

In Search of Tractability for Partial Satisfaction Planning

Michael Katz

IBM Watson Health, Israel
katzm@il.ibm.com

Vitaly Mirkis

Huawei Technologies, Israel
vitaly.mirkis@huawei.com

Abstract

The objective of partial satisfaction planning is to achieve an as valuable as possible state, tacking into account the cost of its achievement. In this work we investigate the computational complexity of restricted fragments of two variants of partial satisfaction: net-benefit and oversubscription planning. In particular, we examine restrictions on the *causal graph* structure and variable domain size of the planning problem, and show that even for the strictest such restrictions, optimal oversubscription planning is hard. In contrast, certain tractability results previously obtained for classical planning also apply to net-benefit planning. We then partially relax these restrictions in order to find the boundary of tractability for both variants of partial satisfaction planning. In addition, for the family of 0-binary value functions we show a strong connection between the complexity of cost-optimal classical and optimal oversubscription planning.

1 Introduction

While in classical planning the concern is solely around the cost of achieving all goals, in partial satisfaction planning the restriction is relaxed, allowing to achieve only some of the goals. This comes not without a price, resulting in an empty plan being a trivial valid solution and requiring to aim at achieving as large as possible value. However, the cost of achieving these values must also be accounted for. In *net-benefit* planning, the cost of achieving the values and the values themselves are assumed to be comparable and the objective is to maximize the achieved goals value minus the cost of its achievement. Net-benefit planning was shown to be PSPACE-complete [Briel *et al.*, 2004] and to have a practical mapping to classical planning [Keyder and Geffner, 2009]. A complexity investigation for restricted fragments of net-benefit planning was recently performed [Aghighi and Jonsson, 2014; Aghighi and Bäckström, 2015], focusing on **PUBS** fragments [Bäckström and Nebel, 1995]. Further, Aghighi and Jonsson (2014) in their discussion suggest that it might be possible to obtain tractable fragments of net-benefit planning by restricting the causal graph structure.

In *oversubscription planning* the cost of achieving the values and the values themselves are not assumed to be comparable. To take the cost into account, a bound on the cost or a *budget* is introduced [Smith, 2004], and the objective is to maximize the achieved goals value, while the cost is constrained by the given budget. While heuristic search has been playing a significant role in the progress of both cost-optimal classical and net-benefit planning, optimal oversubscription planning remained almost completely untouched. A significant performance improvement for optimal oversubscription planning was reported for the first time almost a decade after the oversubscription planning problem was initially introduced [Mirkis and Domshlak, 2013], exploiting a heuristic search approach with admissible heuristics based on the *explicit abstraction* paradigm [Edelkamp, 2001]. The abstract oversubscription planning problems, tractable due to their small size, were then additively composed to derive informative admissible estimates. The reported heuristic performance, compared to the baseline algorithm, in some cases, reduced the search space by three orders of magnitude.

The success of exploiting polynomial complexity fragments for deriving admissible estimates for optimal oversubscription planning was not surprising, as in the classical planning literature the picture was virtually the same. However, in contrast to both classical and net-benefit planning, oversubscription planning remains completely unexplored in terms of complexity analysis. Optimal oversubscription planning is PSPACE-complete even when severely restricting the value function, as a result of a straightforward reduction from cost-optimal classical planning. Even this result, however, to the best of our knowledge does not appear in the literature.

To alleviate this gap, we investigate the tractability of both optimal oversubscription and net-benefit planning for certain classes of problems specified by their causal graph structure and variable domain sizes [Helmert, 2004; Katz and Domshlak, 2010]. Focusing on a most common class of value functions, we show that the most structurally restricted fragments are weakly NP-complete for optimal oversubscription planning and polynomial for optimal net-benefit planning. We continue by relaxing the structural restriction and describe a boundary of tractability for both net-benefit and oversubscription planning. Furthermore, for a restricted class of value functions, we present a generic result relating tractable fragments of oversubscription and classical planning. We

conclude by summarizing our results and presenting future research directions. We find that net-benefit versions of tractable cost-optimal planning problems tend to be tractable as well, while complexity results for oversubscription problems carry over only in the presence of certain limiting assumptions concerning the value function.

