When Local Plans Cannot Be Global: A Sheaf-Theoretic View of Decomposition in Numeric Planning

Jiajia Song, Guang Hu

The University of Melbourne, Australia
jiajia.song1, guang.hu@unimelb.edu.au

Abstract

Numeric planning poses great computational challenges. One way to analyse it is to ask when reasoning can be done locally, i.e., decompose the problem into smaller subproblems, solve them locally, and assemble the local solutions into a global one. However, interactions among numeric variables may prevent valid and synchronised local plans from being lifted to a globally valid solution. We introduce a sheaf-theoretic framework to study the decompositional obstructions and formulate global solvability as a local-to-global consistency condition. Our main result separates global solvability from local compatibility, and proves that overlap consistent covers with no cross-cutting interactions are exact. We further use tree decompositions of the Gaifman graph to induce covers with bag size bounded by treewidth plus one, in which every numeric interaction is represented locally. Finally, we propose Topological CEGAR, a refinement based procedure, to coarsen the numeric cover iteratively when obstructions are detected. This work provides a structural characterisation of sound and exact decomposition in numeric planning.

Introduction

Numeric planning (Fox and Long 2003) extends classical planning with numeric variables, constraints and effects to model resources, time and physical quantities. The expressiveness allows it to be applied in many real-world scenarios (Fox, Long, and Magazzeni 2011; Löhr et al. 2012; Piacentini et al. 2013; Vallati et al. 2016; Kiam and Schulte 2017), but also results in great computational challenges. Numeric planning is undecidable in general (Helmert 2002) and even decidable fragments exhibit high complexity.

Many works address the challenges via heuristic search, relaxation, numeric pattern databases, SAT and SMT encodings, and integer/linear programming (Hoffmann 2003; Aldinger, Mattmüller, and Göbelbecker 2015; Scala et al. 2016a, 2020; Gnad et al. 2025; Shin and Davis 2005; Hoffmann et al. 2007; Scala et al. 2016b; Bofill, Espasa, and Villaret 2016; Leofante et al. 2020; Coles et al. 2008; Scala et al. 2017; Piacentini et al. 2018), and have achieved good performance in many numeric domains.

A natural structural question is whether a numeric planning problem can be decomposed into subproblems whose local plans can be assembled into a valid global plan. Intuitively, local reasoning is attractive since many numeric constraints and effects only involve a small subset of variables. However, decomposition may not be sound. If strongly coupled variables are separated into different subproblems, a constraint or effect involving them may not be enforced by any of the subproblems, and thus locally valid plans may fail to assemble into a globally valid solution due to this structural *obstruction* induced by the decomposition.

In this paper, we study the structural question: when can we perform local search over subproblems and then combine the local plans into a global solution? To the best of our knowledge, this question has received less attention in the literature of numeric planning.

We study this problem using sheaf theory (Lane and Moerdijk 1992). The numeric variable set is decomposed into subsets, with each subset inducing a subproblem. Reachable states in these subproblems are *local sections*, and *compatible* local sections are the ones that agree on overlapping variables. We formulate the local-to-global question as a *sheaf condition*, i.e., compatible local sections can be glued into a globally reachable state, and synchronised compatible local goal states induce a valid global plan.

In our framework, we show that *cross-cutting* interactions can result in structural obstructions to local-to-global consistency. Our main theoretical result is the exact decomposition theorem for numeric covers. We show that if an overlap consistent cover has no cross-cutting interactions, it is *exact*, i.e., the numeric planning problem is solvable if and only if there exists a *synchronised compatible family* of local plans. An overlap consistent cover prevents conflicts on variable updates by subproblems.

Note that this structural characterisation does not reduce numeric planning to an easier problem since it is still computationally hard to find a synchronised compatible family. However, the sheaf-theoretic characterisation provides a structural meaning of solvability and allows us to identify the objects that need to be computed or relaxed for any exact decomposition in numeric planning.

Furthermore, we introduce *Topological CEGAR*, a refinement procedure to repair the local-to-global inconsistency. A cross-cutting interaction becomes a witness of obstruction when a synchronised compatible family of local plans fails to form a valid global plan, and the refinement step merges the subproblems involved in this obstruction. We use Topological CEGAR to demonstrate that our framework can be converted into an operational procedure that is guaranteed to terminate in finitely many steps. It is not intended as a practical algorithm, and we leave the design of efficient and practical refinement algorithms to future work.

The main contributions of this paper are as follows:

- We introduce a sheaf-theoretic formulation of decomposition in numeric planning.
- We define the *reachable constraint presheaf* and the *goal presheaf* whose sections represent locally reachable states and local goal states, respectively.
- We formalise local-to-global obstructions as failures of the sheaf condition, and identify cross-cutting interactions as structural sources of obstructions.
- We present the notions of *sound* and *exact* numeric covers, and prove that overlap consistent covers with no cross-cutting interactions are exact.
- We connect the framework with tree decomposition, and show how the structure of a Gaifman graph can provide covers with controlled local numeric interactions.
- We propose Topological CEGAR, a finite refinement procedure over numeric covers, and use detected obstructions to guide the improvement of covers.

The goal of this paper is to provide a rigorous sheaf-theoretic formulation of decomposition in numeric planning. The focus is on the structure of global solvability, rather than a practical numeric planner.

Numeric Planning

In this section, we provide necessary background for numeric planning. We use a hybrid state representation that separates the assignments for numeric variables and propositional variables. The notation for numeric variables, constraints and effects is based on Hoffmann (2003); Fox and Long (2003). For propositional variables, we allow finite domains as in SAS$^+$ (Bäckström and Nebel 1995).

We denote the domain of a variable v by $\mathrm{dom}(v)$. Let $\mathcal{V}_P$ be a finite set of propositional variables, each with a finite domain. Let $\mathcal{V}_N$ be a finite set of numeric variables. Each numeric variable v has $\mathrm{dom}(v) \subseteq \mathbb{R}$. The total variable set $\mathcal{V}$ is the union of $\mathcal{V}_P$ and $\mathcal{V}_N$, and we assume that $\mathcal{V}_P$ and $\mathcal{V}_N$ are disjoint.

Definition 1 (Hybrid state space)**.** A *global (hybrid) state* s is a tuple $\langle p, x\rangle$ with the propositional component $p \in \mathcal{S}_{\mathcal{V}_P} \triangleq \bigtimes_{v\in\mathcal{V}_P} \mathrm{dom}(v)$, and the numeric component $x \in \mathcal{S}_{\mathcal{V}_N} \triangleq \bigtimes_{v\in\mathcal{V}_N} \mathrm{dom}(v)$. The *global state space* is $\mathcal{S} \triangleq \mathcal{S}_{\mathcal{V}_P} \times \mathcal{S}_{\mathcal{V}_N}$. We denote the value of $v \in \mathcal{V}$ in s by $s[v]$, where $s[v] = p[v]$ if $v \in \mathcal{V}_P$ and $s[v] = x[v]$ if $v \in \mathcal{V}_N$.

Given a numeric variable subset $U \subseteq \mathcal{V}_N$, the *restriction* of the hybrid state $s = \langle p, x\rangle$ to U is $s|_U \triangleq \langle p, x|_U\rangle$, where $x|_U[v] = x[v]$ for all $v \in U$. We denote the set of restricted states over U by $\mathcal{S}_U$.

