
STRUCTURE AND PROBLEM HARDNESS: GOAL ASYMMETRY AND DPLL PROOFS IN SAT-BASED PLANNING

JÖRG HOFFMANN^a, CARLA GOMES^b, AND BART SELMAN^c

^a Digital Enterprise Research Institute, Innsbruck, Austria
e-mail address: joerg.hoffmann@deri.org

^{b,c} Cornell University, Ithaca, NY, USA
e-mail address: {gomes,selman}@cs.cornell.edu

ABSTRACT. In Verification and in (optimal) AI Planning, a successful method is to formulate the application as boolean satisfiability (SAT), and solve it with state-of-the-art DPLL-based procedures. There is a lack of understanding of why this works so well. Focussing on the Planning context, we identify a form of *problem structure* concerned with the symmetrical or asymmetrical nature of the cost of achieving the individual planning goals. We quantify this sort of structure with a simple numeric parameter called *AsymRatio*, ranging between 0 and 1. We run experiments in 10 benchmark domains from the International Planning Competitions since 2000; we show that *AsymRatio* is a good indicator of SAT solver performance in 8 of these domains. We then examine carefully crafted *synthetic planning domains* that allow control of the amount of structure, and that are clean enough for a rigorous analysis of the combinatorial search space. The domains are parameterized by size, and by the amount of structure. The CNFs we examine are unsatisfiable, encoding one planning step less than the length of the optimal plan. We prove upper and lower bounds on the size of the best possible DPLL refutations, under different settings of the amount of structure, as a function of size. We also identify the best possible sets of branching variables (backdoors). With *minimum AsymRatio*, we prove exponential lower bounds, and identify minimal backdoors of size linear in the number of variables. With *maximum AsymRatio*, we identify *logarithmic* DPLL refutations (and backdoors), showing a doubly exponential gap between the two structural extreme cases. The reasons for this behavior – the proof arguments – illuminate the prototypical patterns of structure causing the empirical behavior observed in the competition benchmarks.

2000 ACM Subject Classification: I.2.8, I.2.3.

Key words and phrases: planning, domain-independent planning, planning as SAT, DPLL, backdoors.

^b Research supported by the Intelligent Information Systems Institute, Cornell University (AFOSR grant F49620-01-1-0076).

1. INTRODUCTION

There has been a long interest in a better understanding of what makes combinatorial problems from SAT and CSP hard or easy. The most successful work in this area involves random instance distributions with phase transition characterizations (e.g., [10, 35]). However, the link of these results to more structured instances is less direct. A random unsatisfiable 3-SAT instance from the phase transition region with 1,000 variables is beyond the reach of any current solver. On the other hand, many unsatisfiable formulas from Verification and AI Planning contain well over 100,000 variables and can be proved unsatisfiable within a few minutes (e.g., with Chaff [46]). This raises the question as to whether one can obtain general measures of structure in SAT encodings, and use them to characterize typical case complexity. Herein, we address this question in the context of AI Planning. We view this as an entry point to similar studies in other areas.

In Planning, methods are developed to automatically solve the reachability problem in declaratively specified transition systems. That is, given some formalism to describe system states and transitions (“actions”), the task is to find a solution (“plan”): a sequence of transitions leading from some given start (“initial”) state to a state that satisfies some given non-temporal formula (the “goal”). Herein, we consider the wide-spread formalism known as “STRIPS Planning” [21]. This is a very simple formal framework making use of only Boolean state variables (“facts”), conjunctions of positive atoms, and atomic transition effects; details and notations are given later (Section 3).

We focus on showing infeasibility, or, in terms of Planning, on proving optimality of plans. We consider the difficulty of showing the non-existence of a plan with one step less than the shortest possible – *optimal* – plan. SAT-based search for a plan [39, 40, 41] works by iteratively incrementing a plan length bound b , and testing in each iteration a formula that is satisfiable iff there exists a plan with b steps. So, our focus is on the last unsuccessful iteration in a SAT-based plan search, where the optimality of the plan (found later) is proved. This focus is relevant since proving optimality is precisely what SAT solvers are best for, at the time of writing, in the area of Planning. On the one hand, SAT-based planning won the 1st prize for optimal planners in the 2004 International Planning Competition [28] as well as the 2006 International Planning Competition. On the other hand, finding potentially sub-optimal plans is currently done much faster based on heuristic search (with non-admissible heuristics), e.g. [6, 32, 24, 27, 57].

In our work, we first formulate an intuition about what makes DPLL [15, 14] search hard or easy in Planning. We design a numeric measure of that sort of problem structure, and we show empirically that the measure is a good indicator of search performance in many domains. We also perform a case study: We design synthetic domains that capture the problem structure in a clean form, and we analyze DPLL behavior in detail, within these domains.

1.1. Goal Asymmetry. In STRIPS Planning, the goal formula is a conjunction of goal facts, where each fact requires a Boolean variable to be true. The goal facts are commonly referred to as *goals*, and their conjunction is referred to as a *set* of goals. In most if not all benchmark domains that appear in the literature, the individual goal facts correspond quite naturally to individual “sub-problems” of the problem instance (*task*) to be solved. The sub-problems typically interact with each other – and this is the starting point of our investigation.

Given tasks with the same optimal plan length, our intuitions are these. (1) Proving plan optimality is hard if the optimal plan length arises from complex interactions between many sub-problems. (2) Proving plan optimality is easy if the optimal plan length arises mostly from a single sub-problem. To formalize these intuitions, we take a view based on sub-problem “cost”, serving to express both intuitions with the same notion, and offering the possibility to interpolate between (1) and (2). By “cost”, here, we mean the number of steps needed to solve a (sub-)problem. We distinguish a *symmetrical case* – where the individual sub-problems are all (symmetrically) “cheap” – and an *asymmetrical case* – where a single sub-problem (asymmetrically) “dominates the overall cost”.¹

The asymmetrical case obviously corresponds to intuition (2) above. The symmetrical case corresponds to intuition (1) because of the assumption that the number of steps needed to solve the overall task is the same in both cases. If each single sub-problem is cheap, but their conjunction is costly, then that cost must be the result of some sort of “competition for a resource” – an interaction between the sub-problems. One can interpolate between the symmetrical and asymmetrical cases by measuring to what extent any single sub-problem dominates the overall cost.

It remains to define what “dominating the overall cost” means. Herein, we choose a simple maximization and normalization operation. We select the most costly goal and divide that by the cost of achieving the conjunction of all goals. For a conjunction C of facts let $cost(C)$ be the length of a shortest plan achieving C . Then, for a task with goal set G , $AsymRatio$ is defined as $max_{g \in G} cost(g) / cost(\bigwedge_{g \in G} g)$. $AsymRatio$ ranges between 0 and 1. Values close to 0 correspond to the symmetrical case, values close to 1 correspond to the asymmetrical case. With the intuitions explained above, $AsymRatio$ should be thought of as an indirect measure of the degree of sub-problem interactions. That is, we interpret the value of $AsymRatio$ as an *effect* of those interactions. An important open question is whether more direct measures – syntactic definitions of the *causes* of sub-problem interactions – can be found. Starting points for investigations in this direction may be existing investigations of “sub-goals” and their dependencies [1, 33, 34]; we outline these investigations, and their possible relevance, in Section 8.

Note that $AsymRatio$ is a rather simplistic parameter. Particularly, the use of maximization over individual costs, disregarding any interactions between the facts, can be harmful. The maximally costly goal may be independent of the other goals. In such a case, while taking many steps to achieve, the goal is not the main reason for the length of the optimal plan. An example for this is given in Section 7; we henceforth refer to this as the “red herring”.

To make the above concrete, Figure 1 provides a sketch of one of our synthetic domains, called *MAP*. One moves in an undirected graph and must visit a subset of the nodes. The domain has two parameters: the size, n , and the amount of structure, k . The number of nodes in the graph is $3n - 3$. The value of k changes the set of goal nodes. One always starts in the bottom node, L^0 . If k is 1, then all goal nodes have distance 1 from the start node. For every increase of k by 2, a single one of the goal nodes wanders two steps (edges) further away, and one of the other goal nodes is skipped. As a result, the overall plan length is $2n - 1$ independently of k . Our formulas encode the unsolvable task of finding a plan with at most $2n - 2$ steps; they have $\Theta(n^2)$ variables ($\Theta(n)$ graph nodes at $\Theta(n)$ time steps).

¹Please note that this use of the word “symmetrical” has nothing to do with the wide-spread notion of “problem symmetries” in the sense of problem permutations; the cost of the goals has no implications on whether or not parts of the problem can be permuted.

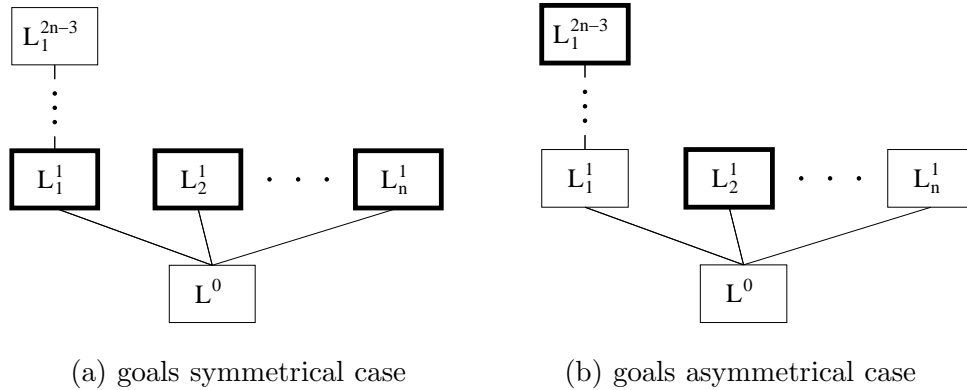


Figure 1: The MAP domain. Goal nodes indicated in bold face.

We have $AsymRatio = k/(2n - 1)$. In particular, for increasing n , $AsymRatio$ converges to 0 for $k = 1$ – symmetrical case – and it converges to 1 for $k = 2n - 3$ – asymmetrical case.² Note how all the goal nodes interact in the symmetrical case, competing for the time steps. Note also that the outlier goal node is not independent of the other goals, but interacts with them in the same competition for the time steps. In our red herring example, Section 7, the main idea is to make the outlier goal node independent by placing it on a separate map, and allowing to move in parallel on both maps. Note finally that, in Figure 8 (a), the graph nodes L_2^1, \dots, L_n^1 can be permuted. This is an artefact of the abstract nature of MAP. In some experiments in competition examples, we did not find a low $AsymRatio$ to be connected with a high amount of problem symmetries (see also Sections 5.1).

Beside MAP, we constructed two domains called SBW and SPH. SBW is a block stacking domain. There are n blocks, and k controls the amount of stacking restrictions. In the symmetrical case, there are no restrictions. In the asymmetrical case, one particular stack of blocks takes many steps to build. SPH is a non-planning domain, namely a structured version of the pigeon hole problem. There are $n + 1$ pigeons and n holes. The parameter k controls how many holes one particular “bad” pigeon needs, and how many “good” pigeons there are, which can share a hole with the “bad” one. The total number of holes needed remains $n + 1$, independently of k . The symmetrical case is the standard pigeon hole. In the asymmetrical case, the bad pigeon needs $n - 1$ holes.

1.2. Results Overview. In our research, the analysis of synthetic domains served to find and explore intuitions about how structure influences search performance. The final results of the analysis are studies of a set of prototypical behaviors. These prototypical behaviors are inherent also to more realistic examples; in particular, we will outline some examples from the competition benchmarks.

We prove upper and lower bounds on the size of the best-case DPLL proof trees. We also investigate the best possible sets of branching variables. Such variable sets were recently coined “backdoors” [61]. In our context, a backdoor is a subset of the variables so that, for every value assignment to these variables, unit propagation (UP) yields an empty clause.³

²When setting $k = 2n - 1$, the formula contains an empty clause.

³In general, a backdoor is defined relative to an arbitrary polynomial time “subsolver” procedure. The subsolver can solve some class of formulas that does not necessarily have a syntactic characterization. Our definition here instantiates the subsolver with the widely used unit propagation procedure.

That is, a smallest possible backdoor encapsulates the best possible branching variables for DPLL, a question of huge practical interest. Identifying backdoors is also a technical device: we obtain our upper bounds as a side effect of the proofs of backdoor properties. In all formula classes we consider, we determine a backdoor subset of variables. We prove that the backdoors are *minimal*: no variable can be removed without losing the backdoor property. In small enough instances, we prove empirically that the backdoors are in fact *optimal* - of minimal size. We conjecture that the latter is true in general.

In the symmetrical case, for MAP and SPH there are exponential (in n) lower bounds on the size of resolution refutations, and thus on DPLL refutations [7, 4]. For SPH, the lower bound is just the known result [8] for the standard Pigeon Hole problem. For MAP, we construct a polynomial reduction of resolution proofs for MAP to resolution proofs for a variant of the Pigeon Hole problem [48]. The reduction is via an intriguing temporal version of the pigeon hole problem, where the holes correspond to the time steps in the planning encoding, in a natural way. This illustrates quite nicely the competition of tasks for time steps that underlies also more realistic examples.

For SBW, it is an open question whether there exists an exponential lower bound on DPLL proof size in the symmetrical case; we conjecture that there is. In all the domains, the backdoor sets in the symmetrical case are linear in the total number of variables: $\Theta(n^2)$ for MAP and SPH, $\Theta(n^3)$ for SBW.

In the asymmetrical case, the DPLL proofs and backdoors become much smaller. In SPH, the minimal backdoors have size $O(n)$. What surprised us most is that in MAP and SBW the backdoors even become logarithmic in n . Considering Figure 1 (b), one would suspect to obtain an $O(n)$ backdoor reasoning along the path to the outlier goal node. However, it turns out that one can pick the branching variables in a way exploiting the need to go back and forth on that path. Going back and forth introduces a factor 2, and so one can double the number of time steps between each pair of variables. The resulting backdoor has size $O(\log n)$. This very nicely complements recent work [61], where several benchmark planning tasks were identified that contain exorbitantly small backdoors in the order of 10 out of 10000 variables. Our scalable examples with provably logarithmic backdoors illustrate the reasons for this phenomenon. Notably, the DPLL proof trees induced by our backdoors in the asymmetrical case degenerate to lines. Thus we get an exponential gap in DPLL proof size for SPH, and a doubly exponential gap for MAP.

To confirm that *AsymRatio* is an indicator of SAT solver performance in more practical domains, we run large-scale experiments in 10 domains from the biennial International Planning Competitions since 2000. The main criterion for domain selection is the availability of an instance generator. These are necessary for our experiments, where we generate and examine thousands of instances in each domain, in order to obtain large enough samples with identical plan length and *AsymRatio* (see below). We use the most successful SAT encoding for Planning: the “Graphplan-based” encoding [40, 41], which has been used in all planning competitions since 1998. We examine the performance of a state-of-the-art SAT solver, namely, ZChaff [46]. We compare the distributions of search tree size for pairs \mathcal{P}^x and \mathcal{P}^y of sets of planning tasks. All the tasks share the same domain, the same instance size parameters (taken from the original competition instances), and the same optimal plan length. The only difference is that *AsymRatio* = x for the tasks in \mathcal{P}^x , and *AsymRatio* = y for the tasks in \mathcal{P}^y , where $x < y$. Very consistently, the mean search tree size is significantly higher in \mathcal{P}_m^x than in \mathcal{P}_m^y . The T test yields confidence level 95% (99.9%, most of the time) in the vast majority of the cases in 8 of the 10 domains. In one domain (Logistics), the T

test fails for almost all pairs \mathcal{P}_m^x and \mathcal{P}_m^y ; in another domain (Satellite), *AsymRatio* does not show any variance so there is nothing to measure.

