

General Policies, Serializations, and Planning Width*

Blai Bonet,¹ Hector Geffner²

¹ Universitat Pompeu Fabra, Barcelona, Spain

² ICREA & Universitat Pompeu Fabra, Barcelona, Spain
bonetblai@gmail.com, hector.geffner@upf.edu

Abstract

It has been observed that in many of the benchmark planning domains, atomic goals can be reached with a simple polynomial exploration procedure, called IW, that runs in time exponential in the problem width. Such problems have indeed a bounded width: a width that does not grow with the number of problem variables and is often no greater than two. Yet, while the notion of width has become part of the state-of-the-art planning algorithms like BFWS, there is still no good explanation for why so many benchmark domains have bounded width. In this work, we address this question by relating bounded width and serialized width to ideas of generalized planning, where general policies aim to solve multiple instances of a planning problem all at once. We show that bounded width is a property of planning domains that admit optimal general policies in terms of features that are explicitly or implicitly represented in the domain encoding. The results are extended to the larger class of domains with bounded serialized width where the general policies do not have to be optimal. The study leads also to a new simple, meaningful, and expressive language for specifying domain serializations in the form of **policy sketches** which can be used for encoding domain control knowledge by hand or for learning it from traces. The use of sketches and the meaning of the theoretical results are all illustrated through a number of examples.

Introduction

Pure width-based search methods exploit the structure of states to enumerate the state space in ways that are different from standard methods like breadth-first, depth-first, or random search (Lipovetzky and Geffner 2012). For this, width-based methods appeal to a notion of novelty to establish a preference for first visiting states that are most novel. Novelty-based methods have also been used in the context of genetic algorithms where a greedy focus on the function to optimize (fitness) often leads to bad local optima (Lehman and Stanley 2011a,b), and in reinforcement learning to guide exploration in large spaces where reward is sparse (Tang et al. 2017; Pathak et al. 2017; Ostrovski et al. 2017).

In classical planning, i.e., planning in factored spaces for achieving a given goal from a known initial state (Geffner and Bonet 2013; Ghallab, Nau, and Traverso

2016), the notion of novelty is now part of state-of-the-art search algorithms like BFWS (Lipovetzky and Geffner 2017b,a) and has been applied successfully in purely exploratory settings where no compact model of the actions or goals is assumed to be known a priori (Francès et al. 2017; Bandres, Bonet, and Geffner 2018). The basic width-based planning algorithms are simple and they all assume that the states are factored, assigning values to a fixed number of boolean features F that in classical planning is given by the atoms in the problem. The procedure IW(1) is indeed a breadth-first search that starts in the given initial state and prunes all the states that do not make a feature from F true for the first time in the search. IW(k) is like IW(1) but using the set of features F^k that stand for conjunctions of up to k features from F . For many benchmark domains, it has been shown that IW(k) for a small and constant value of k like $k = 2$, called the domain width, suffices to compute plans, and indeed optimal plans, for any atomic goal (Lipovetzky and Geffner 2012). Other algorithms like BFWS make use of this property for serializing conjunctive goals into atomic ones; an idea that is also present in algorithms that precompute atomic landmarks (Hoffmann, Porteous, and Sebastia 2004), and in particular, those that use landmark counting heuristics (Richter and Westphal 2010).

A key open question in the area is why these width-based methods are effective at all, and in particular, why so many domains have a small width when atomic goals are considered. Is this a property of the domains? Is it an accident of the manual representations used? In this work, we address these and related questions. For this, we bring the notion of **general policies**; policies that solve multiple instances of a planning domain all at once (Srivastava, Immerman, and Zilberstein 2008; Bonet, Palacios, and Geffner 2009; Hu and De Giacomo 2011; Belle and Levesque 2016; Segovia, Jiménez, and Jonsson 2016), in some cases by appealing to a fixed set Φ of general state features that we restrict to be linear. A number of correspondences are then obtained connecting the notions of *width*, *the size of general policies as measured by the number of features used*, and *the more general and useful notion of serialized width*; i.e., width under a given serialization. The study leads us to formalize the abstract notion of serialization

*This paper is the long version of (Bonet and Geffner 2021).

and to analyze the conditions under which serializations are well-formed. Moreover, from the relations established between general policies and serializations, we obtain a new simple, meaningful, and expressive language for specifying serializations, called **policy sketches**, which can be used for encoding domain control knowledge by hand or for learning it from traces. While exploring these potential uses of sketches is beyond the scope of this paper, the use of sketches and the meaning of the theoretical results are all illustrated through a number of examples.

The paper is organized as follows. We review first the notions of width, representations, and general policies, and relate width with the size of such policies, as measured by the number of features. Then we introduce serializations, the more general notion of serialized width, the relation between general policies and serialized width, and finally, policy sketches and their properties.

Width

IW(1) is a simple search procedure that operates on a rooted directed graph where the nodes represent states, and states assign values v to a given set F of features f (Lipovetzky and Geffner 2012). IW(1) performs a breadth-first search starting at the root but pruning the states that do not make an atom $f=v$ true for the first time in the search. For classical planning problems expressed in languages such as (grounded) STRIPS, the features f are the problem variables which can take the values true or false. In other settings, like the Atari games as supported in ALE (Bellemare et al. 2013), the features and their values are defined in other ways (Lipovetzky, Ramirez, and Geffner 2015; Bandres, Bonet, and Geffner 2018). The procedure IW(k) for $k > 1$ is IW(1) but with a feature set given by F^k .

A finding reported by Lipovetzky and Geffner (2012) and exploited since in width-based algorithms is that the procedure IW(k) with $k = 2$ suffices to solve a wide variety of planning problems. Indeed, Lipovetzky and Geffner consider the 37,921 instances that result from all the domains used in planning competitions until 2012, where each instance with a goal made up of a conjunction of k atoms is split into k instances, each one with a single atomic goal. They report that IW(2) solves more than 88% of such instances, and moreover, 100% of the instances from 26 of the 37 domains considered.

This is remarkable because IW(k) when k is smaller than the number n of problem variables is an *incomplete procedure*. Indeed, while breadth-first search expands a number of nodes that is exponential in n , IW(2) expands a quadratic number of nodes at most. Moreover, the results are not an accident resulting from a lucky choice of the instances, as one can formally prove that IW(2) solves **any** instance of many of these domains when the goal is a single atom.

Underlying the IW algorithms is the notion of *problem width*, which borrows from similar notions developed for constraint satisfaction problems and Bayesian networks, that are intractable but have algorithms that run in time and space that are exponential in the treewidth of the underlying graphs (Freuder 1982; Pearl 1988; Dechter 2013). For planning,

the notion introduced by Lipovetzky and Geffner takes this form:¹

Definition 1 (Based on Lipovetzky and Geffner, 2012). *The width $w(P)$ of problem P is the minimum k for which there is a sequence t_0, t_1, \dots, t_m of atom tuples t_i , each with at most k atoms, such that:*

1. t_0 is true in the initial state of P ,
2. any optimal plan for t_i can be extended into an optimal plan for t_{i+1} by adding a single action, $i = 1, \dots, n - 1$,
3. any optimal plan for t_m is an optimal plan for P .

The width is $w(P) = 0$ iff the initial state of P is a goal state. For convenience, we set $w(P)$ to 0 if the goal of P is reachable in a single step, and to $w(P) = N + 1$ if P has no solution where N is the number of atoms in P .

Chains of tuples $\theta = (t_0, t_1, \dots, t_m)$ that comply with conditions 1–3 are called **admissible**, and the size of the chain is the size $|t_i|$ of the largest tuple in the chain. The width $w(P)$ is thus k if k is the minimum size of an admissible chain for P . Notice that the definition of width does not presuppose a particular language for specifying the actions or goals, which can actually be specified procedurally. It just assumes that states are factored and assign truth values to a pool of atoms. The tuples t_i in an admissible chain can be regarded as subgoals or stepping stones in the way to the goal of P that ensure that it will be reached optimally. The main properties of the IW(k) algorithm can then be expressed as follows:²

Theorem 2 (Lipovetzky and Geffner, 2012). *IW(k) expands up to N^k nodes, generates up to bN^k nodes, and runs in time and space $O(bN^{2k-1})$ and $O(bN^k)$, respectively, where N is the number of atoms and b bounds the branching factor in problem P . IW(k) is guaranteed to solve P optimally (shortest path) if $w(P) \leq k$.*

Proof. Since the number of tuples of size at most k is bounded by N^k , IW(k) expands up to N^k nodes and generates up to bN^k nodes, where each expansion takes time bounded by b . Under the assumption that each transition flips the value of up to M (constant) number of atoms, each dequeued state can make true up to $\binom{N}{k} - \binom{N-M}{k} = O(N^{k-1})$ new tuples of atoms of size at most k . Seen tuples are stored in a *perfect hash* of size $O(N^k)$ which supports operations in constant time. Therefore, the test for expansion for each dequeued state takes time $O(N^{k-1})$, and thus the total time incurred by IW(k) for exploring the search tree is $O(bN^{2k-1})$. From the definition of width, it is not hard to see that if $w(P) \leq k$, then IW(k) generates a node for a goal state, and the path from the initial state to the first goal state that is generated is an optimal (shortest) path. \square

When the width of problem P is not known, the **IW algorithm** can be run instead which calls IW(k) iteratively

¹For convenience, we set the width of problems that can be solved in one step to zero.

²It is assumed that the number of atoms affected by an action is bounded by a constant. When this is not the case, the time bound in Theorem 2 becomes $O(bN^{2k})$.

for $k = 0, 1, \dots, N$ until the problem is solved, or shown to have no solution when $IW(N)$ finds no plan. While these algorithms are not aimed at being practical as planning algorithms, some state-of-the-art planners, make use of these ideas in a slightly different form (Lipovetzky and Geffner 2017b,a).

Example. Let \mathcal{Q}_{clear} be the class of Blocksworld problems P in the standard stack/unstack encoding whose goal is to achieve the atom $clear(x)$ and an empty gripper, starting with $clear(x)$ false and an empty gripper. Let B_1, \dots, B_m for $m > 0$ be the blocks above x from top to bottom in P , and let us consider the chain of subgoals t_0, \dots, t_{2m-1} where $t_0 = \{clear(B_1)\}$, $t_{2i-1} = \{hold(B_i)\}$, and $t_{2i} = \{ontable(B_i)\}$, $i = 1, \dots, m$. It is easy to check that conditions 1–3 above hold for this chain, hence, $w(P) \leq 1$. Since $w(P) > 0$, as the goal cannot be reached in zero or one step in general, $w(P) = 1$. ■

Example. Let \mathcal{Q}_{on} be the class of Blocksworld problems P where the goal is $on(x, y)$. For simplicity, let us assume that x and y are not initially in the same tower, and there are blocks B_1, \dots, B_m above x , and blocks D_1, \dots, D_ℓ above y . It is not difficult to check that the chain $t_0, \dots, t_{2m}, t'_0, \dots, t'_{2\ell}, t''_0, t''_1$ is admissible and of size 2, where t_i for $i = 0, \dots, 2m$ is as in previous example, $t'_{2i} = \{hold(D_i), clear(x)\}$ and $t'_{2i-1} = \{ontable(D_i), clear(y)\}$ for $i = 0, \dots, \ell$, $t''_0 = \{hold(x), clear(y)\}$, and $t''_1 = \{on(x, y)\}$. Then, $w(P) = 2$ since $w(P) \not\leq 1$. ■

Representations

The width of a planning problem is tied to the representation language and the encoding of the problem in the language. For example, the problem of moving a number of packages N from one room to the next, one by one, has a width that grows with N in standard encodings where each package has a name, but width 2 when the packages are indistinguishable from each other and the number of packages is encoded in unary (with one atom per counter value).³

In order to deal with a variety of possible languages and encodings, and since width-based methods rely on the *structure of states* but not on the *structure of actions* (i.e., action preconditions and effects), we consider *first-order languages* for describing states in terms of atoms that represent objects and relations, leaving out from the language the representation of action preconditions and effects. It is assumed that the possible state transitions (s, s') from a state s are a *function of the state* but no particular representation of this function is assumed. In addition, the state language is extended with *features* f whose values $f(s)$ in a state s are determined by the state. The features provide additional expressive power and ways for bridging different state representation languages, although they are logically redundant as their value is determined by the truth value of the atoms in the state. The features extend the notion of *derived predicates* as defined in PDDL, as they do not have to be boolean, and they do not have to be defined in the language of first-order logic or logic programs

³The problem would still not have bounded width if the counter is represented in binary using a logarithmic number of atoms.