Definition 2 (Constraints and effects)**.** A *numeric constraint* over $U \subseteq \mathcal{V}_N$ is a tuple $t = (\mathrm{exp}, \mathrm{comp}, \mathrm{exp}')$ where $\mathrm{exp}, \mathrm{exp}'$ are arithmetic expressions over $U \cup \mathbb{Q}$ using $\{+, -, \times, /\}$ and $\mathrm{comp} \in \{<, \leq, =, \geq, >\}$. A hybrid state $s = \langle p, x\rangle$ *satisfies* t, denoted by $s \models t$, if the values of $\mathrm{exp}, \mathrm{exp}'$ evaluated in x satisfy comp. The *support* of t is $\mathrm{supp}(t) \triangleq \mathit{vars}(\mathrm{exp}) \cup \mathit{vars}(\mathrm{exp}')$.

A *propositional constraint* t is of the form $v = c$, where $v \in \mathcal{V}_P$, $c \in \mathrm{dom}(v)$. A state $s = \langle p, x\rangle$ satisfies t, denoted by $s \models t$, if $p[v] = c$. The *support* of t is $\mathrm{supp}(t) \triangleq \{v\}$.

A *numeric effect* is a tuple $e = (v, \mathrm{assi}, \mathrm{exp})$ where $v \in \mathcal{V}_N$, $\mathrm{assi} \in \{:=, +=, -=, \times=, /=\}$ and exp is an arithmetic expression over $\mathcal{V}_N \cup \mathbb{Q}$. The *support* of e is $\mathrm{supp}(e) \triangleq \mathit{vars}(\mathrm{exp}) \cup \{v\}$.

A *propositional effect* e is in the form of (v, c), where $v \in \mathcal{V}_P$ and $c \in \mathrm{dom}(v)$. The *support* of e is $\mathrm{supp}(e) \triangleq \{v\}$.

In this framework, we focus on numeric constraints and effects. We can also view a numeric effect as a numeric transition constraint that restricts a state pair (s, s'). For example, $v \mathrel{+}= 5$ requires that $s'[v] = s[v] + 5$. Numeric constraints and effects together impose geometric constraints on the reachability space of the planning problem.

Definition 3 (Numeric planning problem)**.** A *numeric planning problem* is a tuple $\Pi = \langle \mathcal{V}_P, \mathcal{V}_N, \mathcal{A}, s_0, \mathcal{G}\rangle$, where $s_0 = \langle p_0, x_0\rangle \in \mathcal{S}$ is the initial state, $\mathcal{A}$ is a finite set of actions and $\mathcal{G} = \mathcal{G}_P \cup \mathcal{G}_N$ is the goal specification, i.e., the union of a finite set of propositional constraints $\mathcal{G}_P$ and a finite set of numeric constraints $\mathcal{G}_N$.

An action a is $\langle \mathit{pre}_P(a), \mathit{pre}_N(a), \mathit{post}_P(a), \mathit{post}_N(a)\rangle$, where $\mathit{pre}_P(a), \mathit{pre}_N(a), \mathit{post}_P(a)$ and $\mathit{post}_N(a)$ are finite sets of propositional constraints, numeric constraints, propositional effects and numeric effects, respectively.

We denote $\mathit{pre}(a) \triangleq \mathit{pre}_P(a) \cup \mathit{pre}_N(a)$ and $\mathit{post}(a) \triangleq \mathit{post}_P(a) \cup \mathit{post}_N(a)$. The *support* of a is $\mathrm{supp}(a) \triangleq (\bigcup_{t\in \mathit{pre}(a)} \mathrm{supp}(t)) \cup (\bigcup_{e\in \mathit{post}(a)} \mathrm{supp}(e))$.

An action a is applicable in a hybrid state $s = \langle p, x\rangle$ if $s \models t$ for all $t \in \mathit{pre}(a)$. The successor state $s' = f(s, a) = \langle p', x'\rangle$ satisfies the following:

- For all $(v, c) \in \mathit{post}_P(a)$, $p'[v] = c$.
- For all $(v, \mathrm{assi}, \mathrm{exp}) \in \mathit{post}_N(a)$, if assi is $:=$, then $x'[v] = \mathrm{exp}(x)$, otherwise, $x'[v] = x[v]\, \mathrm{assi}\, \mathrm{exp}(x)$.
- For all remaining variables v, $s'[v] = s[v]$.

Given a state s and an action sequence $\pi = (a_1, \ldots, a_n)$, a *state trajectory* τ from s is the sequence of states $(s^1, \ldots, s^{n+1})$ where $s^1 = s$, a_i is applicable in s^i and $s^{i+1} = f(s^i, a_i)$ for all $i = 1, \ldots, n$. If for some i, the action a_i is not applicable in s^i, we say that π is not valid.

A state s' is *reachable* from s if there exists a valid action sequence π and the last state in the trajectory is s'. We denote the set of reachable states in $\mathcal{S}$ from s_0 by $\mathsf{Reach}(\Pi)$.

The *goal region* is $G \triangleq \{s \in \mathcal{S} \mid \forall t \in \mathcal{G}, s \models t\}$. An *s-plan* is a valid action sequence from s that reaches some $s' \in G$. An s_0-plan is a *plan* for the planning problem Π and if no such plan exists, we say that Π is *unsolvable*.

Sheaf-Theoretic Framework

In this section, we establish a sheaf-theoretic framework for decomposition in numeric planning. We first define the numeric interaction graph and the local interaction system induced by a numeric cover. Next, we construct the reachable constraint presheaf and the goal presheaf, and introduce the notions of soundness and exactness for numeric covers. Finally, we formulate the local-to-global sheaf condition underlying sound and exact decomposition, and show that cross-cutting interactions are structural sources of decompositional obstructions.

We work over the poset category $(2^{\mathcal{V}_N}, \subseteq)$ (Lane and Moerdijk 1992), whose objects are subsets of numeric variables, with morphisms being set inclusions. A *presheaf* on this category assigns local data to each subset $U \subseteq \mathcal{V}_N$ and assigns restriction maps based on set inclusions. The *sheaf condition* checks whether compatible local data on these subsets can be glued into global data on $\mathcal{V}_N$.

We can distinguish two types of local-to-global inconsistency based on this structure. A *topological obstruction* occurs when the *reachable constraint presheaf* fails the sheaf condition, i.e., locally reachable states in subproblems cannot be glued into a globally reachable state in the original problem. A *goal obstruction* arises when synchronised compatible local goal states cannot induce a valid global plan.

Let $\Pi = \langle \mathcal{V}_P, \mathcal{V}_N, \mathcal{A}, s_0, \mathcal{G} = \mathcal{G}_P \cup \mathcal{G}_N \rangle$ be a numeric planning problem. We denote the set of all numeric constraints and effects in Π as $\mathcal{E} \triangleq \mathcal{E}_{\textit{state}} \cup \mathcal{E}_{\textit{effect}}$, where $\mathcal{E}_{\textit{state}} \triangleq \mathcal{G}_N \cup \{t \in \textit{pre}_N(a) \mid a \in \mathcal{A}\}$ is the set of state-evaluable constraints that are goals and preconditions, and $\mathcal{E}_{\textit{effect}} \triangleq \{e \in \textit{post}_N(a) \mid a \in \mathcal{A}\}$ is the set of numeric effects (transition constraints).

Due to space constraints, most of the proofs are provided in Appendix.

Numeric Interactions and Constraint Presheaf

Definition 4 (Gaifman graph of numeric interactions)**.** The *Gaifman graph* (Gaifman 1982) of Π is $H_\Pi = (\mathcal{V}_N, E_\Pi)$ where $(u, v) \in E_\Pi$ if and only if $u \neq v$ and there exists a numeric precondition, numeric goal constraint or numeric effect e such that $u, v \in \mathrm{supp}(e)$.