Some remarks are in order. Note that *AsymRatio* characterizes a kind of hidden structure. It cannot be computed efficiently even based on the original planning task representation – much less based on the SAT representation. This nicely reflects practical situations, where it is usually impossible to tell a priori how difficult a formula will be to handle. In fact, one interesting feature of our MAP formulas is that their syntax hardly changes between the two structural extreme cases. The content of a single clause makes all the difference between exponential and logarithmic DPLL proofs.

Note further that *AsymRatio* is not a completely impractical parameter. There exists a wealth of well-researched techniques approximating plan length, e.g. [5, 6, 32, 18, 24, 27, 26]. Such techniques can be used to approximate *AsymRatio*. We leave this as a topic for future work.

Note finally that, of course, *AsymRatio* can be fooled. (1) One can easily modify the goal set in a way blinding *AsymRatio*, for example by replacing the goal set with a single goal that has the same semantics. (2) One can also construct examples where a relevant phenomenon is “hidden” behind an irrelevant phenomenon that controls *AsymRatio*, and DPLL tree size grows, rather than decreases, with *AsymRatio*. For (1), we outline in Section 8 how this could be circumvented. As for (2), we explain the construction of such an example in Section 7. As a reply to both, it is normal that heuristics can be fooled. What matters is that *AsymRatio* often is informative in the domains that people actually try to solve.

The paper is organized as follows. Section 2 discusses related work. Section 3 provides notation and some details on the SAT encodings we use. Section 4 contains our experiments with *AsymRatio* in competition benchmarks. Section 5 describes our synthetic domains and our analysis of DPLL proofs; Section 6 briefly demonstrates that state-of-the-art SAT solvers – ZChaff [46] and MiniSat [20, 19] – indeed behave as expected, in these domains. Section 7 describes our red herring example, where *AsymRatio* is *not* an indicator of DPLL proof size. Section 8 concludes and discusses open topics.

For readability, the paper includes only proof sketches. The full proofs, and some other details, can be looked up in a TR [31].

2. RELATED WORK

A huge body of work on structure focusses on phase transition phenomena, e.g., [10, 23, 35, 50]. There are at least two major differences to our work. First, as far as we are aware, all work on phase transitions has to do with transitions between areas of under-constrained instances and over-constrained instances, with the critically constrained area – the phase transition – in the middle. In contrast, the Planning formulas we consider are all just one step short of a solution. In that sense, they are all “critically constrained”. As one increases the plan length bound, for a single planning task, from 0 to ∞ , one naturally moves from an over-constrained into an under-constrained area, where the bounds close to the first satisfiable iteration constitute the most critically constrained region. Put differently, phase transitions are to do with the balance of duties and resources; in our formulas, per definition, the amount of resources is set to a level that is just not enough to fulfill the duties.

A second difference to phase transitions is that these are mostly concerned with random instance distributions. In such distributions, the typical instance hardness is completely

governed by the parameter settings – e.g., the numbers of variables and clauses in a standard k -CNF generation scheme. This is very much not so in more practical instance distributions. If the contents of the clauses are extracted from some practical scenario, rather than from a random number generator, then the numbers of variables and clauses alone are not a good indicator of instance hardness: these numbers may not reflect the semantics of the underlying application. A very good example for this are our MAP formulas. These are syntactically almost identical at both ends of the structure scale; their DPLL proofs exhibit a doubly exponential difference.

Another important strand of work on structure is concerned with the identification of tractable classes, e.g., [9, 17, 11]. Our work obviously differs from this in that we do not try to identify general provable connections between structure and hardness. We identify empirical correlations, and we study particular cases.

Our analysis of synthetic examples is, in spirit, similar to the work in proof complexity, e.g., [13, 25, 8, 38]. There, formula families such as pigeon hole problems have been the key to a better understanding of resolution proofs. Generally speaking, the main difference is that, in proof complexity, one investigates the behavior of different proof calculi in the same example. In contrast, we consider the single proof calculus DPLL, and modify the examples. Major technical differences arise also due to the kinds of formulas considered, and the central goal of the research. Proof complexity considers any kind of synthetic formula provoking a certain behavior, while we consider formulas from Planning. The goal in proof complexity is mostly to obtain lower bounds, separating the power of proof systems. However, to understand structure, and explain the good performance of SAT solvers, interesting formula families with *small* DPLL trees are more revealing.

There is some work on problem structure in the Planning community. Some works [3, 56] investigate structure in the context of causal links in partial-order planning. Howe and Dahlman [36] analyze planner performance from a perspective of syntactic changes and computational environments. Hoffmann [34] investigates topological properties of certain wide-spread heuristic functions. Obviously, these works are quite different from ours. A more closely related piece of work is the aforementioned investigation of backdoors recently done by Williams et al [61]. In particular, this work showed empirically that CNF encodings of many standard Planning benchmarks contain exorbitantly small backdoors in the order of 10 out of 10000 variables. The existence of logarithmic backdoors in our synthetic domains nicely reflects this. In contrast to the previous results, we also explain what these backdoor variables are – what they correspond to in the original planning task – and how their interplay works.

A lot of work on structure can also be found in the Scheduling community, e.g., [59, 58, 60, 55]. To a large extent, this work is to do with properties of the search space surface, and its effect on the performance of local search. A closer relative to our work is the notion of *critical paths* as used in Scheduling, e.g., [42, 62, 45]. Based on efficiently testable inferences, a critical path identifies the (apparently) most critically constrained part of the scheduling task. For example, a critical path may identify a long necessary sequence of job executions implied by the task specification. This closely corresponds to our notion of a high *AsymRatio*, where a single goal fact is almost as costly to solve as the entire planning task. Indeed, such a goal fact tends to ease search by providing a sort of critical path on which a lot of constraint propagation happens. In that sense, our work is an application of critical paths to Planning. Note that Planning and Scheduling are different. In Planning,

there is much more freedom of problem design. Our main observation in here is that our notions of problem structure capture interesting behavior across a *range* of domains.

There is also a large body of work on structure in the constraint reasoning community, e.g., [16, 10, 23, 22, 51, 54, 12, 17, 47, 37, 11]. As far as we are aware, all these works differ considerably from ours. In particular, all works we are aware of define “structure” on the level of the CNF formula/the CSP problem instance; in contrast, we define structure on the level of the modelled application. Further, empirical work on structure is mostly based on random problem distributions, and theoretical analysis is mostly done in a proof complexity sense, or in the context of identifying tractable classes.

One structural concept from the constraint reasoning community is particularly closely related to the concept of a backdoor: *cutsets* [16, 51, 17]. Cutset are defined relative to the *constraint graph*: the undirected graph where nodes are variables and edges indicate common membership in at least one clause. A cutset is a set of variables so that, once these variables are removed from the constraint graph, that graph has a bounded *induced width*; if the bound is 1, then the graph is cycle-free, i.e., can be viewed as a tree. Backdoors are a generalization of cutsets in the sense that any cutset is a backdoor relative to an appropriate subsolver (that exploits properties of the constraint graph). The difference is that cutsets have an “easy” syntactic characterization: one can check in polytime if or not a given set of variables is a cutset. One can, thus, use strategies looking for cutsets to design search algorithms. Indeed, the cutset notion was originally developed with that aim. Backdoors, in contrast, were proposed as a means to characterize phenomena relevant for existing state-of-the-art solvers – which all make use of subsolvers whose capabilities (the solved classes of formulas) have no easy-to-test syntactic characterization. In particular, the effect of unit propagation depends heavily on what values are assigned to the backdoor variables. We will see that, in the formula families considered herein, there are no small cutsets.

3. PRELIMINARIES

We use the STRIPS formalism. States are described as sets of (the currently true) propositional facts. A planning *task* is a tuple of *initial state* (a set of facts), *goal* (also a set of facts), and a set of *actions*. Actions a are fact set triples: the *precondition* $pre(a)$, the *add effect* $add(a)$, and the *delete effect* $del(a)$. The semantics are that an action is applicable to a state (only) if $pre(a)$ is contained in the state. When executing the action, the facts in $add(a)$ are included into the state, and the facts in $del(a)$ are removed from it. The intersection between $add(a)$ and $del(a)$ is assumed empty; executing a non-applicable action results in an undefined state. A *plan* for the task is a sequence of actions that, when executed iteratively, maps the initial state into a state that contains the goal.

As a simple example, consider the task of finding a path from a start node n_0 to a goal node n_g in a directed graph (N, E) . The facts have the form $at-n$ for $n \in N$. The initial state is $\{at-n_0\}$, the goal is $\{at-n_g\}$, and the set of actions is $\{move-x-y = (\{at-x\}, \{at-y\}, \{at-x\}) \mid (x, y) \in E\}$ – precondition $\{at-x\}$, add effect $\{at-y\}$, delete effect $\{at-x\}$. Plans correspond to the paths between n_0 and n_g .

CNF formulas are sets of clauses, where each clause is a set of literals. For a CNF formula ϕ with variable set V , a variable subset $B \subseteq V$, and a truth value assignment a to B , by ϕ_a we denote the CNF that results from inserting the values specified by a , and removing satisfied clauses as well as unsatisfied literals. By $UP(\phi)$, we denote the result of

iterated application of unit propagation to ϕ , where again satisfied clauses and unsatisfied literals are removed. For a CNF formula ϕ with variable set V , a variable subset $B \subseteq V$, and a value assignment a to B , we say that a is *UP-consistent* if $UP(\phi_a)$ does not contain the empty clause. B is a *backdoor* if it has no UP-consistent assignment. We sometimes use the abbreviation *UP* also in informal text.

By a *resolution refutation*, also called *resolution proof*, of a (unsatisfiable) formula ϕ , we mean a sequence C_1, \dots, C_m of clauses C_i so that $C_m = \emptyset$, and each C_i is either an element of ϕ , or derivable by the resolution rule from two clauses C_j and C_k , $j, k < i$. The *size* of the refutation is m . A *DPLL refutation*, also called *DPLL proof* or *DPLL tree*, for ϕ is a tree of partial assignments generated by the DPLL procedure, where the inner nodes are UP-consistent, and the leaf nodes are not. The size of a DPLL refutation is the total number of (inner and leaf) tree nodes.

Planning can be mapped into a sequence of SAT problems, by incrementally increasing a plan length bound b : start with $b = 0$; generate a CNF $\phi(b)$ that is satisfiable iff there is a plan with b steps; if $\phi(b)$ is satisfiable, stop; else, increment b and iterate. This process was first implemented in the *Blackbox* system [39, 40, 41]. There are, of course, different ways to generate the formulas $\phi(b)$, i.e., there are different encoding methods. In our empirical experiments, we use the original *Graphplan-based* encoding used in *Blackbox*. Variants of this encoding have been used by *Blackbox* (more recently named *SATPLAN*) in all international planning competitions since 1998. In our theoretical investigations, we use a somewhat simplified version of the Graphplan-based encoding.⁴

The Graphplan-based encoding is a straightforward translation of a b -step *planning graph* [5] into a CNF. The encoding has b *time steps* $1 \leq t \leq b$. It features variables for facts at time steps, and for actions at time steps. The former encode commitments of the form “fact p is true/false at time t ”, the latter encode commitments of the form “action a is executed/not executed at time t ”. There are artificial *NOOP* actions, i.e. for each fact p there is an action *NOOP- p* whose only precondition is p , and whose only effect is p . The NOOPs are treated just like normal actions in the encoding. Setting a NOOP variable to true means a commitment to “keep fact p true at time t ”. Amongst others, there are clauses to ensure that all action preconditions are satisfied, that the goals are true in the last time step, and that no “mutex” actions are executed in the same time step: actions can be executed in the same time step – in parallel – if their effects and preconditions are not contradictory. The set of fact and action variables at each time step, as well as pairs of “mutex” facts and actions, are read off the planning graph (which is the result of a propagation of binary constraints).

We do not describe the Graphplan-based encoding in detail since that is not necessary to understand our experiments. For the simplified encoding used in our theoretical investigations, some more details are in order. The encoding uses variables only for the actions (including NOOPs), i.e., $a(t)$ is 1 iff action a is to be executed at time t , $1 \leq t \leq b$. A variable $a(t)$ is included in the CNF iff a is *present* at t . An action a is present at $t = 1$ iff a ’s precondition is true in the initial state; a is present at $t > 1$ iff, for every $p \in \text{pre}(a)$, at least one action a' is present at $t - 1$ with $p \in \text{add}(a')$. For each action a present at a time $t > 1$ and for each $p \in \text{pre}(a)$, there is a *precondition* clause of the form $\{-a(t), a_1(t-1), \dots, a_l(t-1)\}$, where a_1, \dots, a_l are all actions present at $t - 1$ with $p \in \text{add}(a_i)$. For each goal fact $g \in G$, there is a *goal* clause $\{a_1(b), \dots, a_l(b)\}$, where

⁴One might wonder whether a different encoding would fundamentally change the results herein. We don’t see a reason why that should be the case. Exploring this is a topic for future work.

a_1, \dots, a_l are all actions present at b that have $g \in \text{add}(a_i)$. Finally, for each *incompatible* pair a and a' of actions present at a time t , there is a *mutex* clause $\{\neg a(t), \neg a'(t)\}$. Here, a pair a, a' of actions is called incompatible iff either both are not NOOPs, or a is a NOOP for fact p and $p \in \text{del}(a')$.




Figure 2: An illustrative example: the planning task is to move from n_0 to n_g within 2 steps.

For example, reconsider the path-finding domain sketched above. Say $(N, E) = (\{n_0, n_1, n_2, n_g\}, \{(n_0, n_1), (n_0, n_2), (n_1, n_g), (n_2, n_g)\})$. There are two paths from n_0 to n_g , one through n_1 , the other through n_2 . An illustration is in Figure 2. The simplified Graphplan-based encoding of this task, for bound $b = 2$, is as follows. The variables for actions present at $t = 1$ are $\text{move-}n_0\text{-}n_1(1)$, $\text{move-}n_0\text{-}n_2(1)$, $\text{NOOP-at-}n_0(1)$. The variables for actions present at $t = 2$ are $\text{move-}n_0\text{-}n_1(2)$, $\text{move-}n_0\text{-}n_2(2)$, $\text{NOOP-at-}n_0(2)$, $\text{move-}n_1\text{-}n_g(2)$, $\text{move-}n_2\text{-}n_g(2)$, $\text{NOOP-at-}n_1(2)$, $\text{NOOP-at-}n_2(2)$. Note here that we can choose to do useless things such as staying at a node, via a NOOP action. We have to insert precondition clauses for all actions at $t = 2$. For readability, we only show the “relevant” clauses, i.e., those that suffice, in our particular example here, to get the correct satisfying assignments. The relevant precondition clauses are those for the actions $\text{move-}n_1\text{-}n_g(2)$ and $\text{move-}n_2\text{-}n_g(2)$, which are $\{\neg \text{move-}n_1\text{-}n_g(2), \text{move-}n_0\text{-}n_1(1)\}$ and $\{\neg \text{move-}n_2\text{-}n_g(2), \text{move-}n_0\text{-}n_2(1)\}$, respectively; in order to move from n_1 to n_g (n_2 to n_g) at time 2, we must move from n_0 to n_1 (n_0 to n_2) at time 1. We get similar clauses for the other (useless) actions, for example $\{\neg \text{move-}n_0\text{-}n_1(2), \text{NOOP-at-}n_0(1)\}$. We get the goal clause $\{\text{move-}n_1\text{-}n_g(2), \text{move-}n_2\text{-}n_g(2)\}$; to achieve our goal, we must move to n_g from either n_1 or n_2 , at time 2.⁵ We finally get mutex clauses. The relevant one is $\{\neg \text{move-}n_0\text{-}n_1(1), \neg \text{move-}n_0\text{-}n_2(1)\}$; we cannot move from n_0 to n_1 and n_2 simultaneously. We also get the clause $\{\neg \text{move-}n_1\text{-}n_g(2), \neg \text{move-}n_2\text{-}n_g(2)\}$, and various mutex clauses between move actions and NOOPs. Now, what the relevant clauses say is: 1. We must move to n_g at time 2, either from n_1 or from n_2 (goal). 2. If we move from n_1 to n_g at time 2, then we must move from n_0 to n_1 at time 1 (precondition). 3. If we move from n_2 to n_g at time 2, then we must move from n_0 to n_2 at time 1 (precondition). 4. We cannot move from n_0 to n_1 and n_2 simultaneously at time 1. Obviously, the satisfying assignments correspond exactly to the solution paths. Please keep in mind that, in contrast to our example here, in general *all* the clauses are necessary to obtain a correct encoding.

