Partitioning, SAT Solvers

Contents

1	Introduction	7
2	Randomization and Simple Distribution	8
3	Partitioning	9
3.1	Unsatisfiable Instances	10
3.2	Satisfiable Instances with Many Solutions	11
3.3	Satisfiable Instances with One Solution	11
4	Composite Approaches	12
5	Implementing a Randomized Partitioning Function	13
6	Experimental Results on Partitioning	13
7	Conclusions	16
	References	16
A	Proofs	18
B	Experimental Results	21

1 INTRODUCTION

In this work¹ we develop distributed techniques for solving challenging instances of the propositional satisfiability problem (SAT). We are interested in using the best available SAT solvers as black-box subroutines or with little modification and in this way take advantage of the rapid development of SAT solver technology.

One of the interesting features in current state-of-the-art SAT solvers is that they use randomization and that their run times can vary significantly for a given instance. This opens up new opportunities for developing distributed solving techniques. The most straightforward idea is to employ a *simple distribution* approach where one just performs a number of independent runs using a randomized solver. This leads to surprising good speed-ups even when used in a grid environment with substantial communication and other delays [9]. The approach could be extended by applying particular restart strategies [13, 12] or using an algorithm portfolio scheme [5, 3]. Another key feature in modern SAT solvers is the use of conflict driven clause learning techniques. This feature can be exploited in the simple distribution approach and it has been shown that combining parallel learning schemes with a simple restart strategy leads to a powerful distributed SAT solving technique [10].

Another approach to developing parallel SAT solving techniques is based on *partitioning* the search space to multiple parts which can be handled in parallel. This can be achieved by constraint-based partitioning where the search space for a SAT instance \mathcal{F} is split to n derived instances $\mathcal{F}_1, \dots, \mathcal{F}_n$ by including additional constraints to \mathcal{F} . Typical implementation techniques include guiding paths [1, 14, 11] and scattering [8].

Both simple distribution and partitioning have their strengths. The former has led to surprisingly good performance but for really challenging SAT instances it provides no mechanism for splitting the search to more manageable portions to be treated in parallel. Search space partitioning techniques offer an approach to achieving this. However, the interaction between partitioning and randomized SAT solvers is poorly understood and the work aims to shed new light on this problem. It studies in detail combination of constraint-based partitioning and randomized SAT solvers, and provides an analysis on how an efficient and robust implementation can be achieved.

The rest of the work is structured as follows. Section 2 reviews briefly relevant key characteristics of modern randomized SAT solvers and the simple distribution approach. Section 3 studies analytically the expected run time of a plain partitioning approach where a SAT instance is partitioned and then a randomized SAT solver is used to solve the resulting instances. The section provides fundamental results for two limiting cases, for ideal and void partitioning functions. Section 4 extends the study to a setting where simple distribution and partitioning are mixed. Section 5 provides an implementation of a randomized partitioning function and Section 6 verifies the results briefly using experiments, and conclusions are given in Section 7.

¹This is an extended version of the paper to appear in the 11th Conference of the Italian Association for Artificial Intelligence (AI*IA 2009) [6] with two new appendices containing proofs for the propositions and additional experimental results.




Figure 1: The run time distributions of two instances for single (*the $q(t)$ plots*) and eight (*the $q_8(t)$ plots*) randomized SAT solvers.

2 RANDOMIZATION AND SIMPLE DISTRIBUTION

Most modern SAT solvers apply search *restarts* and some form of *randomization* to avoid getting stuck at hard subproblems [4]. For instance, MiniSat [2] version 1.14 restarts the search periodically and makes two percent of its branching decisions pseudo-randomly. Despite restarts and randomness, the run times of a SAT solver on an instance \mathcal{F} can vary significantly between some minimum t_{\min} and maximum t_{\max} (we assume that $t_{\min} > 0$ and t_{\max} is finite). Thus, we treat the run time of the solver on the instance as a random variable T and study the associated cumulative run-time distribution $q_T(t) = \Pr(T \leq t)$ (i.e. $q_T(t)$ is the probability that the instance is solved within t seconds) and its expected value $\mathbb{E}(T) = \int_{t_{\min}}^{t_{\max}} tq'(t)dt$. As an example, observe the run-time distribution $q(t)$ (approximated by one hundred sample runs) of an instance given in the left hand side plot of Fig. 1. Depending on the seed given to the pseudo-random number generator of MiniSat v1.14, the run time varies from less than a second to thousands of seconds.

This non-constant run time phenomenon can be exploited in a parallel environment by simply running n SAT solvers on the same instance \mathcal{F} in parallel and terminating the search when one of the solvers reports the solution. We call this approach *Simple Distributed SAT solving* (SDSAT) and denote its run time by the random variable T_{sdsat}^n . The cumulative run time distribution is now improved from $q_T(t)$ of the sequential case to $q_{T_{\text{sdsat}}^n}(t) = 1 - (1 - q_T(t))^n$. This approach can be surprisingly efficient. As an example, for the instance in the left hand side plot of Fig. 1 the expected run-time in the sequential case is $\mathbb{E}(T) \approx 623\text{s}$ while for eight solvers $\mathbb{E}(T_{\text{sdsat}}^8) \approx 31\text{s}$ (that is, around 20 times less). For a more detailed analysis of running SDSAT in a parallel, distributed environment involving communication and other delays, see [9].

Although the SDSAT approach can reduce the expected time to solve an instance, it cannot reduce it below the minimum run time t_{\min} . For an example, observe the sequential run time distribution $q(t)$ of another instance given in the right hand side plot of Fig. 1; the variation of the run time is significantly smaller and the instance seems to have no short run times. Consequently, running eight SAT solvers in parallel does not reduce the expected run time significantly; in numbers, $\mathbb{E}(T) \approx 2,065\text{s}$ while for eight

solvers $\mathbb{E}(T_{\text{sdsat}}^8) \approx 1,334s$ (i.e., only less than two times faster). Even more importantly, the *minimum run time stays the same irrespective of how many parallel solvers are employed*. As a summary, we can establish the following properties for the expected run time of the SDSAT approach²:

Proposition 1 $t_{\min} \leq \mathbb{E}(T_{\text{sdsat}}^n) \leq \mathbb{E}(T)$ for each $n \geq 1$. Furthermore, $\mathbb{E}(T_{\text{sdsat}}^n) \rightarrow t_{\min}$ when $n \rightarrow \infty$.