Definition 5 (Treewidth)**.** (Robertson and Seymour 1986) A *tree decomposition* (T, χ) of the graph $H_\Pi = (\mathcal{V}_N, E_\Pi)$ is a tree T of vertices $V(T)$ and a collection of bags $\chi = (B_t \subseteq \mathcal{V}_N \mid t \in V(T))$, such that

- $\bigcup_{t \in V(T)} (B_t) = \mathcal{V}_N$.
- For each edge $(u, v) \in E_\Pi$, there exists at least one bag B_t such that $u, v \in B_t$.
- If bags B_t and $B_{t'}$ both contain a vertex v, then all bags that are on the path of T between B_t and $B_{t'}$ also contain the vertex v.

The *width* of the decomposition (T, χ) is $\max_{t \in V(T)}(|B_t| - 1)$. The *treewidth* of H_Π, denoted by $\mathrm{tw}(H_\Pi)$, is the minimum width among all tree decompositions.

Similar to the Clique Containment Lemma (Bodlaender and Möhring 1990), we prove that at least one bag of a valid tree decomposition contains each numeric interaction e.

Proposition 1 (Constraint containment)**.** Let (T, χ) be a valid tree decomposition of H_Π. For every $e \in \mathcal{E}$, there exists at least one bag B_{t^*} such that $\mathrm{supp}(e) \subseteq B_{t^*}$.

We will show that treewidth yields covers with bag size bounded by treewidth plus one, in which every numeric interaction is represented locally.

Definition 6 (Feasible regions)**.** Given $U \subseteq \mathcal{V}_N$, the *local interaction system* $\Sigma_U \triangleq \{e \in \mathcal{E} \mid \mathrm{supp}(e) \subseteq U\}$ is the set of numeric constraints and numeric effects whose supports are subsets of U. We denote the subset of state-evaluable interactions as $\Sigma_U^{\textit{state}} \triangleq \Sigma_U \cap \mathcal{E}_{\textit{state}}$.

The *feasible region* over U is $\mathcal{F}(U) \triangleq \{\langle p, x\rangle \in \mathcal{S}_{\mathcal{V}_P} \times \mathcal{S}_U \mid \forall e \in \Sigma_U^{\textit{state}}, \langle p, x\rangle \models e\}$.

The feasible region ensures the static consistency with state-evaluable numeric interactions and we will define the reachable constraint presheaf later to capture the transition constraints from numeric effects.

It is worth noting that our sheaf-theoretic framework is different from the variable projection abstraction (Gnad et al. 2025). In variable projection, a subproblem over U includes all constraints whose supports intersect U. The local feasible region is obtained by existentially quantifying the variables not in U. When U expands, the projection region may shrink as fewer variables can be assigned freely.

However, in our framework, the subproblem over U only considers the constraints whose supports are completely contained in U. When U expands, the local feasible region may shrink as the previously not included constraints are taken into consideration.

Definition 7 (Constraint presheaf)**.** The *constraint presheaf* $\mathcal{F}$ maps each subset $U \subseteq \mathcal{V}_N$ to its feasible region $\mathcal{F}(U)$. For every $V \subseteq U \subseteq \mathcal{V}_N$, the restriction map $\rho_{U,V} : \mathcal{F}(U) \to \mathcal{F}(V)$ is the natural state restriction $\rho_{U,V}(\langle p, x\rangle) \triangleq \langle p, x|_V\rangle$.

A presheaf assigns a space of feasible numeric values to every subset of numeric variables. The restriction map ensures consistency between these local assignments in all overlapping subsets.

Lemma 1. The map $\mathcal{F}$ is a well defined presheaf, i.e., a contravariant functor from the category poset $(2^{\mathcal{V}_N}, \subseteq)$ to the category Set (Lane and Moerdijk 1992).

Reachable Constraint Presheaf

Definition 8 (Subproblem)**.** Given $U \subseteq \mathcal{V}_N$, the *subproblem* over U is a tuple $\Pi_U = \langle \mathcal{V}_P, U, \mathcal{A}_U, s_0|_U, \mathcal{G}_U\rangle$, where:

- $s_0|_U = \langle p_0, x_0|_U\rangle \in \mathcal{S}_{\mathcal{V}_P} \times \mathcal{S}_U$ is the local initial state.
- $\mathcal{G}_U = \mathcal{G}_P \cup (\mathcal{G}_N \cap \Sigma_U^{\textit{state}})$ is the local goal specification.
- $\mathcal{A}_U = \{a|_U \mid a \in \mathcal{A}\}$ is the set of local actions and for each $a \in \mathcal{A}$, $a|_U \triangleq \langle \textit{pre}_P(a), \textit{pre}_N(a) \cap \Sigma_U^{\textit{state}}, \textit{post}_P(a), \textit{post}_N(a) \cap \Sigma_U\rangle$.

A local hybrid state of a subproblem over U is $s = \langle p, x\rangle \in \mathcal{S}_{\mathcal{V}_P} \times \mathcal{S}_U$. It is *reachable* in Π_U if there exists a finite sequence of local actions $\pi|_U = (a_1|_U, \ldots, a_n|_U)$ and a state trajectory $(s^1, \ldots, s^{n+1})$ satisfying the following:

- $s_0|_U = s^1$ and $s = s^{n+1}$.
- For all $i \in \{1, \ldots, n\}$, $s^i \models \textit{pre}(a_i|_U)$.
- For all $i \in \{1, \ldots, n\}$, $s^{i+1} = f_U(s^i, a_i|_U)$ where the local transition function f_U applies the effects in $a_i|_U$ on s^i as defined in the previous section.

The set of all reachable local hybrid states in Π_U is denoted by $\mathsf{Reach}(\Pi_U)$. In sheaf theory (B.R.Tennison 1975), an element assigned to a local domain is called a *section*. We adopt the standard notation and call a local hybrid state over the variable in U a *section*. The set of reachable sections over U is denoted as $\Gamma(U, \mathcal{F}^{\mathsf{Reach}})$.

Definition 9 (Numeric cover)**.** A numeric *cover* of $\mathcal{V}_N$ is a finite collection $\mathcal{U} = \{U_1, \ldots, U_n\}$ such that $U_i \subseteq \mathcal{V}_N$ and $\bigcup_{i=1}^{n} U_i = \mathcal{V}_N$. We denote $U_{ij} \triangleq U_i \cap U_j$.

Definition 10 (Reachable constraint presheaf)**.** Given a numeric cover $\mathcal{U} = \{U_1, \ldots, U_n\}$, the *reachable constraint presheaf* $\mathcal{F}^{\mathsf{Reach}}$ assigns to each U_i the set of locally reachable hybrid states in the subproblem, i.e., $\Gamma(U_i, \mathcal{F}^{\mathsf{Reach}}) \triangleq \mathsf{Reach}(\Pi_{U_i}) = \{s \in \mathcal{S}_{\mathcal{V}_P} \times \mathcal{S}_{U_i} \mid s \text{ is reachable from } s_0|_{U_i} \text{ in } \Pi_{U_i}\}$.

Elements of $\Gamma(U_i, \mathcal{F}^{\mathsf{Reach}})$ are called *reachable local sections* over U_i.

For every U_{ij}, the restriction map preserves the propositional component and restricts the numeric component, i.e., $\langle p, x\rangle \mapsto \langle p, x|_{U_{ij}}\rangle$.

From the definition, we know that $\mathcal{F}^{\mathsf{Reach}}$ is a presheaf relative to the given numeric cover. It assigns reachable local sections to cover elements and its restriction maps are used to compare the sections on overlaps.

Sheaf Condition

Definition 11 (Cross-cutting interaction)**.** Given a numeric cover $\mathcal{U} = \{U_1, \ldots, U_n\}$, a constraint or effect $e \in \mathcal{E}$ is a *cross-cutting interaction* with respect to $\mathcal{U}$ if $\mathrm{supp}(e) \not\subseteq U_i$ for all $i \in \{1, \ldots, n\}$.