(Thiébaux, Hoffmann, and Nebel 2005), but can be defined via procedures. Domains, instances, and states are defined as follows:

Definition 3 (Domains, problems, and states). A *domain* is a pair $D = (R, F)$ where R is a set of **primitive predicate symbols** with their corresponding arities, and F is a set of **features** defined in terms of the primitive predicates with their corresponding range of feature values. A **problem** P over domain $D = (R, F)$ is a tuple $P = (D, O, I, G)$ where O is a set of unique object names c (objects), and I and G are sets of ground atoms that denote the **initial** and **goal** states of P . A ground atom $r(c_1, \dots, c_{a(r)})$ is made of a predicate $r \in R$ and an object tuple in $O^{a(r)}$ for the arity $a(r)$ of r . A **state** s over problem $P = (D, O, I, G)$ is a collection of ground atoms. The state s is a **goal** if $G \subseteq s$, and a **dead end** if a goal state cannot be reached from s in P . A state s denotes a unique valuation for the ground atoms in P : $s \models r(c_1, \dots, c_k)$ iff $r(c_1, \dots, c_k)$ belongs to s .

This is all standard except for the two details mentioned before: there are no action schemas, and there are state features. For the former, it is implicitly assumed that in each problem P , there is a function that maps states s into the set of possible transitions (s, s') . This implies, for example, that the states may contain *static atoms*, like adjacency relations, whose truth value are not affected by any action. For the features, we make the assumption that they are **linear**, in the sense that they can be computed efficiently and only span a linear number of values. More precisely:

Linear features assumption. The features f in F are either boolean or numerical, ranging in the latter case over the non-negative integers. The value of the feature f in a state s for problem P , $f(s)$, can be computed in time bounded by $O(bN)$ where N is the number of atoms and b bounds the branching factor in P . Numerical features can take up to N values. ■

This assumption rules out features like $V^*(s)$ that stands for the optimal cost (distance) from s to a goal which may take a number of values that is not linear in the number of problem atoms, and whose computation may take exponential time. In many cases, the features can be defined in the language of first-order logic but this is not a requirement.

Example. Three state languages for Blocksworld are:

1. \mathcal{L}_{BW}^1 with the binary predicate (symbol) on^2 and the unary $ontable^1$ (superindex indicates arity),
2. \mathcal{L}_{BW}^2 with predicates on^2 , $ontable^1$, $hold^1$, and $clear^1$,
3. \mathcal{L}_{BW}^3 with predicates on^2 and $hold^1$, and boolean features $ontable^1$ and $clear^1$. ■

Example. Four languages for a domain Boxes, where boxes b containing marbles ma must be removed from a table, and for this, the marbles in the box must be removed one by one first:

1. \mathcal{L}_B^1 with predicates $ontable^1(b)$ and $in^2(ma, b)$,
2. \mathcal{L}_B^2 with predicates $ontable^1$, in^2 , and $empty^1(b)$,
3. \mathcal{L}_B^3 with predicates $ontable^1$ and in^2 , and features $n(b)$ that count the number of marbles in b ,

4. \mathcal{L}_B^4 with predicates *ontable*¹ and *in*², and features *m* and *n* counting the number of marbles in a box with the least number of marbles, and the number of boxes left. ■

By abstracting away the details of the domain dynamics and the ability to introduce features, it is simple to move from one state representation to another.

The notion of width and the IW algorithms generalize to state languages containing features in a direct fashion. In both cases, the set of atoms considered is extended to contain the possible feature values $f=v$ where f is a feature and v is one of its possible values. Features are logically redundant but can have drastic effect on the problem width.

The width for class \mathcal{Q} of problems P over some domain D is k , written as $w(\mathcal{Q}) = k$, if $w(P) = k$ for some $P \in \mathcal{Q}$ and $w(P') \leq k$ for every other problem P' in \mathcal{Q} .

Example. The width for the class of problems \mathcal{Q}_{clear} where block x has to be cleared has width 1 in the state languages \mathcal{L}_{BW}^i , $i = 1, 2, 3$, while the class \mathcal{Q}_{on} has width 2 for the three languages. On the other hand, for the class \mathcal{Q}_{B_1} of instances from Boxes with a single box, the width is not bounded as it grows with the number of marbles when encoded in the languages \mathcal{L}_B^1 and \mathcal{L}_B^2 , but it is 1 when encoded in the languages \mathcal{L}_B^3 or \mathcal{L}_B^4 . Likewise, for the class \mathcal{Q}_B of instances from Boxes with arbitrary number of boxes, the encoding in the language \mathcal{L}_B^3 has width that is not bounded as it grows with the number of boxes, but remains bounded and equal to 2 in \mathcal{L}_B^4 . ■

Generalized policies

We want to show that bounded width is a property of domains that admit a certain class of general policies. Different language for expressing general policies have been developed, some of which can deal with relational domains where different instances involve different (ground) actions. Most closely to this work, general policies have been defined in terms of qualitative numerical planning problems (QNPs) (Srivastava et al. 2011; Bonet and Geffner 2018, 2020). We build on this idea but avoid the introduction of QNPs by defining policies directly as mappings from *boolean feature conditions* into *feature value changes*.

A **boolean feature condition** for a set of features Φ is a condition of the form p or $\neg p$ for a boolean feature p in Φ , or $n = 0$ or $n > 0$ for a numerical feature n in Φ . Similarly, a **feature value change** for Φ is an expression of the form p , $\neg p$, or $p?$ for a boolean feature p in Φ , and $n\downarrow$, $n\uparrow$, or $n?$ for a numerical feature n in Φ . General policies are given by a set of rules $C \mapsto E$ where C and E stands for boolean feature conditions and feature changes respectively.

Definition 4 (Policies). A **general policy** π_Φ for a domain D over a set of features Φ is a set of rules of the form $C \mapsto E$, where C is a set of boolean feature conditions and E is a set of feature value changes. The condition $n > 0$ is assumed in rules with effects $n\downarrow$ or $n?$.

The policy π_Φ prescribes the possible actions a to be done in a state s over a problem P indirectly, as the set of state transitions (s, s') that the actions in P make possible and which are *compatible* with the policy:

Definition 5. A transition (s, s') **satisfies** the effect E when:

1. if p (resp. $\neg p$) in E , $p(s') = 1$ (resp. $p(s') = 0$),
2. if $n\downarrow$ (resp. $n\uparrow$) in E , $n(s) > n(s')$ (resp. $n(s) < n(s')$),
3. if p (resp. n) is not mentioned at all in E , $p(s) = p(s')$ (resp. $n(s) = n(s')$).

The transition (s, s') is **compatible** with policy π_Φ (or is a π_Φ -transition) if there is a policy rule $C \mapsto E$ such that s makes true C and (s, s') satisfies E .

Policy rules provide a description of how the value of the features must change along the state trajectories that are compatible with the policy. Every transition (s, s') in such trajectories have to be compatible with a policy rule $C \mapsto E$. The expressions $p?$ and $n?$ in E stand for uncertain effects, meaning that p and n may change in any way or not change at all. Features not mentioned in E , on the other hand, must keep their values unchanged in a transition compatible with the rule $C \mapsto E$.

The definition does not exclude the presence of multiple policy rules $C \mapsto E$ with conditions C that are all true in a state s . In such a case, for a state transition (s, s') to be compatible with the policy, the definition requires (s, s') to satisfy one of the effect expressions E . On the other hand, if the body C of a single rule $C \mapsto E$ is true in s , for the transition to be compatible with the policy, it must satisfy the effect E of that rule.

Example. A policy for solving the class \mathcal{Q}_{clear} can be expressed in terms of the features $\Phi = \{H, n\}$, where H is true if a block is being held, and n counts the number of blocks above x . The policy can be expressed with two rules:

$$\{\neg H, n > 0\} \mapsto \{H, n\downarrow\} ; \{H, n > 0\} \mapsto \{\neg H\}. \quad (1)$$

The first rule says that when the gripper is empty and there are blocks above x , an action that decreases n and makes H true must be chosen, while the second rule says that when the gripper holds a block and there are blocks above x , an action that makes H false and does not affect n must be selected. ■

The conditions under which a general policy π_Φ solves an instance P and class \mathcal{Q} are:

Definition 6 (Trajectories and solutions). A **state trajectory** s_0, \dots, s_n for problem P is **compatible** with policy π_Φ over features Φ (or is π_Φ -trajectory), iff s_0 is the initial state of P , no state s_i is goal, $0 \leq i < n$, and each pair (s_i, s_{i+1}) is a possible state transition in P that is compatible with π_Φ . It is **maximal** if either s_n is a goal state, no transition (s_n, s_{n+1}) in P is compatible with π_Φ , or the trajectory is infinite (i.e., $n = \infty$). A policy π_Φ **solves** a problem P if all maximal state trajectories s_0, \dots, s_n compatible with π_Φ are goal reaching (i.e., s_n is a goal state in P). π_Φ solves a collection \mathcal{Q} of problems if it solves each problem in \mathcal{Q} .

It is easy to show that the policy captured by the rules in (1) solves \mathcal{Q}_{clear} . The verification and synthesis of general policies of this type have been addressed by Bonet and Geffner (2020) in the context of qualitative numerical planning, and by Bonet, Frances, and Geffner (2019) where the set of features Φ and QNP model are learned from traces.

Generalized Policies and Width

The first result establishes a relation between classes of problems that are solvable by certain type of policies and their width. Namely, if there is a “Markovian” policy π_Φ that generates optimal plans for a class \mathcal{Q} , the problems in \mathcal{Q} can be encoded to have width bounded by $|\Phi|$.

A policy π_Φ over the features Φ is Markovian in \mathcal{Q} when the features provide a suitable abstraction of the states; i.e., when the possible next feature valuations are not just a function of the current state, but of the feature valuation in the state. More precisely, if we say that a state s is *optimal for a feature valuation* f in a problem P when $f = f(s)$ and there is no state s' with the same feature valuation $f = f(s')$ that is reachable in P in less number of steps than s , the Markovian property is defined as follows:

Definition 7 (Markovian). *A policy π_Φ is **Markovian** for a problem P iff the existence of a transition (s, s') compatible with π_Φ with feature valuations $f = f(s)$ and $f' = f(s')$, implies the existence of transitions (s_1, s'_1) with the same feature valuations $f = f(s_1)$ and $f' = f(s'_1)$, in all states s_1 that are optimal for f in P . The policy is **Markovian** for a class of problems \mathcal{Q} if it is so for each P in \mathcal{Q} .*

The states s_1 on which the Markovian condition is required in the definition are not necessarily among those which are reachable with the policy. The definition captures a *weak* Markovian assumption. The standard, but stronger, Markovian assumption requires the same condition on all states s_1 that reach the feature valuation f and not just those that reach f optimally. A sufficient condition that ensures the (strong) Markovian property is for the policy to be **deterministic** over the feature valuations, meaning that if f is followed by f' in some transition compatible with the policy, f can always be followed by f' in every state transition (s, s') compatible with the policy where $f(s) = f$, and moreover, that f can only be followed by f' then. This form of determinism results when the bodies C of the policy rules $C \mapsto E$ are logically inconsistent with each other, the heads E do not have uncertain effects $p?$ or $n?$, the domain is such that all increments and decrements $n\uparrow$ and $n\downarrow$ are by fixed amounts, and there is a transition (s, s') compatible with E in the reachable states where C holds.