By making every pair of non-NOOPs incompatible in our simplified Graphplan-based encoding, we allow at most one (non-NOOP, i.e., “real”) action to execute per time step. This is a restriction in domains where there exist actions that can be applied in parallel; in our synthetic domains, no parallel actions are possible anyway, so the mutex clauses have an

⁵Note that, if b was 3, then we would also have the option to simply *stay* at n_g , i.e., the goal clause would be $\{\text{move-}n_1\text{-}n_g(3), \text{move-}n_2\text{-}n_g(3), \text{NOOP-at-}n_g(3)\}$. At $t = 2$, this NOOP is not present.

effect only on the power of UP (unit propagation). We also investigated backdoor size, in our synthetic domains, with a weaker definition of incompatible action pairs, allowing actions to be applied in parallel unless their effects and preconditions directly contradict each other. We omit the results for the sake of readability. In a nutshell, one obtains the same DPLL lower bounds and backdoors in the symmetrical case, but larger ($O(n)$) backdoors and DPLL trees in the asymmetrical case. Note that, thus, the restrictive definition of incompatible action pairs we use in our simplified encoding gives us an exponential efficiency advantage. This is in line with the use of the Graphplan-based encoding in our experiments: planning graphs discover the linear nature of our synthetic domains, including all the mutex clauses present in our simplified encoding.

4. GOAL ASYMMETRY IN PLANNING BENCHMARKS

In this section, we explore empirically how *AsymRatio* behaves in a range of Planning benchmarks. We address two main questions:

- (1) What is the distribution of *AsymRatio*?
- (2) How does *AsymRatio* behave compared to search performance?

Section 4.1 gives details on the experiment setup. Sections 4.2 and 4.3 address questions (1) and (2), respectively. Before we start, we reiterate how *AsymRatio* is defined.

Definition 4.1. Let P be a planning task with goal G . For a conjunction C of facts, let $cost(C)$ be the length of a shortest plan achieving C . The *asymmetry ratio* of P is:

$$AsymRatio(P) := \frac{\max_{g \in G} cost(g)}{cost(\bigwedge_{g \in G} g)}$$

Note that $cost(\bigwedge_{g \in G} g)$, in this definition, is the optimal plan length; to simplify notation, we will henceforth denote this with m . Note also that such a simple definition can not be fool-proof. Imagine replacing G with a single goal g , and an additional action with precondition G and add effect $\{g\}$; the (new) goal is then no longer a set of “sub-problems”. However, in the benchmark domains that are used by researchers to evaluate their algorithms, G is almost always composed of several goal facts resembling sub-problems. A more stable approach to define *AsymRatio* is a topic for future work, outlined in Section 8; for now, we will see that *AsymRatio* often works quite well.

4.1. Experiment Setup. Denote by $\phi(P, b)$, for a planning task P and integer b , the Graphplan-based CNF encoding of b action steps. Our general intuition is that $\phi(P, m - 1)$ is easier to prove unsatisfiable for tasks P with higher *AsymRatio*, than for tasks with lower *AsymRatio*, provided “all other circumstances are equal”. By this, we mean that the tasks all share the same *domain*, the same *size parameters*, and the same *optimal plan length*. We will not formalize the notions of “domain” and “size parameters”; doing so is cumbersome and not necessary to understand our experiments. An informal explanation is as follows:

- **Domain.** A domain, or *STRIPS domain*, is a (typically infinite) set of related planning tasks. More technically, a domain is defined through a finite set of logical predicates (relations), a finite set of *operators*, and an infinite set of *instances*. Operators are functions

from *object tuples* into STRIPS actions;⁶ they are described using predicates and variables. Applying an operator to an object tuple just means to instantiate the variables, resulting in a propositional STRIPS action description. Each instance defines a finite set of objects, an initial state, and a goal condition. Both initial state and goal condition are sets of instantiated predicates, i.e., propositional facts.

An example is the MAP domain previewed in Section 1, Figure 1. The predicates are “edge(x,y)”, “at(x)”, and “visited(x)” relations; the single operator has the form “move(x,y)”; the instances define the graph topology and the initial/goal nodes. A more general example would be a transportation domain with several vehicles and transportable objects, and additional operators loading/unloading objects to/from vehicles.

- **Size parameters.** An instance generator for a domain is usually parameterized in several ways. In particular, one needs to specify how many objects of each type (e.g., vehicles, transportable objects) there will be. For example, n is the size parameter in the MAP domain. If two instances of a transportation domain both contain the same numbers of locations, vehicles, and transportable objects, then they share the same size parameters.

In our experiments, we always fix a domain \mathcal{D} , and a setting \mathcal{S} of the size parameters of an instance generator for \mathcal{D} . We then generate thousands of (randomized) instances, and divide those into classes \mathcal{P}_m with identical optimal plan length m . Within each \mathcal{P}_m , we then examine the distribution of *AsymRatio*, and its behavior compared to performance, in the formulas $\{\phi(P, m - 1) \mid P \in \mathcal{P}_m\}$. Note here that the size parameters and the optimal plan length determine the size of the formula. So the restriction we make by staying inside the classes \mathcal{P}_m basically comes down to fixing the domain, and fixing the formula size.

We run experiments in a range of 10 STRIPS domains taken from the biennial International Planning Competition (IPC) since 2000. The main criterion for domain selection is the availability of an instance generator:

- **IPC 2000.** From this IPC, we select the Blocksworld, Logistics, and Miconic-ADL domains. Blocksworld is the classical block-stacking domain using a robot arm to rearrange the positions of blocks on a table; we use the instance generator by Slaney and Thiebaux [53]. Logistics is a basic transportation domain involving cities, places, trucks, airplanes, and packages. Each city contains a certain number of places; trucks move within cities, airplanes between them; packages must be transported. Moves are instantaneous, i.e., a truck/airplane can move in one action between any two locations. We implemented a straightforward instance generator ourselves. Miconic-ADL is an elevator transport domain coming from a real-world application [43]. It involves all sorts of interesting side constraints regarding, e.g., the prioritized transportation of VIPs, and constraints on which people have access to which floors. The 2000 IPC instance generator, which we use, was written by one of the authors. The domain makes use of first-order logic in action preconditions; we used the “adl2strips” software [29] to translate this into STRIPS.
- **IPC 2002.** Here, instance generators are provided by the organizers [44], so we select all the domains. They are named Depots, Driverlog, Freecell, Rovers, Satellite, and Zenotravel. Depots is a mixture between Blocksworld and Logistics, where blocks must be transported *and* arranged in stacks. Driverlog is a version of Logistics where trucks need drivers; apart from this, the main difference to the classical Logistics domain is that drivers and trucks move on directed graph road maps, rather than having instant access from any location to any other location. Freecell encodes the well-known solitaire card game

⁶In the planning community, constants are commonly referred to as “objects”; we adopt this terminology.

where the task is to re-order a random arrangement of cards, following certain stacking rules, using a number of “free cells” for intermediate storage. Rovers and Satellite are simplistic encodings of NASA space-applications. In Rovers, rovers move along individual road maps, and have to gather data about rock or soil samples, take images, and transfer the data to a lander. In Satellite, satellites must take images of objects, which involves calibrating cameras, turning the right direction, etc. Zenotravel is a version of Logistics where moving a vehicle consumes fuel that can be re-plenished using a “refuel” operator.

- **IPC 2004.** From these domains, the only one for which a random generator exists is called PSR, where one must reconfigure a faulty power supply network. PSR is formulated in a complex language involving first-order formulas and predicates whose value is derived as an effect of the values of other predicates. (Namely, the flow of power through the network is determined by the setting of the switches.) One can translate this to STRIPS, but only through the use of significant simplification methods [29]. In the thus simplified domain, every goal can be achieved in a single step, making *AsymRatio* devoid of information. We emphasize that this is not the case for the more natural original domain formulation.
- **IPC 2006.** From these domains, the only one that we can use is called TPP; for all others, instance generators were not available at the time of writing. TPP is short for Travelling Purchase Problem. Given sets of products and markets, and a demand for each product, one must select a subset of the markets so that routing cost and purchasing cost are minimized. The STRIPS version of this problem involves unit-cost products, and a road map graph for the markets.

We finally run experiments in a domain of purely random instances generated using Rintanen’s [50] “Model A”. We consider this an interesting contrast to the IPC domains. We will refer to this domain with the name “Random”.

To choose the size parameters for our experiments, we simply rely on the size parameters of the original IPC instances. This is justified since the IPC instances are the main benchmark used in the field. Precisely, we use the following method. For each domain, we consider the original IPC test suite, containing a set of instances scaling in size. For each instance, we generate a few random instances with its size parameters, and test how fast we can compute *AsymRatio*. (That computation is done by a combination of calls to Blackbox.) We select the largest instance for which each test run is completed within a minute. For example, in Driverlog we select the instance indexed 10 out of 20, and, accordingly, generate random instances with 6 road junctions, 2 drivers, 6 transportable objects, and 3 trucks. While the instances generated thus are not at the very limit of the performance of SAT-based planners, they are reasonably close to that limit. The runtime curves of Blackbox undergo a sharp exponential increase in all the considered domains. There is not much room between the last instance solved in a minute, and the first instance not solved at all. For example, in Driverlog we use instance number 10, while the performance limit of Blackbox is reached around instance number 12.⁷ In the random domain, we use “40 state variables” which is reasonable according to Rintanen’s [50] results. We also run a number

⁷In Logistics, we use IPC instance number 18 of 50, and the performance limit is around instance number 20. In Miconic-ADL, we use IPC instance 7 out of 30, and the performance limit is around instance 9. In Depots, these numbers are 8, 20, and 10; in Rovers, they are 8, 20, and 10; in Satellite, they are 7, 20, and 9; in Zenotravel, they are 13, 20, and 15; in TPP, they are 15, 30, and 20. In Blocksworld, we use 11 blocks and the performance limit is at around 13 blocks. In Freecell, Blackbox is very inefficient, reaching its limit around IPC instance 2 (of 20), which is what we use.

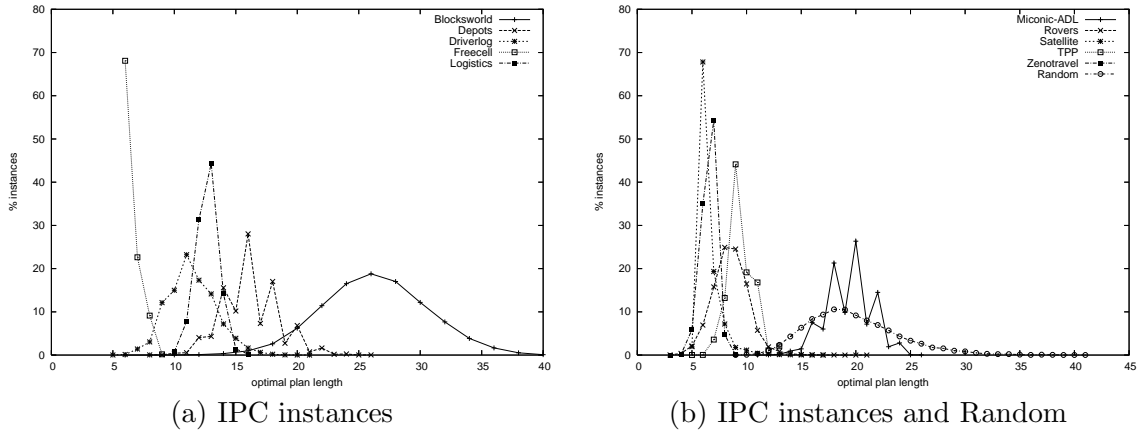


Figure 3: Distribution of optimal plan length. Split across two graphs for readability. of experiments where we modify the size parameters in certain domains, with the aim to obtain a picture of what happens as the parameters change. This will be detailed below.

In most cases, the distribution of optimal plan length m is rather broad, for given domain \mathcal{D} and size parameters \mathcal{S} ; see Figure 3.⁸ Thus we need many instances in order to obtain reasonably large classes \mathcal{P}_m . Further, we split each class \mathcal{P}_m into subsets, *bins*, with identical *AsymRatio*. To make the bins reasonably large, we need even more instances. In initial experiments, we found that between 5000 and 20000 instances were usually sufficient to identify the overall behavior of the domain and size parameter setting (\mathcal{D} and \mathcal{S}). To be conservative, we decided to fix the number of instances to 50000, per \mathcal{D} and \mathcal{S} . To avoid noise, we skip bins with less than 100 elements; the remaining bins each contain around a few 100 up to a few 1000 instances.

4.2. AsymRatio Distribution. What interests us most in examining the distribution of *AsymRatio* is how “spread out” the distribution is, i.e., how many different values we obtain within the \mathcal{P}_m classes, and how far they are apart. The more values we obtain, the more cases can we distinguish based on *AsymRatio*; the farther the values are apart, the more clearly will those cases be distinct. Figure 4 shows some of the data; for all the settings of \mathcal{D} and \mathcal{S} that we explored, it shows the *AsymRatio* distribution in the most populated class \mathcal{P}_m .

In Figure 4, the x axes show *AsymRatio*, and the y axes show percentage of instances within \mathcal{P}_m . Let us first consider Figure 4 (a) and (b), which show the plots for the IPC settings \mathcal{S} of the size parameters. In Figure 4 (a), we see vaguely normal distributions for Blocksworld, Depots, and Driverlog. Freecell is unusual in its large weight at high *AsymRatio* values, and also in having relatively few different *AsymRatio* values. This can be attributed to the unusually small plan length in Freecell.⁹ For Logistics, we note that there are only 2 different values of *AsymRatio*. This can be attributed to the trivial road maps in this domain, where one can instantaneously move between any two locations. The

⁸We see that the plan length is mostly normal distributed. The only notable exception is the Freecell domain, Figure 3 (a), where the only plan lengths occurring are 6, 7, 8, and 9, and \mathcal{P}_6 is by far the most populated class.

⁹Remember that our optimal planner scales only to instance number 2 out of 20 in the IPC 2002 test suite.




Figure 4: Distribution of $AsymRatio$, for the respective most populated class \mathcal{P}_m as indicated. (a) and (b), for IPC settings of \mathcal{S} . (c), for various settings of \mathcal{S} in Logistics. (d), for various settings of \mathcal{S} in Driverlog. (e), for two variants of Miconic-ADL. Explanations see text.

instance shown in Figure 4 has 6 cities and 2 airplanes, and thus there is not much variance in the cost of achieving a single goal (transporting a single package). This will be explored further below.

In Figure 4 (b), at first glance one sees that there are quite a few odd-looking distributions. The distributions for Rovers and Random are fairly spread out. Miconic-ADL has 77% of its weight at $AsymRatio = 0.20$, 18% of its weight at $AsymRatio = 0.30$, and 2.4% at $AsymRatio = 0.4$; a lot of other $AsymRatio$ values have non-zero weights of less than 1%. This distribution can be attributed to the structure of the domain, where the elevator can instantaneously move between any two floors; most of the time, each goal (serving a passenger) thus takes 4 steps to achieve, corresponding to $AsymRatio = 0.25$ in our picture. Due to the side constraints, however, sometimes serving a passenger involves more effort (e.g., some passengers need to be attended in the lift) – hence the instances with higher $AsymRatio$ values. We will consider below a Miconic-ADL variant with non-instantaneous lift moves.