A cross-cutting interaction is not fully represented in any local constraint system Σ_{U_i}. Although it may be enforced in the global problem, no local subproblem induced by the numeric cover enforces it.

Definition 12 (Compatible family)**.** Given a numeric cover $\mathcal{U} = \{U_1, \ldots, U_n\}$, a family $(s_i \in \Gamma(U_i, \mathcal{F}^{\mathsf{Reach}}))_{i=1}^{n}$ is *compatible* if for all $i, j \in \{1, \ldots, n\}$, $s_i|_{U_{ij}} = s_j|_{U_{ij}}$. We denote the set of all compatible families over $\mathcal{U}$ as $\Gamma^0(\mathcal{U}, \mathcal{F}^{\mathsf{Reach}})$.

Proposition 2 (Global witness)**.** For any compatible family $(s_i)_{i=1}^{n} \in \Gamma^0(\mathcal{U}, \mathcal{F}^{\mathsf{Reach}})$, there exists a uniquely determined global hybrid state $s^* \in \mathcal{S}$, called the *global witness* of $(s_i)_{i=1}^{n}$ such that $s^*|_{U_i} = s_i$ for all i.

Intuitively, if we glue the local numeric assignments, we may obtain a global witness with the same propositional component as the local states. For example, let $U_1 = \{x, y\}$ and $U_2 = \{y, z\}$ with $s_1 = (x = 1, y = 2, p = \mathsf{true})$ and $s_2 = (y = 2, z = 3, p = \mathsf{true})$. Since s_1 and s_2 agree on the overlap $\{y\}$ and have the same propositional component $p = \mathsf{true}$, they determine the global witness $s = (x = 1, y = 2, z = 3, p = \mathsf{true})$. However, this global witness may not be reachable in the original planning problem. Whether every such compatible family can be glued into a globally reachable state is exactly what the sheaf condition tests.

Definition 13 (Topological obstruction)**.** Given a numeric cover $\mathcal{U} = \{U_i \mid i = 1, \ldots, n\}$, the *gluing map* $\gamma_{\mathcal{U}} \colon \Gamma(\mathcal{V}_N, \mathcal{F}^{\mathsf{Reach}}) \to \Gamma^0(\mathcal{U}, \mathcal{F}^{\mathsf{Reach}})$ assigns each global section s to its uniquely induced compatible family $(s|_{U_i})_{i=1}^{n}$. The reachable constraint presheaf $\mathcal{F}^{\mathsf{Reach}}$ satisfies the *sheaf condition* over the cover $\mathcal{U}$ if $\gamma_{\mathcal{U}}$ is surjective, i.e., every compatible family lifts to a global section.

If the sheaf condition fails, the cover introduces a *topological obstruction*.

In sheaf theory, the sheaf condition requires the gluing map to be a bijection from the global sections to matching families, which allows every compatible family to lift to a global section. The gluing map $\gamma_{\mathcal{U}}$ is always injective due to Proposition 2. Therefore, the sheaf condition reduces to the surjectivity of $\gamma_{\mathcal{U}}$.

Proposition 3. Given a numeric cover $\mathcal{U} = \{U_1, \ldots, U_n\}$, the reachable constraint presheaf over $\mathcal{U}$ fails the sheaf condition if and only if there exists a compatible family $(s_i)_i \in \Gamma^0(\mathcal{U}, \mathcal{F}^{\mathsf{Reach}})$ whose unique global witness s^* satisfies $s^* \notin \Gamma(\mathcal{V}_N, \mathcal{F}^{\mathsf{Reach}})$.

Goal Presheaf

In order to analyse local-to-global consistency of the plans, we consider the *goal presheaf*. Similar to $\mathcal{F}^{\mathsf{Reach}}$, the goal presheaf is relative to the given numeric cover.

Definition 14 (Goal presheaf)**.** Given a numeric cover $\mathcal{U} = \{U_1, \ldots, U_n\}$, the *goal presheaf* $\mathcal{F}^{\mathcal{G}}$ maps each numeric subset U_i to its local goal-satisfying sections, i.e., $\Gamma(U_i, \mathcal{F}^{\mathcal{G}}) \triangleq \{s \in \Gamma(U_i, \mathcal{F}^{\mathsf{Reach}}) \mid s \models \mathcal{G}_{U_i}\}$.

Given a numeric cover $\mathcal{U} = \{U_i \mid i = 1, \ldots, n\}$, a subproblem Π_{U_i} is *locally solvable* if and only if it can reach at least one valid state in the local goal region from $s_0|_{U_i}$, i.e., $\Gamma(U_i, \mathcal{F}^{\mathcal{G}}) \neq \emptyset$. However, static consistency in the overlap is insufficient to ensure the global solvability as the plans for the subproblems must be synchronised.

Definition 15 (Synchronised compatible family)**.** Given a numeric cover $\mathcal{U} = \{U_i \mid i = 1, \ldots, n\}$, a compatible family of local goal states $(s_i \in \Gamma(U_i, \mathcal{F}^{\mathcal{G}}))_{i=1}^{n}$ is *synchronised* if there exists an action sequence $\pi = (a_1, \ldots, a_m)$ such that for all i, the local restriction $\pi|_{U_i}$ is a valid plan to reach s_i from $s_0|_{U_i}$ in Π_{U_i}. Note that π needs not satisfy the global preconditions or effects of Π.

We call such an action sequence a *shared synchronised local plan* for $\mathcal{U}$. We denote the set of all synchronised compatible families over $\mathcal{U}$ as $\Gamma^0_{sync}(\mathcal{U}, \mathcal{F}^{\mathcal{G}})$.

Definition 16 (Goal obstruction)**.** Given a numeric cover $\mathcal{U} = \{U_1, \ldots, U_n\}$, the goal presheaf $\mathcal{F}^{\mathcal{G}}$ satisfies the *sheaf condition* over the cover $\mathcal{U}$ if the gluing map $\gamma^{\mathcal{G}}_{\mathcal{U}} \colon \Gamma(\mathcal{V}_N, \mathcal{F}^{\mathcal{G}}) \to \Gamma^0_{sync}(\mathcal{U}, \mathcal{F}^{\mathcal{G}})$ is surjective.

If the sheaf condition fails, then a synchronised compatible family in $\Gamma^0_{sync}(\mathcal{U}, \mathcal{F}^{\mathcal{G}})$ that has no preimage in $\Gamma(\mathcal{V}_N, \mathcal{F}^{\mathcal{G}})$ introduces a *goal obstruction*.

The subproblem Π_U contains all propositional variables and all associated propositional preconditions and effects. The shared propositional component in every local section ensures the global consistency in these variables globally without a separate decomposition structure. We leave the simultaneous decomposition of propositional and numeric variables for future work.

Definition 17 (Sound and exact covers)**.** A numeric cover $\mathcal{U}$ of $\mathcal{V}_N$ is *sound* for Π if every synchronised compatible family $(s_i)_i \in \Gamma^0_{sync}(\mathcal{U}, \mathcal{F}^{\mathcal{G}})$ induces a valid global plan.

The cover $\mathcal{U}$ is *exact* if Π is solvable if and only if $\Gamma^0_{sync}(\mathcal{U}, \mathcal{F}^{\mathcal{G}}) \neq \emptyset$.

Soundness ensures that synchronised local plans are not spurious global witnesses. Exactness means that solving the problem globally is equivalent to finding a synchronised compatible family over the cover.