As we have seen, SDSAT can allow super-linear speedup (meaning $\mathbb{E}(T_{\text{sdsat}}^n) < \mathbb{E}(T)/n$) for instances with strongly varying run time. However, as the maximum speedup obtainable with SDSAT is $\mathbb{E}(T)/t_{\min}$, this can only happen for “smallish” values of n and for more than $\mathbb{E}(T)/t_{\min}$ solvers the speedup is guaranteed to be sub-linear.

3 PARTITIONING

The basic idea in the form of partitioning we use in this paper is quite simple: given a SAT instance \mathcal{F} and a positive integer n , use a *partitioning function* to compute a set $\mathcal{F}_1, \dots, \mathcal{F}_n$ of *derived SAT instances* such that

$$\mathcal{F} \equiv \mathcal{F}_1 \vee \dots \vee \mathcal{F}_n. \quad (1)$$

Now, in order to find whether \mathcal{F} is satisfiable or not, we solve, in parallel, all $\mathcal{F}_1, \dots, \mathcal{F}_n$ and deduce that \mathcal{F} is satisfiable if at least one of $\mathcal{F}_1, \dots, \mathcal{F}_n$ is. This method is called the *plain partitioning approach* in order to distinguish it from the composite approaches in Sect. 4. One way to implement partitioning functions is described in [8] (also see Sect. 5), where each \mathcal{F}_i is obtained from \mathcal{F} by conjoining it with a set of additional partitioning constraints.³ In addition to the requirement (1), partitioning functions often ensure that the models of $\mathcal{F}_1, \dots, \mathcal{F}_n$ are mutually disjoint.

Intuitively, the ideal case is that the partitioning function can partition the instance \mathcal{F} into n new instances $\mathcal{F}_1, \dots, \mathcal{F}_n$ so that each new instance \mathcal{F}_i is n times easier to solve than the original. That is, if the original instance \mathcal{F} has the cumulative run time distribution $q_T(t)$, then the distribution of each \mathcal{F}_i is $q_{T_i}(t) = q_T(nt)$. In this case we say that the partition function is *ideal* for the instance. As obtaining ideal partitioning functions can be difficult, we also consider the case of a *void* partitioning function where the partitioning fails totally, resulting in new instances which are as hard to solve as the original, i.e. have the same distribution $q_{T_i}(t) = q_T(t)$. This is a realistic scenario because modern, clause-learning SAT solvers, such as MiniSat, use sophisticated heuristics in the search: it is possible that values of certain variables are practically never considered. If the partition function constrains only these irrelevant variables, the difficulty of the instance does not decrease, and thus such a function is void.

In this section, we give an analytic study of the efficiency of the plain partitioning approach, under both ideal and void functions, when the fact that

²For proofs, see Appendix A.

³As explained in [7], guiding paths [1, 14] can also be interpreted as partitioning constraints.

the SAT solver is randomized is taken into account. As the efficiency depends heavily on the satisfiability of the instance, we consider three cases: unsatisfiable instance, a satisfiable instance with many solutions, and a satisfiable instance with a unique solution. We have also simulated the plain partitioning approach on run time distributions of some real SAT instances; some results are given later in Sect. 4 after some composite approaches mixing simple distribution and plain partitioning have been described. A real partitioning function is considered in Sect. 5.

3.1 Unsatisfiable Instances

Assume that an unsatisfiable instance \mathcal{F} is partitioned into n new instances $\mathcal{F}_1, \dots, \mathcal{F}_n$ fulfilling Eq. (1). All new instances need to be shown unsatisfiable to deduce that \mathcal{F} is unsatisfiable. When performed in parallel, this corresponds to waiting for the termination of the “unluckiest” run.

In the case of ideal partitioning function, each new instance \mathcal{F}_i is n times easier to solve than the original \mathcal{F} , having run time distributions $q_{T_i}(t) = q_T(nt)$. We denote the random variable capturing the run time of the resulting plain partitioning approach under an ideal partitioning function by $T_{\text{part(ideal)}}^n$. As all the new instances have to be solved (in parallel), the corresponding run time distribution is $q_{T_{\text{part(ideal)}}^n}(t) = q(nt)^n$. Based on this, we have the following interesting results⁴. First, ideal partitioning functions can provide *at most linear expected speed-up* on unsatisfiable instances:

Proposition 2 $\mathbb{E}(T_{\text{part(ideal)}}^n) \geq \mathbb{E}(T)/n$ for each $n \geq 1$.

In fact, it can be shown that linear speed-up can only be obtained on instances that have a constant run time distribution, i.e. when $t_{\min} = t_{\max}$. However, the expected run time is never worse than that of solving the original instance with one solver:

Proposition 3 $\mathbb{E}(T_{\text{part(ideal)}}^n) \leq \mathbb{E}(T)$ for each $n \geq 1$.

When the number n of SAT solvers run in parallel is increased, the expected run time $\mathbb{E}(T_{\text{part(ideal)}}^n)$ approaches t_{\max}/n , i.e., linear speed-up w.r.t. the *maximum* run time. Plain partitioning with ideal partitioning functions and simple distribution cannot be totally ordered; there are distributions for which $\mathbb{E}(T_{\text{sdsat}}^n) < \mathbb{E}(T_{\text{part(ideal)}}^n)$ and others for which $\mathbb{E}(T_{\text{part(ideal)}}^n) < \mathbb{E}(T_{\text{sdsat}}^n)$ when n is smallish⁵. However, as $\mathbb{E}(T_{\text{part(ideal)}}^n) \leq t_{\max}/n$ and $\mathbb{E}(T_{\text{sdsat}}^n) \geq t_{\min}$, we have that $\mathbb{E}(T_{\text{part(ideal)}}^n) < \mathbb{E}(T_{\text{sdsat}}^n)$ for sufficiently large n .

Let us next consider the case of a void partitioning function, i.e. the case when the partitioning fails so that the run time distribution $q_{T_i}(t)$ of each new instance \mathcal{F}_i is equal to $q_T(t)$ of the original instance \mathcal{F} . We denote by $T_{\text{part(void)}}^n$ the run time of the resulting plain partitioning approach. As all \mathcal{F}_i have to be solved, the run time distribution of $T_{\text{part(void)}}^n$ is $q_{T_{\text{part(void)}}^n}(t) = q_{T_i}(t)^n = q_T(t)^n$. From this it follows that for unsatisfiable instances *it is not possible to obtain any speedup with void functions*:

⁴For proofs, see Appendix A.

⁵For the case $\mathbb{E}(T_{\text{part(ideal)}}^n) < \mathbb{E}(T_{\text{sdsat}}^n)$, see Fig. 2. For the case $\mathbb{E}(T_{\text{sdsat}}^n) < \mathbb{E}(T_{\text{part(ideal)}}^n)$ see Appendix B.

