

Model Counting in Product Configuration

Andreas J. Kübler Christoph Zengler

Wolfgang Küchlin

Symbolic Computation Group,
Wilhelm Schickard Institute for Computer Science,
Universität Tübingen, Germany

<http://www-sr.informatik.uni-tuebingen.de>

{kuebler,zengler,kuechlin}@informatik.uni-tuebingen.de

We describe how to use propositional model counting for a quantitative analysis of product configuration data. Our approach computes valuable meta information such as the total number of valid configurations or the relative frequency of components. This information can be used to assess the severity of documentation errors or to measure documentation quality. As an application example we show how we apply these methods to product documentation formulas of the Mercedes-Benz line of vehicles. In order to process these large formulas we developed and implemented a new model counter for non-CNF formulas. Our model counter can process formulas, whose CNF representations could not be processed up till now.

1 Introduction

Since R1/XCON [15] was used by DEC to support computer system configuration and assembly, product configuration systems have been among the most prominent and successful applications of AI methods in practice [19]. As a result computer aided configuration systems have been used in managing complex software, hardware or network settings. Another application area of these configuration systems is the automotive industry. Here they helped to realize the transition from the mass production paradigm to present-day mass customization.

Model counting is a technique to count the number of different satisfying variable assignments of a formula in propositional logic. Up to now model counting has mostly been used within bayesian networks [22, 1] and planning problems [17, 14]. In contrast, our work aims at using propositional model counting for a quantitative analysis of comprehensive product configuration setups.

To demonstrate the applicability of our novel methods, we provide examples within a well-studied car configuration context. The German automobile industry follows a build-to-order strategy based on an exceedingly large product variety. This makes it possible, especially for the manufacturers of luxury cars, to offer each customer their unique tailor-made car, and it differentiates the model lines from the mass market. Pil and Holweg [18] have discussed the interconnection of product variety and order-fulfillment strategies.

In this context it is interesting for the management to know precisely how much variety there is in the product line. For the automobile industry, this question is surprisingly difficult to answer. In particular, sales options cannot simply be multiplied out because of complex interdependencies (e.g. automatic transmission being standard on U.S. exports, and not available on European cars with small engines). Pil and Holweg published data for the year 2002, based on an analysis of company material, about customer selectable variations. In summary, numbers range from the order of 10^3 for the smallest Peugeot and Nissan models to the order of 10^8 or 10^9 for the GM Astra and Corsa or the Ford Focus. In the luxury sector, the BMW 3-Series was estimated at $6 \cdot 10^{16}$, and the Mercedes C-Class at 10^{21} . Finally, for the

Mercedes E-Class, 2 body styles, 15 power trains, 285 paint-and-trim options, and 70 other factory-fitted options could be combined to variations on the order of $3 \cdot 10^{24}$. It gets even worse: Variations induced by sales to countries outside Europe (due to different emission standards, fuel grades etc.) apparently were not considered by Pil and Holweg. And customer visible variations are only a subset of the variations afforded by technical and legal restrictions.

New technology options, such as hybridization, are a continuous source of new variability, and therefore the issue of rising complexity must have the ongoing attention of management. Since we work with our industrial partner Mercedes-Benz on the symbolic verification of configuration problems [13], such as whether a given configuration is valid or not, or which part is required in which configurations, we were able to investigate *how many valid configurations there are for one specific car line*. Obviously, there are a number of useful questions associated with this number. Some of these are: How does the complexity evolve over time? Which change in complexity would result from adding or canceling an equipment option? What is the complexity of each body style? Of each engine or power train configuration? In each country on Earth? For which percentage of the total variations is a particular part needed? In practice, detailed questions such as the above can only be answered precisely if the number of variations can be computed automatically from the manufacturer’s configuration data. This was not possible up to now, which also means that the above numbers were not verified independently so far.

This paper covers only the first steps towards using model counting in configuration and manufacturing problems. Clearly, some kind of additional information about the expected or observed frequency of chosen options in the real product should be included in the model counts. E.g. an option can be allowed in 90% of all valid configurations, but is chosen only in 10% of actual orders. Still, to the best of our knowledge our results represent the first cases where exact model counts were obtained directly and automatically from the internal manufacturer configuration data of a current highly complex product line. We are certain that more ways will be found in which our methods can be used for computing management information. Also, more extensive data sets will certainly be produced in the future.