Zenotravel has an extreme $AsymRatio$ distribution, with 99% of the weight at $AsymRatio = 0.57$. Again, this can be attributed to the lack of a road map graph, i.e., to instantaneous vehicle moves. In TPP, 94% of the time there is a single goal that takes as many steps to achieve as the entire goal set. It is unclear to us what the reason for that is (the goal sets are large, containing 10 facts in all cases). The most extreme $AsymRatio$ distribution is exhibited by Satellite: in all classes \mathcal{P}_m in our experiments, all instances have the same $AsymRatio$ value ($AsymRatio = 0.67$ in Figure 4 (b)). Partly, once again this is because of instantaneous “moves” – here, changing the direction a satellite observes. Partly it seems to be because of the way the 2002 IPC Satellite instance generator works, most of the time including at least one goal with the same maximal cost.

To sum the above up, $AsymRatio$ does show a considerable spread of values across all our domains except Logistics, Satellite, TPP, and Zenotravel. For Logistics, Satellite, and Zenotravel, this can be attributed to the lack of road map graphs in these domains. This is a shortcoming of the domains, rather than a shortcoming of $AsymRatio$; in reality one does not move instantaneously (though sometimes one would wish to ...). We ran additional experiments in Logistics to explore this further. We generated instances with 8 instead of 6 cities, and with just a single airplane instead of 2 airplanes; we had to decrease the number of packages from 18 to 8 to make this experiment feasible. We then repeated the experiment, but with the number of airplanes set to 2 and 3, respectively. Figure 4 (c) shows the results, including the IPC \mathcal{S} setting for comparison. “Logistics- xa ” denotes our respective new experiment with x airplanes. Considering these plots, we see clearly that the $AsymRatio$ distribution shifts to the right, and becomes more dense, as we increase the number of airplanes; eventually (with more and more airplanes) all weight will be in a single spot.

In Driverlog, we run experiments to explore what happens as we change the relative numbers of vehicles and transportable objects. Figure 4 (d) shows the results. The IPC instance has 6 road junctions, 2 drivers, 6 transportable objects, and 3 trucks. For “Driverlog-1t-10p” in Figure 4 (d), we set this to 6 road junctions, 1 driver, 10 transportable objects, and 1 truck. For “Driverlog-1t-20p” and “Driverlog-1t-30p”, we increased the number of transportable objects to 20 and 30, respectively. “Driverlog-5t-5p”, on the other hand, has 6 road junctions, 5 drivers, 5 transportable objects, and 5 trucks. The intuition is that the position of the $AsymRatio$ distribution depends on the ratio between the number of transportable objects and the number of means for transport. The higher that ratio is, the

lower do we expect *AsymRatio* to be – if there are many objects for a single vehicle then it is unlikely that a single object will take all the time. Also, for very low and very high values of the ratio we expect the distribution of *AsymRatio* to be dense – if every object takes almost all the time/no object takes a significant part of the time, then there should not be much variance. Figure 4 (d) clearly shows this tendency for “Driverlog-5t” (ratio = 5/5), “Driverlog-IPC” (ratio = 6/3), “Driverlog-10p” (ratio = 10/1), and “Driverlog-20p” (ratio = 20/1). Interestingly, the step from “Driverlog-20p” to “Driverlog-30p” does not make as much of a difference. Still, if we keep increasing the number of transportable objects then eventually the distribution will trivialize. Note here that 30 objects in a road map with only 6 nodes already constitute a rather dense transportation problem. Further, transport problems with 30 objects really push the limits of the capabilities of current SAT solvers.

Let us finally consider Figure 4 (e). “Miconic-ADL Modified” denotes a version of Miconic-ADL where the elevator is constrained to take one move action for every move between adjacent floors, rather than moving between any two floors instantaneously. As one would expect, this changes the distribution of *AsymRatio* considerably. We get a distribution spreading out from *AsymRatio* = 0.34 to *AsymRatio* = 0.75, and further – with very low percentages – to *AsymRatio* = 0.97.

4.3. AsymRatio and Search Performance. Herein, we examine whether *AsymRatio* is indeed an indicator for search performance. Let us first state our hypothesis more precisely. Fixing a domain \mathcal{D} , a size parameter setting \mathcal{S} , and a number $\delta \in [0; 1]$, our hypothesis is:

Hypothesis 1 ($\mathcal{D}, \mathcal{S}, \delta$). *Let m be an integer, and let $x, y \in [0; 1]$, where $y - x > \delta$. Let \mathcal{P}_m be the set of instances in \mathcal{D} with size \mathcal{S} and optimal plan length m . Let $\mathcal{P}_m^x = \{P \in \mathcal{P}_m \mid \text{AsymRatio}(P) = x\}$ and $\mathcal{P}_m^y = \{P \in \mathcal{P}_m \mid \text{AsymRatio}(P) = y\}$.*

If both \mathcal{P}_m^x and \mathcal{P}_m^y are non-empty, then the mean DPLL search tree size is significantly higher for $\{\phi(P, m - 1) \mid P \in \mathcal{P}_m^x\}$ than for $\{\phi(P, m - 1) \mid P \in \mathcal{P}_m^y\}$.

The hypothesis is parameterized by \mathcal{D} , \mathcal{S} , and by an *AsymRatio difference threshold* δ . Regarding \mathcal{D} , we do not mean to claim that the stated correlation will be true for every domain; we construct a counter example in Section 7, and we will see that Hypothesis 1 is not well supported by two of the domains in our experiments. Regarding \mathcal{S} , we have already seen that the size parameter setting influences the distribution of *AsymRatio*; we will see that this is also true for the behavior of *AsymRatio* compared to performance. Regarding the *AsymRatio* difference threshold δ , this serves to trade-off the strength of the claim vs its applicability. In other words, one can make Hypothesis 1 weaker by increasing δ , at the cost of discriminating between less classes of instances. We will explore different settings of δ below.

For every planning task P generated in our experiments, with optimal plan length m , we run ZChaff[46] on the formula $\phi(P, m - 1)$. We measure the search tree size (number of backtracks) as our indication of how hard it is to prove a formula unsatisfiable. We compare the distributions of search tree size between pairs of *AsymRatio* bins within a class \mathcal{P}_m . Some quantitative results will be discussed below. First, we show qualitative results summarizing how often the distributions differ significantly. For every \mathcal{D} and \mathcal{S} , we distinguish between three different values of δ ; these will be explained below. For every class \mathcal{P}_m , and for every pair $\mathcal{P}_m^x, \mathcal{P}_m^y$ of *AsymRatio* bins with $y - x > \delta$, we run a T-test to determine whether or not the mean search tree sizes differ significantly. Namely, we run

Experiment	$\delta = 0$		$\delta_{95}(\mathcal{D}, \mathcal{S})$				$\delta_{100}(\mathcal{D}, \mathcal{S})$			
	95	99.9	δ	#	95	99.9	δ	#	95	99.9
Blocksworld	94	87	0.07	84	95	90	0.15	55	100	97
Depots	84	78	0.08	70	95	90	0.16	39	100	98
Driverlog-IPC	97	94	0.00	100	97	94	0.06	92	100	98
Freecell	100	100	0.00	100	100	100	0.00	100	100	100
Logistics-IPC	16	16	0.28	0			0.28	0		
Miconic-ADL Org	85	70	0.15	29	100	87	0.15	29	100	87
Rovers	96	96	0.00	100	96	96	0.06	96	100	100
TPP	100	100	0.00	100	100	100	0.00	100	100	100
Zenotravel	100	50	0.00	100	100	50	0.00	100	100	50
Miconic-ADL Mod	86	74	0.04	77	97	91	0.09	49	100	100
Random	40	27	0.24	5	100	71	0.24	5	100	71
Logistics-1a	91	83	0.08	62	100	100	0.08	62	100	100
Logistics-2a	77	77	0.12	55	100	100	0.12	55	100	100
Logistics-3a	66	66	0.12	33	100	100	0.12	33	100	100
Driverlog-5t-5p	90	90	0.15	50	100	100	0.15	50	100	100
Driverlog-1t-10p	93	86	0.06	72	96	94	0.11	53	100	98
Driverlog-1t-20p	88	84	0.10	49	96	96	0.13	33	100	100
Driverlog-1t-30p	80	75	0.11	37	95	90	0.13	28	100	94

Table 1: Summary of T-test results comparing the distributions of ZChaff’s search tree size for every pair $\mathcal{P}_m^x, \mathcal{P}_m^y$ of *AsymRatio* bins within every class \mathcal{P}_m , for every \mathcal{D} and \mathcal{S} , and different settings of δ ; explanation see text. The “ δ ” columns give the value of δ ; the “#” columns give the percentage of pairs with $y - x > \delta$; the “95” and “99.9” columns give the percentage of pairs whose distribution means differ significantly as hypothesized, at the respective level of confidence.

the Student’s T-test for unequal sample sizes, and check whether the means differ with a confidence of 95% and/or a confidence of 99.9%. Table 1 summarizes the results.

The three different values of δ distinguished for every \mathcal{D} and \mathcal{S} are $\delta = 0$, and two values called $\delta_{95}(\mathcal{D}, \mathcal{S})$ and $\delta_{100}(\mathcal{D}, \mathcal{S})$. $\delta = 0$ serves to show the situation for all pairs; $\delta_{95}(\mathcal{D}, \mathcal{S})$ and $\delta_{100}(\mathcal{D}, \mathcal{S})$ show how far one has to increase δ in order to obtain 95% and 100% accuracy of Hypothesis 1, respectively. Precisely, $\delta_{95}(\mathcal{D}, \mathcal{S})$ is the smallest number $\delta \in \{0, 0.01, 0.02, \dots, 1.0\}$ so that the T-test succeeds for at least 95% of the pairs with $y - x > \delta$, with confidence level 95%. Similarly, $\delta_{100}(\mathcal{D}, \mathcal{S})$ is the smallest number $\delta \in \{0, 0.01, 0.02, \dots, 1.0\}$ so that the T-test succeeds for all the pairs with $y - x > \delta$, with confidence level 95%. Note that, for some \mathcal{D} and \mathcal{S} , $\delta_{95}(\mathcal{D}, \mathcal{S}) = \delta_{100}(\mathcal{D}, \mathcal{S})$, or $0 = \delta_{95}(\mathcal{D}, \mathcal{S})$, or even $0 = \delta_{95}(\mathcal{D}, \mathcal{S}) = \delta_{100}(\mathcal{D}, \mathcal{S})$.

The upper most part of Table 1 is for the IPC domains and the respective settings of \mathcal{S} . Satellite is left out because not a single class \mathcal{P}_m in this domain contained more than one *AsymRatio* bin. Then there are separate parts for the modified version of Miconic-ADL, for the Random domain, and for our exploration of different \mathcal{S} settings in Logistics and Driverlog.

One quick way to look at the data is to just examine the leftmost columns, where $\delta = 0$. We see that there is good support for Hypothesis 1, with the T-test giving a 95% confidence level in the vast majority of cases (pairs of *AsymRatio* bins). Note that most of the T-tests succeed with a confidence level of 99.9%. Logistics-IPC and Random, and to some extent

Logistics-3a, are the only experiments (rows of Table 1) that behave very differently. In Freecell, TPP, and Zenotravel, there are only few *AsymRatio* pairs, c.f. Figure 4; precisely, we got 9 pairs in Freecell, 4 pairs in TPP, and only 2 pairs in Zenotravel. For all of these pairs, the T-test succeeds, with 99.9% confidence in all cases except one of the pairs in Zenotravel.

Another way to look at the data is to consider the values of $\delta_{95}(\mathcal{D}, \mathcal{S})$ and $\delta_{100}(\mathcal{D}, \mathcal{S})$, for each experiment. The smaller these values are, the more support do we have for Hypothesis 1. In particular, at $\delta_{100}(\mathcal{D}, \mathcal{S})$, every pair of *AsymRatio* bins encountered within 50000 random instances behaves as hypothesized. Except in Logistics-IPC and Random, the maximum $\delta_{100}(\mathcal{D}, \mathcal{S})$ we get in any of our experiments is the $\delta = 0.16$ needed for Depots. For some of the other experiments, $\delta_{100}(\mathcal{D}, \mathcal{S})$ is considerably lower; the mean over the IPC experiments is 9.55, the mean over all experiments is 11.27. The mean value of $\delta_{95}(\mathcal{D}, \mathcal{S})$ over the IPC experiments is 6.44 (3.75 without Logistics-IPC), the mean over all experiments is 8.88.

When considering $\delta_{95}(\mathcal{D}, \mathcal{S})$ and $\delta_{100}(\mathcal{D}, \mathcal{S})$, we must also consider the relative numbers of *AsymRatio* pairs that actually remain given these δ values. This information is provided in the “#” columns in Table 1. We can nicely observe how increasing δ trades off the accuracy of Hypothesis 1 vs its applicability. In particular, we see that no pairs remain in Logistics-IPC, and that hardly any pairs remain in Random; for these settings of \mathcal{D} and \mathcal{S} , the only way to make Hypothesis 1 “accurate” is by raising δ so high that it excludes almost all pairs. So here *AsymRatio* is useless. The mean percentage of pairs remaining at $\delta_{95}(\mathcal{D}, \mathcal{S})$ is 75.88 for the IPC experiments (85.37 without Logistics-IPC), and 60.16 for all experiments. The mean percentage of pairs remaining at $\delta_{100}(\mathcal{D}, \mathcal{S})$ is 67.88 for the IPC experiments (76.37 without Logistics-IPC), and 54.38 for all experiments.

As expected, in comparison to the original version of Miconic-ADL, *AsymRatio* is more reliable in our modified version, where lift moves are not instantaneous between any pair of floors. It is unclear to us what the reason for the unusually bad results in Logistics and Random is. A general speculation is that, in these domains, even if there is a goal that takes relatively many steps to achieve, this does not imply tight constraints on the solution. More concretely, regarding Logistics, consider a transportation domain, and a goal that takes many steps to achieve because it involves travelling a long way on a road map. Then the number of options to achieve the goal corresponds to the number of optimal paths on the map. Unless the map is densely connected, this will constrain the search considerably. In Logistics, however, the map is actually fully connected. So, (A), there is not much variance in how many steps a goal needs; and, (B), it is not a tight constraint to force one of the airplanes to move from one location to another at a certain time step, particularly if there are many airplanes. Given (A), it is surprising that in Logistics-1a *AsymRatio* is as good an indicator of performance as in most other experiments. Given (B), it is understandable why this phenomenon gets weaker for Logistics-2a and Logistics-3a. Logistics-IPC is even tougher than Logistics-3a, presumably due to the larger ratio between number of packages and road map size.

In comparison to Logistics, the results for Driverlog-5t-5p show the effect of a non-trivial road map. Even with 5 trucks for just 5 transportable objects, *AsymRatio* is a good performance indicator. Increasing the number of transportable objects over Driverlog-1t-10p, Driverlog-1t-20p, and Driverlog-1t-30p, *AsymRatio* becomes less reliable. Note, however, that stepping from Driverlog-1t-10p to Driverlog-1t-20p affects the behavior more than stepping from Driverlog-1t-20p to Driverlog-1t-30p. In the former step, $\delta_{95}(\mathcal{D}, \mathcal{S})$ and

$\delta_{100}(\mathcal{D}, \mathcal{S})$ increase by 0.04 and 0.02, respectively; in the latter step, $\delta_{95}(\mathcal{D}, \mathcal{S})$ increases by 0.01 and $\delta_{100}(\mathcal{D}, \mathcal{S})$ remains the same. Also, we reiterate that 30 objects on a map with 6 nodes – Driverlog-1t-30p – already constitute a very dense transportation problem, and that optimal planners do not scale beyond 30 objects anyway.¹⁰

To sum our observations up, on the negative side there are domains like Random where Hypothesis 1 is mostly wrong, and there are domains like Logistics where it holds only for very restricted settings of \mathcal{S} ; in other domains, *AsymRatio* is probably devoid in extreme \mathcal{S} settings. On the positive side, Hypothesis 1 holds in almost all cases – all pairs of *AsymRatio* bins – we encountered in the IPC domains (other than Logistics) and IPC parameter size settings. With increasing δ , the bad cases quickly disappear; in our experiments, all bad cases are filtered out at $\delta = 0.28$, and all bad cases except those in Logistics and Random are filtered out at $\delta = 0.16$.