Proposition 4 $\mathbb{E}(T_{\text{part}(\text{void})}^n) \geq \mathbb{E}(T)$ for each $n \geq 1$.

In fact, *the more resources one uses, the closer to the maximum run time one gets*: $\mathbb{E}(T_{\text{part}(\text{void})}^n) \rightarrow t_{\max}$ when $n \rightarrow \infty$.

3.2 Satisfiable Instances with Many Solutions

We next consider the case when a satisfiable SAT instance \mathcal{F} is partitioned into n new instances $\mathcal{F}_1, \dots, \mathcal{F}_n$ fulfilling Eq. (1). In order to deduce that \mathcal{F} is satisfiable, it is enough to show that *any* of the new instances is satisfiable. In this section we assume that each new instance \mathcal{F}_i is satisfiable, postponing the case where only one is satisfiable to the next section.

Let us consider the case of ideal partitioning function first. Again, we denote the random variable describing the run time of the resulting plain partitioning approach by $T_{\text{part}(\text{ideal})}^n$. As the probability that none of the n solvers has solved the associated new instance within time t is $(1 - q_T(nt))^n$, run time distribution of $T_{\text{part}(\text{ideal})}^n$ is $q_{T_{\text{part}(\text{ideal})}^n}(t) = 1 - (1 - q_T(nt))^n$. Several interesting properties follow from this. First, with n parallel solvers, *the expected run time is n times smaller than that of the Simple Distributed SAT*: $\mathbb{E}(T_{\text{part}(\text{ideal})}^n) = \mathbb{E}(T_{\text{sdsat}}^n)/n$. Therefore, when compared to solving the original instance, we notice that *on satisfiable instances with many solutions we may expect at least linear speed-up*:

Proposition 5 $\mathbb{E}(T_{\text{part}(\text{ideal})}^n) \leq \mathbb{E}(T)/n$ for each $n \geq 1$.

When the number n of parallel SAT solvers is increased, $\mathbb{E}(T_{\text{part}(\text{ideal})}^n)$ approaches t_{\min}/n . Thus, one can in principle obtain almost linear speed-up w.r.t. *the minimum run time*.

In the case of a void partitioning function the run time of each new instance is the same as that of the original. As each new instance is assumed to be satisfiable, solving any of them is enough to deduce the satisfiability of the original instance. Therefore, *for satisfiable instances with many solutions, the plain partitioning approach with a void partitioning function effectively reduces to Simple Distributed SAT*:

Proposition 6 $\mathbb{E}(T_{\text{part}(\text{void})}^n) = \mathbb{E}(T_{\text{sdsat}}^n)$.

3.3 Satisfiable Instances with One Solution

When a satisfiable instance \mathcal{F} with only one satisfying truth assignment is partitioned into n new instances $\mathcal{F}_1, \dots, \mathcal{F}_n$, it is likely that only one of the new instances is satisfiable while the others are unsatisfiable. Therefore, the satisfiable new instance has to be solved to deduce that \mathcal{F} is satisfiable. The problem is that it is not known which of the new instances this is.

In the case of ideal partitioning function, the run time of the satisfiable new instance is n times smaller than that of the original instance. Therefore, if all the n new instances are solved in parallel and the solving is terminated as soon as the satisfiable new instance is solved, *linear speed-up is obtained with an ideal partitioning function*:

Proposition 7 $\mathbb{E}(T_{\text{part(ideal)}}^n) = \mathbb{E}(T)/n$.

In the case of a void partitioning function, the run time of the satisfiable new instance is the same as that of the original instance. Thus, *using a void partitioning function results neither in speed-up nor in loss of efficiency*:

Proposition 8 $\mathbb{E}(T_{\text{part(void)}}^n) = \mathbb{E}(T)$.

4 COMPOSITE APPROACHES

The analysis of the previous section shows that the plain partitioning approach can potentially obtain even super-linear speed-ups, whereas an improper, void implementation can by Prop. 4 result in worse expected run time than that of one solver. The two approaches presented here aim at being at least as efficient as solving the instance with one solver. Assume that we have resources to run n SAT solvers in parallel and consider the following approaches that mix simple distribution and plain partitioning.

- *Repeated Partitioning.* In this approach, we run in parallel $k = \lfloor \sqrt{n} \rfloor$ copies of the plain partitioning approach, each copy splitting the instance \mathcal{F} into k new instances $\mathcal{F}_{i,1}, \dots, \mathcal{F}_{i,k}$ and solving each $\mathcal{F}_{i,j}$ once. We denote the random variable describing the run time of this approach by $T_{\text{rep-part}}^n$.
- *Safe Partitioning.* This approach reverses the order of SDSAT and partitioning compared to repeated partitioning: the instance \mathcal{F} is partitioned into $k = \lfloor \sqrt{n} \rfloor$ new instances $\mathcal{F}_1, \dots, \mathcal{F}_k$ and each new instances \mathcal{F}_i is solved with k SAT solvers in parallel. The run time of this is denoted by the random variable $T_{\text{safe-part}}^n$.

Unfortunately, when using repeated partitioning on an unsatisfiable instance and the partitioning function is void, the experiments suggest that $\mathbb{E}(T) \leq \mathbb{E}(T_{\text{rep-part}}^n)$ and that $\mathbb{E}(T_{\text{rep-part}}^n) \rightarrow t_{\max}$ when $n \rightarrow \infty$. However, safe partitioning (i) is as good as repeated partitioning (i.e. $\mathbb{E}(T_{\text{safe-part}}^n) = \mathbb{E}(T_{\text{rep-part}}^n)$) on satisfiable instances⁶ and we conjecture that safe partitioning is at least as good (i.e. $\mathbb{E}(T_{\text{safe-part}}^n) \leq \mathbb{E}(T_{\text{rep-part}}^n)$) on unsatisfiable ones when the same partition function is applied; (ii) is equal to SDSAT on satisfiable instances with many solutions when the partition function is void⁷; and (iii) seems experimentally at least as fast as solving the original instance with one solver, i.e., $\mathbb{E}(T_{\text{safe-part}}^n) \leq \mathbb{E}(T)$ even when the instance is unsatisfiable and the partitioning function is void.

To illustrate the approaches and results presented in Sects. 2, 3, and 4, Fig. 2 shows the expected run times of different approaches when applied to the same instances as in Fig. 1. As the left hand side instance is satisfiable with many solutions, the “sdsat+others” line depicts the behavior of the SDSAT approach as well as all the considered partitioning approaches when a void partitioning function is applied. The instance at the right hand side is unsatisfiable.