The plan of this paper as follows: In Section 2 we summarize the theory of model counting and give a short overview of current algorithmic approaches. Section 3 presents possible application fields for model counting in configuration problems. In Section 4 we discuss the specific application of our new methods in the customer order process at Mercedes-Benz. We give a short description of our new non-CNF model counter `ncnt` and present benchmarks of configuration problems used in the current product line. Section 5 finally summarizes our contributions and points at some future research directions.

2 Model Counting

Let φ be a formula in propositional logic, and let $\text{vars}(\varphi)$ denote the finite set of variables occurring in the formula φ . An *assignment* α for φ is a partial function $\alpha : \text{vars}(\varphi) \rightarrow \{\top, \perp\}$ mapping variables to truth values \top and \perp . We follow the convention to write $\alpha \models \varphi$ when some formula φ holds with respect to α . *Propositional model counting* or #SAT is the problem of computing the number of all full assignments α for which φ holds, i.e. $|\{\alpha \mid \alpha \models \varphi\}|$.

Analogous to SAT, which is the canonical NP-complete problem, #SAT is the canonical #P-complete problem. The complexity class #P is the class of all problems p for which there exists a non-deterministic polynomial-time bound Turing machine $M(p)$ such that for each instance $I(p)$ of p there exist exactly as many computation paths of $M(p)$ as solutions for $I(p)$. Intuitively #P is the class of counting problems for polynomial-time decidable problems. According to [25] even the counting problems for polynomial-time solvable problems like 2-SAT, Horn-SAT, or DNF-SAT can be #P-complete.

In this paper we only deal with *exact model counting* (in contrast to approximative counting). We distinguish between two different approaches for exact counting: (1) DPLL-like exhaustive search and (2) knowledge compilation.

The vast majority of successful SAT solvers uses the DPLL approach [10, 9]. DPLL is basically a complete search in the search space of all 2^n variable assignments with early cuts in the search tree when an unsatisfiable branch is detected. DPLL-style model counters like CDP [4], Re1Sat [2], or Cachet [20] are extensions to existing SAT solvers and require an input formula in CNF. If a formula φ with n variables is not satisfiable, the output is 0. If a satisfying (and possible partial) assignment α is found, the number of models for this α is computed with $2^{n-|\alpha|}$ and the algorithm proceeds to explore the rest of the search tree. There are two important improvements of this DPLL-based approach. The first one is *component analysis* [2] where one identifies different components C_1, \dots, C_n in the constraint graph G of a CNF formula φ . Let $\varphi_1, \dots, \varphi_n$ be the sub-formulas of φ corresponding to the components C_1, \dots, C_n . Then the model count $\#SAT(\varphi)$ is equal to $\#SAT(\varphi_1) \times \dots \times \#SAT(\varphi_n)$, thus we can calculate the model count of each component independently. This identification of components can be performed dynamically while descending into the search tree. The second improvement is the $\#SAT$ correspondence to clause learning in SAT: *component caching* [20, 24]. Since during the counting process we often compute counts for the same sub-formulas multiple times, we cache signatures of sub-formulas and their model count according to certain caching schemes. Variable selection heuristics as known from SAT have to be adjusted for $\#SAT$ wrt. component analysis and caching: While in SAT one tries to narrow down the search to one specific solution by intelligently choosing the branching variables, in $\#SAT$ we try to choose variables where the according constraint graph is decomposed in various components [21].

In the knowledge compilation based approach we convert the formula φ into another logical representation such that $\#SAT(\varphi)$ can be computed in polynomial time. One well known approach for this is the compilation of φ into a binary decision diagram (BDD) [5]. Once we have the BDD, we can count all paths from the root node to the \top labeled node to get the model count of the formula at hand. Narodytska and Walsh discussed this approach for configuration problems [16]. However our own experiments with the formulas presented in Section 4 showed that for these large formulas, which emerge in our industrial automotive project, the compilation into BDDs is not feasible. Another logical representation for propositional formulas, which is used e.g. in the system c2d [8], is the *deterministic decomposable negation normal form* (d-DNNF) [7]. A DNNF is an extension of a negation normal form where for each conjunction $\bigwedge_{i=0}^n f_i$ with sub-formulas f_0, \dots, f_n it must hold that $\text{vars}(f_i) \cap \text{vars}(f_j) = \emptyset$ for all $0 \leq i < j \leq n$. In a d-DNNF we have the additional condition that for each disjunction $\bigvee_{i=0}^n g_i$ with sub-formulas $G = \{g_0, \dots, g_n\}$ it must hold that for each α we have $|\{g_i \in G \mid \alpha \models g_i\}| \leq 1$, i.e. no sub-formulas share the same satisfying assignment. Once we have the d-DNNF φ' of a formula φ we can count all models by these two rules:

$$\#SAT\left(\bigwedge_{i=0}^n f_i\right) = \prod_{i=0}^n \#SAT(f_i), \quad \#SAT\left(\bigvee_{i=0}^n g_i\right) = \sum_{i=0}^n \#SAT(g_i).$$

It turns out that a d-DNNF representation is closer to the original CNF formula and therefore is easier to compute.

3 Applications in Product Configuration

In this section we will point out possible fields of applications for model counting in product configuration. We use the definition of a configuration problem as given in [11, Definition 1]: a configuration

problem is a triple (\mathcal{V}, D, Ψ) where \mathcal{V} is a set of variables x_1, x_2, \dots, x_n , D is a set of their finite domains D_1, D_2, \dots, D_n and $\Psi = \{\psi_1, \psi_2, \dots, \psi_m\}$ is a set of propositional formulas over atomic propositions $x_i = v$ where $v \in D_i$, specifying conditions that the variable assignments have to satisfy. For each formula $\psi \in \Psi$ we have $\text{vars}(\psi) \subseteq \mathcal{V}$. A *valid configuration* is an assignment α with $\text{dom}(\alpha) = \mathcal{V}$ such that $\alpha \models \bigwedge_{\psi \in \Psi} \psi$.

In this paper we consider the special case where we have only propositional variables in \mathcal{V} and hence $D_i = \{\top, \perp\}$ for all $1 \leq i \leq n$. The set \mathcal{O} is the finite set of all configuration options for a product. Each variable $x_o \in \mathcal{V}$ represents a configuration option $o \in \mathcal{O}$. The variable x_o is assigned to \top if the option o is chosen, otherwise it is assigned to \perp . Following this course, the resulting formulas $\psi \in \Psi$ are propositional formulas and hence $\varphi = \bigwedge_{\psi \in \Psi} \psi$ is a propositional formula describing all valid configurations. We will also refer to φ as *product overview formula (POF)* [13].

Remark. The restriction of the variables $x \in \mathcal{V}$ to propositional variables does not limit the expressiveness of our problem description. Since the domains D_i are finite and we only allow atomic propositions of the form $x = v$, we can use a reduction [6] from equality logic to propositional logic.

3.1 Number of Valid Configurations

The first question which naturally arises is the total number of valid product configurations $v = \#SAT(\varphi)$. Obviously $2^{|\mathcal{V}|}$ is an upper bound for v , but in most cases $v \ll 2^{|\mathcal{V}|}$. Nevertheless this number v can often demonstrate the sheer complexity of a given product. Subsection 4.2 supports this claim by presenting some of these numbers for the car lines of Mercedes-Benz.

One can also count valid configurations of a product under certain preconditions. E.g. we can force options $P \subset \mathcal{O}$ to be chosen before performing $\#SAT(\varphi)$. This can be achieved by computing $\#SAT(\varphi \wedge \bigwedge_{p \in P} x_p)$. This method can yield important information about the influence of certain options o on the number of valid configurations. This information again can be used for special domain-specific variable heuristics for the SAT solving process of formulas of the application area at hand. The main idea of state-of-the-art SAT solvers is to narrow down the search space as fast as possible. Thus when we know a set of variables $X \subseteq \mathcal{V}$ representing configuration options $P \subset \mathcal{O}$ which reduce the number of models of a given formula to a great extent, we can give these variables $x \in X$ high activity in the SAT solving process of formulas representing configuration problems of the same product line.

Preselecting a certain option can also be used to compute the relative frequency of this option in valid configurations. We compute the frequency $f(o)$ of a given option o with

$$f(o) = \frac{\#SAT(\varphi \wedge x_o)}{\#SAT(\varphi)}.$$

$f(o)$ can be used as additional information to statistical data about the frequency of options for demand estimation or process optimization.