It would be interesting to explore how the behavior of *AsymRatio* changes as a function of all the parameters of the domains, i.e., to extend our above observations regarding Logistics and Driverlog, and to perform similar studies for all the other domains. Given the number of parameters the domains have, such an investigation is beyond the scope of this paper, and we leave it as a topic for future work.

Figures 5 and 6 provide some quantitative results. In Figure 5, we take the mean value of each *AsymRatio* bin, and plot that over *AsymRatio*. We do so for a selection of domains, with IPC size parameter settings, and for a selection of classes \mathcal{P}_m . We select Blocksworld, Depots, Driverlog, Miconic-ADL, Rovers, and TPP; we show data for the 3 most populated \mathcal{P}_m classes. The decrease of mean search tree size over *AsymRatio* is very consistent, with some minor perturbations primarily in Blocksworld and Depots. The most remarkable behavior is obtained in Rovers, where the mean search tree size decreases exponentially over *AsymRatio*; note the logarithmic scale of the y axis in Figure 5 (e).¹¹ From the relative positions of the different curves, one can also nicely see the influence of optimal plan length/formula size – the longer the optimal plan, the larger the search tree.

Figure 6 provides some concrete examples of the search tree size distributions. They are plotted in terms of their survivor functions. In these plots, search tree size increases on the x axis, and the y axis shows the percentage of instances that require more than x search nodes. Figure 6 shows data for the same domains and size settings as Figure 5. For each domain, the maximum m of the 3 most populated \mathcal{P}_m classes is selected – in other words, we select the maximum m classes from Figure 5. Each graph contains a separate survivor function for every distinct value of *AsymRatio*, within the respective domain and \mathcal{P}_m . The graphs thus show how the survivor function changes with *AsymRatio*. The y axes are log-scaled to improve readability; without the log-scale, the outliers, i.e., the few instances with a very large search tree size, cannot be seen.

The behavior we expect to see, according to Hypothesis 1, is that the survivor functions shift to the left – to lower search tree sizes – as *AsymRatio* increases. Indeed, with only

¹⁰In fact, we were surprised that Driverlog instances with 30 objects can be solved. If one increases the number of trucks (state space branching factor) or the number of map nodes (branching factor and plan length) only slightly, the instances become extremely challenging; even with 20 packages and 10 map nodes, ZChaff often takes hours.

¹¹We can only speculate what the reason for this extraordinarily strong behavior in Rovers is. High *AsymRatio* values arise mainly due to long paths to be travelled on the road map. Perhaps there is less need to go back and forth in Rovers than in domains like Driverlog; one can transmit data from many locations. This might have an effect on how tightly the search gets constrained.

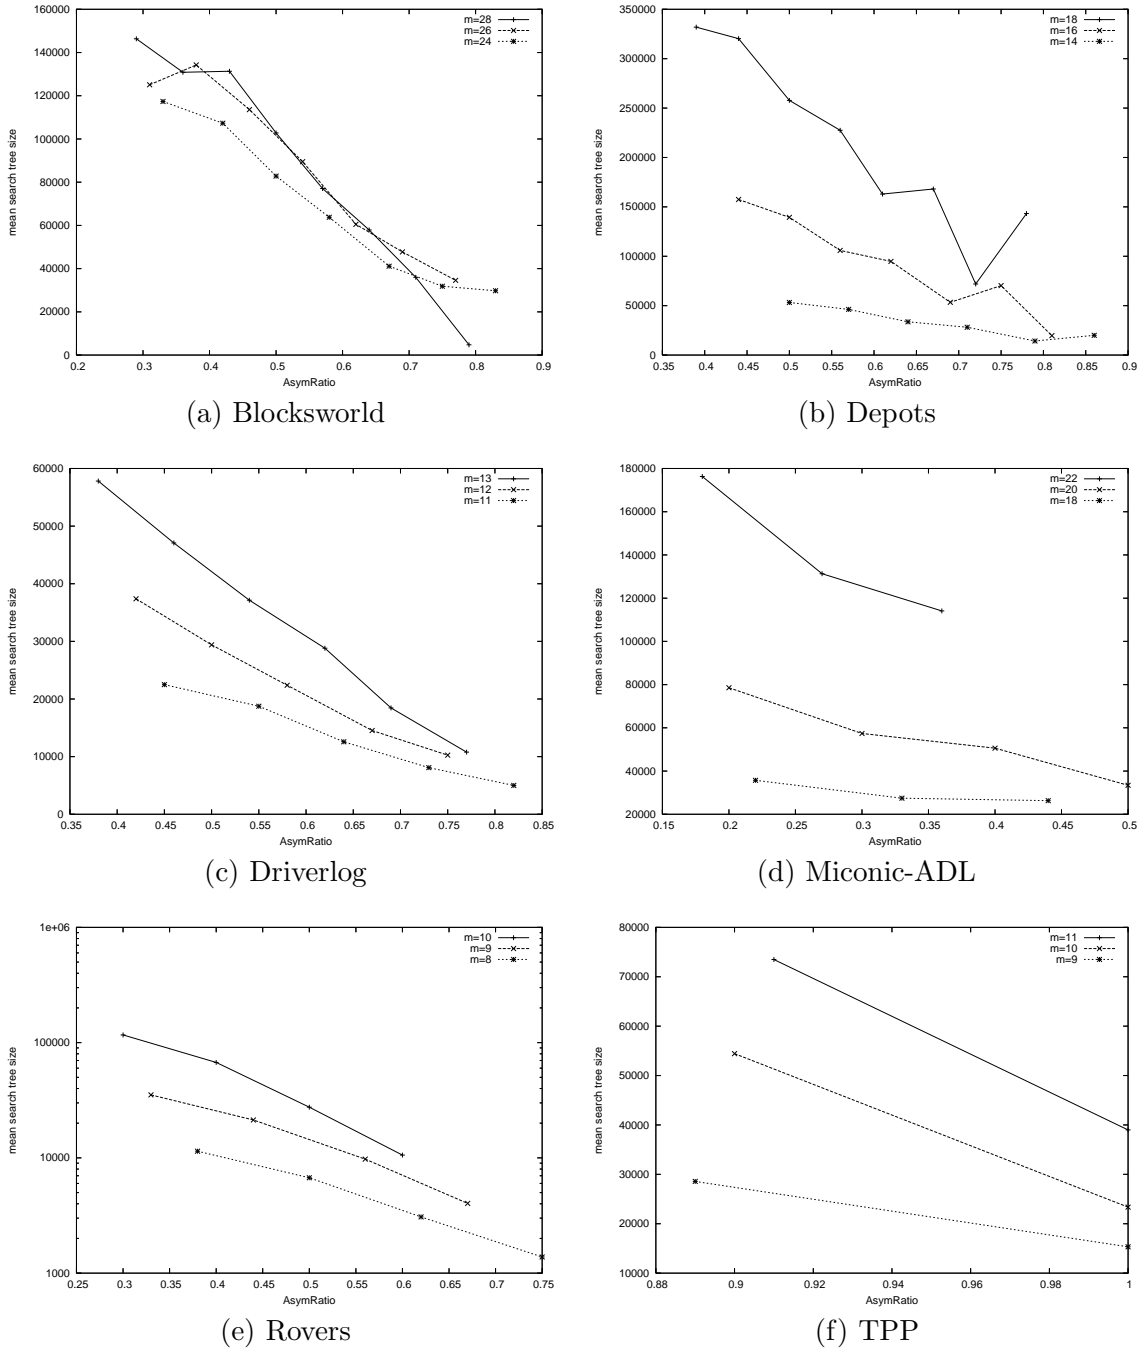


Figure 5: Mean search tree size of ZChaff, plotted against *AsymRatio* (log-plotted in (e)). Shown for the 3 most populated \mathcal{P}_m classes of the respective IPC settings of \mathcal{S} .




Figure 6: Search tree size distributions for different settings of $AsymRatio$, in terms of the survivor functions over search tree size; log-scaled in y . Shown for the maximum m class \mathcal{P}_m for each of the domains from Figure 5, with the IPC settings of \mathcal{S} .




Figure 7: Size/depth ratio of ZChaff’s search trees in Rovers. Left hand side: mean values plotted against *AsymRatio*. Right hand side: survivor functions for $m = 10$.

few exceptions, this is what happens. Consider the upper halves of the graphs, containing 99% of the instances in each domain. Almost without exception, the survivor functions are parallel curves shifting to the left over increasing *AsymRatio*. In the lower halves of the graphs, the picture is a little more varied; the curves sometimes cross in Blocksworld, Depots, and Driverlog. In Driverlog, for example, with *AsymRatio* = 0.37, the maximum costly instance takes 353953 nodes; with *AsymRatio* = 0.46, 0.1% of the instances require more search nodes than that. So it seems that, sometimes, *AsymRatio* is not as good an indicator on outliers. Still, the bulk of the distributions behaves according to Hypothesis 1.

Beside the search tree size of ZChaff, we also measured the (maximum) search tree depth, the size of the identified backdoors (the sets of variables branched upon), and the ratio between size and depth of the search tree. The latter gives an indication of how “broad” or “thin” the shape of the search tree is. Denoting depth with d , if the tree is full binary then the ratio is $(2^{d+1} - 1)/d$; if the tree is degenerated to a line then the ratio is $(2d + 1)/d$. Plotting these parameters over *AsymRatio*, in most domains we obtained behavior very similar to what is shown in Figures 5 and 6. As an example, Figure 7 shows the size/depth ratio data for the Rovers domain. We find the results regarding size/depth ratio particularly interesting. They nicely reflect the intuition that, as problem structure increases, UP can prune many branches early on and so makes the search tree grow thinner.

5. ANALYZING GOAL ASYMMETRY IN SYNTHETIC DOMAINS

In this section, we perform a number of case studies. We analyze synthetic domains constructed explicitly to provoke interesting behavior regarding *AsymRatio*. The aim of the analysis is to obtain a better understanding of how this sort of problem structure affects the behavior of SAT solvers; in fact, the definition of *AsymRatio* was motivated in the first place by observations we made in synthetic examples.

The analytical results we obtain in our case studies are, of course, specific to the studied domains. We do, however, identify a set of prototypical patterns of structure that also appear in the planning competition examples; we will point this out in the text.

We analyze three classes of synthetic domains/CNF formulas, called MAP, SBW, and SPH. MAP is a simple transportation kind of domain, SBW is a block stacking domain. SPH is a structured version of the pigeon hole problem. Each of the domains/CNF classes

is parameterized by size n and structure k . In the planning domains, we use the simplified Graphplan-based encoding described in Section 3, and consider CNFs that are one step short of a solution. We denote the CNFs with MAP_n^k , SBW_n^k , and SPH_n^k , respectively.

We choose the MAP and SBW domains because they are related to Logistics and Blocksworld, two of the most classical Planning benchmarks. We chose SPH for its close relation to the formulas considered in proof complexity. The reader will notice that the synthetic domains are *very* simple. The reasons for this are threefold. First, we wanted to capture the intended intuitive problem structure in as clean a form as possible, without “noise”. Second, even though the Planning tasks are quite simple, the resulting CNF formulas are complicated – e.g., much more complicated than the pigeon hole formulas often considered in proof complexity. Third, we identify provably minimal backdoors. To do so, one has to take account of every tiny detail of the effects of unit propagation. The respective proofs are already quite involved for our simple domains – for MAP, e.g., they occupy 9 pages, featuring myriads of interleaved case distinctions. To analyze more complex domains, one probably has to sacrifice precision.

For the sake of readability, herein we discuss only MAP in detail, and we replace the proofs with proof sketches. The details for SBW and SPH, and the proofs, are in the TR [31].

5.1. MAP.

5.1.1. *Domain Definition.* In this domain, one moves on the undirected graph shown in Figure 8 (a) and (b). The available actions take the form *move-x-y*, where x is connected to y with an edge in the graph. The precondition is $\{at-x\}$, the add effect is $\{at-y, visited-y\}$, and the delete effect is $\{at-x\}$.

The number of nodes in the graph is $3n - 3$. Initially one is located at L^0 . The goal is to visit a number of locations. Which locations must be visited depends on the value of $k \in \{1, 3, \dots, 2n - 3\}$. If $k = 1$ then the goal is to visit each of $\{L_1^1, \dots, L_n^1\}$. For each increase of k by 2, the goal on the L_1 -branch goes up by two steps, and one of the other goals is skipped. For $k = 2n - 3$ the goal is $\{L_1^{2n-3}, L_2^1\}$. (For $k = 2n - 1$, MAP_n^k contains an empty clause: no supporting action for the goal is present at the last time step.) We refer to $k = 1$ as the *symmetrical case*, and to $k = 2n - 3$ as the *asymmetrical case*, see Figure 8 (a) and Figure 8 (b), respectively.¹²

The length of a shortest plan is $2n - 1$ independently of k : one first visits all goal locations on the right branches (i.e., the branches except the L_1 -branch), going forth and back from L^0 ; then one descends into the L_1 -branch. Our CNFs encode $2n - 2$ steps. *AsymRatio* is $k/(2n - 1)$ because achieving the goal on the L_1 -branch takes k steps. In the symmetrical case, *AsymRatio* = $1/(2n - 1)$ which converges to 0; in the asymmetrical case, *AsymRatio* = $(2n - 3)/(2n - 1)$ which converges to 1.

Figure 8 (c) and (d) illustrate that the setting of k has only very little impact on the size and shape of the constraint graph. As mentioned in Section 2, the constraint graph

¹²In Figure 8 (a), the graph nodes L_2^1, \dots, L_n^1 can be permuted. Running SymChaff [52], a version of ZChaff extended to exploit symmetries, we could solve the MAP formulas relatively easily. This is an artefact of the simplified structure of the MAP domain. In some experiments we ran on low *AsymRatio* examples from Driverlog and Rovers, exploiting symmetries did not make a discernible difference. Intuitively, like in MAP, all goals are cheap to achieve; unlike in MAP, they are not completely symmetric – this is a side-effect of MAP’s overly abstract nature. Exploring this in more depth is a topic for future work.




Figure 8: Goals and constraint graphs in MAP. In (a) and (b), goal locations are indicated in bold face. In (c) and (d), $n = 4$; the variables at growing time steps lie on circles with growing radius, edges indicate common membership in at least one clause.

is the undirected graph where the nodes are the variables, and the edges indicate common membership in at least one clause. Constraint graphs form the basis for many notions of structure that have been investigated in the area of constraint reasoning, e.g., [16, 51, 17].

Clearly, the constraint graphs do not capture the difference between the symmetrical and asymmetrical cases in MAP. When stepping from Figure 8 (c) to Figure 8 (d), one new edge within the outmost circle is added, and three edges within the outmost circle disappear (one of these is visible on the left side of the pictures, just below the middle). More generally, between formulas MAP_n^k and $MAP_n^{k'}$, $k' > k$, there is no difference except that $k' - k$ goal clauses are skipped, and that the content of the goal clause for the L_1 -branch changes. Thus, the problem structure here cannot be detected based on simple examinations of CNF syntax. In particular, it is easy to see that there are no small cutsets in the MAP formulas, irrespectively of the setting of k . The reason are large cliques of variables present in the constraint graphs for all these formulas. Details on this are in the TR [31].

5.1.2. *Symmetrical Case.* The structure in our formulas does not affect the formula syntax much, but it does affect the size of DPLL refutations, and backdoors. First, we proved that, in the symmetrical case, the DPLL trees are large.