⁶See Proposition 9 in Appendix A.

⁷See Proposition 10 in Appendix A; also compare to Proposition 11 in Appendix A concerning instances with one solution.

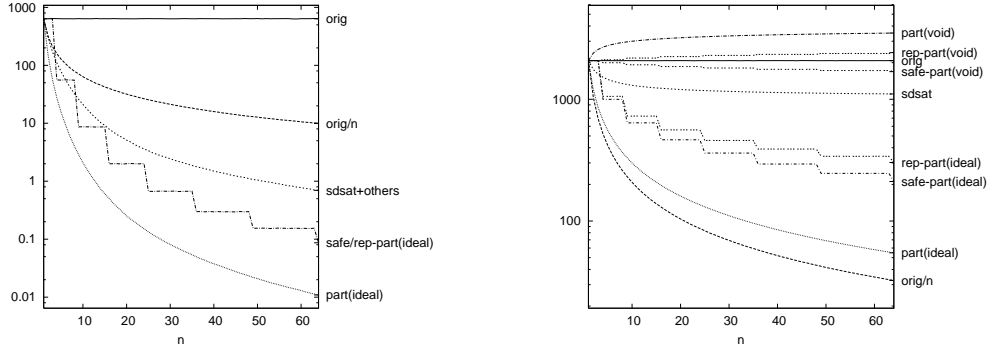


Figure 2: Expected run times (in seconds) of different approaches on the instances in Fig. 1 when the number n of SAT solvers run in parallel is varied.

5 IMPLEMENTING A RANDOMIZED PARTITIONING FUNCTION

This section briefly describes how we implemented a partitioning function called *scattering* [8]. The implementation is based on MiniSat 1.14, constructs partitions having pairwise disjoint models, and is randomized so that the partitions differ depending on the random seed given as input to the function.

The function works in two phases. First, the function simply runs MiniSat as is to obtain heuristic values for the Boolean variables in the instance. The first phase ends when the instance is solved or a fixed time limit (currently 300 seconds) is reached. If the instance was not solved, the function enters the second phase, where the derived instances $\mathcal{F}_1, \dots, \mathcal{F}_n$ are constructed from \mathcal{F} by adding constraints. For each $1 \leq i \leq n - 1$, the scattering function uses the obtained heuristic values to select a conjunction of d_i literals, $C_i = l_{i,1} \wedge \dots \wedge l_{i,d_i}$, and conjoins it to the derived instance \mathcal{F}_i . The function ensures that no derived instances share models by inserting the negation ($\neg l_{i,1} \vee \dots \vee \neg l_{i,d_i}$) of the conjunction C_i in \mathcal{F}_i to each derived instance \mathcal{F}_j with $i < j \leq n$. Scattering is a partition function, since the last instance contains only the negations of the conjunctions corresponding to the “remaining” search space. Finally, the randomization of the scattering function follows naturally from that of MiniSat: the derived instances depend on the random seed passed to MiniSat.

6 EXPERIMENTAL RESULTS ON PARTITIONING

The following experiments study the behavior of the presented approaches both under ideal and real (scattering) partitioning functions on some real-world SAT instances. A summary of the results is presented in Table 1. The instances represent hard SAT formulas in the sense that their randomized run times often exceed one hour. Furthermore, `cube-11-h14-sat` is a satisfiable instance where the scattering function always resulted in a unique satisfiable derived instance, `dated-10-13-s` is a satisfiable instance where the scattering function always resulted in several satisfiable derived instances (unless solved by the function), and `AProVE07-09` is an unsatisfiable instance. The first column, labeled T , reports the sequential run-time distribution of the

instances. The next two columns, labeled "simple distribution" in the table, report the results for Simple Distributed SAT solving for eight and 64 resources. The next four columns, labeled "ideal partitioning", report run-time distributions when an ideal partitioning function is used: first for the plain partitioning approach with eight and 64 resources, and then for safe and repeated partitioning approaches. The last four columns, labeled "scattering", report the run time distributions obtained when scattering (recall Sect. 5) is used as the partitioning function; first for plain partitioning with eight and 64 resources, and then for safe and repeated partitioning approaches with 64 resources. The rows of the table report the expected, minimum, median and maximum run times together with the first and third quartile (the values of t such that $q(t) \leq 0.25$ and $q(t) \leq 0.75$).

The run-time distributions for SDSAT and "ideal partitioning" approaches were obtained by solving the instance one hundred times with MiniSat 1.14. The resulting distributions were then used to compute the results analytically. None of the results include delays associated with parallel environments.

The distributions in columns T_{scatter}^8 and T_{scatter}^{64} are obtained by running the plain partitioning approach with the scattering function fifty times using different random seeds. The resulting distribution was directly used to compute the values for the repeated partitioning approach ($T_{\text{rep-part}}^{64}$) under "scattering". To compute the results for the column $T_{\text{safe-part}}^{64}$ under "scattering", each derived instance was solved seven more times, thereby directly simulating an implementation of safe partitioning with scattering. The run times do not include the time required to run the scattering function. If the instance was solved while scattering, the run time is reported as zero.

The results in "simple distribution" columns show good scalability for dated-10-13-s and moderate scalability for other instances, as predicted by analytical results when t_{\min} is close to $\mathbb{E}(T)$. The columns under "ideal partitioning" show that partitioning can in theory result in even better speed-up for these instances. Surprisingly, in the actual implementation ($T_{\text{scatter}}^8, T_{\text{scatter}}^{64}$) we see that plain scattering results in higher expected run-times than simple distribution for these instances. This reflects the difficulty of obtaining ideal partitioning functions.