Definition 18 (Overlap consistent cover)**.** A numeric cover $\mathcal{U}$ is *overlap consistent* if for every pair $U_i, U_j \in \mathcal{U}$, every variable $v \in U_i \cap U_j$, and every numeric effect e that updates v, we have $\mathrm{supp}(e) \subseteq U_i$ and $\mathrm{supp}(e) \subseteq U_j$.

An overlap consistent cover ensures that local subproblems do not conflict in variable updates.

Theorem 1 is one of the main results in this framework. When every numeric interaction is represented locally and local subproblems agree on the updates to shared variables, synchronised compatibility is sufficient to solve the planning problem globally.

Theorem 1. If a numeric cover $\mathcal{U}$ is overlap consistent and has no cross-cutting interactions, then $\mathcal{U}$ is exact for Π, i.e., Π is solvable if and only if $\Gamma^0_{sync}(\mathcal{U}, \mathcal{F}^{\mathcal{G}}) \neq \emptyset$.

Proof sketch. If Π is solvable, the restriction of any valid global plan π over U_i is a valid local plan $\pi|_{U_i}$ for Π_{U_i}. The goal state s^* reached by π is restricted to $s^*|_{U_i}$, which is a local goal state for Π_{U_i} since every local goal constraint is contained in U_i. The restrictions agree on overlaps by definition, thus $(s^*|_{U_i})_i$ is a synchronised compatible family.

Let $(s_i)_i \in \Gamma^0_{sync}(\mathcal{U}, \mathcal{F}^{\mathcal{G}})$ with shared action sequence $\pi = (a_1, \dots, a_m)$. At each step ℓ, every numeric precondition of a_ℓ is contained in some subset U_i in the cover, and is satisfied by local trajectory over U_i. Since the glued global state agrees with the local state on U_i, the precondition is satisfied globally. Propositional preconditions are also satisfied globally since they are shared by all subproblems. Moreover, every numeric effect is contained in some cover element, and overlap consistency ensures that there is no conflict on the updates of shared variables. Therefore, the local successors glue into a valid unique global successor. By induction, π induces a valid global trajectory. Finally, since every numeric goal constraint is contained in some U_i of the cover and propositional goals are shared, the glued final state satisfies the global goal, thus π is valid globally. $\square$

Theorem 2. Let H_Π be the Gaifman graph of the numeric interactions of Π. If $\mathrm{tw}(H_\Pi) = k$, then there exists a numeric cover $\mathcal{U}$ such that every $U \in \mathcal{U}$ has size no more than $k+1$ and every numeric interaction is represented locally.

If H_Π has treewidth k, we use the bags from the tree decomposition as the cover. By Proposition 1, every numeric interaction is contained in some bag, therefore, no cross-cutting interactions exist with respect to the cover. If this cover is also overlap consistent, then we can apply Theorem 1 to show that the cover is exact.

Example

We illustrate cross-cutting interactions with a simple numeric planning task.

Assume that a rover has two numeric variables e (the remaining energy) and d (the stored data). The action $sample$ increases stored data, i.e., $d \mathrel{+}= \delta$ for some $\delta > 0$, and the action $transmit$ sends all stored data to a base station with precondition $e \geq d$ and effect $d := 0$. We can $transmit$ data only if the remaining energy level is no less than the data.

Let the numeric cover $\mathcal{U} = \{U_1 = \{e\}, U_2 = \{d\}\}$. The precondition $e \geq d$ is a cross-cutting interaction since $\mathrm{supp}(e \geq d) = \{e, d\}$ is not contained in either U_1 or U_2.

The subproblem Π_{U_1} has no numeric precondition for $transmit$ and thus it is not restricted by any local numeric precondition. Similarly, the subproblem Π_{U_2} does not enforce the numeric precondition $e \geq d$, but still has the local effect $d := 0$. Therefore, the action sequence $(sample, transmit)$ is locally valid for both subproblems and is a shared synchronised local plan.

However, this sequence may not be globally valid. After $sample$ increases d, the global state may satisfy $e < d$, thus $transmit$ cannot be applied in the original problem. The violated constraint $e \geq d$ witnesses an obstruction, since it is a cross-cutting interaction invisible to both local subproblems, and thus prevents the local plans from being glued globally.

Not all cross-cutting interaction can cause obstruction. For example, if the subproblem over U_1 enforces the invariant $e \geq 10$, and the subproblem over U_2 enforces $d \leq 5$, then every compatible family over $\mathcal{U}$ can be glued into a global state such that $e \geq 10$ and $d \leq 5$, therefore, every such global witness also satisfies $e \geq d$. So, the cross-cutting interaction $e \geq d$ does not witness an obstruction even though it is not contained in any single subproblem.

This example motivates us to consider the refinement strategy in the next section. We do not refine a numeric cover simply because there is a cross-cutting interaction. Refinement is triggered only when a cross-cutting interaction actually witnesses an obstruction.

Topological CEGAR

In the previous section, we showed that an overlap consistent numeric cover with no cross-cutting interaction is exact, with any synchronised compatible family inducing a valid global plan via the sheaf condition. However, an arbitrary numeric cover may result in an invalid global plan from a synchronised local plan as some interaction is invisible to any single subproblem.

In this section, inspired by the Counterexample-Guided Abstraction Refinement (CEGAR) technique (Clarke et al. 2003; Seipp and Helmert 2018), we present *Topological CEGAR*, a refinement process over the space of numeric covers. When a synchronised compatible family does not induce a valid global plan, a cross-cutting interaction witnesses this obstruction. We merge the subproblems related to this witness to make the interaction local.

Given a planning problem, we may naively construct a cover by letting $U_e = \mathrm{supp}(e)$ for every $e \in \mathcal{E}$ and it can guarantee that there is no cross-cutting interaction. However, this construction has two limitations.

First, the subproblems induced by this cover may overlap cyclically as several constraints may share variables in closed loops. Obtaining a synchronised compatible family over the cyclic cover is computationally difficult. It is hard to ensure simultaneous agreement on the overlaps and fixing an assignment on one overlap may conflict with previous assignments. This results in excessive backtracking similar to solving a cyclic CSP (Dechter and Pearl 1989).

Second, not every cross-cutting interaction results in obstructions in a planning problem. As shown in the previous example, local solutions may already satisfy the cross-cutting constraints implicitly. Therefore, the naive construction may cause unnecessary overhead and make the subproblems become too large to solve or just be the global one.

Topological CEGAR addresses these issues by initialising an acyclic cover and refining it only when an obstruction is detected. The initial cover acts as a structural approximation of the numeric interactions.

For example, we may use a tree decomposition of the Gaifman graph (Definition 4). By Theorem 2, if $\mathrm{tw}(H_\Pi) \leq k$ and $\mathcal{U}_0$ is the set of bags each with size no greater than $k+1$, then every numeric interaction is represented locally. However, if $\mathrm{tw}(H_\Pi) > k$, we may construct a relaxed Gaifman graph that ignores some interactions and use the bags from the relaxed tree decomposition as the initial cover.

Algorithm 1 explains the high level refinement step. It searches for a synchronised compatible family from the current cover and if the synchronised local plan is not valid globally, it identifies the obstruction witness and coarsens the cover by merging subsets to make the witness local. It is worth noting that the algorithm does not include the method to compute synchronised compatible family and we leave it for future work.