Theorem 5.1 (MAP symmetrical case, Resolution LB). *Every resolution refutation of MAP_n^1 must have size exponential in n .*

Corollary 5.2 (MAP symmetrical case, DPLL LB). *Every DPLL refutation of MAP_n^1 must have size exponential in n .*

The proof of Theorem 5.1 proceeds by a “reduction” of MAP_n^1 to a variant of the pigeon hole problem. A reduction here is a function that transforms a resolution refutation of MAP_n^1 into a resolution refutation of the pigeon hole. Given a reduction function from formula class A into formula class B, a lower bound on the size of resolution refutations of B is also valid for A, modulo the maximal size increase induced by the reduction.

We define a reduction function from MAP_n^1 into the *onto functional pigeon hole problem*, of PHP_n . This is the standard pigeon hole – where n pigeons must be assigned to $n - 1$ holes – plus “onto” clauses saying that at least one pigeon is assigned to each hole, and “functional” clauses saying that every pigeon is assigned to at most one hole. Razborov [48] proved that every resolution refutation of $ofPHP_n$ must have size $\exp(\Omega(n/(\log(n+1))^2))$.

Our reduction proceeds by first setting many variables in MAP_n^1 to 0 or 1, and identifying other variables (renaming x and y to a new variable z).¹³ By some rather technical (but essentially simple) arguments, we prove that such operations do not increase the size of a resolution refutation. The reduced formula is a “temporal” version of the onto pigeon hole problem; we call it $oTPHP_n$. We will discuss this in detail below. We prove that, from a resolution refutation of $oTPHP_n$, one can construct a resolution refutation of $ofPHP_n$ by replacing each resolution step with at most $n^2 + n$ new resolution steps. This proves Theorem 5.1. Corollary 5.2 follows immediately since DPLL corresponds to a restricted form of resolution. The same is true for DPLL with clause learning [4], as done in the ZChaff solver we use in our experiments.

The temporal pigeon hole problem is similar to the standard pigeon hole problem except that now the “holes” are time steps. This nicely reflects what typically goes on in Planning encodings, where the available time steps are the main resource. So it is worth having a closer look at this formula. We skip the “onto” clauses since these are identical in both the standard and the temporal version. We denote the standard pigeon hole with PHP_n , and the temporal version with $TPHP_n$. Both make use of variables x_y , meaning that pigeon x goes into hole y . Both have clauses of the form $\{\neg x_y, \neg x'_y\}$, saying that no pair of pigeons can go into the same hole. PHP_n has clauses of the form $\{x_{-1}, \dots, x_{-(n-1)}\}$, saying that each pigeon needs to go into at least one hole. In $TPHP_n$, these clauses are replaced with the following constructions of clauses:

- (1) $\{\neg px_{-2}, x_{-1}\}$
- (2) $\{\neg px_{-3}, x_{-2}, px_{-2}\} \dots \{\neg px_{-(n-1)}, x_{-(n-2)}, px_{-(n-2)}\}$
- (3) $\{x_{-(n-1)}, px_{-(n-1)}\}$

Here, the px_y are new variables whose intended meaning is that px_y is set to 1 iff x is assigned to some hole $y' < y$ – some earlier time step than y . The px_y variables correspond

¹³For example, we set all *NOOP-at* variables to 0. Such a variable will never be set to 1 in an optimal plan, since that would mean a commitment to not move at all in a time step. Similar (more complicated) intuitions are behind all the operations performed.

to the NOOP actions in Graphplan-based encodings. Clause (1) says that if we decide to put x into a time step earlier than 2, we must put it into step 1. This is an action precondition clause. Clauses (2) say that if we decide to put x into a step earlier than some y , then we must put x into either $y - 1$ or earlier than $y - 1$. These also are action precondition clauses. Clause (3) is a goal clause; it says that we must have put x somewhere by step $n - 1$ at the latest.¹⁴

We believe that the temporal pigeon hole problem is quite typical for planning situations; at least, it is clearly contained in some Planning benchmarks more involved than MAP. A simple example is the Gripper domain, where the task is to move n balls from one room into another, two at a time. With similar reduction steps as we use for MAP, this problem can be transformed into the *TPHP*. The same is true for various transportation domains with trivial road maps, e.g., the so-called Ferry and Miconic-STRIPS domains. Likewise, Logistics can be transformed into *TPHP* if there is only one road map (i.e., only one city, or only airports); in the general case, one obtains a kind of sequenced *TPHP* formula, where each pigeon must be assigned a sequence of holes (corresponding to truck, airplane, truck). In a similar fashion, assigning samples to time steps in the Rovers domain is a temporal pigeon hole problem. Formulated very generally, situations similar to the *TPHP* arise when there is not enough time to perform a number of duties. If each duty takes only few steps, then there are many possible distributions of the duties over the time steps. In its most extreme version, this situation is the *TPHP* as defined above.

The proof of Theorem 5.1 by reduction to the pigeon hole is intriguing, but not particularly concrete about what is actually going on inside a DPLL procedure run on MAP_n^1 . To shed more light on this, we now investigate the best choices of branching variables for such a procedure. The backdoor we identify in the symmetrical case, called MAP_n^1B , is shown in Figure 9 for $n = 5$. MAP_n^1B contains, at every time step with an odd index, $move-L^0-L_i^1$ and $NOOP-visited-L_i^1$ variables for all branches i on the MAP, with some exceptions. A few additional variables are needed, including $NOOP-at-L^0$ at the first time step of the encoding. Full details are in the TR [31]. The size of MAP_n^1B is $\Theta(n^2)$: $\Theta(n)$ time steps with $\Theta(n)$ variables each. Remember that the total number of variables is also $\Theta(n^2)$, so the backdoor is a linear-size variable subset.

Theorem 5.3 (MAP symmetrical case, BD). *MAP_n^1B is a backdoor for MAP_n^1 .*

For the proof, first note that, in the encoding, *any pair of move actions is incompatible*. So if one move action is set to 1 at a time step, then all other move actions at that step are forced out by UP over the mutex clauses – the time step is “occupied” (this is relevant also in the asymmetrical case below). Now, to see the basic proof argument, assume for the moment that MAP_n^1B contains all $move-L^0-L_i^1$ and $NOOP-visited-L_i^1$ variables, at each odd time step. Assigning values to all these variables results, by UP, in a sort of goal regression. In the last time step of the encoding, $t = 2n - 2$, the goal clauses form n constraints requiring to either visit a location L_i^1 , or to have visited it earlier already (i.e., to achieve it via a NOOP). Examining the interactions between *moves* and *NOOPs* at $t = 2n - 3$, one sees that, if all these are set, then at least $n - 1$ goal constraints will be transported down to $t = 2n - 4$. Iterating the argument over the $n - 2$ odd time steps,

¹⁴Note that PHP_n can be obtained from $TPHP_n$ within few resolution steps resolving on the $px-y$ variables. So it is easy to turn a refutation of PHP_n into a refutation of $TPHP_n$. The inverse direction, which we need for our proof, is less trivial; one replaces each variable $px-y$ with its meaning $\{x-1, \dots, x-(y-1)\}$, and then reasons about how to repair the resolution steps.

t=1	Nat-L ⁰				
t=2					
t=3	mv-L ⁰ -L ₁ ¹	mv-L ⁰ -L ₂ ¹	mv-L ⁰ -L ₃ ¹ Nv-L ₃ ¹	mv-L ⁰ -L ₄ ¹ Nv-L ₄ ¹	mv-L ⁰ -L ₅ ¹ Nv-L ₅ ¹
t=4					
t=5		mv-L ⁰ -L ₂ ¹	mv-L ⁰ -L ₃ ¹ Nv-L ₃ ¹	mv-L ⁰ -L ₄ ¹ Nv-L ₄ ¹	mv-L ⁰ -L ₅ ¹ Nv-L ₅ ¹
t=6					
t=7		mv-L ⁰ -L ₂ ¹	mv-L ⁰ -L ₃ ¹ Nv-L ₃ ¹	mv-L ⁰ -L ₄ ¹ Nv-L ₄ ¹	mv-L ⁰ -L ₅ ¹ Nv-L ₅ ¹
t=8					

Figure 9: The MAP_n^1B variables for $n = 5$. The vertical axis corresponds to time steps t , the horizontal axis corresponds to branches on the map. “NOOP-at” is abbreviated as “Nat”, “NOOP-visited” is abbreviated as “Nv”, “move” is abbreviated as “mv”.

one gets two goal constraints at $t = 2$: two nodes L_i^1 must be visited within the first two time steps. It is easy to see, then, that branching over $NOOP-at-L^0$ at time 1 yields an empty clause in either case. What makes identifying a non-redundant (minimal) backdoor difficult is that UP is slightly more powerful than just performing the outlined “regression”. MAP_n^1B contains hardly any variables for branch $i = 1$. So at $t = 2$ one gets only a single goal constraint, achieving which isn’t a problem. We perform an intricate case distinction about the precise pattern of time steps that are occupied after the regression, taking account of, e.g., such subtleties as the possibility to achieve $visited-L_1^1$ by moving in from L_1^2 above. In the end, one can show that UP enforces commitments to accommodate also the two $move-L^0-L_i^1$ actions that weren’t accommodated in the regression. For this, there is not enough room left.

We conjecture that the backdoor identified in Theorem 5.3 is also a minimum size (i.e., an optimal) backdoor; for $n \leq 4$ we verified this empirically, by enumerating all smaller variable sets. (Enumerating variable sets in small enough examples was also our method to find the backdoors in the first place.) We proved that the backdoor is minimal.

Theorem 5.4 (MAP symmetrical case, BD minimality). *Let B' be a subset of MAP_n^1B obtained by removing one variable. Then the number of UP-consistent assignments to the variables in B' is always greater than 0, and at least $(n - 3)!$ for $n \geq 3$.*

To prove this theorem, one figures out how wrong things can go when a variable is missing in the proof of Theorem 5.3.

t=1	$mv-L^0-L_1^1$
t=2	
t=3	$mv-L_1^2-L_1^3$
t=4	
t=5	
t=6	
t=7	$mv-L_1^6-L_1^7$
t=8	

Figure 10: The $MAP_n^{2n-3}B$ variables for $n = 5$. The vertical axis corresponds to time steps t , the horizontal axis corresponds to branches on the map (of which only one, the L_1 -branch, is relevant for the backdoor). The variables stay the same for $n = 6, 7, 8$. Compare to Figure 9.

5.1.3. *Asymmetrical Case.* When starting to investigate the backdoors in the asymmetrical case, our expectation was to obtain $\Theta(n)$ backdoors involving only the map branch to the outlier goal node. We were surprised to find that one can actually do much better. The backdoor we identify, called MAP_n^{2n-3} , is shown in Figure 9 for $n = 5$. The general form is:

- $move-L^0-L_1^1$ at step 1
- $move-L_1^2-L_1^3$ at step 3
- $move-L_1^6-L_1^7$ at step 7
- $move-L_1^{14}-L_1^{15}$ at step 15
- ...

That is, starting with $t = 2$ one has $move-L_1^{t-2}-L_1^{t-1}(t-1)$ variables where the value of t is doubled between each two variables. The size of $MAP_n^{2n-3}B$ is $\lceil \log_2 n \rceil$.

Theorem 5.5 (MAP asymmetrical case, BD). $MAP_n^{2n-3}B$ is a backdoor for MAP_n^{2n-3} .

We again conjecture that this is also a minimum size backdoor. For $n \leq 8$ we verified this empirically. Note that, while the size of the backdoor is $O(\log n)$, n itself is asymptotic to the square root of the number of variables in the formula. We can show that the backdoor is minimal. Precisely, we have:

Theorem 5.6 (MAP asymmetrical case, BD minimality). *Let B' be a subset of $MAP_n^{2n-3}B$ obtained by removing one variable. Then there is exactly one UP-consistent assignment to the variables in B' .*

The proof of Theorem 5.5 is explained with an example below. Proving Theorem 5.6 is a matter of figuring out what can go wrong in the proof to Theorem 5.5, after removing one variable.

We consider it particularly interesting that the MAP_n^{2n-3} formulas have *logarithmic* backdoors. This shows, on the one hand, that these formulas are (potentially) easy for Davis

Putnam procedures, having polynomial-size refutations.¹⁵ On the other hand, the formulas are non-trivial, in two important respects. First, they do have non-constant backdoors and are not just solved by unit propagation. Second, finding the logarithmic backdoors involves, at least, a non-trivial branching heuristic – the worst-case DPLL refutations of MAP_n^{2n-3} are still exponential in n . For example, UP will not cause any propagation if the choice of branching variables is $\{NOOP\text{-visited-}L_i^1(t) \mid 2 \leq i \leq n\}$ for one time step $3 \leq t \leq 2n - 3$.

More generally, our identification of $O(\log n)$ backdoors rather than $O(n)$ backdoors is interesting since it may serve to explain recent findings in more practical examples. Williams et al [61] have empirically found backdoors in the order of 10 out of 10000 variables in CNF encodings of many standard Planning benchmarks (as well as Verification benchmarks). In this context, it is instructive to have a closer look at how the logarithmic backdoors in the MAP formulas arise. We do so in detail below. A high-level intuition is that one can pick the branching variables in a way exploiting the need to go back and forth on the path to the outlier goal node. Going back and forth introduces a factor 2, and so one can double the number of time steps between each pair of variables. Similar phenomena arise in other domains, in the presence of “outlier goal nodes”, i.e., goals that necessitate a long sequence of steps. Good examples for this are transportation of a package to a far-away destination, or the taking of a far-away sample in Rovers. Note that this corresponds to a distorted version of the *TPHP*, where one pigeon must be assigned to an entire sequence of time steps, and hence the number of different distributions of pigeons over time steps is small.

We now examine an example formula, MAP_n^{2n-3} for $n = 8$, in detail. The proof of Theorem 5.5 uses the following two properties of UP, in MAP_n^{2n-3} :

- (1) If one sets a variable $move\text{-}L_1^{i-1}\text{-}L_1^i(i)$ to 1, then at all time steps $j < i$ a move variable is set to 1 by UP.
- (2) If one sets a variable $move\text{-}L_1^{i-1}\text{-}L_1^i(i)$ to 0, then at all time steps $j > i$ a move variable is set to 1 by UP.

Both properties are caused by the “tightness” of branch 1, i.e., by UP over the precondition clauses of the actions moving along that branch. Other than what one may think at first sight, the two properties by themselves are *not* enough to determine the log-sized backdoor. The properties just form the foundation of a subtle interplay between the different settings of the backdoor variables, exploiting exponentially growing UP implication chains on branch 1. For $n = 8$, the backdoor is $\{move\text{-}L^0\text{-}L_1^1(1), move\text{-}L_1^2\text{-}L_1^3(3), move\text{-}L_1^6\text{-}L_1^7(7)\}$. Figure 11 contains an illustration.

Consider the first (lowest) variable in the backdoor, $move\text{-}L^0\text{-}L_1^1(1)$. If one sets this to 0, then property (2) applies: only 13 of the 14 available steps are left to move towards the goal location L_1^{13} ; UP recognizes this, and forces moves towards L_1^{13} at all steps $2 \leq t \leq 14$. Since $t = 1$ is the only remaining time step not occupied by a move action, UP over the L_2^1 goal clause sets $move\text{-}L^0\text{-}L_2^1(1)$ to 1, yielding a contradiction to the precondition clause of the move set to 1 at time 2. So $move\text{-}L^0\text{-}L_1^1(1)$ must be set to 1.

Consider the second variable in the backdoor, $move\text{-}L_1^2\text{-}L_1^3(3)$. Say one sets this to 0. By property (2) this forces moves at all steps $4 \leq t \leq 14$. So the goal for L_2^1 must be achieved by an action at step 3. But we have committed to $move\text{-}L^0\text{-}L_1^1$ at step 1. This forces us to move back to L^0 at step 2 and to move to L_2^1 at step 3. But then the move forced in earlier at 4 becomes impossible. It follows that we must assign $move\text{-}L_1^2\text{-}L_1^3(3)$ to

¹⁵We will see below that the DPLL refutations in fact degenerate to lines, having the same size as the backdoors themselves.