Comparison of the "ideal partitioning" approaches confirms the discussion in Sect. 4. In particular, safe partitioning results in lower expected run time than repeated partitioning for unsatisfiable instances. However, the results under "scattering" show the opposite; repeated partitioning has consistently lower expected run time than safe partitioning. For example, observe the expected run times for safe and repeated partitioning approaches for the instance cube-11-h14-sat with unique satisfiable derived instance: in "ideal partitioning" they are equal, whereas the scattering-based safe partitioning is significantly worse than the scattering-based repeated partitioning approach. To study this, we computed run-time distributions for some of the satisfiable derived instances (not shown in the table), and it turns out that their expected run times varied between 109.1 and 4,773 seconds. Thus the hardness (expected run time) of a derived instance produced by scattering is also a random variable with possibly a very large range, and running the scattering function independently several times increases the probability of finding derived instances with low expected run times. This explains

Table 1: Comparing approaches for parallel search

T	simple distrib.		ideal partitioning				scattering				
	T_{sdsat}^8	T_{sdsat}^{64}	$T_{\text{part(ideal)}}^8$	$T_{\text{part(ideal)}}^{64}$	$T_{\text{safe-part}}^{64}$	$T_{\text{rep-part}}^{64}$	T_{scatter}^8	T_{scatter}^{64}	$T_{\text{safe-part}}^{64}$	$T_{\text{rep-part}}^{64}$	
				cube-11-h14-sat							
Exp	3,110	2,685	604,0	75.50	388.7	388.7	3,537	3,378	2,265	839.6	
Min	2,629	2,629	328.6	41.07	328.6	328.6	117.5	69.49	13.93	117.5	
Q_1	2,748	2,629	455.1	56.89	343.5	343.5	2,291	1,193	1,665	141.7	
Med	3,009	2,640	572.7	72.83	376.1	376.1	3,459	3,473	2,381	338.2	
Q_3	3,362	2,741	716.3	89.53	420.2	420.2	4,662	5,288	3,182	1,719	
Max	10,050	10,050	1,256	157.0	1,256	1,256	11,500	7,199	4,405	11,500	
				dated-10-13-s							
Exp	128.2	16.45	16.02	0.2570	2.056	2.056	261.2	317.3	21.22	0.2566	
Min	10.09	10.09	1.262	0.1577	1.262	1.262	0	0	0	0	
Q_1	28.00	10.09	3.499	0.1577	1.262	1.262	0	0	0	0	
Med	784.7	109.4	13.67	0.1653	1.322	1.322	19.86	9.858	7.037	0	
Q_3	2,093	181.4	22.68	0.2755	2.204	2.204	227.7	84.70	17.77	0	
Max	37,930	37,930	4,741	529.6	4,741	4,741	6,449	10,095	172.0	6,449	
				AProVE07-09							
Exp	4,016	2,361	759.7	117.7	392.5	586.7	2,598	1,719	1,632	1,559	
Min	1,552	1,552	194.1	24.26	194.1	194.1	1,261	486.3	723.9	1,261	
Q_1	3,086	2,033	668.0	106.2	352.3	554.2	1,757	1,122	1,271	1,466	
Med	3,905	2,389	770.6	110.3	396.4	581.4	2,267	1,437	1,637	1,567	
Q_3	4,732	2,666	843.7	129.8	414.3	606.4	3,550	1,894	1,912	1,694	
Max	9,302	9,302	1,163	145.3	1,163	1,163	4,539	6,617	2,968	4,539	

the good speed-up obtained by repeated partitioning when compared to safe partitioning.

7 CONCLUSIONS

The paper investigates distributed techniques for solving challenging SAT instances and focuses on combining constraint-based search space partitioning with randomized SAT solving techniques. The paper studies first analytically the expected run time of a plain partitioning approach where a SAT instance is partitioned and then a randomized SAT solver is used to solve the resulting instances. Analytical results are derived for two limiting cases, for ideal and void partitioning functions. The investigation is then extended to a setting where simple distribution and partitioning are mixed. Finally the paper proposes a randomized partitioning function and compares the function against the ideal case.

The analytical results show that partitioning can potentially lead to catastrophic failures where an increase in computing resources leads to a decrease in solving efficiency for unsatisfiable instances. The empirical results show in part that a good implementation is usually able to avoid this failure, but plain partitioning can nevertheless be worse than an approach based on simple distributed SAT solving (SDSAT). Both problems are avoided in practice with safe and repeated partitioning. The experimental and analytical comparisons show an interesting relationship between the safe and repeated partitioning approaches, suggesting that an ideal partitioning function would profit from safe partitioning whereas randomness in the partitioning function can be better exploited with repeated partitioning.

ACKNOWLEDGMENTS

The authors are grateful for the financial support of the Academy of Finland (projects 122399 and 112016), Helsinki Graduate School in Computer Science and Engineering, Jenny and Antti Wihuri Foundation, Emil Aaltosen Säätiö, Finnish Foundation for Technology Promotion, and Technology Industries of Finland Centennial Foundation.

REFERENCES

- [1] Max Böhm and Ewald Speckenmeyer. A fast parallel SAT-solver: Efficient workload balancing. *Annals of Mathematics and Artificial Intelligence*, 17(4–3):381–400, 1996.
- [2] Niklas Eén and Niklas Sörensson. An extensible SAT-solver. In *Proc. SAT’03*, volume 2919 of *LNCS*, pages 502–518. Springer, 2004.
- [3] Carla P. Gomes and Bart Selman. Algorithm portfolios. *Artificial Intelligence*, 126(1–2):43–62, 2001.

- [4] Carla P. Gomes, Bart Selman, Nuno Crato, and Henry A. Kautz. Heavy-tailed phenomena in satisfiability and constraint satisfaction problems. *Journal of Automated Reasoning*, 24(1/2):67–100, 2000.
- [5] Bernardo A. Huberman, Rajan M. Lukose, and Tad Hogg. An economics approach to hard computational problems. *Science*, 275(5296):51–54, 1997.
- [6] Antti Hyvärinen, Tommi Junttila, and Ilkka Niemelä. Partitioning search spaces of a randomized search. In *Proc. AI*IA, LNCS*. Springer, 2009. To appear.
- [7] Antti E. J. Hyvärinen. Approaches to grid-based SAT solving. Research Report TKK-ICS-R16, Helsinki University of Technology, June 2009.
- [8] Antti E. J. Hyvärinen, Tommi Junttila, and Ilkka Niemelä. A distribution method for solving SAT in grids. In *Proc. SAT'06*, volume 4121 of *LNCS*, pages 430–435. Springer, 2006.
- [9] Antti E. J. Hyvärinen, Tommi Junttila, and Ilkka Niemelä. Strategies for solving SAT in grids by randomized search. In *Proc. AISC'08*, volume 5144 of *LNAI*, pages 125–140. Springer, 2008.
- [10] Antti E. J. Hyvärinen, Tommi Junttila, and Ilkka Niemelä. Incorporating clause learning in grid-based randomized SAT solving. *Journal on Satisfiability, Boolean Modeling and Computation*, 6:223–244, 2009.
- [11] B. Jurkowiak, C. Li, and G. Utard. A parallelization scheme based on work stealing for a class of SAT solvers. *Journal of Automated Reasoning*, 34(1):73–101, 2005.
- [12] Michael Luby and Wolfgang Ertel. Optimal parallelization of Las Vegas algorithms. In *Proc. STACS'94*, volume 775 of *LNCS*, pages 463–474. Springer, 1994.
- [13] Michael Luby, Alistair Sinclair, and David Zuckerman. Optimal speedup of Las Vegas algorithms. *Information Processing Letters*, 47(4):173–180, 1993.
- [14] H. Zhang, M. Bonacina, and J. Hsiang. PSATO: A distributed propositional prover and its application to quasigroup problems. *Journal of Symbolic Computation*, 21(4):543–560, 1996.