Algorithm 1: Topological CEGAR

Require: Numeric planning problem Π, initial numeric cover $\mathcal{U}_0$
$\mathcal{U} \leftarrow \mathcal{U}_0$
loop
    Search for a synchronised compatible family $(s_i)_i \in \Gamma^0_{\mathit{sync}}(\mathcal{U}, \mathcal{F}^{\mathcal{G}})$
    if no such family exists **then**
        return not found
    end if
    Let π be a shared synchronised local plan for $(s_i)_i$
    Construct the induced global trajectory of π
    if π is a valid plan for Π **then**
        return π
    end if
    Let e witness the obstruction of the induced trajectory
    $\mathcal{U}_{\mathit{obs}}(e) \leftarrow \{U \in \mathcal{U} \mid U \cap \mathrm{supp}(e) \neq \emptyset\}$
    $\mathcal{U} \leftarrow (\mathcal{U} \setminus \mathcal{U}_{\mathit{obs}}(e)) \cup \{\bigcup_{U \in \mathcal{U}_{\mathit{obs}}(e)} U\}$
end loop

Cover Coarsening

The repair in Topological CEGAR is done only on the cover via subproblem merging as the obstruction is caused by the decomposition structure. To formalise this, we equip the space of numeric covers with a coarsening preorder.

Definition 19 (Cover coarsening)**.** Let $\mathcal{U}$ and $\mathcal{U}'$ be two numeric covers of $\mathcal{V}_N$. The cover $\mathcal{U}'$ is a *coarsening* of $\mathcal{U}$, denoted by $\mathcal{U} \preceq \mathcal{U}'$, if for every $U \in \mathcal{U}$, there exists $U' \in \mathcal{U}'$ such that $U \subseteq U'$.

It is straightforward to show that $\preceq$ is a preorder on the space of all numeric covers. Intuitively, if $\mathcal{U} \preceq \mathcal{U}'$, then every subproblem of $\mathcal{U}$ is "wrapped" in some subproblem of $\mathcal{U}'$. Therefore, coarsening the cover cannot make more cross-cutting interactions and can make previously hidden cross-cutting interactions visible to local subproblems.

Refinement is triggered only by the detection of an obstruction. Given a synchronised compatible family $(s_i)_i \in \Gamma^0_{\mathit{sync}}(\mathcal{U}, \mathcal{F}^{\mathcal{G}})$ with shared synchronised local plan $\pi = (a_1, \ldots, a_m)$, let s^ℓ denote the induced global state at step ℓ. An obstruction is detected if either a state-evaluable interaction $e \in \mathcal{E}_{\mathit{state}}$ is violated at some state s^ℓ, or a transition interaction $e \in \mathcal{E}_{\mathit{effect}}$ is violated by some transition $(s^\ell, s^{\ell+1})$. We call such e an *obstruction witness*.

Lemma 2. Let $(s_i)_i$ be a synchronised compatible family of $\mathcal{U}$. If an interaction $e \in \mathcal{E}$ is an obstruction witness, then it is a cross-cutting interaction with respect to $\mathcal{U}$.

Definition 20 (Cutting subfamily)**.** Given a numeric cover $\mathcal{U}$, let $e \in \mathcal{E}$ be an obstruction witness. The *cutting subfamily* induced by e is $\mathcal{U}_{\mathit{obs}}(e) \triangleq \{U \in \mathcal{U} \mid U \cap \mathrm{supp}(e) \neq \emptyset\}$.

The *merge* with respect to the obstruction witness e results in the cover $\mathcal{U}' \triangleq (\mathcal{U} \setminus \mathcal{U}_{\mathit{obs}}(e)) \cup \{\bigcup_{U \in \mathcal{U}_{\mathit{obs}}(e)} U\}$. Let $U^* = \bigcup_{U \in \mathcal{U}_{\mathit{obs}}(e)} U$, thus the merge ensures that e is no longer a cross-cutting interaction in $\mathcal{U}'$. Since e is cross-cutting, thus $|\mathcal{U}'| < |\mathcal{U}|$ and $\mathcal{U} \preceq \mathcal{U}'$.

A critical property of Topological CEGAR is that the refinement reduces the number of cross-cutting interactions with no backtracking. Given a numeric cover $\mathcal{U}$, we define the quantity $\mu(\mathcal{U})$ to be the number of cross-cutting interactions with respect to $\mathcal{U}$.

Lemma 3. If $e \in \mathcal{E}$ has its support $\mathrm{supp}(e)$ contained in some $U \in \mathcal{U}$, then for any coarsening $\mathcal{U} \preceq \mathcal{U}'$, e is not a cross-cutting interaction in $\mathcal{U}'$. Furthermore, $\mu(\mathcal{U}') \leq \mu(\mathcal{U})$.

Termination and Completeness

A cover $\mathcal{U}$ is *terminal* for Topological CEGAR if the procedure stops at $\mathcal{U}$ and either finds a valid global plan, or no synchronised compatible family exists over $\mathcal{U}$, or no obstruction witness is detected.

Theorem 3. [Finite termination] Given any initial numeric cover $\mathcal{U}_0$ which contains finitely many subsets, any sequence of merges triggered by obstruction witnesses terminates after at most $|\mathcal{U}_0| - 1$ steps. Moreover, if each merge eliminates at least its witness, then it also terminates after at most $\mu(\mathcal{U}_0)$ witness-eliminating steps.

Since the CEGAR refinement will terminate for any initial cover, we explain the relationship between the terminal cover and the solvability of the numeric planning instance in the following theorem.

Theorem 4. Given Π, let $\mathcal{U}^*$ be a terminal cover of a Topological CEGAR sequence, then

1. If Π is solvable, then every global plan induces a synchronised compatible family over $\mathcal{U}^*$ and thus $\Gamma^0_{sync}(\mathcal{U}^*, \mathcal{F}^{\mathcal{G}}) \neq \emptyset$.
2. Assume $\mu(\mathcal{U}^*) = 0$ and $\mathcal{U}^*$ is overlap consistent, then Π is solvable if and only if $\Gamma^0_{sync}(\mathcal{U}^*, \mathcal{F}^{\mathcal{G}}) \neq \emptyset$.

From above, we know that when $\mu(\mathcal{U}^*) = 0$ and $\mathcal{U}^*$ is overlap consistent, the terminal cover is exact by Theorem 1, therefore, Π is solvable if and only if $\Gamma^0_{sync}(\mathcal{U}^*, \mathcal{F}^{\mathcal{G}}) \neq \emptyset$.

When $\mu(\mathcal{U}^*) > 0$, the framework can only guarantee on one side. A synchronised compatible family may exist but induce an invalid global plan from the shared synchronised local plan, since the cross-cutting interaction is not enforced in any subproblem.

This asymmetry is not due to the refinement strategy as any decomposition may encounter this issue as long as the cross-cutting interaction exists. The advantage of Topological CEGAR is that it does not remove all such constraints initially, but rather refines the cover only when needed. When it terminates, it either produces a valid global plan or witnesses unsolvability relative to the terminating cover by showing that no synchronised compatible family exists.

A refinement example

We illustrate Topological CEGAR with a small delivery instance. Consider four numeric variables f, b, w, t, which represent fuel, battery, package weight and the number of deliveries. Let $s_0 = (f = 2, b = 2, w = 0, t = 0)$ and the numeric goal $\mathcal{G}_N = \{w = 0, t \geq 1\}$. The actions $load$ and $recharge$ have no numeric precondition. The numeric effects of $load$ and $recharge$ are $w \mathrel{+}= 5$ and $f \mathrel{+}= 2$, respectively. The action $deliver$ has numeric precondition $f + b \geq w$, and effect $w := 0$ and $t \mathrel{+}= 1$.