3.2 Rating Errors

Quite important issues arise when reporting errors. Observations from formal methods in software verification [3] tell us that the more bugs you report, the smaller the probability gets that they will eventually be fixed. Developers as well as product documentation engineers tend to get overwhelmed quite quickly by extensive error reports leaving them uncertain where to start correcting defects.

Model counting might help classifying errors according to their severity. We consider scenarios in which satisfiability of the input formula φ indicates error situations — hence any satisfying assignment may be interpreted as a counterexample. In configuration problems such situations mostly arise

when checking mutually exclusive component inclusion. Let $c_1, c_2 \in \text{vars}(\varphi)$ denote binary flags for the inclusion of two mutually exclusive components. Each assignment satisfying $\varphi \wedge c_1 \wedge c_2$ may be interpreted as a product featuring both mutually exclusive components at the same time, thus computing $\#SAT(\varphi \wedge c_1 \wedge c_2)$ results in the total number of invalid configurations w.r.t. c_1, c_2 . Formulas yielding many such “counterexamples” intuitively tend to fail more likely in practice than formulas yielding only negligible numbers. However, if there are very few models for a given error, we assume that this error is more intricate than others.

Measuring error severity using model counters as proposed may however turn out not to be applicable in every domain. Think of Boolean encodings of configuration options where each option o is assigned to a value of a finite integer domain D and a consistency assertion $a(o)$ fixing o to a specific integer value i . If checking the Boolean encoding w.r.t. $a(o)$ yields that o is not restricted to i but may instead take an arbitrary value of D , model counting will consequently return some large number of invalid configurations. In case o is some exotic configuration option, using the computed number as a measure of severity might be misleading.

3.3 Measuring Documentation Quality

Often the individual constraints $\psi \in \Psi$ are reflected in a product documentation. This product documentation is either automatically produced or manually maintained by experts. For complex products with thousands of options this documentation evolves over years. This fact can be observed in the development of the numbers of valid configurations of different versions of the documentation for the same product. We made observations where for the same product the number of valid configurations went down from a magnitude of 10^{34} to a magnitude of 10^{10} due to a better and more detailed documentation (c.f. Subsection 4.2).

Model counting can also be used interactively by the documentation engineer. One can immediately see the impact of adding or changing certain constraints in terms of valid configurations. E.g. when adding a new constraint does not change the number of valid configurations it can be redundant and therefore can be omitted.

4 Example: Automotive Product Configuration

In this section we give examples how to apply propositional model counting to formulas as used in constructibility checking of customer orders at Mercedes-Benz. Therefore we first give a brief introduction of the notions used in the product documentation and of the configuration system used in the mass production of individually configured, personalized cars.

4.1 Formula Semantics

For the Mercedes lines of cars, product documentation is done based on propositional logic. For each vehicle class, there are about 1.500 variables which represent the configuration options. Processing of customer orders at Mercedes-Benz is organized as a three-staged process [13]:

1. The customer’s choice of configuration options $P \subset \mathcal{O}$ with $P = \{x_{o_1}, \dots, x_{o_{|P|}}\}$ is compiled into an initial assignment $\alpha = \{x_{o_1} \leftarrow \top, \dots, x_{o_{|P|}} \leftarrow \top\}$. We refer to propositional variables representing an option $o \in \mathcal{O}$ also as *codes*.

2. The assignment α computed in Step 1 is iteratively extended using so called supplementary rules. A supplementary rule $S(x) \rightarrow x$ consists of a conditional part $S(x)$ in propositional logic and a supplementary code x . If $\alpha \models S(x)$ for the current assignment α we extend $\alpha = \alpha \cup \{x \leftarrow \top\}$.
3. For each code x there is a constructibility condition in propositional logic of the form $x \rightarrow C(x)$. The initial assignment α of Step 1, supplemented in Step 2, is finally checked for constructibility in this third step. If $\alpha \models \bigwedge_{x \in \text{dom}(\alpha)} C(x)$ holds, the customer order is constructible otherwise it will be refused.