Figure 11: The workings of the optimal backdoor for MAP_8^{13} . Arrows indicate moves on the L_1 -branch forced to 1 by UP. Direction \rightarrow means towards L_1^{13} , \leftarrow means towards L^0 . When only a single open step is left, $move-L^0-L_2^1$ is forced to 1 at that step by UP, yielding an empty clause.

1. With property (1), this implies that, by UP, all time steps below 3 get occupied with move actions. (Precisely, in our case here, $move-L_1^1-L_1^2(2)$ is also set to 1.)

Consider the third variable in the backdoor, $move-L_1^6-L_1^7(7)$. If we set this to 0, then by property (2) moves are forced in by UP at the time steps $8 \leq t \leq 14$. So, to achieve the L_2^1 goal at step 7, we have to take **three steps to move back from L_1^3 to L^0** : steps 4, 5, and 6. A move to L_2^1 is forced in at step 7, in contradiction to the move at 8 forced in earlier. Finally, if we assign $move-L_1^6-L_1^7(7)$ to 1, then by property (1) moves are forced in by UP at all steps below 7. We need **seven steps to move back from L_1^7 to L^0** , and an eighth step to get to L_2^1 . But we have only the 7 steps $8, \dots, 14$ available, so the goal for L_2^1 is unachievable.

The key to the logarithmic backdoor size is that, to achieve the L_2^1 goal, we have to move back from L_1^t locations we committed to earlier (as indicated in bold face above for $t = 3$ and $t = 7$). We committed to move to L_1^t , and the UP propagations force us to move back, thereby occupying $2 * t$ steps in the encoding. This yields the possibility to double the value of t between variables.

The DPLL tree for MAP_n^{2n-3} degenerates to a line:

Corollary 5.7 (MAP asymmetrical case, DPLL UB). *For MAP_n^{2n-3} , there is a DPLL refutation of size $2 * \lceil \log_2 n \rceil + 1$.*

Proof. Follows directly from the proof to Theorem 5.5: If one processes the $MAP_n^{2n-3}B$ variables from $t = 1$ upwards, then, for every variable, assigning the value 0 yields a contradiction by UP. There are $\lceil \log_2 n \rceil$ non-failed search nodes. Adding the failed search nodes, this shows the claim. \square

This is an interesting result since it reflects the intuition that, in practical examples, constraint propagation can cut out many search branches early on, yielding a nearly degenerated search tree. This phenomenon is also visible in our empirical data regarding search tree size/depth ratio, c.f. Section 4, Figure 7.

Note that we have now shown a doubly exponential gap between the sizes of the best-case DPLL refutations in the symmetrical case and the asymmetrical case. It would be interesting to determine what the optimal backdoors are in general, i.e., in MAP_n^k , particularly at what point the backdoors become logarithmic. Such an investigation turns out to be extremely difficult. For interesting combinations of n and k , it is practically impossible to find the optimal backdoors empirically, and so get a start into the theoretical investigation. We developed an enumeration program that exploits symmetries in the planning task to cut down on the number of variable sets to be enumerated. Even with that, the enumeration didn't scale up far enough. We leave this topic for future work.

5.2. SBW. This is a block-stacking domain, with stacking restrictions on what blocks can be stacked onto what other blocks. The blocks are initially all located side-by-side on a table t_1 . The goal is to bring all blocks onto another table t_2 , that has only space for a single block; so the n blocks must be arranged in a single stack on top of t_2 . The parameter k , $0 \leq k \leq n$, defines the amount of restrictions. There are k “bad” blocks b_1, \dots, b_k and $n - k$ “good” blocks g_1, \dots, g_{n-k} . Each b_i , $i > 1$, can only be stacked onto b_{i-1} ; b_1 can be stacked onto t_2 and any g_i . The g_i can be stacked any g_j , and onto t_2 .

Independently of k , the optimal plan length is n : move actions stack one block onto another block or a table. *AsymRatio* is $1/n$ if $k = 0$, and k/n otherwise. In the symmetrical case, $k = 0$, we identify backdoors of size $\Theta(n^3)$ – linear in the total number of variables. In the asymmetrical case, $k = n - 2$, there are $O(\log n)$ DPLL refutations and backdoors. It is an open question whether there is an exponential lower bound in the symmetrical case.


5.3. SPH. For this domain, we modified the pigeon hole problem. In our SPH_n^k formulas, as usual the task is to assign $n + 1$ pigeons to n holes. The new feature is that there is one “bad” pigeon that requires k holes, and $k - 1$ “good” pigeons that can share a hole with the bad pigeon. The remaining $n - k + 1$ pigeons are normal, i.e., need exactly one hole each. The range of k is between 1 and $n - 1$. Independently of k , $n + 1$ holes are needed overall.

We identify minimal backdoors for all combinations of k and n ; their size is $(n - k) * (n - 1)$. For $k = n - 1$ we identify an $O(n)$ DPLL refutation. With results by Buss and Pitassi [8], this implies an exponential DPLL complexity gap to $k = 1$.

6. EMPIRICAL BEHAVIOR IN SYNTHETIC DOMAINS


Regarding the analysis of synthetic domains, it remains to verify whether the theoretical observations carry over to practice – i.e., to verify whether state-of-the-art SAT solvers do indeed behave as expected. To confirm this, we ran ZChaff [46] and MiniSat [20, 19] on our synthetic CNFs. The results are in Figure 12.

The axes labelled “X” in Figure 12 refer to *AsymRatio*. The axes labelled “Y” refer to values of n . The range of *AsymRatio* is 0 to 1 in all pictures, growing on a linear scale from left to right. For each value of n , the data points shown on the X axis correspond to the *AsymRatio* values for all possible settings of k . The range of n depends on domain and SAT solver; in all pictures, n grows on a linear scale from front to back. We started each n range at the lowest value yielding non-zero runtimes, and scaled until the first time-out occurred, which we set to two hours (we used a PC running at 1.2GHz with 1GB main memory and 1024KB cache running Linux). The plots in Figure 12 include all values of n for which *no* time-out occurred. With this strategy, in MAP, for ZChaff we got data for




(a) ZChaff in MAP

(b) MiniSat in MAP



(c) ZChaff in SBW

(d) MiniSat in SBW



(e) ZChaff in SPH

(f) MiniSat in SPH

Figure 12: Runtime of ZChaff and MiniSat, log-plotted against *AsymRatio* (axes labelled “X”) and n (axes labelled “Y”), in our synthetic CNFs. The range of n depends on domain and SAT solver, see text; the value grows from front to back. The range of *AsymRatio* is $[0; 1]$ in all plots, growing from left to right. For each value of n , the data points shown on the X axis correspond to the *AsymRatio* values for all possible settings of k .

$n = 5, \dots, n = 9$, and for MiniSat we got data for $n = 5, \dots, n = 10$. In SBW, for ZChaff we got data for $n = 6, \dots, n = 10$, and for MiniSat we got data for $n = 6, \dots, n = 11$. In SPH, for ZChaff we got data for $n = 7, \dots, n = 12$, and for MiniSat we got data for $n = 7, \dots, n = 11$. (Note that MiniSat outperforms ZChaff in MAP and SBW, while ZChaff is better in SPH.)

All in all, the empirical results comply very well with our analytical results, meaning that the solvers do succeed, at least to some extent, in exploiting the problem structure and finding short proofs. When walking towards the back on a parallel to the Y axis – when increasing n – as expected runtime grows exponentially in all but the most asymmetric (rightmost) instances, for which we proved the existence of polynomial (logarithmic, in MAP and SBW) proofs. Also, when walking to the left on a parallel to the X axis – when decreasing *AsymRatio* – as expected runtime grows exponentially. There are a few remarkable exceptions to the latter, particularly for ZChaff and MiniSat in MAP, and for ZChaff in SPH: there, the most symmetric (leftmost) instances are solved faster than their direct neighbors to the right. There actually also is such a phenomenon in SBW. This is not visible in Figure 12 since $k = 0$ and $k = 1$ lead to the same *AsymRatio* value in SBW, and Figure 12 uses $k = 0$; the $k = 1$ instances take about 20% – 30% more runtime than the $k = 0$ instances. So, the SAT solvers often find the completely symmetric cases easier than the only slightly asymmetric ones. We can only speculate what the reason for that is. Maybe the phenomenon is to do with the way these SAT solvers perform clause learning; maybe the branching heuristics become confused if there is only a very small amount of asymmetry to focus on. It seems an interesting topic to explore this issue – in the “slightly asymmetric” cases, significant runtime could be saved.

7. A RED HERRING

Being defined as a simple cost ratio – even one that involves computing optimal plan lengths – *AsymRatio* cannot be fool-proof. We mentioned in Section 4 already that one can replace G with a single goal g , and an additional action with precondition G and add effect $\{g\}$, thereby making *AsymRatio* devoid of information. Such “tricks” could be dealt with by defining *AsymRatio* over “necessary sub-goals” instead, as explained in the next section. More importantly perhaps, of course there are examples of (parameterized) planning tasks where *AsymRatio* does *not* correlate with DPLL proof size. One way to construct such examples is by “hiding” a relevant phenomenon behind an irrelevant phenomenon that controls *AsymRatio*. Figure 13 provides the construction of such a case, based on the MAP domain.

The actions in Figure 13 are the same as in MAP, moving along edges in the graph, precondition $\{at-x\}$, add effect $\{at-y, visited-y\}$, delete effect $\{at-x\}$. The difference is that we now have *two* maps (graphs), and we assume that we can move on both in parallel, i.e., our CNFs allow parallel move actions on separate graphs (which, for example, the Graphplan-based encoding indeed does). Initially one is located at L^0 and R^0 . The goal is to visit the locations shown in bold face in Figure 13: L_1^k and all of R_1^1, \dots, R_n^1 , where k ranges between 1 and $2n - 2$. Independently of k , the optimal plan length is $2n - 1$, and our CNF encodes $2n - 2$ plan steps. The left map is solvable within this number of steps, so unsatisfiability must be proved on the right map, which does not change over k and requires exponentially long resolution proofs, c.f. Theorem 5.1. *AsymRatio*, however, is $k/(2n - 1)$.




Figure 13: The red herring example, consisting of two parallel maps on which nodes must be visited. Parallel moves on separate maps are allowed, thus plan length is $2n - 1$ irrespectively of k . Plan length bound for the CNFs is $2n - 2$, the left map is solvable within $2n - 2$ steps, the right map is not. *AsymRatio* is controlled by the left map, but to prove unsatisfiability one has to deal with the right map, which does not change over k and demands exponentially long resolution proofs.

In the red herring example, *AsymRatio* scales as the ratio between k and n , as before, but the best-case DPLL proof size remains constant. One can make the situation even “worse” by making the right map more complicated. Say we introduce additional locations R_1^2, \dots, R_n^2 , each R_i^2 being linked to R_i^1 . We further introduce the location R_1^3 linked to R_1^2 . The idea is to introduce a varying number of goal nodes that lie “further out”. The larger the number of such further out goal nodes, the easier we expect it to be (from our experience with the MAP domain) to prove the right map unsatisfiable. We can always “balance” the goal nodes in a way keeping the optimal plan length constant at $2n - 1$. We then *invert* the correlation between *AsymRatio* and DPLL performance by introducing a lot of further out goals when k is small, and less further out goals when k is large. Concretely, with maximum k we take the goal to be the same as before, to visit R_1^1, \dots, R_n^1 . With minimum k , for simplicity say $n = 2m$. Then our goal will be to visit $R_1^3, R_2^2, \dots, R_m^2$. The optimal plan here first visits all of R_2^2, \dots, R_m^2 , taking $4 * (m - 1) = 4 * (n/2 - 1) = 2n - 4$ steps (4 steps each since one needs to go forth and back). Then R_1^3 is visited, taking another 3 steps. All in all, with this construction, *AsymRatio* is still dominated by the left map. But the right map becomes *harder* to prove unsatisfiable as k increases.

The example contains *completely separate* sub-problems, the left and right maps. The complete separation invalidates the connection between *AsymRatio* and our intuitions about the relevant “problem structure”. As explained in the introduction, the general intuition is that (1) in the symmetrical case (low *AsymRatio*) there is a fierce competition of many sub-problems (goals) for the available resources, and (2) in the asymmetrical case (high *AsymRatio*) a single sub-problem uses most of the resources on its own. In the red herring example, (1) is still valid, but (2) is not since the asymmetric sub-problem in question has access to its own resources – the parallel moves on the left map. So, in this

case, *AsymRatio* fails to characterize the intuitive “degree of sub-problem interactions”. The mistake lies in the maximization over individual goal fact costs, which disregards to what extent the respective sub-problems are actually interconnected. In this sense, the red herring is a suggestion to research more direct measures of sub-problem interactions, identifying their causes rather than their effects.

8. CONCLUSIONS AND FURTHER DIRECTIONS

We defined a concrete notion of what problem structure is in Planning, and we revealed empirically that this structure often governs SAT solver performance. Our analytical results provide a detailed case study of how this phenomenon arises. In particular, we show that the phenomenon can make a doubly exponential difference, and we identify some prototypical behaviors that appear also in planning competition benchmarks.

Our parameterized CNFs can be used as a set of benchmarks with a very “controlled” nature. This is interesting since it allows very clean and detailed tests of how good a SAT solver is at exploiting this particular kind of problem structure. For example, our experiments from Section 6 suggest that ZChaff and MiniSat have difficulties with “slightly asymmetric” instances.

There are some open topics to do with the definition of *AsymRatio*. First, one should try to define a version of *AsymRatio* that is more stable with respect to the “sub-goal structure”. As mentioned, our current definition, $AsymRatio = \max_{g \in G} cost(g) / cost(\bigwedge_{g \in G} g)$, can be fooled by replacing G with a single goal g , and introducing an additional action with precondition G and add effect $\{g\}$. A more stable approach would be to identify a hierarchy of layers of “necessary sub-goals” – facts that must be achieved along any solution path. Such facts were termed “landmarks” in recent research [33]. In a nutshell, one could proceed as follows. Build a sequence of landmark “layers” G_0, \dots, G_m . Start with $G_0 := G$. Iteratively, set $G_{i+1} := \bigcup_{g \in G_i} \bigcap_{a: g \in add(a)} pre(a)$, until G_{i+1} is contained in the previous layers. Set $m := i$, i.e., consider the layers up to the last one that brought a new fact. It has been shown that such a process results in informative problem decompositions in many benchmark domains [33]. For an alternative definition of *AsymRatio*, the idea would now be to select one layer G_i , and define *AsymRatio* based on that, for example as $\max_{g \in G_i} cost(g) + i$ divided by $cost(\bigwedge_{g \in G} g)$. A good heuristic may be to select the layer G_i that contains the largest number of facts. This approach could not be fooled by replacing G with a single goal. It remains an open question in what kinds of domains the approach would deliver better (or worse) information than our current definition of *AsymRatio*. We remark that the approach does not solve the red herring example. There, $G_1 = \{at-L_1^{k-1}, at-R^0\}$, and $G_i = \{at-L_1^{k-i}\}$ for $2 \leq i \leq k$. So, for all i , $\max_{g \in G_i} cost(g) + i$ is the same as $\max_{g \in G} cost(g)$. Motivated by this, one could try to come up with more complex definitions of *AsymRatio*, taking into account to what extent the sub-problems corresponding to the individual goals are interconnected. Interconnection could, e.g., be approximated using shared landmarks. A hard sub-problem that is tightly interconnected would have more weight than one that is not.

A potentially more important topic is to explore whether it is possible to define more direct measures of “the degree of sub-problem interactions”. A first option would be to explore this in the context of the above landmarks. Previous work [33] has already defined what “interactions” between landmarks could be taken to mean. A landmark L is said to interact with a landmark L' if it is known that L has to be true at some point after L' , and

that achieving L' involves deleting L . This means that, if L is achieved before L' , then L must be re-achieved after achieving L' – hence, L should be achieved *before* L' . One can thus define an order over the landmarks, based on their interactions. It is unclear, however, how this sort of interactions could be turned into a number that stably correlates with performance across a range of domains.