A PROOFS

Proposition 1 $t_{\min} \leq \mathbb{E}(T_{\text{sdsat}}^n) \leq \mathbb{E}(T)$ for each $n \geq 1$. Furthermore, $\mathbb{E}(T_{\text{sdsat}}^n) \rightarrow t_{\min}$ when $n \rightarrow \infty$.

Proof Recall that by definition $q_T(t) = 0$ for all $t < t_{\min}$ and $q_T(t) = 1$ for all $t \geq t_{\max}$. By applying integration by parts, i.e. $\int_a^b f(t)g'(t)dt = f(b)g(b) - f(a)g(a) - \int_a^b f'(t)g(t)dt$, the expected value given by $\mathbb{E}(T) = \int_0^{t_{\max}} tq_T'(t)dt$ can be rewritten as

$$\mathbb{E}(T) = t_{\max} - \int_0^{t_{\max}} q_T(t)dt. \quad (2)$$

Recall that $q_{T_{\text{sdsat}}^n}(t) = 1 - (1 - q_T(t))^n$ and thus $q_{T_{\text{sdsat}}^n}(t) = 0$ if $t < t_{\min}$, $0 < q_{T_{\text{sdsat}}^n}(t) \leq 1$ if $t_{\min} \leq t < t_{\max}$, and $q_{T_{\text{sdsat}}^n}(t) = 1$ if $t \geq t_{\max}$. Now $\int_0^{t_{\max}} q_{T_{\text{sdsat}}^n}(t)dt \leq t_{\max} - t_{\min}$ and therefore $\mathbb{E}(T_{\text{sdsat}}^n) = t_{\max} - \int_0^{t_{\max}} q_{T_{\text{sdsat}}^n}(t)dt \geq t_{\min}$.

As $0 \leq q_T(t) \leq 1$ and $(1 - q_T(t))^n \leq (1 - q_T(t))$ for all $n \geq 1$, we have that $q_{T_{\text{sdsat}}^n}(t) = 1 - (1 - q_T(t))^n \geq q_T(t)$. This implies that $\int_0^{t_{\max}} q_{T_{\text{sdsat}}^n}(t)dt \geq \int_0^{t_{\max}} q_T(t)dt$ and thus $\mathbb{E}(T_{\text{sdsat}}^n) \leq \mathbb{E}(T)$.

When $n \rightarrow \infty$, we have that $1 - (1 - q_T(t))^n \rightarrow 1$ for all $t_{\min} \leq t \leq t_{\max}$. Thus $\int_0^{t_{\max}} q_{T_{\text{sdsat}}^n}(t)dt \rightarrow t_{\max} - t_{\min}$ and $\mathbb{E}(T_{\text{sdsat}}^n) = t_{\max} - \int_0^{t_{\max}} q_{T_{\text{sdsat}}^n}(t)dt \rightarrow t_{\min}$. ■

Proposition 2 $\mathbb{E}(T_{\text{part(ideal)}}^n) \geq \mathbb{E}(T)/n$ for each $n \geq 1$.

Proof By definition, the probability that an ideally partitioned instance is solved in time t or less is $q(t) = q_T(nt)$. Since the instance is unsatisfiable, all n instances need to be solved by the partitioning approach. Therefore, for unsatisfiable instances, the probability that the plain partitioning approach has solved the instance is $q_{T_{\text{part(ideal)}}^n}(t) = q_T(nt)^n$. By (2), the expected run time of the ideal partitioning approach is

$$\mathbb{E}(T_{\text{part(ideal)}}^n) = \frac{t_{\max}}{n} - \int_0^{t_{\max}/n} q_T(nt)^n dt = \frac{t_{\max}}{n} - \int_0^{t_{\max}} \frac{1}{n} q_T(\tau)^n d\tau,$$

where the last equality is obtained by substituting $t = \tau/n$ and $dt = \frac{d\tau}{n}$ in the middle equation.

Since $q_T(t) \leq 1$ for all $t \geq 0$, we have $q_T(t)^n \leq q_T(t)$ when $n \geq 1$. Thus

$$\mathbb{E}(T) = t_{\max} - \int_0^{t_{\max}} q_T(t)dt \leq t_{\max} - \int_0^{t_{\max}} q_T(t)^n dt = n\mathbb{E}(T_{\text{part(ideal)}}^n).$$

Note that the probability that one of the parallel solvers ends up running an execution corresponding to the highest run time t_{\max}/n approaches 1 as $n \rightarrow \infty$. Therefore, $\mathbb{E}(T_{\text{part(ideal)}}^n) \rightarrow \frac{t_{\max}}{n}$ when $n \rightarrow \infty$, that is, speed-up of the expected run time of the ideal partitioning approach tends to t_{\max}/n . Hence, linear speed-up of the expected run time is only obtained when $t_{\max} = t_{\min}$. ■

Proposition 3 $\mathbb{E}(T_{\text{part(ideal)}}^n) \leq \mathbb{E}(T)$ for each $n \geq 1$

Proof The proposition can be proved by induction on n . When $n = 1$, $\mathbb{E}(T_{\text{part(ideal)}}^n) = \mathbb{E}(T)$. Assume that the claim holds for some $n \geq 1$. To prove that the claim holds for $n + 1$, we express the expected run time using the derivative of the cumulative run time distribution. As before, $q_T(t)$ denotes the probability that the original instance is solved at or before time t . Then $q_T(nt)$ denotes the probability that a partition produced by the ideal partitioning function from the original instance is solved at or before time t . By definition, the expected run time of the plain partitioning approach for unsatisfiable instances with ideal partitioning function is

$$\mathbb{E}(T_{\text{part(ideal)}}^n) = \int_0^{t_{\max}} t q'_{T_{\text{part(ideal)}}^n}(t) dt,$$

where $q'_{T_{\text{part(ideal)}}^n}(t) = \frac{d}{dt} q_T(nt)^n = n^2 q_T(nt)^{n-1} q'_T(nt)$. Using this and noting that $q'_T(nt) = 0$ when $t \geq t_{\max}/n$, we get

$$\mathbb{E}(T_{\text{part(ideal)}}^n) = \int_0^{t_{\max}/n} t n^2 q_T(nt)^{n-1} q'_T(nt) dt.$$