2 Background

In line with the SAS⁺ formalism for deterministic planning [Bäckström and Klein, 1991; Bäckström and Nebel, 1995], a *planning task structure* is given by a pair $\langle V, O \rangle$, where V is a set of n finite-domain *state variables*, and O is a finite set of *operators*. Each complete assignment to V is called a *state*, and $S = \text{dom}(v_1) \times \dots \times \text{dom}(v_n)$ is the *state space* of the structure $\langle V, O \rangle$. A partial assignment to V is called a *partial state*. Each operator o is a pair $\langle \text{pre}(o), \text{eff}(o) \rangle$ of partial states called *preconditions* and *effects*, respectively. Denoting by $\mathcal{V}(p) \subseteq V$ the subset of variables instantiated by a partial state p , operator o is applicable in a state s iff $s[v] = \text{pre}(o)[v]$ for all $v \in \mathcal{V}(\text{pre}(o))$. Applying o changes the value of each $v \in \mathcal{V}(\text{eff}(o))$ to $\text{eff}(o)[v]$. The resulting state is denoted by $s\llbracket o \rrbracket$; by $s\llbracket \langle o_1, \dots, o_k \rangle \rrbracket$ we denote the state obtained from sequential application of the (applicable in turn) operators o_1, \dots, o_k starting at state s .

The *causal graph* of a planning task structure $\langle V, O \rangle$ is a digraph $\text{CG} = \langle V, E \rangle$ over the set of nodes V that contains an arc (v, v') iff $v \neq v'$ and both $v \in \mathcal{V}(\text{pre}(o)) \cup \mathcal{V}(\text{eff}(o))$ and $v' \in \mathcal{V}(\text{eff}(o))$ for some $o \in O$. The *domain transition graph* $\text{DTG}(v)$ of a variable $v \in V$ is an arc-labeled digraph with nodes $\text{dom}(v)$ that contains an arc (ϑ, ϑ') labeled with $\text{pre}(o) \setminus \text{pre}(o)[v]$ iff $\text{eff}(o)[v] = \vartheta'$ and either $\text{pre}(o)[v] = \vartheta$ or $v \notin \mathcal{V}(\text{pre}(o))$.

In classical planning, a planning task $\Pi = \langle V, O; s_0, G, \mathcal{C} \rangle$ extends its structure with an *initial state* $s_0 \in S$, a *goal specification* G , typically modeled as a partial state, and an *operator cost function* $\mathcal{C} : O \rightarrow \mathbb{R}^{0+}$. An operator sequence π is called a *plan* if it is applicable in s_0 , and $G \subseteq s\llbracket \pi \rrbracket$. A plan is *optimal* if the sum of its operator costs is minimal among all plans. The objective in classical planning is to find a plan of as low cost as possible, with optimal classical planning being devoted to searching for optimal plans only.

The causal graph of the planning task Π is the causal graph of its structure. In this paper we investigate several previously studied causal graph structures, exemplified in Figure 1. First, *fork* and *inverted fork* structures are directed graphs $G = \langle N, E \rangle$ such that there exists a node $r \in N$ for which $(u, v) \in E \iff u = r$, if the structure is a fork, and $(u, v) \in E \iff v = r$, if the structure is an inverted fork. We refer to planning problems whose causal graphs are (inverted) forks as (inverted) fork structured planning problems. Optimal planning has been shown to be tractable for fork structured planning problems if $|\text{dom}(r)| = 2$ and NP-complete for $|\text{dom}(r)| > 2$. For inverted fork structured planning problems optimal planning is tractable for any $|\text{dom}(r)| \in O(1)$ [Katz and Domshlak, 2010; Katz and Keyder, 2012]. Another previously studied structure is a *polytree*. A polytree is a directed graph whose underlying undirected graph is a tree. Plan generation for problems with




Figure 1: Causal graph structures: (a) fork, (b) inverted fork, (c) polytree, (d) semifork, and (e) hourglass.