For reasoning about the policy by reasoning about the features, the features however must also distinguish goal from non-goal states:

Definition 8 (Separation). *The features Φ **separate goals from non-goals** in \mathcal{Q} iff there is a set of boolean feature valuations κ such that for any problem P in \mathcal{Q} and any reachable state s in P , s is a goal state iff $f(s)$ is in κ . The valuations in κ are called **goal valuations**.*

The **boolean feature valuations** determined by state s refer to the truth valuations of the expressions p and $n = 0$ for the boolean and numerical features p and n in Φ , respectively. While the number of feature valuations is not bounded as the size of the instances in \mathcal{Q} is not bounded in general, the number of boolean feature valuations is always $2^{|\Phi|}$.

The last notion needed for relating general policies and the width of the instances is the notion of optimality:

Definition 9 (Optimal policies). *A policy π_Φ that solves a class of problems \mathcal{Q} is **optimal** if any plan ρ induced by π_Φ over a problem P in \mathcal{Q} is optimal for P .*

If the policy π_Φ solves a problem P , the plans induced by the policy are the action sequences ρ that yield the goal-reaching state trajectories s_0, \dots, s_n that are compatible with π_Φ . These plans are optimal for P if there are no shorter plans for P .

It can be shown that if π_Φ is a Markovian policy that solves the class of problems \mathcal{Q} optimally with the features Φ separating goals from non-goals, then any feature valuation f reached in the way to the goal in an instance P in \mathcal{Q} is reached optimally; i.e., no action sequence in P can reach a state s with the same feature valuation $f(s) = f$ in a smaller number of steps.

Theorem 10. *Let π_Φ be a **Markovian** policy that solves a class of problems \mathcal{Q} **optimally**, and where the features Φ **separate goals** in \mathcal{Q} . Then, π_Φ is optimal relative to feature valuations. That is, if ρ is a sequence of actions induced by π_Φ over an instance $P \in \mathcal{Q}$ that reaches the state s_i , ρ is an **optimal plan** in P for the feature valuation $f(s_i)$.*

Proof. Let $\tau^* = (s_0, \dots, s_n)$ be an optimal trajectory for P compatible with π_Φ , and let τ be an optimal trajectory that ends in state s with $f(s) = f(s_i)$, for $0 \leq i \leq n$. We want to show that $|\tau| = i$. By the Markovian property, there is a π_Φ -transition (s, s') with $f(s') = f(s_{i+1})$. If τ extended with the state s' is an optimal trajectory for $f(s_{i+1})$, we can extend it again with a π_Φ - (s', s'') into a trajectory for $f(s_{i+2})$. Otherwise, there is an optimal trajectory τ' for $f(s_{i+1})$ which can then be extended using the Markovian property into a trajectory for $f(s_{i+2})$. Thus, repeating the argument, there is an optimal trajectory τ_n for $f(s_n)$ of length at most $|\tau| + n - i$ (where the length of a trajectory is the number of transitions in it). On the other hand, since π_Φ is optimal and the features separate the goals, τ_n is a goal-reaching trajectory and thus $n \leq |\tau| + n - i$ which implies $i \leq |\tau|$. Since τ^* contains a trajectory that reaches $f(s_i)$ on i steps, $|\tau| \leq i$ and thus $|\tau| = i$. \square

As a result, under these conditions, the width of the instances in \mathcal{Q} can be bounded by the number of features $|\Phi|$ in the policy π_Φ , provided that the features are represented explicitly in the instances:

Theorem 11. *Let π_Φ be a **Markovian** policy that solves a class of problems \mathcal{Q} **optimally**, where the features Φ **separate goals**. If the features in Φ are **explicitly represented** in the instances P in \mathcal{Q} , $w(P) \leq |\Phi|$.*

Proof. Let $k = |\Phi|$, let $\tau^* = (s_0, \dots, s_n)$ be a goal-reaching π_Φ -trajectory in P , and let $t_i = f(s_i)$, $i = 0, 1, \dots, n$. Clearly, $|t_i| = k$ and we only need to show that the chain $\theta = (t_0, t_1, \dots, t_n)$ is admissible; i.e., it satisfies the 3 conditions in Definition 1. The first condition is direct since $t_0 = f(s_0)$.

For the second condition, let ρ be an optimal plan that achieves tuple t_i at state s . We need to show that there is an action a in P such that (ρ, a) is an optimal plan for t_{i+1} .

Since the transition (s_i, s_{i+1}) is reachable by π_Φ (i.e., belongs to τ^*), the Markovian property implies there must be a π_Φ -transition (s, s') associated to an action a such that $f(s') = f(s_{i+1}) = t_{i+1}$. By Theorem 10, $|\rho| = i$ and any optimal plan for t_{i+1} is of length $i + 1$. Therefore, the plan (ρ, a) is a plan for t_{i+1} of length $i + 1$ that is optimal.

For the last condition, we show that any optimal plan that achieves t_n is an optimal plan for problem P . This is direct since t_n is a goal valuation as Φ separates goals: a trajectory is goal reaching iff it ends in a state that makes t_n true. \square

Indeed, if a policy π_Φ solves \mathcal{Q} optimally under the given conditions, an *admissible chain* t_0, t_1, \dots, t_n of size $k = |\Phi|$ can be formed for solving each problem P in \mathcal{Q} optimally, where t_i is the valuation of the features in Φ at the i -th state of any state trajectory that results from applying the policy.

Related to this result, if we let IW_Φ refer to the variant of IW that replaces tuples of atoms by feature valuations and thus deems a state s novel in the search (unpruned) when $f(s)$ has not been seen before, one can show:

Theorem 12. *Let π_Φ be a Markovian policy that solves a class of problems \mathcal{Q} optimally with features Φ that separate the goals. The procedure IW_Φ solves any instance P in \mathcal{Q} optimally in time $O(N^{|\Phi|})$ where N is the number of atoms in P .*

IW_Φ can be thought as a standard breadth-first search that treats states with the same feature valuations as “duplicate” states. It runs in time $O(N^{|\Phi|})$ as the number of features in Φ is fixed and does not grow with the instance size, as opposed to N that stands for the number of atoms in the instance.

Proof. Let s_0, s_1, \dots, s_n be a goal-reaching trajectory generated by π_Φ on problem P in \mathcal{Q} , and let f_0, f_1, \dots, f_n be the sequence of feature valuations for the states in the trajectory. We prove by induction on k that IW_Φ finds a node n'_k for state s'_k such that $f(s'_k) = f_k$, the path from the root node n_0 to n'_k is **optimal**, and the node n'_k is not pruned. The base case $k = 0$ is direct since the empty path leads to s_0 and it is optimal. Let us assume that the claim holds for $0 \leq k < n$. By inductive hypothesis, IW_Φ finds a node n'_k for state s'_k such that $f(s'_k) = f_k$, the path from n_0 to n'_k is optimal, and n'_k is not pruned. Since the transition (s_k, s_{k+1}) exists in P , the Markovian property implies that there is a transition (s'_k, s'_{k+1}) with $f(s'_{k+1}) = f_{k+1}$, and thus IW_Φ generates a node n'_{k+1} for state s'_{k+1} . If n'_{k+1} is the first node where f_{k+1} holds, it is not pruned. Otherwise, there is another node n'' that makes f_{k+1} true and the length of the path from n_0 to n'' is less than or equal to the length of the path from n_0 to n'_{k+1} . In either case, since the length of an optimal path that achieves f_{k+1} is $k + 1$ by Theorem 10, IW_Φ finds an optimal path to a node that achieves f_{k+1} and is not pruned.

We have just shown that IW_Φ finds an optimal path to a state that makes f_n true. Since the features separate the goals, such a path is an optimal path for P . \square

Example. A general policy for Boxes is given by the rules: $\{m > 0\} \mapsto \{m \downarrow\}$ and $\{m = 0, n > 0\} \mapsto \{n \downarrow, m?\}$ where m and n are two features that count the number of marbles

in a box with a least number of marbles, and the number of boxes left on the table. Since the policy complies with the conditions in Theorem 11 and the two features are represented explicitly in the language \mathcal{L}_B^2 , it follows that the width of instances of Boxes in such encoding is 2. \blacksquare

Example. Similarly, the policy π_Φ with features $\Phi = \{H, n\}$ for \mathcal{Q}_{clear} given by the two rules in (1) is Markovian, solves \mathcal{Q}_{clear} optimally, and the features separate goals. Yet there are no atoms representing the counter n explicitly. Still, Theorem 12 implies that IW_Φ solves the instances in \mathcal{Q}_{clear} optimally in quadratic time. \blacksquare

Theorem 11 relates the number of features in a general policy π_Φ that solves \mathcal{Q} with the width of \mathcal{Q} provided that the features are part of the problem encodings. This, however, is not strictly necessary:

Theorem 13. *Let π_Φ be an optimal and Markovian policy that solves a class \mathcal{Q} of problems over some domain D for which Φ separates the goals. The width of the problems P is bounded by k if for any sequence of feature valuations $\{f_i\}_i$ generated by the policy π_Φ in the way to the goal, there is a sequence of sets of atoms $\{t_i\}_i$ in P of size at most k such that the optimal plans for t_i and the optimal plans for f_i coincide.*

Proof. Let f_0, \dots, f_n be the sequence of feature valuations generated by the policy π_Φ in the way to the goal in some instance P in \mathcal{Q} , and let t_0, \dots, t_n be the sequence of sets of atoms that comply with the conditions in the theorem, where $|t_i| \leq k$ for $i = 0, \dots, n$. We show that t_0, \dots, t_n is an admissible chain for P .

First, since any optimal plan for f_0 is an optimal plan for t_0 and the empty plan is optimal plan for f_0 , $t_0 \subseteq s_0$. Second, if ρ is an optimal plan for t_n , ρ is an optimal plan for f_n , and therefore an optimal plan for P as the features separate the goals. Finally, if ρ is an optimal plan for t_i we must show that there is an action a in P such that (ρ, a) is an optimal plan for t_{i+1} , $i < n$. Let ρ be an optimal plan for t_i that ends in state s . Since π_Φ is Markovian and ρ is optimal for f_i , there is a transition (s, s') in P with $f(s') = f_{i+1}$. That is, there is an action b such that the plan (ρ, b) reaches f_{i+1} . By Theorem 10, the optimal plans for f_{i+1} are of length $1 + |\rho|$, and thus (ρ, b) is an optimal plan for f_{i+1} . \square

Example. Consider an instance P in \mathcal{Q}_{clear} where B_1, \dots, B_m are the blocks above x initially, from top to bottom, $m > 0$. The feature valuations in the way to the goal following the Markovian policy π_Φ are $f_i = \{H, n = m - i\}$, $i = 1, \dots, m$, and $g_i = \{\neg H, n = m - i\}$, $i = 0, \dots, m - 1$. The policy is optimal, Markovian, and the features separate the goals. The tuples of atoms in P that capture these valuations as expressed in Theorem 13 are $t_{f_i} = \{\text{hold}(B_i)\}$ and $t_{g_i} = \{\text{ontable}(B_i)\}$ for $i > 0$, and $t_{g_0} = \{\text{clear}(B_1)\}$. Since these tuples have size 1, the widths $w(P)$ and $w(\mathcal{Q}_{clear})$ are both equal to 1. \blacksquare

Admissible Chains and Projected Policies

Theorem 13 relates the width of a class \mathcal{Q} to the size of the atom tuples t_i in the instances of \mathcal{Q} that capture the values

of the features f_i , following an optimal policy for \mathcal{Q} . We extend this result now by showing that it is often sufficient if the tuples t_i capture the value of the features f_i in some of those trajectories only. The motivation is to explain all proofs of bounded width for infinite classes of problems \mathcal{Q} that we are aware of in terms of general policies. For this, we start with the notion of *feasible chains*:

Definition 14 (Feasible chain). *Let $\theta = (t_0, t_1, \dots, t_n)$ be a chain of tuples of atoms from P . The chain θ is **feasible** in problem P if t_0 is true in the initial state, the optimal plans for t_n have length n , and they are all optimal for P .*

An **admissible** chain, as used in the definition of width, is a feasible chain that satisfies an extra condition; namely, that every optimal plan ρ for the tuple t_i in the chain can be extended with a single action into an optimal plan for the tuple t_{i+1} , $i = 0, 1, \dots, n-1$. In order to account for this key property, we map feasible chains into features and policies as follows:

Definition 15. *Let $\theta = (t_0, t_1, \dots, t_n)$ be a chain of atom tuples from P , and let $\tilde{t}_i(s)$ denote the **boolean state feature** that is true in s when t_i is true in s and t_j is false for all $i < j \leq n$, $i = 1, \dots, n$. The chain defines a **policy** π_θ over P with rules $\{\tilde{t}_i\} \mapsto \{t_{i+1}, \neg t_i\}$, $i = 0, \dots, n-1$.*

The first result gives necessary and sufficient conditions for a chain θ to be admissible.