Let X be the set of all allowed codes for a given line of vehicles, then the afore mentioned POF is defined as follows

$$\text{POF}(X) = \bigwedge_{x \in X} (S(x) \rightarrow x) \wedge (x \rightarrow C(x)).$$

Example. Consider as a toy example a vehicle where the configuration options are three different engines with codes e_1, e_2, e_3 , two different gearboxes g_1, g_2 and three additional features a_1, a_2, a_3 . Engine e_1 must be combined with gearbox g_1 (sr_1), e_2 must be combined with g_2 (sr_2). In a car with e_3 and g_1 also a_2 has to be chosen (sr_3), in a car with e_3 and g_2 , we must choose a_3 (sr_4). The resulting supplementary rules are

$$\text{sr}_1 = e_1 \rightarrow g_1, \quad \text{sr}_2 = e_2 \rightarrow g_2, \quad \text{sr}_3 = e_3 \wedge g_1 \rightarrow a_2, \quad \text{sr}_4 = e_3 \wedge g_2 \rightarrow a_3.$$

There is exactly one engine in a vehicle (cc_1 – cc_4) and exactly one gearbox (cc_5 – cc_7). There is also the condition, that feature a_2 must not be combined with a_3 (cc_8). The resulting constructibility conditions are

$$\begin{aligned} \text{cc}_1 &= \top \rightarrow e_1 \vee e_2 \vee e_3, & \text{cc}_2 &= e_1 \rightarrow \neg e_2 \wedge \neg e_3, & \text{cc}_3 &= e_2 \rightarrow \neg e_1 \wedge \neg e_3, & \text{cc}_4 &= e_3 \rightarrow \neg e_1 \wedge \neg e_2 \\ \text{cc}_5 &= \top \rightarrow g_1 \vee g_2, & \text{cc}_6 &= g_1 \rightarrow \neg g_2, & \text{cc}_7 &= g_2 \rightarrow \neg g_1, & \text{cc}_8 &= a_2 \rightarrow \neg a_3 \end{aligned}$$

The POF is the conjunction of all constructibility conditions and supplementary rules:

$$\text{POF} = \bigwedge_{i=1}^8 \text{cc}_i \wedge \bigwedge_{j=1}^4 \text{sr}_j.$$

The solutions of the POF represent exactly those customer orders (as a set of configuration options) which can be built under the configuration constraints. Hence any model α of the POF (i.e. $\alpha \models \text{POF}$) describes a single valid configuration. Based on the notion of the POF, several consistency tests such as finding necessary codes, detecting redundant parts or intersections may be performed quite naturally using SAT solvers [13].

Given the variability described above, it is no longer possible to store separate parts lists (bills of material – BOM) for each possible order. Therefore, there is a single BOM for the model line (e.g. E-Class), where each part p is associated with a selection condition $P(p)$ in propositional logic. For each customer order, given as the supplemented assignment α , every part whose selection condition evaluates to true under α is pulled into the order specific BOM. Hence the set of all parts p occurring in the BOM is defined as $\{p \mid \alpha \models P(p)\}$.

The BOM is actually structured into so-called positions, each of them grouping together a set of parts which may be alternatively selected for installation at the same geometric position of the car (e.g. the

set of available radios). In addition, there is one fictitious (null-)part which represents the ”‘nothing-to-install’” case. Counting the models of the selection formulas for each part may give an indication of the relative frequency of the parts, and of the frequency of the null-case, which often represents some kind of exception.

4.2 Model Counting in the Mercedes-Benz Scenario

In the rest of this paper, we will present some results as described in Section 3 gathered by model counting product documentation formulas in use at Mercedes-Benz. As those formulas are originally non-CNF, we developed a new exact propositional model counter that operates directly upon the raw formulas without prior normalization. It is based upon a reimplementaion of the non-CNF SAT solver introduced by Thiffault et al. [23] and uses the DPLL-style approach to model counting with connected component detection and caching.

Remark. When converting a non-CNF formula φ into a CNF formula φ' in order to utilize a CNF model counter, it is important to use a transformation method which preserves the number of models of φ , i.e. $|\{\alpha \mid \alpha \models \varphi\}| = |\{\alpha \mid \alpha \models \varphi'\}|$. Consider e.g. the original formula $\varphi = a \vee (b \wedge c)$. To get an equisatisfiable formula φ' in CNF we can use the Tseitin transformation with either implications ($T^{\rightarrow}(\varphi)$) or equivalences ($T^{\leftrightarrow}(\varphi)$):