By substituting $nt = \tau$ above, we get

$$\begin{aligned} \mathbb{E}(T_{\text{part(ideal)}}^n) &= \int_0^{t_{\max}} \frac{\tau}{n} n^2 q_T(\tau)^{n-1} q'_T(\tau) \frac{1}{n} d\tau = \int_0^{t_{\max}} \tau q_T(\tau)^{n-1} q'_T(\tau) d\tau \\ &\geq \int_0^{t_{\max}} \tau q_T(\tau)^n q'_T(\tau) d\tau = \mathbb{E}(T_{\text{part(ideal)}}^{n+1}), \end{aligned}$$

where the inequality follows from $q_T(\tau) \leq 1$ for all $0 \leq \tau \leq t_{\max}$. \blacksquare

Proposition 9 $\mathbb{E}(T_{\text{safe-part}}^n) = \mathbb{E}(T_{\text{rep-part}}^n)$ for satisfiable instances and ideal partitioning function.

Proof Recall that when given n computing resources and letting $k = \lfloor \sqrt{n} \rfloor$,

- the safe partitioning approach partitions the original instance into k instances and solves each resulting instance with k solvers in parallel, and
- the repeated partitioning approach performs, in parallel, k independent plain partitioning approaches over k resources.

To clarify the proofs, we introduce two operators, D_{sdsat}^k and D_{part}^k , which describe respectively the effects of the SDSAT and plain partitioning approaches to a run time distribution $q(t)$. Note that the definition covering the partitioning approach is more general than what is strictly required for the proofs, since it includes a function $\epsilon(k)$ capturing the effect of the partitioning to the run time: $\epsilon(k) = k$ for ideal partitioning function, and $\epsilon(k) = 1$ for void partitioning function. The operators D_{sdsat}^k and D_{part}^k map a distribution function to another distribution function so that

$$D_{\text{sdsat}}^k(q(t)) = 1 - (1 - q(t))^k$$

is the run time distribution resulting from applying the SDSAT approach to a solving process with run time distribution $q(t)$, and

$$D_{\text{part}}^k(q(t)) = \begin{cases} q(\epsilon(k)t)^k & \text{if all partitions are unsatisfiable, and} \\ 1 - (1 - q(\epsilon(k)t))^r & \text{if } r \text{ of the partitions are satisfiable} \end{cases}$$

is the run time distribution resulting from applying the ideal partitioning approach to a solving process with run time distribution $q(t)$.

We are now ready to give the proofs. In safe partitioning, an instance is partitioned once and all resulting partitions are solved with the SDSAT approach. Therefore the distribution for the unsatisfiable case is

$$D_{\text{part}}^k(D_{\text{sdsat}}^k(q(t))) = (1 - (1 - q(\epsilon(k)t))^k)^k,$$

and for the case where r of the k partitions are satisfiable

$$D_{\text{part}}^k(D_{\text{sdsat}}^k(q(t))) = 1 - (1 - (1 - (1 - q(\epsilon(k)t))^k))^r.$$

When the satisfiable case is simplified, we have

$$D_{\text{part}}^k(D_{\text{sdsat}}^k(q(t))) = \begin{cases} (1 - (1 - q(\epsilon(k)t))^k)^k & \text{if unsatisfiable, and} \\ 1 - (1 - q(\epsilon(k)t))^{kr} & \text{if satisfiable,} \end{cases} \quad (3)$$

where r is the number of satisfiable partitions resulting from the partitioning ($1 \leq r \leq k$). In repeated partitioning an instance is partitioned k times and each instance is solved with a single solver. Therefore the distribution for unsatisfiable instances is

$$D_{\text{sdsat}}^k(D_{\text{part}}^k(q(t))) = 1 - (1 - q(\epsilon(k)t))^k$$

and for satisfiable instances with r satisfiable partitions

$$D_{\text{sdsat}}^k(D_{\text{part}}^k(q(t))) = 1 - (1 - (1 - (1 - q(\epsilon(k)t))^r))^k.$$

When the satisfiable case is simplified, we have as a summary

$$D_{\text{sdsat}}^k(D_{\text{part}}^k(q(t))) = \begin{cases} 1 - (1 - q(\epsilon(k)t))^k & \text{if unsatisfiable, and} \\ 1 - (1 - q(\epsilon(k)t))^{rk} & \text{if satisfiable,} \end{cases} \quad (4)$$

where again $1 \leq r \leq k$. Therefore, if the instance is satisfiable, we see directly from (3) and (4) that $D_{\text{sdsat}}^k(D_{\text{part}}^k(q(t))) = D_{\text{part}}^k(D_{\text{sdsat}}^k(q(t)))$, which also means that $\mathbb{E}(T_{\text{safe-part}}^n) = \mathbb{E}(T_{\text{rep-part}}^n)$ follows.

Proposition 10 $\mathbb{E}(T_{\text{safe-part}}^n) = \mathbb{E}(T_{\text{rep-part}}^n)$ for satisfiable instances, when all partitions are satisfiable and the partition function is void. Furthermore, if $n = k^2$ for some $k \in \mathbb{N}$, $\mathbb{E}(T_{\text{safe-part}}^n) = \mathbb{E}(T_{\text{rep-part}}^n) = \mathbb{E}(T_{\text{sdsat}}^n)$.

Proof This follows immediately from (3) and (4) by setting $\epsilon(k) = 1$ and $r = k$, and noting that $q_{T_{\text{sdsat}}^n}(t) = 1 - (1 - q(t))^n$. ■

Proposition 11 $\mathbb{E}(T_{\text{rep-part}}^n) = \mathbb{E}(T_{\text{safe-part}}^n)$ for satisfiable instances, when exactly one partition is satisfiable and partitioning function is void. Furthermore, $\mathbb{E}(T_{\text{rep-part}}^n) = \mathbb{E}(T_{\text{safe-part}}^n) = \mathbb{E}(T_{\text{sdsat}}^k)$, where $k = \lfloor \sqrt{n} \rfloor$.

Proof This follows directly from (3) and (4) by setting $r = 1$ and $\epsilon(k) = 1$, and noting that $q_{T_{\text{sdsat}}^k}(t) = 1 - (1 - q(t))^k$. ■




Figure 3: An (artificial) distribution of an unsatisfiable instance, described by Eq. (5), where SDSAT outperforms partitioning for quite large n .