Theorem 16. *Let $\theta = (t_0, t_1, \dots, t_n)$ be a chain of tuples for a problem P . θ is **admissible** in P if and only if θ is **feasible** and the policy π_θ solves P **optimally** and is **Markovian**.*

Proof. Let s_0 be the initial state of problem P . As always, the length $|\tau|$ of a (state) trajectory $\tau = (s_0, \dots, s_n)$ is n . We first establish some claims:

C1. If θ is admissible, there is no optimal trajectory τ for t_i that also reaches t_j for $0 \leq i < j \leq n$.

Proof. Any optimal trajectory for t_i can be extended with $j-i$ transitions into an optimal trajectory for t_j , thus no optimal trajectory for t_i can reach t_j , $j > i$.

C2. If θ is admissible, the optimal trajectories for t_i have length i like the optimal trajectories for \tilde{t}_i , $0 \leq i \leq n$.

Proof. Let τ be an optimal trajectory for t_i , and let $\tilde{\tau}$ be an optimal trajectory for \tilde{t}_i . By definition, $|\tau| \leq |\tilde{\tau}|$, by C1, $|\tilde{\tau}| \leq |\tau|$, and by admissibility of θ , $|\tau| = i$.

C3. If θ is feasible, π_θ is optimal and Markovian for P , and τ is a trajectory for t_i , there is an optimal trajectory for P of length at most $|\tau| + n - i$, $0 \leq i \leq n$.

Proof. Let τ^* be a goal-reaching π_θ -trajectory for P , and let τ be a trajectory for t_i . τ reaches some feature \tilde{t}_j for $i \leq j \leq n$. Then, there is an optimal trajectory τ_j for \tilde{t}_j , and thus for t_j , of length $|\tau_j| \leq |\tau|$. Using the Markovian property with τ^* , the trajectory τ_j can be extended with a π_θ -transition that reaches \tilde{t}_{j+1} . Repeating the argument, we find an optimal trajectory τ_n for \tilde{t}_n of length at most $|\tau| + n - i$.

C4. If θ is feasible, π_θ is optimal and Markovian for P , and τ is an optimal trajectory for t_i , $|\tau| = i$, $0 \leq i \leq n$.

Proof. Let τ be an optimal trajectory for t_i . By C3, there is an optimal trajectory for t_n of length at most $|\tau| + n - i$. By feasibility of θ , the length of any optimal trajectory for t_n is exactly n . Therefore, $n \leq |\tau| + n - i$ which implies $i \leq |\tau|$. On the other hand, any goal-reaching trajectory τ^* induced by π_θ contains a subtrajectory of length i for t_i , and thus $|\tau| \leq i$.

C5. If θ is feasible, π_θ is optimal and Markovian for P , and τ is an optimal trajectory for t_i , τ is also optimal for \tilde{t}_i , $0 \leq i \leq n$.

Proof. In the proof of C3, if the feature \tilde{t}_j is for $j > i$, we can find an optimal trajectory for P of length at most $|\tau| + n - j \leq n - 1$ since $|\tau| = i$ by C4. Then, τ reaches \tilde{t}_i . On the other hand, if τ' is an optimal trajectory for \tilde{t}_i with $|\tau'| < |\tau|$, using the Markovian property we can construct a goal-reaching trajectory of length strictly less than n . Therefore, τ is optimal for \tilde{t}_i .

Forward implication. Let us assume that θ is admissible for P . Then, by definition, θ is **feasible**. Let us now consider a maximal π_θ -trajectory τ in P . Since the initial state s_0 makes true t_0 , by admissibility and C1 and C2, the trajectory τ is of length n and ends in a state s_n that makes true \tilde{t}_n . In particular, τ is an optimal trajectory for t_n and thus it is an optimal goal-reaching trajectory for P . This shows that π_θ is **optimal**. Finally, to see that π_θ is Markovian for P , let (s_1, s'_1) be a π_θ -reachable transition and let s_2 be a state closest to s_0 with $f(s_2) = f(s_1)$. On one hand, By definition, there is one and only one policy rule in π_θ that is compatible with (s_1, s'_1) , and thus the states s_1 and s'_1 make true the features t_i and t_{i+1} respectively. On the other hand, by C2, s_2 is a closest state to s_0 that makes t_i true. Hence, by admissibility, the trajectory τ that leads to s_2 can be extended with one transition (s_2, s'_2) into an optimal trajectory for t_{i+1} . Therefore, by C1, the state s'_2 makes true \tilde{t}_{i+1} , and the transition (s_2, s'_2) is compatible with π_θ . Hence, π_θ is **Markovian**.

Backward implication. Let us assume that θ is feasible, and that π_θ is optimal and Markovian for P . We need to show that θ is an admissible chain; i.e., 1) t_0 is true in the initial state s_0 , 2) any optimal plan for t_i can be extended into an optimal plan for t_{i+1} by adding a single action, and 3) any optimal plan for t_n is an optimal plan for P . Conditions 1 and 3 follow directly from the definition of feasible chains. For 2, let $\tau = (s_0, \dots, s_i)$ be an optimal trajectory that reaches t_i ; its length is i by C4. We want to show that τ can be extended with a transition (s_i, s_{i+1}) into an optimal trajectory for t_{i+1} . By C5, τ is an optimal trajectory for \tilde{t}_i , and thus, by the Markovian property, there is a transition (s_i, s_{i+1}) that is compatible with π_θ . Therefore, the state s_{i+1} makes true \tilde{t}_{i+1} and t_i . The extended trajectory $\tau' = (s_0, \dots, s_{i+1})$ is optimal for t_{i+1} by C4. \square

This result connects admissible chains with policies that are optimal and Markovian, but it does not connect admissible chains with general policies. This is done next. We first define when a policy π_1 can be regarded as a **projection** of another policy π_2 :

Definition 17 (Projection). Let π_Φ be a policy over a class \mathcal{Q} and let $\pi_{\Phi'}$ be a policy for a problem P in \mathcal{Q} . The policy $\pi_{\Phi'}$ is a **projection** of π_Φ in P if every **maximal state trajectory** compatible with $\pi_{\Phi'}$ in P is a **maximal state trajectory** in P compatible with π_Φ .

Notice that it is not enough for the state trajectories s_0, \dots, s_i compatible with $\pi_{\Phi'}$ to be state trajectories compatible with π_Φ ; it is also required that if s_i is a final state in the first trajectory that it is also a final state in the second one. This rules out the possibility that there is a continuation of the first trajectory that is compatible with π_Φ but not with $\pi_{\Phi'}$. A result of this is that if π_Φ is optimal for \mathcal{Q} , the projected policy $\pi_{\Phi'}$ must be optimal for P . From this, the main theorem of this section follows:

Theorem 18. Let π_Φ be an **optimal** policy for a class \mathcal{Q} of problems. If for any problem P in \mathcal{Q} , there is a **feasible** chain θ of size at most k such that π_θ is a **projection** of π_Φ in P that is **Markovian**, then $w(\mathcal{Q}) \leq k$.

Proof. Theorem 16 implies that if θ is feasible and π_θ is Markovian and a projection of an optimal policy π_Φ in $P \in \mathcal{Q}$, then θ is admissible, and hence that $w(P)$ is bounded by the size of θ (i.e., max size of a tuple in θ). If this size is bounded by k for all P in \mathcal{Q} , $w(\mathcal{Q})$ is bounded by k . \square

While the proof of this theorem is a direct consequence of Theorem 16, the result is important as it renders explicit the logic underlying all proofs of bounded width for meaningful classes of problems \mathcal{Q} that we are aware of. In all such cases, the proofs have been constructed by finding feasible chains with tuples t_i that are a function of the instance $P \in \mathcal{Q}$, and whose role is to capture suitable projections of **some general optimal policy** π_Φ .

Example. Let \mathcal{Q}_G be a grid navigation domain where an agent at some initial cell has to move to a target cell with atoms of the form $x(i)$ and $y(j)$ that specify the coordinates of the agent along the two axes. An optimal policy for \mathcal{Q}_G can be obtained with the singleton $\Phi = \{d\}$ where the linear feature d measures the distance to the target cell, $d=0$ identifies the goal states, and there is a single policy rule $\{d>0\} \mapsto \{d\downarrow\}$. Let P be a problem in \mathcal{Q}_G where the agent is initially located at cell (i_0, j_0) and the goal cell is (i_n, j_m) , which for simplicity satisfies $i_0 \leq i_n$ and $j_0 \leq j_m$. Further, let $\theta = (t_0, t_1, \dots, t_{n+m})$ be the chain of tuples where $t_k = \{x(i_0 + k)\}$ for $0 \leq k < n$, and $t_k = \{x(i_n), y(j_0 + k - n)\}$ for $n \leq k \leq n+m$. Then, since θ is feasible, and π_θ is the projection of the optimal policy π_Φ while being Markovian, it follows from Theorem 18 that $w(\mathcal{Q}_G) \leq 2$. Notice that the tuples t_i in θ do not capture the values d_i of the distance feature d in **all** the trajectories that are compatible with the policy π_Φ in P (an exponential number of such trajectories), but in **one** such trajectory only; namely, the one that is compatible with the projection π_θ of π_Φ , with states $s_i, i = 0, \dots, n+m$, where the tuple t_i is true iff the feature \tilde{t}_i used in π_θ is true. \blacksquare

Example. In the Delivery (D) domain an agent moves in a grid to pick up packages and deliver them to a target cell, one by one; Delivery-1 (D_1) is the version with

one package. A general policy π_Φ for D and D_1 is given in terms of four rules that capture: move to the (nearest) package, pick it up, move to the target cell, and drop the package, in a cycle, until no more packages are left. The rules can be expressed in terms of the features $\Phi = \{H, p, t, n\}$ that express holding, distance to the nearest package (zero if agent is holding a package or no package to be delivered remains), distance to the target cell, and the number of undelivered packages respectively. Hence, $n=0$ identifies the goal states for the problems in the classes \mathcal{Q}_D and \mathcal{Q}_{D_1} for the problems D and D_1 respectively. The rules that define the policy π_Φ are $\{\neg H, p > 0\} \mapsto \{p\downarrow, t?\}$, $\{\neg H, p = 0\} \mapsto \{H\}$, $\{H, t > 0\} \mapsto \{t\downarrow\}$, and $\{H, n > 0, t = 0\} \mapsto \{\neg H, n\downarrow, p?\}$.⁴

Let us consider a problem P in the class \mathcal{Q}_{D_1} . If the encoding of P contains atoms like $at(cell_i)$, $atp(pkg_i, cell_j)$, $hold(pkg_i)$, and $empty$, it can be shown that \mathcal{Q}_{D_1} has width 2. Indeed, without loss of generality, let us assume that in P , the package is initially at $cell_i$ and has to be delivered at $cell_t$, and the agent is initially at $cell_0$, and let $\theta = (t_0, t_1, \dots, t_n)$ be the chain made of tuples of the form $\{at(cell_k)\}$ for the cells on a shortest path from $cell_0$ to $cell_i$, followed by tuples of the form $\{at(cell_k), hold(pkg_j)\}$, with $cell_k$ now ranging over a shortest path from $cell_i$ up to $cell_t$, and a last tuple t_n of the form $\{at(pkg_j, cell_t)\}$. It is easy to show that the chain θ is feasible, and π_θ is the projection of the general policy π_Φ that in D_1 is both optimal and Markovian (although not in D). Then, by Theorems 16 and 18, it follows that the chain θ is admissible, and $w(\mathcal{Q}_{D_1}) \leq 2$. \blacksquare

Serialized Width

Theorem 18 suggests that bounded width is most often the result of general policies that can be projected on suitable tuples of atoms. We extend now these relations to the much larger and interesting class of problems that have bounded *serialized width*, where the general policies do not have to be optimal or Markovian. The notion of serialized width is based on the decomposition of problems into subproblems, and explicates and generalizes the logic and scope of the Serialized IW (SIW) algorithm of Lipovetzky and Geffner (2012). There are indeed many domains with large or unbounded widths that are simple and admit simple policies, like the Delivery domain above.