$$\begin{aligned} T^{\rightarrow}(\varphi) &= f_{or} \wedge (f_{or} \rightarrow a \vee f_{and}) \wedge (f_{and} \rightarrow b \wedge c) = && f_{or} \wedge (\neg f_{or} \vee a \vee f_{and}) \wedge (\neg f_{and} \vee b) (\neg f_{and} \vee c) \\ T^{\leftrightarrow}(\varphi) &= f_{or} \wedge (f_{or} \leftrightarrow a \vee f_{and}) \wedge (f_{and} \leftrightarrow b \wedge c) = && f_{or} \wedge (\neg f_{or} \vee a \vee f_{and}) \wedge (f_{or} \vee \neg a) \wedge (f_{or} \vee \neg f_{and}) \wedge \\ &&& (\neg f_{and} \vee b) \wedge (\neg f_{and} \vee c) \wedge (f_{and} \vee \neg b \vee \neg c) \end{aligned}$$

We can now compare the number of models for each formula: $\#SAT(\varphi) = 5$, $\#SAT(T^{\rightarrow}(\varphi)) = 6$, and $\#SAT(T^{\leftrightarrow}(\varphi)) = 5$. We see that only $T^{\leftrightarrow}(\varphi)$ preserves the number of models.

4.2.1 Number of Valid Configurations

Model counting pure PO formulas without any extensions (i.e. computing $\#SAT(POF)$, the number of different valid orders) demonstrates the vast multitude of constructible vehicles due to customer choices: We have been able to compute results of 5.9×10^{10} up to 9.9×10^{103} constructible orders for the E-class line of Mercedes-Benz cars (cf. Tab. 1). C212_A, C212_B, and C212_C indicate different body styles of the same car line. Those impressive numbers additionally provide an invaluable aid in arguing for formal methods use in large configuration settings. Our non-CNF model counter has directly been applied to the non-normalized POFs whereas Cachet and c2d have been given Tseitin transformed (T^{\leftrightarrow}) CNF versions as input. Our comparison states that preserving formula structure by directly reasoning over the input without prior normalization pays off in scenarios featuring large input: Cachet timed out after 6 hours on C212_C respectively exited abnormally on C212_B, c2d has been manually aborted (as it is lacking a timeout option) after 6 hours returning no results.

In table 2 we summarize some computations we performed in order to analyze the number of valid orders including individual parts of the BOM, i.e. $\#SAT(POF \wedge P(p))$. Each row identifies a position variant (e.g. a certain device in the set of available radios), i.e. a part that might be installed at one single position in the car. As simultaneously installing different parts at the same position poses a physical impossibility, position variants need to be mutually exclusive in order to preserve product documentation consistency. Using table 2 one may compute inclusion probabilities for each part depending on the

¹Environment: Linux (2.6.23), 3.4 GHz Intel Pentium D, 2 GB memory

Table 1: Computing the total number of orders¹

problem	#vars	#clauses	ncnt		Cachet		c2d	
			result	time	result	time	result	time
C212_A	4898	10289	5.9×10^{10}	2.57	5.9×10^{10}	0.05	5.9×10^{10}	51.82
C212_B	9574	40809	9.4×10^{60}	4709.51	bus error	—	timeout	—
C212_C	10637	45990	9.9×10^{103}	7915.00	timeout	—	timeout	—

Table 2: Computing the number of orders including specific parts using ncnt²

part p	C212_B		C212_C	
	result	time	result	time
10	3.6×10^{57}	5160.25	4.1×10^{100}	6528.81
20	2.7×10^{58}	7686.87	2.3×10^{101}	7900.62
30	5.3×10^{56}	4626.25	5.5×10^{99}	7620.96
40	1.8×10^{57}	5530.34	1.9×10^{100}	10933.19
60	6.0×10^{56}	5607.98	6.1×10^{99}	8251.84
70	6.0×10^{55}	1833.41	4.8×10^{98}	3840.32
999	1.2×10^{45}	111.87	4.4×10^{88}	136.41

corresponding POF. One can clearly see that for the few variants we computed, variant 20 dominates the others. These numbers might serve as a basis in applications estimating future requirements in purchasing parts or for speeding up SAT solving.

4.2.2 Rating Errors

In practice it is highly desirable to rank the findings of an automated verifier. In the real world, nothing is perfect, and every corrective action has a price. If we have an error condition, which is represented by a formula φ being satisfiable, then $\#SAT(\varphi)$ may give an indication, how many cars are potentially affected by the error.