Let the initial cover be $\mathcal{U}_0 = \{U_1 = \{f, b\}, U_2 = \{w\}, U_3 = \{t\}\}$. The constraint $f + b \geq w$ is cross-cutting since $\{f, b, w\}$ is not contained in any subset. Let the shared action sequence $\pi = (load, deliver)$. For U_1, the local trajectory is $(f = 2, b = 2) \rightarrow (f = 2, b = 2) \rightarrow (f = 2, b = 2)$ since $load$ and $deliver$ have no visible numeric precondition or effect locally. For U_2, the action $load$ changes w from 0 to 5, and the effect of $deliver$ makes $w = 0$, thus we have a trajectory $(w = 0) \rightarrow (w = 5) \rightarrow (w = 0)$. For U_3, the trajectory is $(t = 0) \rightarrow (t = 0) \rightarrow (t = 1)$. Therefore, π is locally valid for each subproblem. However, π is not valid globally. After taking the action $load$, the global state becomes $(f = 2, b = 2, w = 5, t = 0)$, since $f + b = 4 < 5 = w$, the precondition of $f + b \geq w$ fails, thus $deliver$ cannot be applicable globally. The constraint $f + b \geq w$ witnesses an obstruction.

Topological CEGAR merges the subsets U_1 and U_2. Now the refined cover is $\mathcal{U}' = \{U_1' = \{f, b, w\}, U_3 = \{t\}\}$. Let a shared action sequence $\pi' = (load, recharge, deliver)$. After $load$ and $recharge$, $f + b = 6 \geq 5 = w$, the action $deliver$ is applicable. The refined cover has no cross-cutting interactions and is overlap consistent since the cover is a partition, so by Theorem 1, the synchronised compatible family induced by π' can be glued into a valid global plan.

Related Work

Since our contribution is primarily a structural and mathematical formulation rather than a new planning algorithm, we focus on related work in applied sheaf theory, topological methods and decomposition-based planning.

Sheaf-theoretic and topological methods The mathematical foundation of this paper follows the standard textbook (Lane and Moerdijk 1992). The application of topological ideas to AI planning is inspired by the book (Ghrist 2014). In recent years, sheaf-theoretic methods have been applied in many areas such as sensor networks (Robinson 2017), distributed task solvability (Felber, Galeana, and Flores 2025), multi-agent reinforcement learning (Schmid 2025) and deep learning (Ayzenberg et al. 2025).

In the seminal work of Abramsky and Brandenburger (2011), the authors used sheaf theory to describe nonlocality and contextuality in quantum systems. They characterised these phenomena as obstructions to the existence of global sections for compatible families of local measurement outcomes. Later, Abramsky et al. (2015) introduced cohomological obstructions as formal witnesses to failures of local-to-global gluing. We follow their insight and adopt the local-to-global viewpoint in numeric planning.

Decomposition-based planning Amir and Engelhardt (2003) factorised classical domains into interacting subdomains via graph decomposition techniques and used dynamic programming to obtain a global plan without global backtracking. The compatibility between subdomains was handled via discrete message passing. Kelareva et al. (2007) proposed the factoring algorithm dTreePlan with backtracking. They recursively partitioned a hypergraph to create a decomposition tree with minimal variable sharing between subtrees and used a SAT solver to resolve the inserted abstract actions between subproblems.

Brafman and Domshlak (2006) demonstrated that the treewidth of a domain's causal graph indicates the effectiveness of decomposition-based methods in classical planning. They introduced a local iterative deepening algorithm to compile the search into a discrete Constraint Satisfaction Problem, which ensured agreement on boundaries between subdomains. Their use of the treewidth of the causal graph to analyse domain complexity is conceptually similar to our use of the Gaifman graph. However, in this paper our focus is the formal characterisation of structural compatibility rather than developing an algorithm to exploit this concept.

Fabre et al. (2010) framed the cost-optimal factored planning problem as the selection of compatible local plans for interacting components. This formulation is conceptually similar to our notion of synchronised compatible families, though we have a stricter synchronisation requirement. In order to handle the problems with unbounded local plan lengths, they represented valid local plans as regular languages recognised by weighted finite automata and facilitated global compatibility via a message passing algorithm.

In (Gnad and Hoffmann 2018), a particular structural profile in decomposition, a star topology, was studied for classical planning. The authors showed that if a star topology decoupling is exhibited, exponential reduction in state-space size can be achieved by decoupled search. Later, Gnad, Álvaro Torralba, and Fišer (2022) generalised decoupled search to arbitrary variable partitions without a strict star centre. Unlike their works, our framework provides a formal obstruction theory for numeric variable decomposition which does not require a particular decomposition topology.

Conclusion

We introduce a sheaf-theoretic framework to analyse decomposition in numeric planning. In this framework, states of subproblems are viewed as local sections, and global solvability is formulated as a local-to-global gluing problem. Rather than providing an immediate practical planner, our goal is to give a rigorous structural description of when local reasoning over numeric subsets can be assembled into valid global reasoning.

We defined sound and exact covers, and showed that overlap consistent covers with no cross-cutting interaction are exact. We also connected numeric covers with tree decompositions of the Gaifman graph. Finally, we presented Topological CEGAR, a refinement procedure over numeric covers. When an obstruction is detected, the refinement step coarsens the cover. This shows that our framework can be converted into a finite mechanism.

We believe this work can open a path towards modeling planning problems via sheaf theory, and we suggest the following future theoretical and practical directions.

First of all, cohomology is often used in sheaf theory to detect and classify obstructions to gluing. However, our framework only provides a set-valued obstruction theory. The presheaves assign to each subset of numeric variables a set of local sections such as reachable states or goal satisfying states. These objects do not automatically carry the algebraic structure required for standard Čech cohomology, such as abelian group, module and vector-space operations. In order to apply cohomology to planning, we need to equip the planning objects with suitable structures and represent cochains, cocycles and coboundaries accordingly. It will be interesting to use cohomology to distinguish different kinds of local-to-global inconsistencies and to explain how refinement of covers can eliminate obstruction classes.

Second, we only consider numeric covers in this paper. This allows us to keep the formalism clean. However, if we only decompose the numeric variable set, the local subproblems may not be small enough to be solved efficiently and thus a hybrid cover is more desirable. In order to construct a hybrid cover, we may need to consider dependencies among numeric variables themselves, propositional variables themselves and numeric and propositional variables. This will require a more careful treatment to handle the interaction structure and a corresponding notion of obstructions.

Third, synchronised compatible family plays a central role in the structural characterisation of global solvability. However, it is non-trivial to extract these families. It will be interesting to develop practical algorithms to compute symbolic representations of locally reachable goal states and action sequences that reach them. When the cover is acyclic or has low-width overlap structure, we may use message-passing or dynamic programming methods. Local subproblems can send summaries of feasible boundary states and compatible action prefixes to neighbour subproblems, and global synchronisation can be checked by propagating these summaries along the cover structure.

Fourth, the cross-cutting number is not by itself a solvability criterion. We may consider other quantities such as the maximum support size of cross-cutting interactions, the treewidth of the Gaifman graph, and the overlap complexity of the cover as better indicators of the structural difficulty for numeric planning. Studying these parameters can help us understand when certain classes of numeric planning instance can have useful decompositions.

Another direction is to use numeric covers to derive admissible heuristics. When every global plan restricts to a valid local plan for each subproblem, the maximum of the optimal costs is an admissible heuristic. It will be interesting to study the trade-off between cover refinement and heuristic value, i.e., a coarser cover may have better estimate, but may increase the cost of solving local subproblems.

Finally, a practical direction is to design refinement strategies to repair detected obstructions while restricting the growth of local subproblems. The initial cover provides a tractable approximation and a good selection may reduce the number of merges. It will be important to develop practical methods that balance bag size and the expected number of cross-cutting interactions.

Acknowledgments

This research was supported by the Commonwealth through an Australian Government Research Training Program Scholarship [DOI: https://doi.org/10.82133/C42F-K220] and also by the University of Melbourne Faculty of Engineering and Information Technology Ingenium Scholarship Award.