In SIW, problems are serialized by using one feature, a *goal counter* denoted by $\#g$, that tracks the number of unachieved goal atoms, assuming that goals are conjunctions of atoms. Other serializations have appealed to other counters and also to heuristics (Lipovetzky and Geffner 2017a).⁵

⁴A different formulation involves packages that need to be delivered to target cells that depend on the package. The set of features is the same $\Phi = \{H, p, t, n\}$ except that t is defined as the distance to the *current target cell* (zero if agent holds nothing or there are no more packages to deliver). A policy for this formulation has the rules $\{\neg H, p > 0\} \mapsto \{p\downarrow\}$, $\{\neg H, p = 0\} \mapsto \{H, t\uparrow\}$, $\{H, t > 0\} \mapsto \{t\downarrow\}$, and $\{H, n > 0, t = 0\} \mapsto \{\neg H, n\downarrow, p?\}$.

⁵These serializations appear in state-of-the-art algorithms like BFWS that use novelty measures inside complete best-first proce-

We take a more general and abstract approach that replaces the goal counter by a *strict partial order* (an irreflexive and transitive binary relation), over the feature valuations, also referred to as Φ -tuples. A serialization for a class \mathcal{Q} of problems is defined as follows:

Definition 19 (Serializations). *Let \mathcal{Q} be a class of problems, let Φ be a set of goal-separating features for \mathcal{Q} , and let \prec be a **strict partial order** over Φ -tuples. The pair (Φ, \prec) is a **serialization** over \mathcal{Q} if 1) the ordering \prec is well-founded; i.e. there is no infinite descending chain $f_1 \succ f_2 \succ f_3 \succ \dots$ where $f_i \succ f_{i+1}$ stands for $f_{i+1} \prec f_i$, and 2) the goal feature valuations f are \prec -minimal; i.e., no $f' \prec f$ for any feature valuation f' .*

A serialization over \mathcal{Q} decomposes each problem P into subproblems which define the width of the serialization or serialized width:

Definition 20 (Subproblems). *Let (Φ, \prec) be a serialization. The subproblem $P[s, \prec]$ is the problem of finding a state s' reachable from s such that s' is goal in P or $f(s') \prec f(s)$; i.e., the goal states in $P[s, \prec]$ are the goal states in P and the states s' such that $f(s') \prec f(s)$. The collection $P[\prec]$ of subproblems for problem P is the smallest subset of problems $P[s, \prec]$ that comply with:*

1. if the initial state s_0 in P is not goal, $P[s_0, \prec] \in P[\prec]$,
2. $P[s', \prec] \in P[\prec]$ if $P[s, \prec] \in P[\prec]$ and s' is a non-goal state in P that is at **shortest distance** from s such that $f(s') \prec f(s)$, and no goal state of P is strictly closer from s than s'

Definition 21 (Serialized width). *Let (Φ, \prec) be a serialization for a collection \mathcal{Q} of problems. The **serialized width** of problem P relative to (Φ, \prec) is k , written as $w_\Phi(P) = k$ with the ordering “ \prec ” left implicit, if there is a subproblem $P[s, \prec]$ in $P[\prec]$ that has width k , and every other subproblem in $P[\prec]$ has width at most k . The **serialized width** for \mathcal{Q} is k , written as $w_\Phi(\mathcal{Q}) = k$, if $w_\Phi(P) = k$ for some problem $P \in \mathcal{Q}$, and $w_\Phi(P) \leq k$ every other problem $P' \in \mathcal{Q}$.*

If a class of problems \mathcal{Q} has bounded serialized width and the ordering $f \prec f'$ can be tested in polynomial time, each problem P in \mathcal{Q} can be solved in polynomial time using a variant SIW_Φ of the SIW algorithm: starting at the state $s = s_0$, SIW_Φ performs an IW search from s to find a state s' that is a goal state or that renders the precedence constraint $f(s') \prec f(s)$ true. If s' is not a goal state, s is set to s' , $s := s'$, and the loop repeats until a goal state is reached. The result below follows from the observations by Lipovetzky and Geffner (2012) once the goal counter is replaced by a partial order, and the notion of serialized width is suitably formalized:

Theorem 22. *Let \mathcal{Q} be a collection of problems and let (Φ, \prec) be a serialization for \mathcal{Q} with width $w_\Phi(\mathcal{Q}) \leq k$. Any problem P in \mathcal{Q} is solved by SIW_Φ in time and space bounded by $O(bN^{|\Phi|+2k-1}\Lambda)$ and $O(bN^k + N^{|\Phi|+k})$ respectively, where b bounds the branching factor in P , N is the number of atoms in P , and Λ bounds the time to test the order \prec for the Φ -tuples that arise from the states in P .*

dures.

Proof. We show something more general. If P is a problem in \mathcal{Q} , SIW_Φ finds a plan ρ for P and state sequence s_0, s_1, \dots, s_n such that

1. the state s_0 is the initial state in P ,
2. the state s_{i+1} , reachable from state s_i , can be found by $\text{IW}(k)$ on the problem $P[s_i, \prec]$, for $i = 0, 1, \dots, n-1$,
3. $f(s_{i+1}) \prec f(s_i)$, for $i = 0, 1, \dots, n-1$,
4. s_n is a goal state, and
5. $|\rho| \leq N^{|\Phi|+k}$.

SIW_Φ achieves this in time and space bounded by $O(bN^{|\Phi|+2k-1}\Lambda)$ and $O(bN^k + N^{|\Phi|+k})$ respectively.

By definition, SIW_Φ calls IW iteratively to find such a sequence. In the first call, over subproblem $P[s_0, \prec]$, IW finds a state s_1 such that $f(s_1) \prec f(s_0)$. Then, iteratively, after s_i is found, IW is called to solve the subproblem $P[s_i, \prec]$. The process continues until IW finds a goal state s_n .

Let us assume that such a sequence is found by SIW_Φ . A run of IW involves runs of $\text{IW}(1), \text{IW}(2), \dots$ until the problem is solved (guaranteed to succeed before the call to $\text{IW}(k+1)$ since $w_\Phi(\mathcal{Q}) \leq k$). By Theorem 2 and the assumption on the cost of testing \prec , each call to IW requires time and space bounded by $O(bN^{2k-1}\Lambda)$ and $O(bN^k)$ respectively. The plan ρ_i found by IW that maps s_i into s_{i+1} has length bounded by N^k . All these plans are concatenated into a single plan ρ that solves P . Since the number of Φ -tuples is bounded by $N^{|\Phi|}$, SIW_Φ invests total time and space bounded by $O(bN^{|\Phi|+2k-1})$ and $O(bN^k + N^{|\Phi|+k})$ respectively, and $|\rho| \leq N^{|\Phi|+k}$.

We now show that the state sequence is indeed found. The first requirement is clear since s_0 is the initial state in P . We reason inductively to show that the subproblem $P[s_i, \prec]$ belongs to the collection $P[\prec]$ of subproblems, and that the state s_{i+1} is at minimum distance from s_i such that $f(s_{i+1}) \prec f(s_i)$.

For the base case, $P[s_0, \prec]$ belongs to $P[\prec]$ by definition of $P[\prec]$ when s_0 is not a goal state; if s_0 is a goal state, the sequence with the single state s_0 satisfies the requirements. Since $w_\Phi(P) \leq k$, $\text{IW}(k)$ is guaranteed to find a state s_1 at minimum distance from s_0 that is a goal or $f(s_1) \prec f(s_0)$. If s_1 is a goal state, the sequence s_0, s_1 satisfies the requirements. Otherwise, by definition of $P[\prec]$, $P[s_1, \prec]$ belongs to $P[\prec]$. Assume now that s_i is not a goal state and $P[s_i, \prec]$ belongs to $P[\prec]$. A similar argument shows that $\text{IW}(k)$ finds a state s_{i+1} at minimum distance from s_i that is a goal or $f(s_{i+1}) \prec f(s_i)$. If s_{i+1} is a goal state, SIW_Φ has found the required state sequence. Otherwise, the problem $P[s_{i+1}, \prec]$ belongs to $P[\prec]$. Finally, since the number of Φ -tuples is finite, this process terminate with a goal state s_n . \square

The class \mathcal{Q}_{VA} of instances for the problem VisitAll, where all cells in a grid must be visited, has serialized width $w_\Phi(\mathcal{Q}_{VA}) = 1$ for $\Phi = \{\#g\}$ where $\#g$ counts the number of unvisited cells (unachieved goals), with the obvious ordering ($'\prec'$ set to $'<'$). The class \mathcal{Q}_{BW} of Blocksworld instances cannot be serialized into subproblems of bounded width with $\Phi' = \{\#g\}$ because achieved goals may have to

be undone, yet with $\Phi = \{\#m\}$ where $\#m$ counts the number of misplaced blocks (i.e., not on their targets or above one such block), $w_\Phi(Q_{BW}) = 2$. The result above implies that SIW_Φ solves these problems in polynomial time.

From General Policies to Serializations

We show next how serializations can be inferred from general policies, a result that paves the way to introduce a convenient language for defining serializations. Given a general policy π_Φ that solves a class of problems \mathcal{Q} for features Φ that separates the goals, the intuition is to consider serializations (Φ, \prec) where $f' \prec f$ if there is a state transition (s, s') compatible with the policy π_Φ in some instance P in \mathcal{Q} with $f' = f(s')$ and $f = f(s)$. The problem with this ordering, however, is that it does not necessarily define a strict partial order as required, as it is possible that there is another transition (s_1, s_2) compatible with the policy in another instance $P' \in \mathcal{Q}$ where $f(s_1) = f'$ and $f(s_2) = f$. In other words, even if the policy orders the feature valuations in a strict partial order in an instance P , it is possible that the orderings over two or more instances in \mathcal{Q} cannot be combined into a single strict partial order.

Thus in order to obtain a serialization (Φ, \prec) over a class of problems \mathcal{Q} from a policy π_Φ certain additional restrictions are required. For this, we focus on the policies π_Φ that solve \mathcal{Q} **structurally**; i.e., by virtue of the effects of the actions on the features as expressed in the policy rules in π_Φ . A policy solves a class of problems \mathcal{Q} structurally when this can be verified from the form of the policy rules without having to consider the instances in \mathcal{Q} . Fortunately, we do not have to formulate the conditions under which a policy π_Φ solves a class of problems \mathcal{Q} structurally from scratch, we can use ideas developed in the context of QNPs (Srivastava et al. 2011; Bonet and Geffner 2020).