As an example consider physically overlapping parts p_1, p_2 : If, due to the POF, there is a constructible order featuring both, p_1 and p_2 , any assignment α with $\alpha \models POF \wedge P(p_1) \wedge P(p_2)$ describes an erroneously constructible order. Thus computing $\#SAT(POF \wedge P(p_1) \wedge P(p_2))$ will return the total number of erroneously constructible orders wrt. p_1 and p_2 . Using those numbers retrieved by model counting one may intuitively classify errors as follows: Overlaps leading to a high number of constructible orders (thus being more likely to actually occur in production) are intuitively more severe than overlaps featuring a negligible number. Experts concerned with fixing documentary flaws may thus prioritize their work using results produced by model counting.

4.2.3 Measuring Documentation Quality

Exact model counts also give rise to questions about the appropriate product documentation language and method. We are currently engaged in introducing symbolic verification methods to detect and help

²Environment: Mac OS X (10.5), 2×4 Core Intel Xeon (2.8 GHz each), 32 GB memory, using 1 core

avoid documentation errors. Knowing the exact model count gives some indication (at least on the level of management decisions) of the complexity of detecting errors. E.g. in order to prove that a part is no longer needed in a 2002 model line, it must be verified that none of the 10^{24} cars needs the part.

Our results furthermore indicate that model counting POFs might serve as a measure of product documentation maturity (cf. Subsection 3.3): Model counting an early developer's version of C212.A resulted in 2.9×10^{34} models whereas the version in productive use at Mercedes-Benz yields only 5.9×10^{10} .

5 Conclusion & Future Work

In this paper we introduced new methods for the quantitative analysis of configuration formulas using model counting. We described methods for gaining additional information like the total number of valid configurations, the relative frequency of options in valid configurations, or the severity of errors. This information can be used to speed up future SAT solving processes of similar formulas, illustrate the complexity of the problem, or measuring the documentation quality. Using our non-CNF model counter `ncnt` we have been able to compute results for formulas whose CNF representations are too large for recent model counters to cope with. Our results give proof of the applicability of our newly introduced analysis methods to industrial-scale configuration problems.

As our benchmarks show, model counting for large propositional formulas is still a quite time-consuming job to do. In d-DNNF compilation static heuristics based upon hypergraph decomposition have been successfully used to generate good variable orders for decomposing large and complex formulas [12]. To our knowledge this approach has not yet been used in non-compilation based settings. Such heuristics should be introduced in DPLL-style model counting to speed up decomposition.

Moreover, decomposing problems into independent connected components quite naturally suggests to harness multithreaded CPU architectures by parallelizing propositional model counting.

6 Acknowledgement

We would like to thank Matthias Sauter (formerly STZ OIT) for providing us with helpful remarks, advice, and support.

References

- [1] Fahiem Bacchus, Shannon Dalmao & Toniann Pitassi (2003): *Algorithms and Complexity Results for #SAT and Bayesian Inference*. In: *Proceedings of the FOCS '03*, IEEE Computer Society, pp. 340–352.
- [2] Roberto J. Bayardo, Jr. & Joseph Daniel Pehoushek (2000): *Counting Models Using Connected Components*. In: *Proceedings of the 17th National Conference on Artificial Intelligence and 12th Conference on Innovative Applications of Artificial Intelligence*, AAAI Press / The MIT Press, pp. 157–162.
- [3] Al Bessey, Ken Block, Ben Chelf, Andy Chou, Bryan Fulton, Seth Halleem, Charles Henri-Gros, Asya Kamsky, Scott McPeak & Dawson Engler (2010): *A Few Billion Lines Of Code Later: Using Static Analysis to Find Bugs in the Real World*. *Communications of the ACM* 53(2), pp. 66–75.
- [4] Elazar Birnbaum & Eliezer L. Lozinskii (1999): *The Good Old Davis-Putnam Procedure Helps Counting Models*. *Journal of Artificial Intelligence Research* 10(1), pp. 457–577.
- [5] Randal E. Bryant (1986): *Graph-Based Algorithms for Boolean Function Manipulation*. *IEEE Transactions on Computers* 35(8), pp. 677–691.