B EXPERIMENTAL RESULTS

Table 2 lists statistics similar to those presented in Table 1 for four additional instances from the SAT 2007 competition. All four instances are unsatisfiable. The results support the conclusions made based on the smaller set: especially the variance in separate invocations of the randomized partitioning function results in better performance of the repeated partitioning approach compared to that of the safe partitioning approach.

Figure 3 demonstrates an artificial unsatisfiable instance with a distribution where the SDSAT approach (the *sdsat* graph) performs better than the partitioning approach (the *part(ideal)* graph) for large n . The distribution is artificially generated, having

$$q_T(t) = \begin{cases} 0 & \text{for } 0 \leq t < 1, \\ 0.8 & \text{for } 1 \leq t < 1000000, \text{ and} \\ 1 & \text{for } t \geq 1000000. \end{cases} \quad (5)$$

The distribution is also useful in demonstrating some anomalies that can be observed with repeated partitioning. For example, repeated partitioning with void partitioning function (the *rep-part(void)* graph) may result in an expected run time which is higher than the sequential expected run time (the *orig* graph) for large n , and its performance may be decreasing for some range of n even with ideal partitioning function (the *rep-part(ideal)* graph). For this particular function, a “local minimum” for the expected run time of repeated partitioning with an ideal partitioning function is reached when $n \approx 35$, and is only reached again when $n \approx 1600$.

Table 2: Comparing approaches for parallel search

		ideal partitioning				scattering					
		simple distrib.		ideal partitioning		scattering					
T	$T_{\text{sd sat}}^{T^8}$	$T_{\text{sd sat}}^{T^{64}}$	$T_{\text{part (ideal)}}^{T^8}$	$T_{\text{part (ideal)}}^{T^{64}}$	$T_{\text{sat-part}}^{T^{64}}$	$T_{\text{rep-part}}^{T^{64}}$	$T_{\text{scatter}}^{T^8}$	$T_{\text{scatter}}^{T^{64}}$	$T_{\text{sat-part}}^{T^{64}}$	$T_{\text{rep-part}}^{T^{64}}$	
1invr1nv5											
Exp	2829	2440	2207	417.3	60.11	331.9	367.1	1354	1311	1107	986.4
Min	2168	2168	2168	271.1	33.88	271.1	271.1	895.9	815.5	738.5	895.9
Q_1	2644	2297	2168	377.7	56.33	323.4	359.2	1068	1088	873.8	904.6
Med	2799	2481	2172	414.2	62.11	331	366.9	1232	1229	1018	1000
Q_3	2943	2548	2225	443.8	62.93	339.2	372.3	1707	1604	1359	1033
Max	4046	4046	4046	505.7	63.22	505.7	505.7	2682	1896	2301	2682
contest03-SGI_30_50_30_20_3-dir											
Exp	1433	1221	1116	210.6	29.24	165.4	191.3	1324	1107	1111	1060
Min	1076	1076	1076	134.6	16.82	134.6	134.6	925.8	650.9	889.3	925.8
Q_1	1309	1186	1076	200.8	28.36	161.4	185.4	1174	944.7	1027	955.1
Med	1421	1222	1083	206.6	29.23	164.5	192.2	1332	1115	1124	1071
Q_3	1550	1273	1169	220.4	30.32	169.1	196.2	1461	1272	1197	1142
Max	1993	1993	1993	249.2	31.14	249.2	249.2	1698	1588	1284	1698
hwb-r28-02-s818962541											
Exp	4654	4031	3624	667.5	93.19	561.3	597.3	2247	2303	1881	1807
Min	3596	3596	3596	449.6	56.19	449.6	449.6	1285	1867	1279	1285
Q_1	4506	3685	3596	610.7	92.36	556.9	588.2	2065	2087	1881	1601
Med	4660	4062	3598	665	93.22	563.8	595.4	2227	2413	1921	1949
Q_3	4767	4299	3615	718.3	95.19	572.5	601.6	2495	2509	1952	1986
Max	6191	6191	6191	773.9	96.74	773.9	773.9	2634	2973	2107	2634
unsat-set-b-fclqcolor-10-07-09											
Exp	2027	1446	1133	333.7	51.19	212.7	285.8	1656	1487	1149	1108
Min	1012	1012	1012	126.5	15.81	126.5	126.5	913.4	387	558.9	913.4
Q_1	1630	1346	1012	303.7	48.43	196.5	277.9	1362	1024	941.2	945.7
Med	2090	1481	1033	316.5	51.59	205.3	286.7	1536	1460	1259	1056
Q_3	2304	1553	1284	364.8	54.72	218.4	294.5	2050	1670	1385	1235
Max	3652	3652	3652	456.5	57.06	456.5	456.5	3347	4626	1631	3347

TKK REPORTS IN INFORMATION AND COMPUTER SCIENCE

- TKK-ICS-R12 Markus Koskela, Jorma Laaksonen
Specification of information interfaces in PinView. November 2008.
- TKK-ICS-R13 Jorma Laaksonen
Definition of enriched relevance feedback in PicSOM. November 2008.
- TKK-ICS-R14 Jori Dubrovin
Checking Bounded Reachability in Asynchronous Systems by Symbolic Event Tracing.
April 2009.
- TKK-ICS-R15 Eerika Savia, Kai Puolamäki, Samuel Kaski
On Two-Way Grouping by One-Way Topic Models. May 2009.
- TKK-ICS-R16 Antti E. J. Hyvärinen
Approaches to Grid-Based SAT Solving. June 2009.
- TKK-ICS-R17 Tuomas Launiainen
Model checking PSL safety properties. August 2009.
- TKK-ICS-R18 Roland Kindermann
Testing a Java Card applet using the LIME Interface Test Bench: A case study.
September 2009.
- TKK-ICS-R19 Kalle J. Palomäki, Ulpu Remes, Mikko Kurimo (Eds.)
Studies on Noise Robust Automatic Speech Recognition. September 2009.
- TKK-ICS-R20 Kristian Nybo, Juuso Parkkinen, Samuel Kaski
Graph Visualization With Latent Variable Models. September 2009.
- TKK-ICS-R21 Sami Hanhijärvi, Kai Puolamäki, Gemma C. Garriga
Multiple Hypothesis Testing in Pattern Discovery. November 2009.

ISBN 978-952-248-230-3 (Print)

ISBN 978-952-248-231-0 (Online)

ISSN 1797-5034 (Print)

ISSN 1797-5042 (Online)