The key idea is the notion of **terminating policy graph**. Recall that every feature valuation f defines a projected boolean valuation $b(f)$ over the expressions p and $n = 0$ for the boolean and numerical features p and n in Φ , where p and $n = 0$ are true in $b(f)$ iff p and $n = 0$ are true in f respectively. The policy graph for policy π_Φ is:

Definition 23 (Policy graph). *The policy graph $G(\pi_\Phi)$ for policy π_Φ has nodes b , one for each of the $2^{|\Phi|}$ boolean feature valuations over Φ , and edges $b \rightarrow b'$ labeled with E if b is not a goal valuation and (b, b') is compatible with a rule $C \rightarrow E$ in the policy.*

The edge $b \rightarrow b'$ is compatible with a rule $C \rightarrow E$ if C is true in b , and (b, b') is compatible with E ; namely, if p (resp. $\neg p$) is in E , p (resp. $\neg p$) must be true in b' , and if $n\uparrow$ is in E , $n = 0$ must be false in E . Effects $p?$, $n?$, and $n\downarrow$ in E put no constraints on b' . In particular, in the latter case, $n = 0$ can be either true or false in b' , meaning that after a decrement, a numerical feature may remain positive or have value 0. Notice that the policy graph associated with a policy π_Φ does not take the class of problems \mathcal{Q} into account. Indeed, the policy graph may have an edge $b \rightarrow b'$ even if there is no state transition (s, s') in an instance in \mathcal{Q} where this transition between boolean feature valuations arises. The policy

graph and the notion of **termination** below are purely **structural** as they depend on the form of the policy rules only:

Definition 24 (Termination). *A policy π_Φ and a policy graph $G(\pi_\Phi)$ are **terminating** if for every cycle in the graph $G(\pi_\Phi)$, i.e., any sequence of edges $b_i \rightarrow b_{i+1}$ with labels E_i that start and end in the same node, there is a numerical feature n in Φ that is decreased along some edge and increased in none. That is, $n\downarrow \in E_k$ for some k -th edge in the cycle, and $n\uparrow \notin E_j$ and $n? \notin E_j$ for all others edges in the cycle.*

The termination condition can be checked in time that is polynomial in the size of the policy graph by a procedure called SIEVE that looks at the strongly connected components in the graph (Srivastava et al. 2011; Bonet and Geffner 2020).⁶ A key property of terminating policies is that they give rise to state trajectories that are finite.

Theorem 25. *Let π_Φ be a terminating policy with features Φ over \mathcal{Q} that separate the goals. Then, π_Φ cannot give rise to infinite state trajectories in instances P in \mathcal{Q} .*

Proof. For a proof by contradiction, let us assume that there is an infinite state trajectory s_0, s_1, \dots that is compatible with π_Φ . Since the number of states in P is finite, there must be indices $i < j$ such that $s_i = s_j$. The subsequence s_i, \dots, s_j defines a cycle $b(f(s_i)), \dots, b(f(s_j))$ in the policy graph for π_Φ . The termination condition implies that there is a numerical feature n in Φ that is decreased by some edge in the cycle but not increased by any edge in the cycle. However, this is impossible since $s_i = s_j$ implies that the feature n has the same value on both states. Therefore, such infinite state trajectory cannot exist. \square

Since terminating policies induce state trajectories with a final state, a sufficient condition for a terminating policy to solve a class \mathcal{Q} is to be closed:

Definition 26 (Closed policy). *A policy π_Φ is **closed** over a class of problems \mathcal{Q} if for any non-goal state s in a problem P in \mathcal{Q} that is π_Φ -reachable, there is a transition (s, s') that is compatible with π_Φ .*

Theorem 27. *If π_Φ is closed over \mathcal{Q} , the features in Φ separate goals in \mathcal{Q} , and π_Φ is terminating, π_Φ solves \mathcal{Q} .*

Proof. If π_Φ does not solve instance P in \mathcal{Q} , there is a maximal non-goal reaching trajectory in P that is compatible with π_Φ . This trajectory is either infinite or terminates at state s_n where no transition (s_n, s) compatible with π_Φ exists in P . The former is impossible by Theorem 25, and the latter is impossible since π_Φ is closed. \square

We are interested in the conditions under which general policies determine serializations, and in particular, serializations with bounded width. For the first part, we do not need the policy to be closed or solve \mathcal{Q} , but just to be terminating:

⁶See Bonet and Geffner (2020) for more details on the formal properties of termination in QNPs. Our setting is slightly different as our numerical features are linear and cannot grow without bound as in QNPs, but we do not use this property to define termination.

Theorem 28. A terminating policy π_Φ with features Φ over \mathcal{Q} that separate goals determines a serialization (Φ, \prec) of \mathcal{Q} where ‘ \prec ’ is the minimal strict partial relation (i.e., the transitive closure) that satisfies $f' \prec f$ for a non-goal valuation f , if (f, f') is compatible with a policy rule in π_Φ .

Here a pair of feature valuations (f, f') is compatible with a rule $C \rightarrow E$, if C is true in f and the change in values from f to f' is compatible with E . Notice that if a state transition (s, s') is compatible with $C \rightarrow E$, the pair of feature valuations $(f(s), f(s'))$ is compatible with the rule, but the existence of such a state transition is not a requirement for the pair (f, f') to be compatible with the policy rule; they can be arbitrary feature valuations over Φ .

Proof. We must show that \prec is irreflexive and well founded. For the first, assume that there is a sequence of feature valuations $f_1 \prec f_2 \prec \dots \prec f_n$ with $n > 1$ and $f_1 = f_n$. Without loss of generality, we may assume that the pair (f_i, f_{i+1}) is compatible with a policy rule (i.e., $f_{i+1} \prec f_i$ is not the result of the transitive closure). In such a case, the boolean valuations b_i associated with the feature valuations f_i must form a cycle in the policy graph, and therefore from the termination condition, some variable must be decreased and not increased in the cycle, which contradicts $f_1 = f_n$.

For well-foundedness, let us assume that there is an infinite chain $f_1 \succ f_2 \succ \dots$ where $f_i \succ f_{i+1}$ stands for $f_{i+1} \prec f_i$, and let us assume, without loss of generality, that (f_i, f_{i+1}) satisfies a policy rule. Consider the sequence boolean valuations $\{b_i\}_i$ associated with the feature valuations $\{f_i\}_i$, let b be a boolean valuation that appears infinitely often in the sequence, and let B be the set of edges $b \rightarrow b'$ in the policy graph that are traversed infinitely often in the sequence. Then, by definition of B , there is an infinite subsequence $\{i_j\}_{j \geq 0}$ of indices such that $b_{i_j} = b$ and all the edges traversed in the sequence $b_{i_j}, \dots, b_{i_{j+1}}$ belong to B , for $j \geq 0$. By definition of termination, there is a numerical feature n that is decremented in some edge in B but not incremented by any edge in B . Yet this is impossible, since the numerical feature n is integer valued, non-negative, it is decremented infinitely often, and never incremented. The contradiction comes for supposing the existence of the infinite chain $f_1 \succ f_2 \succ \dots$. Therefore, \prec is a well founded strict partial order. \square

There are simple syntactic conditions that ensure that a set of policy rules and the graph that they define are terminating (Bonet and Geffner 2020). For example: a) each policy rule decreases some feature and increases none, b) there is an ordering n_1, \dots, n_k of the numerical features such that each policy rule decreases a feature n_i , while possibly increasing features n_j , $j > i$, and c) a further generalization where b) holds except in some policy rules that affect boolean features only, which cannot produce a cycle of truth valuations over the boolean features without involving a policy rule that affects some numerical variable. If the set of policy rules complies with the conditions a) or b), the rules are said to be **regular**; if they comply with c), the rules are said to be **weakly regular**. In either case, these are sufficient syntactic conditions that ensure termination; unlike the conditions

captured by Definition 24, they are not necessary.

Example. The policy π_Φ for Boxes with rules $\{m > 0\} \mapsto \{m \downarrow\}$ and $\{m = 0, n > 0\} \mapsto \{n \downarrow, m?\}$ and goal $n = 0$ is **closed** and **regular**. It is closed since for any state s where $n > 0$, there is a transition (s, s') that is compatible with π_Φ : if $m = 0$, s' results from an action that puts an empty box away, while if $m > 0$, s' results from an action that puts a marble away from a box with a least number of marbles. Theorem 27 implies that π_Φ solves \mathcal{Q}_B . \blacksquare

If a terminating policy for the class \mathcal{Q} defines a serialization over \mathcal{Q} , a terminating policy that solves \mathcal{Q} should define a serialization over \mathcal{Q} with 0 width. Two conditions are needed for this though. The first is the notion of **goal connectedness**:

Definition 29 (Goal connected). A policy π_Φ for \mathcal{Q} with goal separating features and its policy graph are said to be **goal connected** when all nodes $b(f(s_0))$ associated with the initial states s_0 of instances P in \mathcal{Q} are connected only to nodes b that are connected with goal nodes.

Clearly, a policy π_Φ for \mathcal{Q} is not closed if its policy graph is not goal-connected, but goal-connectedness does not imply that the policy is closed. For this, we need a condition that goes beyond the structure of the policy graph.⁷

Definition 30 (Sound policy). A policy π_Φ over \mathcal{Q} is **sound** if for any reachable non-goal state s in an instance P in \mathcal{Q} where the policy rule $C \mapsto E$ is applicable (i.e., where C holds), there is a transition (s, s') in P that is compatible with π_Φ .

Soundness and goal connectedness imply that a policy is closed, and both are critical for establishing the conditions under which the serialization induced by a terminating policy has zero width:

Theorem 31. If π_Φ is sound and goal-connected in \mathcal{Q} , then π_Φ is closed in \mathcal{Q} .

From Theorem 27, it follows that a terminating policy π_Φ that is sound and goal-connected in \mathcal{Q} , solves \mathcal{Q} . In such a case, we say that the policy π_Φ solves the class of problems \mathcal{Q} **structurally**, as two of the conditions, termination and goal-connectedness, can be tested on the policy graph. Soundness, on the other hand, is the condition that ensures that only sink nodes in the policy graph over any instance $P \in \mathcal{Q}$ are the goal nodes.

Proof. If π_Φ is not closed on some instance P in \mathcal{Q} , there is a non-goal state s in P that is reachable from the initial state s_0 using π_Φ , but there is no transition (s, s') that is compatible with π_Φ . The node $b(f(s_0))$ is connected in the policy graph to $b(f(s))$, which from the definition of goal connectedness, must be connected to a goal node. This implies that there is an edge $b(f(s)) \rightarrow b'$ in the policy graph, and therefore there is some policy rule $C \mapsto E$ that is applicable at s . Soundness of π_Φ then implies that there is a transition (s, s') that is compatible with π_Φ , a contradiction. \square

⁷The notion of soundness is similar to action soundness in QNPs when QNPs are used to abstract classes of problems (Bonet and Geffner 2018; Bonet, Frances, and Geffner 2019).