- [6] Randal E. Bryant & Miroslav N. Velev (2000): *Boolean Satisfiability with Transitivity Constraints*. In: *Computer Aided Verification, LNCS 1855*, Springer, Berlin, Heidelberg, Germany, pp. 85–98.
- [7] Adnan Darwiche (2001): *On the Tractable Counting of Theory Models and its Application to Truth Maintenance and Belief Revision*. *Journal of Applied Non-Classical Logics* 11(1-2), pp. 11–34.
- [8] Adnan Darwiche (2004): *New Advances in Compiling CNF to Decomposable Negational Normal Form*. In: *Proceedings of the ECAI 2004*, IOS Press.
- [9] Martin Davis, George Logemann & Donald Loveland (1962): *A Machine Program for Theorem-Proving*. *Communications of the ACM* 5(7), pp. 394–397.
- [10] Martin Davis & Hilary Putnam (1960): *A Computing Procedure for Quantification Theory*. *Journal of the ACM* 7(3), pp. 201–215.
- [11] Tarik Hadzic, Subbarayan Sathiamoorthy, Rune M. Jensen, Henrik R. Andersen, Jesper Møller & Henrik Hulgaard (2004): *Fast Backtrack Free Product Configuration using Precompiled Solution Space Representations*. In: *Proceedings of the International Conference on Economic, Technical and Organisational aspects of Product Configuration Systems*.
- [12] Jinbo Huang & Adnan Darwiche (2003): *A Structure-Based Variable Ordering Heuristic for SAT*. In: *Proceedings of the IJCAI'03*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, pp. 1167–1172.
- [13] Wolfgang Küchlin & Carsten Sinz (2000): *Proving Consistency Assertions for Automotive Product Data Management*. *Journal of Automated Reasoning* 24(1-2), pp. 145–163.
- [14] T. K. Satish Kumar (2002): *A Model Counting Characterization of Diagnoses*. In: *Proceedings of the DX-02*, pp. 70–76.
- [15] John McDermott (1982): *RI: A Rule-Based Configurer of Computer Systems*. *Artificial Intelligence* 19(1), pp. 39–88.
- [16] Nina Narodytska & Toby Walsh (2007): *Constraint and Variable Ordering Heuristics for Compiling Configuration Problems*. In: *Proceedings of the IJCAI'07*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, pp. 149–154.
- [17] Héctor Palacios, Blai Bonet, Adnan Darwiche & Héctor Geffner (2005): *Pruning Conformant Plans by Counting Models on Compiled D-DNNF Representations*. In: Susanne Biundo, Karen Myers & Kanna Rajan, editors: *Proceedings of the ICAPS 2005*, AAAI Press, pp. 141–150.
- [18] Frits K. Pil & Matthias Holweg (2004): *Linking Product Variety to Order-Fulfillment Strategies*. *Interfaces* 34(5), pp. 394–403.
- [19] Daniel Sabin & Rainer Weigel (1998): *Product Configuration Frameworks-A Survey*. *IEEE Intelligent Systems* 13(4), pp. 42–49.
- [20] Tian Sang, Fahiem Bacchus, Paul Beame, Henry Kautz & Toniann Pitassi (2004): *Combining Component Caching and Clause Learning for Effective Model Counting*. In: *Proceedings of the SAT 2004*.
- [21] Tian Sang, Paul Beame & Henry Kautz (2005): *Heuristics for Fast Exact Model Counting*. In: Fahiem Bacchus & Toby Walsh, editors: *Theory and Applications of Satisfiability Testing - SAT 2005, LNCS 3569*, Springer, Berlin, Heidelberg, Germany, pp. 226–240.
- [22] Tian Sang, Paul Beame & Henry Kautz (2005): *Performing Bayesian Inference by Weighted Model Counting*. In: *Proceedings of the AAAI'05*, AAAI Press, Menlo Park, CA, USA, pp. 475–481.
- [23] Christian Thiffault, Fahiem Bacchus & Toby Walsh (2004): *Solving Non-Clausal Formulas With DPLL Search*. In: *Principles and Practice of Constraint Programming – CP 2004, LNCS 3258*, Springer, pp. 663–678.
- [24] Marc Thurley (2006): *sharpSAT — Counting Models with Advanced Component Caching and Implicit BCP*. In: *Theory and Applications of Satisfiability Testing - SAT 2006, LNCS 4121*, Springer, Berlin, Heidelberg, Germany, pp. 424–429.
- [25] Leslie G. Valiant (1979): *The Complexity of Computing the Permanent*. *Theoretical Computer Science* 8(2), pp. 189–201.