References

Abramsky, S.; Barbosa, R. S.; Kishida, K.; Lal, R.; and Mansfield, S. 2015. Contextuality, Cohomology and Paradox. In *24th EACSL Annual Conference on Computer Science Logic*, LIPIcs, 211–228.

Abramsky, S.; and Brandenburger, A. 2011. The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics*, 13.

Aldinger, J.; Mattmüller, R.; and Göbelbecker, M. 2015. Complexity of Interval Relaxed Numeric Planning. In *KI 2015: Advances in Artificial Intelligence*, 19–31.

Amir, E.; and Engelhardt, B. 2003. Factored Planning. *Proc. of IJCAI*.

Ayzenberg, A. A.; Gebhart, T.; Magai, G.; and Solomadin, G. 2025. Sheaf theory: from deep geometry to deep learning. *arXiv preprint arXiv.2502.15476*.

Bäckström, C.; and Nebel, B. 1995. Complexity results for SAS$^+$ planning. *Computational Intelligence*, 11: 625–655.

Bodlaender, H. L.; and Möhring, R. H. 1990. The Pathwidth and Treewidth of Cographs. In *Proceedings of the 2nd Scandinavian Workshop on Algorithm Theory*, 301 – 309.

Bofill, M.; Espasa, J.; and Villaret, M. 2016. The RANTANPLAN planner: system description. In *The Knowledge Engineering Review*, volume 31, 452–464.

Brafman, R. I.; and Domshlak, C. 2006. Factored Planning: How, When, and When Not. *Proc. of AAAI*, 809–814.

B.R.Tennison. 1975. *Sheaf Theory*. Cambridge: Cambridge University Press.

Clarke, E.; Grumberg, O.; Jha, S.; Lu, Y.; and Veith, H. 2003. Counterexample-guided abstraction refinement for symbolic model checking. In *Journal of the ACM*, 752 – 794.

Coles, A.; Fox, M.; Long, D.; and Smith, A. 2008. A Hybrid Relaxed Planning Graph-LP Heuristic for Numeric Planning Domains. In *Proc. of ICAPS*, 52–59.

Dechter, R.; and Pearl, J. 1989. Tree clustering for constraint networks. In *Artificial Intelligence*, volume 38, 353–366.

Fabre, E.; Jezequel, L.; Haslum, P.; and Thiébaux, S. 2010. Cost-Optimal Factored Planning: Promises and Pitfalls. *Proc. of ICAPS*, 65–72.

Felber, S.; Galeana, H. R.; and Flores, B. H. 2025. A Sheaf-Theoretic Characterization of Tasks in Distributed Systems. *Structural Information and Communication Complexity*.

Fox, M.; and Long, D. 2003. PDDL2.1 : An Extension to PDDL for Expressing Temporal Planning Domains. *Journal of Artificial Intelligence Research (JAIR)*, 20: 61–124.

Fox, M.; Long, D.; and Magazzeni, D. 2011. Automatic Construction of Efficient Multiple Battery Usage Policies. In *Proc. of IJCAI*, 2620 – 2625.

Gaifman, H. 1982. On Local and Non-Local Properties. In *Studies in Logic and the Foundations of Mathematics*, volume 107, 105–135.

Ghrist, R. 2014. *Elementary Applied Topology*. CreateSpace Independent Publishing Platform.

Gnad, D.; Álvaro Torralba; and Fišer, D. 2022. Beyond Stars – Generalized Topologies for Decoupled Search. *Proc. of ICAPS*, 110–118.

Gnad, D.; and Hoffmann, J. 2018. Star-topology decoupled state space search. *Artificial Intelligence*, 257: 24–60.

Gnad, D.; or Alon, L.; Weiss, E.; and Shleyfman, A. 2025. PDBs go numeric: pattern-database heuristics for simple numeric planning. In *Proc. of AAAI*, 26507–26515.

Helmert, M. 2002. Decidability and Undecidability Results for Planning with Numerical State Variables. In *Proc. of the International Conference on AI Planning & Scheduling (AIPS)*, 44–53.

Hoffmann, J. 2003. The Metric-FF Planning System: Translating "Ignoring Delete Lists" to Numeric State Variables. *Journal of Artificial Intelligence Research (JAIR)*.

Hoffmann, J.; Gomes, C.; Selman, B.; and Kautz, H. 2007. SAT encodings of state-space reachability problems in numeric domains. In *Proc. of IJCAI*, 1918 – 1923.

Kelareva, E.; Buffet, O.; Huang, J.; and Thiébaux, S. 2007. Factored Planning Using Decomposition Trees. *Proc. of IJCAI*.

Kiam, J. J.; and Schulte, A. 2017. Multilateral quality mission planning for solar-powered long-endurance UAV. In *Aeropsace Conference*, 1–10.

Lane, S. M.; and Moerdijk, I. 1992. *Sheaves in Geometry and Logic*. New York: Springer-Verlag.

Leofante, F.; Giunchiglia, E.; Ábrahám, E.; and Tacchella, A. 2020. Optimal Planning Modulo Theories. In *Proc. of IJCAI*.

Löhr, J.; Eyerich, P.; Keller, T.; and Nebel, B. 2012. A Planning Based Framework for Controlling Hybrid Systems. In *Proc. of ICAPS*, 164–171.

Piacentini, C.; Alimisis, V.; Fox, M.; and Long, D. 2013. Combining a Temporal Planner with an External Solver for the Power Balancing Problem in an Electricity Network. In *Proc. of ICAPS*, volume 23, 398–406.

Piacentini, C.; Castro, M. P.; Cire, A. A.; and Beck, J. C. 2018. Linear and Integer Programming-based Heuristics for Cost-optimal Numeric Planning. In *Proc. of AAAI*.

Robertson, N.; and Seymour, P. 1986. Graph minors. V. Excluding a planar graph. In *Journal of Combinatorial Theory*, volume 41, 92 – 114.

Robinson, M. 2017. Sheaves are the canonical data structure for sensor integration. *Information Fusion*, 36: 208–224.

Scala, E.; Haslum, P.; Magazzeni, D.; and Thiébaux, S. 2017. Landmarks for Numeric Planning Problems. *Proc. of IJCAI*, 4384–4390.

Scala, E.; Haslum, P.; Thiebaux, S.; and Ramírez, M. 2016a. Interval-based relaxation for general numeric planning. In *Proc. of ECAI*, 655 – 663.

Scala, E.; Ramírez, M.; Haslum, P.; and Thiebaux, S. 2016b. Numeric Planning with Disjunctive Global Constraints via SMT. In *Proc. of ICAPS*, 276–284.

Scala, E.; Saetti, A.; Serina, I.; and Gerevini, A. E. 2020. Search-Guidance Mechanisms for Numeric Planning Through Subgoaling Relaxation. In *Proc. of ICAPS*, 226 – 234.

Schmid, E. 2025. Applied Sheaf Theory For Multi-agent Artificial Intelligence (Reinforcement Learning) Systems: A Prospectus. *arXiv preprint arXiv:2504.17700*.

Seipp, J.; and Helmert, M. 2018. Counterexample-Guided Cartesian Abstraction Refinement for Classical Planning. In *Journal of Artificial Intelligence Research (JAIR)*, volume 62.

Shin, J.-A.; and Davis, E. 2005. Processes and continuous change in a SAT-based planner. In *Artificial Intelligence*, volume 166, 194 – 253.

Vallati, M.; Magazzeni, D.; Schutter, B. D.; Chrpa, L.; and McCluskey, T. 2016. Efficient Macroscopic Urban Traffic Models for Reducing Congestion: A PDDL+ Planning Approach. In *Proc. of AAAI*.