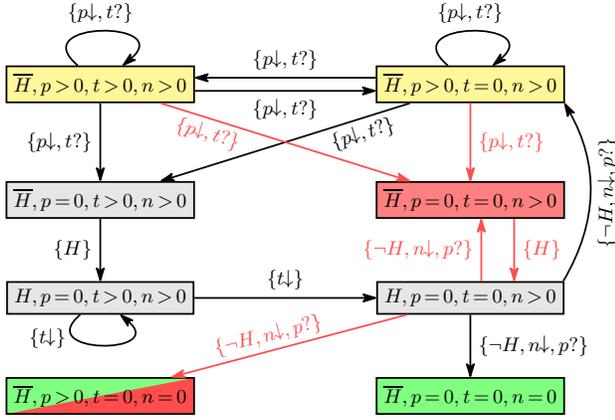


Figure 1: Policy graph for Delivery for the policy defined by the rules $\{\neg H, p > 0\} \mapsto \{p \downarrow, t?\}$, $\{\neg H, p = 0\} \mapsto \{H\}$, $\{H, t > 0\} \mapsto \{t \downarrow\}$, and $\{H, n > 0, t = 0\} \mapsto \{\neg H, n \downarrow, p?\}$. Yellow and green nodes denote initial and goal nodes respectively. Red nodes and edges stand for nodes and transitions in the policy graph that do not arise in the instances. The graph is terminating and goal connected, and the policy is closed and sound for the classes \mathcal{Q}_D and \mathcal{Q}_{D_1} .

As expected, if a terminating policy π_Φ solves \mathcal{Q} , the width of \mathcal{Q} under the serialization determined by the policy (Theorem 28) is 0, provided however, that certain structural conditions hold:

Theorem 32. *Let π_Φ be a policy that solves \mathcal{Q} structurally and let (Φ, \prec) be the serialization over \mathcal{Q} determined by π_Φ . If π_Φ is sound and goal connected, $w_\Phi(\mathcal{Q}) = 0$.*

Proof. By definition of serialized width, for any problem P in \mathcal{Q} , we need to show that the width of any subproblem $P[s, \prec]$ in $P[\prec]$ is zero. For this, it is enough to show that for any such state s , there is a transition (s, s') that is compatible with π_Φ because then, by definition, $f(s') \prec f(s)$ and $P[s, \prec]$ has width equal to zero.

If $P[s_0, \prec]$ is in $P[\prec]$ for the initial state s_0 of P , s_0 is not a goal state and there is some transition (s_0, s') that is compatible with π_Φ since π_Φ solves P . Let us now consider the subproblem $P[s', \prec]$ for $s' \neq s_0$. By definition, there is a subproblem $P[s, \prec]$ in $P[\prec]$ such that s' is a non-goal state at shortest distance from s with $f(s') \prec f(s)$. By definition of \prec and reasoning inductively, there is a path in the policy graph for π_Φ from the boolean feature valuation $b(f(s_0))$ to $b(f(s'))$. Since the features in Φ separate goals in \mathcal{Q} and π_Φ is goal connected, there is a policy rule $C \mapsto E$ that is applicable at s' . Thus, by soundness, there must be a transition (s', s'') in P that is compatible with π_Φ . \square

Example. The policy for Boxes in the previous example is closed, goal connected, and sound. By Theorem 32, it determines a serialization (Φ, \prec) of width zero, where $f(s) = [n(s), m(s)] \prec f(s') = [n(s'), m(s')]$ iff $n(s) < n(s')$ or $n(s) = n(s')$ and $m(s) < m(s')$. \blacksquare

Example. In Delivery, we use the features $\Phi = \{H, p, t, n\}$ and the policy π_Φ defined by the rules $\{\neg H, p > 0\} \mapsto$

$\{p \downarrow, t?\}$, $\{\neg H, p = 0\} \mapsto \{H\}$, $\{H, t > 0\} \mapsto \{t \downarrow\}$, and $\{H, n > 0, t = 0\} \mapsto \{\neg H, n \downarrow, p?\}$. It is easy to check that π_Φ is sound for the classes \mathcal{Q}_D and \mathcal{Q}_{D_1} since for any reachable non-goal state s in a problem P , there is a transition (s, s') that is compatible with π_Φ . On the other hand, the policy graph for π_Φ , depicted in Fig. 1, is clearly terminating and goal connected. Since Φ separates goals for the classes \mathcal{Q}_D and \mathcal{Q}_{D_1} , by Theorem 27, π_Φ solves both classes structurally, and the induced serialization (Φ, \prec) has width zero for the classes \mathcal{Q}_D and \mathcal{Q}_{D_1} by Theorem 32. \blacksquare

Sketches: A Language for Serializations

The notion of serialization plays a key role in the connection between general policies and serialized width, and since width is a special case of serialized width,⁸ serializations emerge as a central concept in the study. Indeed, serializations can be used to reduce the width of a problem, and if this width is reduced to zero over \mathcal{Q} , a general policy for \mathcal{Q} results. We focus next on a specific **language** for specifying serializations in a compact manner. The language can be used either to **encode serializations** by hand, as a form of domain-specific knowledge for extending the scope of width-based algorithms such as SIW, or for **learning serializations** from traces. Such uses, however, are beyond the scope of this paper.

A **policy sketch** or simply **sketch** R_Φ , for a class of problems \mathcal{Q} , is a set of policy rules over the features Φ that distinguish the goals of \mathcal{Q} . The sketch R_Φ can be a full fledged policy over \mathcal{Q} , part of it, or just set of policy rules $C \rightarrow E$, including the empty set. By interpreting R_Φ as a policy, we can transfer previous results that involve policy graphs and termination. We call the rules in a sketch R_Φ , **sketch rules** because their **semantics** is different from the semantics of policy rules.

Definition 33 (Sketch). *A sketch for \mathcal{Q} is a set R_Φ of sketch rules $C \rightarrow E$ over features Φ that **separate goals** in \mathcal{Q} . The sketch R_Φ is **well-formed** if the set of rules R_Φ interpreted as a policy is **terminating**.*

Notice that the definition of terminating policies does not require the policy to be closed or even to solve \mathcal{Q} . Theorem 28 directly yields:

Theorem 34. *A well-formed sketch R_Φ for \mathcal{Q} defines a serialization (Φ, \prec) over \mathcal{Q} where ‘ \prec ’ is the smallest strict partial order that satisfies $f' \prec f$ if the pair of feature valuations (f, f') is compatible with a sketch rule in R_Φ .*

The distinction between policy and sketch rules is **semantic**, not syntactical. A policy π_Φ defines a **filter on state transitions**: a state transition is compatible with the policy if it is compatible with one of its rules. A sketch R_Φ , on the other hand, is not to be used in this way: a well-formed sketch defines a serialization, and a sketch rule $C \mapsto E$ defines **subproblems of the serialization**: the subproblem of going from a state s where C holds in $f = f(s)$ to a state s' with feature valuation $f' = f(s')$ such that the pair (f, f') is compatible with the sketch rule. The key difference

⁸ $w(P) = w_\Phi(P)$ for the empty serialization (Φ, \prec) ; i.e., the one for which $f \prec f'$ is always false.

is that this subproblem is not solvable in a single step in general. Theorems 25 and 28 about policy rules that are terminating and that induce well-formed serializations are valid for sketch rules, because in both cases, the relevant notions, like sketch graphs, are defined structurally, without considering the state transitions that are possible in the target class \mathcal{Q} . Another way to see the difference between policies and sketches is that (terminating) policies π_Φ play **two different roles** in our analysis: they specify **control**, i.e., the possible actions to be done in a given instance P of \mathcal{Q} , and they define **serializations**. Sketches, on the other hand, play the latter role only.

Sketches provide a language for decomposing a problem into subproblems and thus for reducing its width, which goes well beyond the language of goal counters and variations, as the language for sketches includes the language of general policies.

Example. The serialization given by the single feature $\#g$ that counts the number of unachieved (top) goals is captured with the sketch that only contains the rule $\{\#g > 0\} \mapsto \{\#g \downarrow\}$ when there are no other features, and the rule $\{\#g > 0\} \mapsto \{\#g \downarrow, p?, n?\}$ when p and n are other features. The rules say that it is “good” to decrease the goal counter independently of the effects on other features. In problems such as Blocksworld, this serialization does not work (has unbounded width), but serializing instead with the single feature $\#m$ that counts the number of misplaced blocks with the sketch rule $\{\#m > 0\} \mapsto \{\#m \downarrow\}$, yields a serialization of width 2. A block is misplaced when it is on a wrong block or is above a misplaced block. The sketch thus decomposes any Blocksworld problem into subproblems whose goals are to put a block away or to put them at the right place. The width of the first subproblem is 1 while for second is 2. \blacksquare

Our last results are about the width of the serializations defined by sketches, and the modifications in the SIW_Φ algorithm to work with sketches:

Definition 35 (Sketch width). *Let R_Φ be a well-formed sketch for a class of problems \mathcal{Q} such that Φ separates the goals, and let s be a reachable state in some instance P of \mathcal{Q} . The width of the sketch R_Φ at state s of problem P , $w_R(P[s])$, is the width of the subproblem $P[s]$ that is like P but with initial state s and goal states s' such that s' is a goal state of P , or the pair $(f(s), f(s'))$ is compatible with a sketch rule $C \mapsto E$. The **width of the sketch** R_Φ , $w_R(\mathcal{Q})$, is the maximum width $w_R(P[s])$ for any reachable state s in any problem P in \mathcal{Q} .*

Theorem 36. *Let R_Φ be a well-formed sketch for a class \mathcal{Q} of problems, and let (Φ, \prec) be the serialization determined by R_Φ from Theorem 34. The width $w_\Phi(\mathcal{Q})$ of the serialization is bounded by the width $w_R(\mathcal{Q})$ of the sketch.*

Proof. By definition, $w_\Phi(\mathcal{Q}) \leq k$ if the width $w(P[s, \prec])$ of the subproblems $P[s, \prec]$ in $P[\prec]$ is bounded by k , for any P in \mathcal{Q} . The goals of subproblem $P[s, \prec]$ are the states s' that are a goal of P , or $f(s') \prec f(s)$. In particular, if the pair $(f(s), f(s'))$ is compatible with a sketch rule,

then $f(s') \prec f(s)$, but the converse does not hold in general. That is, for any subproblem $P[s, \prec]$ in $P[\prec]$, the goals of the subproblem $P[s]$ are also goals of $P[s, \prec]$. Hence, $w_\Phi(P[s, \prec]) \leq w_R(P[s])$ and $w_\Phi(\mathcal{Q}) \leq w_R(\mathcal{Q})$. \square

If a well-formed sketch R_Φ has bounded width for a class of problems \mathcal{Q} , then the problems in \mathcal{Q} can be solved in polynomial time by the algorithm SIW_R that is like SIW_Φ , with the difference that the precedence test $f \prec f'$ among pairs of feature valuations f and f' is replaced by the test of whether the feature valuation pair (f, f') is compatible with a rule in R_Φ . In other words, SIW_R start at the state $s := s_0$, where s_0 is the initial state of P , and then performs an IW search from s to find a state s' that is a goal state or such the pair $(f(s), f(s'))$ is compatible with a sketch rule in R_Φ . Then if s' is not a goal state, s is set to s' , $s := s'$, and the loop repeats until a goal state is reached. The precedence test in R_Φ can be done in constant time unlike the general test $f' \prec f$ in SIW_Φ .⁹ The runtime properties of SIW_R are thus similar to those of SIW_Φ , as captured in Theorem 22, with precedence tests that can be done in constant time:

Theorem 37. *Let R_Φ be a well-formed sketch for a class \mathcal{Q} of problems. If the sketch width $w_R(\mathcal{Q})$ is bounded by k , SIW_R solves any problem P in \mathcal{Q} in $O(bN^{|\Phi|+2k-1})$ time and $O(bN^k + N^{|\Phi|+k})$ space, where b and N bound the branching factor and number of atoms in P respectively.*

Proof. Like the proof of Theorem 22 but with the test $f \prec f'$ in SIW_Φ replaced by checking if the pair of feature valuations (f, f') is compatible with a rule in R_Φ . \square

Example. Different and interesting sketches are given in Table 1 for the two classes of problems for Delivery: the class \mathcal{Q}_D of problems with an arbitrary number of packages and the class \mathcal{Q}_{D_1} of problems with a single package. In the table, the entries in the columns \mathcal{Q}_{D_1} and \mathcal{Q}_D upper bound the width of the different sketches in the table for the two classes of Delivery problems. The entries *unb* and ‘—’ stand respectively for unbounded width and ill-defined (non-terminating) sketch. The features used are: (boolean) H for holding a package, p is distance to nearest package (zero if holding a package or no package to be delivered remains), t is distance to current target cell (zero if holding nothing), and n is number of packages still to be delivered.