Alternative approaches to design more direct measures of sub-problem interaction could be based on (1) the “fact generation trees” examined previously to draw conclusions about the quality of certain heuristic functions [34], or (2) the “criticality” measures proposed previously to design problem abstractions [1]. As for (1), the absence of a certain kind of interactions in a fact generation tree allows us to conclude that a certain heuristic function returns the precise goal distance (as one would expect, this is not the case in any but the most primitive examples). Like for the landmarks, it is unclear how this sort of observation could be turned into a number estimating the “degree” of interaction. As for (2), this appears promising since, in contrast to the previous two ideas, “criticalities” already are numbers – estimating how “critical” a fact is for a given set of actions. The basic idea is to estimate how many alternative ways there are to achieve a fact, and, recursively, how critical the prerequisites of these alternatives are. The challenge is that, in particular, the proposed computation of criticalities disregards initial state, goal, and delete effects – which is fine in the original context, but not in our context here. All three (initial state, goal, and delete effects) must be integrated into the computation in order to obtain a measure of “sub-problem interaction” in our sense. It is also unclear how the resulting measure for each individual fact would be turned into a measure of interactions between (several) facts.

We emphasize that our synthetic domains are relevant no matter how one chooses to characterize “the degree of sub-problem interactions”. The intuition in the domains is that, on one side, we have a CNF whose unsatisfiability arises from interactions between many sub-problems, while on the other side we have a CNF whose unsatisfiability arises mostly from the high degree of constrainedness of a single sub-problem. As long as it is this sort of structure one wants to capture, the domains are representative examples. In fact, they can be used to test quickly whether or not a candidate definition is sensible – if the definition does not capture the transition on our k scales, then it can probably be discarded.

In some ways, our empirical and analytical observations could be used to tailor SAT solvers for planning. A few insights obtained from our analysis of backdoors are these. First, one should branch on actions serving, recursively over add effects and preconditions, to support the goals; such actions can be found by simple approximative backward chaining. Second, the number of branching decisions related to the individual goal facts should have a similar distribution as the cost of these facts. Third, the branching should be distributed evenly across the time steps, with not much branching in directly neighbored time steps (since decisions at t typically affect $t - 1$ and $t + 1$ a lot anyway). To the best of our knowledge, such techniques have hardly been scratched so far in the (few) existing planning-specific backtracking solvers [49, 2, 30]. The only similar technique we know of is used in “BBG” [30]; this system branches on actions occurring in a “relaxed plan”, which is similar to the requirement to support the goals.

A perhaps more original topic is to approximate *AsymRatio*, and take decisions based on that. Due to the wide-spread use of heuristic functions in Planning, the literature contains a wealth of well-researched techniques for approximating the number of steps needed to achieve a goal, e.g. [5, 6, 32, 18, 24, 27, 26]. We expect that (a combination of)

these techniques could, with some experimentation, be used to design a useful approximation of *AsymRatio*.

There are various possible uses of approximated *AsymRatio*. First, one could design specialized branching heuristics, and choose one depending on the (approximated) value of *AsymRatio*. The very different forms of our identified backdoors in the symmetrical and asymmetrical cases offer a rich source of ideas for such heuristics. For example, it seems that a high *AsymRatio* allows for large “holes” in the time steps branched upon, while a low *AsymRatio* asks for more close decisions. Also, it seems that NOOP variables are not important with high *AsymRatio*, but do play a crucial role with low *AsymRatio*; this makes intuitive sense since, in the presence of subtle interactions, it may be more important to preserve the truth value of some fact through time steps. Perhaps most importantly, with some more work approximated *AsymRatio* could probably be made part of a successful runtime predictor. If not in general – across domains – it is reasonable to expect that this will work within some fixed domain (or class of domains) of interest.

Last but not least, we would like to point out that our work is merely a first step towards understanding practical problem hardness in optimal planning. Some important limitations of the work, in its current state, are:

- The experiments reported in Section 4 are performed only for a narrow range of instance size parameters, for each domain. It is important to determine how the behavior of *AsymRatio* changes as a function of changing parameter settings. We have merely scratched this direction for Driverlog and Logistics; future work should address it in full detail for some selected domains. The problem in performing such a research lies in the large numbers of domain parameters, and the exorbitant amount of computation time needed to explore even a single parameter setting – in our experiments, often a single case required several months of computation time.
- Step-optimality, which we explore herein, is often not the most relevant optimization criterion. In practice, one typically wants to minimize makespan – the end time point of a concurrent durational plan – or the summed up cost of all performed actions.¹⁶ While there are obvious extensions of *AsymRatio* to these settings, it is unclear if the phenomena observed herein will carry over.
- The analyzed synthetic domains are extremely simplistic; while they provide clear intuitions, their relevance for complex practical scenarios – in particular for scenarios going beyond the planning competition benchmarks – is unclear.

To sum up, we are far from claiming that our research “solves” the addressed problems in a coherent way. Rather, we hope that our work will inspire other researchers to work on similar topics. Results on practical problem structure are difficult to obtain, but they are of vital importance for understanding combinatorial search and its behavior in practice. We believe that this kind of research should be given more attention.

Acknowledgements. We thank Ashish Sabharwal for advice on proving resolution lower bounds. We thank Rina Dechter for discussions on the relation between backdoors and cutsets. We thank Toby Walsh for discussions on the intuition behind *AsymRatio*, and for the suggestion to construct a red herring. We thank Jussi Rintanen for providing the code of his random instance generator, and we thank Martin Gütlein for implementing the

¹⁶Note that this is a problem of present-day optimal AI planning in general, not only of our research. Planning research has only quite recently started to develop techniques addressing makespan and, in particular, cost optimization.

enumeration of CNF variable sets under exploitation of problem symmetries. We finally thank the anonymous reviewers for their comments, which contributed a lot to improve the paper.

REFERENCES

- [1] R. Sebastiani A. Bundy, F. Giunchiglia and T. Walsh. Calculating criticalities. *Artificial Intelligence*, 88(1–2):39–67, 1996.
- [2] M. Baiocchi, S. Marcugini, and A. Milani. DPPlan: An algorithm for fast solution extraction from a planning graph. In *Proc. AIPS’00*, 2000.
- [3] A. Barrett and D. Weld. Partial order planning: Evaluating possible efficiency gains. *Artificial Intelligence*, 67:71–112, 1994.
- [4] P. Beame, H. Kautz, and A. Sabharwal. Towards understanding and harnessing the potential of clause learning. *Journal of Artificial Intelligence Research*, 22:319–351, 2004.
- [5] A. Blum and M. Furst. Fast planning through planning graph analysis. *Artificial Intelligence*, 90(1–2):279–298, 1997.
- [6] B. Bonet and H. Geffner. Planning as heuristic search. *Artificial Intelligence*, 129(1–2):5–33, 2001.
- [7] M. Bonet, J. Esteban, N. Galesi, and J. Johansen. On the relative complexity of resolution refinements and cutting planes proof systems. *SIAM Journal of Computing*, 30(5):1462–1484, 2001.
- [8] S. Buss and T. Pitassi. Resolution and the weak pigeon-hole principle. In *Proc. CSL’97*, pages 149–156, 1997.
- [9] Tom Bylander. The computational complexity of propositional STRIPS planning. *Artificial Intelligence*, 69(1–2):165–204, 1994.
- [10] P. Cheeseman, B. Kanefsky, and W. Taylor. Where the *really* hard problems are. In *Proc. IJCAI’91*, pages 331–337, 1991.
- [11] H. Chen and V. Dalmau. Beyond hypertree width: Decomposition methods without decompositions. In *Proc. CP’05*, 2005.
- [12] S. Climer and W. Zhang. Searching for backbones and fat: A limit-crossing approach with applications. In *Proc. AAAI-02*, 2002.
- [13] S. Cook and R. Reckhow. The relative efficiency of propositional proof systems. *Journal of Symbolic Logic*, 44:26–50, 1979.
- [14] M. Davis, G. Logemann, and D. Loveland. A machine program for theorem-proving. *Communications of the ACM*, 5:394–397, 1962.
- [15] M. Davis and H. Putnam. A computing procedure for quantification theory. *J. ACM*, 7(3):201–215, 1960.
- [16] R. Dechter. Enhancement schemes for constraint processing: Backjumping, learning and cutset decomposition. *Artificial Intelligence*, 41(3):273–312, 1990.
- [17] R. Dechter. *Constraint Processing*. Morgan-Kaufmann, 2003.
- [18] S. Edelkamp. Planning with pattern databases. In *Proc. ECP’01*, pages 13–24, 2001.
- [19] N. Een and A. Biere. Effective preprocessing in SAT through variable and clause elimination. In *Proc. SAT-05*, 2005.
- [20] N. Een and N. Sörensson. An extensible SAT solver. In *Proc. SAT-03*, 2003.
- [21] R. Fikes and N. Nilsson. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2:189–208, 1971.
- [22] J. Frank, P. Cheeseman, and J. Stutz. When gravity fails: Local search topology. *Journal of Artificial Intelligence Research*, 7:249–281, 1997.
- [23] I. Gent and T. Walsh. The SAT phase transition. In *Proc. ECAI’94*, pages 105–109, 1994.
- [24] Alfonso Gerevini, Alessandro Saetti, and Ivan Serina. Planning through stochastic local search and temporal action graphs. *Journal of Artificial Intelligence Research*, 20:239–290, 2003.
- [25] A. Haken. The intractability of resolution. *Theoretical Computer Science*, 39:297–308, 1985.
- [26] P. Haslum, B. Bonet, and H. Geffner. New admissible heuristics for domain-independent planning. In *Proc. AAAI’05*, 2005.
- [27] Malte Helmert. A planning heuristic based on causal graph analysis. In *Proc. ICAPS’04*, pages 161–170, 2004.

- [28] J. Hoffmann and S. Edelkamp. The deterministic part of IPC-4: An overview. *Journal of Artificial Intelligence Research*, 24:519–579, 2005.
- [29] J. Hoffmann, S. Edelkamp, S. Thiebaut, R. Englert, F. Liporace, and S. Trüg. Engineering benchmarks for planning: the domains used in the deterministic part of IPC-4. *Journal of Artificial Intelligence Research*, 2006. Submitted.
- [30] J. Hoffmann and H. Geffner. Branching matters: Alternative branching in graphplan. In *Proc. ICAPS'03*, pages 22–31, 2003.
- [31] J. Hoffmann, C. Gomes, and B. Selman. Structure and problem hardness: Goal asymmetry and DPLL proofs in SAT-based planning. Technical Report, 2007. Available at <http://members.deri.at/~joergh/papers/tr-lmcs07.ps.gz>.
- [32] J. Hoffmann and B. Nebel. The FF planning system: Fast plan generation through heuristic search. *Journal of Artificial Intelligence Research*, 14:253–302, 2001.
- [33] J. Hoffmann, J. Porteous, and L. Sebastia. Ordered landmarks in planning. *Journal of Artificial Intelligence Research*, 22:215–278, 2004.
- [34] Jörg Hoffmann. Where ‘ignoring delete lists’ works: Local search topology in planning benchmarks. *Journal of Artificial Intelligence Research*, 24:685–758, 2005.
- [35] T. Hogg, B. Huberman, and C. Williams. Phase Transitions and Complexity. *Artificial Intelligence*, 81, 1996.
- [36] A. Howe and E. Dahlman. A critical assessment of benchmark comparison in planning. *Journal of Artificial Intelligence Research*, 17:1–33, 2002.
- [37] T. Hulubei and B. Sullivan. Optimal refutations for constraint satisfaction problems. In *Proc. IJCAI05*, 2005.
- [38] K. Iwama and S. Miyazaki. Tree-like resolution is super-polynomially slower than DAG-like resolution for the pigeonhole principle. In *Proc. Algorithms and Computation*, pages 133–142, 1999.
- [39] H. Kautz and B. Selman. Planning as satisfiability. In *Proc. ECAI'92*, pages 359–363, 1992.
- [40] H. Kautz and B. Selman. Pushing the envelope: Planning, propositional logic, and stochastic search. In *Proc. AAAI'96*, pages 1194–1201, 1996.
- [41] H. Kautz and B. Selman. Unifying SAT-based and graph-based planning. In *Proc. IJCAI'99*, pages 318–325, 1999.
- [42] T. Kirkpatrick and N. Clark. PERT as an aid to logic design. *IBM Journal on Research Development*, 10(2):135–141, 1966.
- [43] J. Koehler and K. Schuster. Elevator control as a planning problem. In *Proc. AIPS'00*, pages 331–338, 2000.
- [44] D. Long and M. Fox. The 3rd international planning competition: Results and analysis. *Journal of Artificial Intelligence Research*, 20:1–59, 2003.
- [45] M. Marek-Sadowska and T. Xiao. Functional correlation analysis in crosstalk induced critical paths identification. In *Proc. DAC'01*, pages 653–656, 2001.
- [46] M. Moskewicz, C. Madigan, Y. Zhao, L. Zhang, and S. Malik. Chaff: Engineering an efficient SAT solver. In *Proc. DAC'01*, pages 530–535, 2001.
- [47] E. Nudelman, K. Leyton-Brown, H. Hoos, A. Devkar, and Y. Shoham. Understanding random SAT: beyond the clauses-to-variable ratio. In *Proc. CP-04*, pages 438–452, 2004.
- [48] A. Razborov. Resolution lower bounds for perfect matching principles. *J. Computer and Systems Sciences*, 69(1):3–27, 2004.
- [49] J. Rintanen. A planning algorithm not based on directional search. In *Proceedings KR'98*, pages 617–624, 1998.
- [50] J. Rintanen. Phase transitions in classical planning: An experimental study. In *Proc. ICAPS'04*, pages 101–110, 2004.
- [51] I. Rish and R. Dechter. Resolution versus search: Two strategies for SAT. *Journal of Automated Reasoning*, 24(1/2):225–275, 2000.
- [52] A. Sabharwal. Symchaff: A structure-aware satisfiability solver. In *Proc. AAAI'05*, pages 467–474, 2005.
- [53] J. Slaney and S. Thiebaut. Blocks world revisited. *Artificial Intelligence*, 125:119–153, 2001.
- [54] J. Slaney and T. Walsh. Backbones in optimization and approximation. In *Proc. IJCAI01*, 2001.
- [55] M. Streeter and S. Smith. Characterizing the distribution of low-makespan schedules in the job shop scheduling problem. In *Proc. ICAPS'05*, pages 61–70, 2005.

- [56] M. Veloso and J. Blythe. Linkability: Examining causal link commitments in partial order planning. In *Proc. AIPS'94*, pages 170–175, 1994.
- [57] Benjamin Wah and Yixin Chen. Subgoal partitioning and global search for solving temporal planning problems in mixed space. *International Journal of Artificial Intelligence Tools*, 13(4):767–790, 2004.
- [58] J. Watson, L. Barbulescu, L. Whitley, and A. Howe. Contrasting structured and random permutation flow-shop scheduling problems: Search space topology and algorithm performance. *INFORMS Journal on Computing*, 14(2):98–123, 2002.
- [59] J. Watson, J. Beck, L. Barbulescu, and L. Whitley. Toward a descriptive model of local search cost in job-shop scheduling. In *Proc. ECP'01*, 2001.
- [60] J. Watson, J. Beck, A. Howe, and D. Whitley. Problem difficulty for tabu search in job-shop scheduling. *Artificial Intelligence*, 143(2):189–217, 2003.
- [61] R. Williams, C. Gomes, and B. Selman. Backdoors to typical case complexity. In *Proc. IJCAI'03*, 2003.
- [62] S. Yen, D. du, and S. Ghanta. Efficient algorithms for extracting the k most critical paths in timing analysis. In *Proc. DAC'89*, pages 649– 654, 1989.