We briefly explain the entries in the table without providing formal proofs (such proofs can be obtained with Theorem 18). σ_0 is the empty sketch whose width is the

⁹For testing $f' \prec f$, one needs to check if there is a sequence $\{f_i\}_{i=0}^n$ of feature valuations such that $f_0 = f$, $f_n = f'$, and each pair (f_i, f_{i+1}) is compatible with a sketch rule, $i = 0, 1, \dots, n-1$. However, this test can be done in constant time too, provided that the binary relation ‘ \prec ’ for each instance P is precompiled in a boolean hash table with N^k rows and N^k columns where N is the number of atoms in P , k is the number of features, and N^k is the number of feature valuations in P . Unlike the procedure SIW_R , this precompilation is not practical in general. The efficiency of SIW_R comes at a price: by testing “progress” with the sketch rules directly and not with the serializations that results from such rules, SIW_R is not using the serialization at full, as it ignores the transitive closure of the precedence relations.

Policy sketch	\mathcal{Q}_{D_1}	\mathcal{Q}_D
$\sigma_0 = \{\}$	2	<i>unb</i>
$\sigma_1 = \{\{H\} \mapsto \{\neg H, p?, t?\}\}$	2	<i>unb</i>
$\sigma_2 = \{\{\neg H\} \mapsto \{H, p?, t?\}\}$	1	<i>unb</i>
$\sigma_3 = \sigma_1 \cup \sigma_2$	—	—
$\sigma_4 = \{\{n > 0\} \mapsto \{n\downarrow, H?, p?, t?\}\}$	2	2
$\sigma_5 = \sigma_2 \cup \sigma_4$	1	1
$\sigma_6 = \{\{\neg H, p > 0\} \mapsto \{p\downarrow, t?\}\}$	2	<i>unb</i>
$\sigma_7 = \{\{H, t > 0\} \mapsto \{t\downarrow, p?\}\}$	2	<i>unb</i>
$\sigma_8 = \sigma_2 \cup \sigma_4 \cup \sigma_6 \cup \sigma_7$	0	0

Table 1: Upper bounds on the width of different sketches for the classes \mathcal{Q}_{D_1} and \mathcal{Q}_D of Delivery problems. The entries *unb* and ‘—’ mean, respectively, unbounded width and ill-defined sketch. For sketches of bounded width, SIW_R solves any instance in the class in polynomial time.

same as the plain width, 2 for D_1 and unbounded for D , as no problem P is decomposed into subproblems. The rule $\{H\} \mapsto \{\neg H, p?, t?\}$ in σ_1 does not help in initial states that do not satisfy H , and hurts a bit in states that do satisfy H . Indeed, in the latter states s , there is a state s' at one step ahead from s with $f(s') \prec f(s)$ (obtained from s by dropping the package) that gets chosen by any algorithm like SIW_R. However, in the resulting subproblem with initial state s' there is no state s'' for which $f(s'') \prec f(s')$ holds as there is no sketch rule $C \mapsto E$ where C is true in s' . As a result, the goal of the subproblem rooted as s' is the true goal of the problem, which may actually involve going back from s' to s and from s to the goal. This is indeed what SIW_R does given σ_1 . Since this subproblem has width 2, the serialized width of D_1 remains 2. For σ_2 , the rule $\{\neg H\} \mapsto \{H, p?, t?\}$ says that a state s where $\neg H$ holds can be “improved” by finding a state s' where H holds, while possibly affecting p , t , or both. Hence, any problem P in D_1 is split in two subproblems: achieve H first and then the goal, reducing the sketch width of D_1 to 1 but not the sketch width of D . The sketch σ_3 is not well-formed as it is not terminating: indeed the resulting ordering is not a strict partial order: a state s where the agent is at the same location as a package but does not hold it is “improved” into a state s' by picking the package, which is improved again by dropping the package without affecting any numerical feature. For σ_4 , that subgoals on the top goal captured by the feature n , that counts the number of undelivered packages, the sketch width of D reduces to 2 but that of D_1 remains unchanged. The sketch σ_5 combines the rules in σ_2 and σ_4 , and as a result, decomposes D into width 2 subproblems, the first of which, that corresponds to D_1 , is further decomposed into two subproblems. The sketch width of both D and D_1 becomes then 1: the first subproblem is to collect the nearest package, the second one to drop it at the target cell, and the same sequence of subproblems gets repeated. The sketch σ_6 renders the subproblem of getting close to the nearest package with width 0, but the remaining subproblem in D_1 that involves picking up the package and delivering it at the target cell, still has width 2. The sketch σ_7 does not decompose

Thm	Notes
2	Performance and guarantees of IW(k).
11	Optimal and Markovian policies bound width if features encoded.
12	Markovian policies guarantee optimal solutions with IW _{Φ} .
13	Bounded width also when tuples capture the features in all optimal trajectories.
18	Bounded width when tuples capture features in a projection.
22	Performance and guarantees of SIW _{Φ} .
25	Conditions for termination of policies.
27	Closed and terminating policies define structural solutions.
28	Terminating policies define serializations.
32	Structural solutions that are sound and closed define serializations of zero width.
34	Well-formed sketches define serializations.
36	Sketch width bounds width of induced serialization.
37	Performance and guarantees of SIW _R given sketches.

Table 2: Summary of main formal results.

the problem when H is initially false. Finally, the combination in σ_8 yields a full policy, and thus a serialization of width 0 where each subproblem is solved in a single step. ■

Conclusions

We have established a number of connections between the notions of width, as developed and used in classical planning, and the notion of generalized plans, which are summarized in Table 2. The results suggest a deep connection between these two notions and that bounded width for infinite collections of problems \mathcal{Q} is often the result of simple general policies that solve \mathcal{Q} optimally in terms of features that are partially represented in the problem encodings. When this is not the case, we have shed light on the representations that deliver such properties, and hence, polynomial-time searches. We have also formalized and generalized the notion of serialized width by appealing to an explicit and abstract notion of serializations, and established connections between generalized policies and serialized width by borrowing notions from QNPs. Moreover, from this connection, we introduced policy sketches which make use of the language of policy rules but with a different semantics for providing a convenient and powerful language for specifying subproblems and serializations that can be exploited by algorithms such as SIW. The language can be used for encoding domain-specific knowledge by hand, or as a target language for learning domain serializations automatically from traces. These are interesting challenges that we would like to address in the future.

Acknowledgments

The research is partially funded by an ERC Advanced Grant (No 885107), by grant TIN-2015-67959-P from MINECO, Spain, and by the Knut and Alice Wallenberg (KAW) Foundation through the WASP program. H. Geffner is also a Wallenberg Guest Professor at Linköping University, Sweden.

References

- Bandres, W.; Bonet, B.; and Geffner, H. 2018. Planning with Pixels in (Almost) Real Time. In *AAAI*.
- Belle, V.; and Levesque, H. J. 2016. Foundations for Generalized Planning in Unbounded Stochastic Domains. In *Proc. KR*, 380–389.
- Bellemare, M. G.; Naddaf, Y.; Veness, J.; and Bowling, M. 2013. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research* 47: 253–279.
- Bonet, B.; Frances, G.; and Geffner, H. 2019. Learning features and abstract actions for computing generalized plans. In *Proc. AAAI*, 2703–2710.
- Bonet, B.; and Geffner, H. 2018. Features, Projections, and Representation Change for Generalized Planning. In *Proc. IJCAI*, 4667–4673.
- Bonet, B.; and Geffner, H. 2020. Qualitative Numeric Planning: Reductions and Complexity. *Journal of Artificial Intelligence Research (JAIR)* 69: 923–961.
- Bonet, B.; and Geffner, H. 2021. General Policies, Representations, and Planning Width. In *Proc. AAAI*.
- Bonet, B.; Palacios, H.; and Geffner, H. 2009. Automatic Derivation of Memoryless Policies and Finite-State Controllers Using Classical Planners. In *Proc. ICAPS-09*, 34–41.
- Dechter, R. 2013. *Reasoning with probabilistic and deterministic graphical models: Exact algorithms*. Morgan & Claypool Publishers.
- Francès, G.; Ramírez, M.; Lipovetzky, N.; and Geffner, H. 2017. Purely Declarative Action Representations are Overrated: Classical Planning with Simulators. In *Proc. IJCAI*.
- Freuder, E. C. 1982. A sufficient condition for backtrack-free search. *Journal of the ACM* 29(1): 24–32.
- Geffner, H.; and Bonet, B. 2013. *A Concise Introduction to Models and Methods for Automated Planning*. Morgan & Claypool Publishers.
- Ghallab, M.; Nau, D.; and Traverso, P. 2016. *Automated planning and acting*. Cambridge U.P.
- Hoffmann, J.; Porteous, J.; and Sebastia, L. 2004. Ordered Landmarks in Planning. *Journal of Artificial Intelligence Research* 22: 215–278.
- Hu, Y.; and De Giacomo, G. 2011. Generalized planning: Synthesizing plans that work for multiple environments. In *Proc. IJCAI*, 918–923.
- Lehman, J.; and Stanley, K. O. 2011a. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation* 19(2): 189–223.
- Lehman, J.; and Stanley, K. O. 2011b. Evolving a diversity of virtual creatures through novelty search and local competition. In *Proc. of the 13th annual conference on Genetic and evolutionary computation*, 211–218.
- Lipovetzky, N.; and Geffner, H. 2012. Width and serialization of classical planning problems. In *Proc. ECAI*, 540–545.
- Lipovetzky, N.; and Geffner, H. 2017a. Best-First Width Search: Exploration and Exploitation in Classical Planning. In *Proc. AAAI*.
- Lipovetzky, N.; and Geffner, H. 2017b. A polynomial planning algorithm that beats LAMA and FF. *Proc. ICAPS*.
- Lipovetzky, N.; Ramirez, M.; and Geffner, H. 2015. Classical Planning with Simulators: Results on the Atari Video Games. In *Proc. IJCAI*.
- Ostrovski, G.; Bellemare, M. G.; Oord, A.; and Munos, R. 2017. Count-Based Exploration with Neural Density Models. In *Proc. ICML*, 2721–2730.
- Pathak, D.; Agrawal, P.; Efros, A. A.; and Darrell, T. 2017. Curiosity-driven exploration by self-supervised prediction. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 16–17.
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann.
- Richter, S.; and Westphal, M. 2010. The LAMA planner: Guiding cost-based anytime planning with landmarks. *Journal of Artificial Intelligence Research* 39(1): 127–177.
- Segovia, J.; Jiménez, S.; and Jonsson, A. 2016. Generalized planning with procedural domain control knowledge. In *Proc. ICAPS*, 285–293.
- Srivastava, S.; Immerman, N.; and Zilberstein, S. 2008. Learning generalized plans using abstract counting. In *Proc. AAAI*, 991–997.
- Srivastava, S.; Zilberstein, S.; Immerman, N.; and Geffner, H. 2011. Qualitative Numeric Planning. In *AAAI*.
- Tang, H.; Houthoofd, R.; Foote, D.; Stooke, A.; Chen, O. X.; Duan, Y.; Schulman, J.; DeTurck, F.; and Abbeel, P. 2017. # exploration: A study of count-based exploration for deep reinforcement learning. In *Advances in neural information processing systems*, 2753–2762.
- Thiébaux, S.; Hoffmann, J.; and Nebel, B. 2005. In defense of PDDL axioms. *Artif. Intell.* 168(1-2): 38–69.