polytree causal graphs and binary variable domains is NP-hard, but becomes tractable when imposing a constant bound on causal graph in-degree [Brafman and Domshlak, 2003; Giménez and Jonsson, 2008]. Moreover, imposing this restriction makes even optimal planning tractable [Katz and Domshlak, 2008a]. Alternatively, restricting the operators to have at most one precondition also makes optimal planning for binary polytrees tractable [Katz and Domshlak, 2008a]. Lastly, a *semifork* structure consists of a fork structure, and the remaining nodes, called a *hat*, have edges either among themselves or to the root of the fork (called a *center*), and *hourglass* structure consists of both a fork and an inverted fork rooted at the same variable (also called a *center*). Semiforks with binary center variable domain and constantly bounded hat and hourglasses with binary center variable domain and operators with at most one precondition are tractable for optimal planning [Katz and Keyder, 2012].

2.1 Oversubscription Planning

In contrast to classical planning, an *oversubscription planning task* $\Pi = \langle V, O; s_0, \mathcal{C}, u, b \rangle$ extends its structure with four components: an initial state $s_0 \in S$ and an operator cost function $\mathcal{C} : O \rightarrow \mathbb{R}^{0+}$ as above, plus a succinctly represented and efficiently computable *state value function* $u : S \rightarrow \mathbb{R}^{0+}$, and a *cost budget* $b \in \mathbb{R}^{0+}$. An operator sequence π is called a plan if it is applicable in s_0 , and $\sum_{o \in \pi} \mathcal{C}(o) \leq b$; by $\hat{u}(\pi)$ we refer to the value of the end-state of π , that is, $\hat{u}(\pi) = u(s_0\llbracket \pi \rrbracket)$. While an empty operator sequence is always a plan, the objective in oversubscription planning is to find a plan that achieves an as valuable state as possible, and *optimal oversubscription planning* is devoted to searching for optimal plans only: A plan π is *optimal* if $\hat{u}(\pi)$ is maximal among all the plans.

In what follows, we mostly focus on *additive state value functions*, computed as a sum over the state variables $u(s) \stackrel{\text{def}}{=} \sum_{v \in V} u(s[v])$. However, in some cases a value function can be given explicitly. One such special case is the so-called *0-binary* value function, with its image restricted to at most 2 values [Domshlak and Mirkis, 2015]. Originally, 0-binary

value functions were used in the context of explicit abstractions, and were defined on a small set of explicitly stated abstract states. We extend the definition to general oversubscription planning tasks, however restricting ourselves to 0-binary value functions compactly represented by a set of partial states with non-zero values.

2.2 Net-Benefit Planning

Similarly, a *net-benefit planning task* $\Pi = \langle V, O; s_0, \mathcal{C}, u \rangle$ extends its structure with three components: an initial state $s_0 \in S$, an operator cost function $\mathcal{C} : O \rightarrow \mathbb{R}^{0+}$, and a succinctly represented and efficiently computable state value function $u : S \rightarrow \mathbb{R}^{0+}$, as above. An operator sequence π is a plan if it is applicable in s_0 . The cost of a plan is defined as a sum of its operators costs, and the value is defined by the value of the reached state. The objective of net-benefit planning is to find a plan that maximizes the difference between its value and its cost.

2.3 The Knapsack Problem

In what follows we will exploit two variants of the so-called *knapsack problem*. The first variant is the *0-1 knapsack problem*. Given n items to pack in some knapsack of *capacity* c , each with a *profit* p and a *weight* w , the problem is to choose which items to pack, such that the profit sum is maximized without having the weight sum exceed c . The second variant is the *multiple choice knapsack problem*. Given k classes N_1, \dots, N_k of items to pack in some knapsack of *capacity* c . Each item $j \in N_i$ has a *profit* p_{ij} and a *weight* w_{ij} , the problem is to choose one item from each class such that the profit sum is maximized without having the weight sum exceed c . Both 0-1 and multiple choice knapsack are NP-hard, but both can be solved through dynamic programming in time polynomial in the problem description size and in the unary representation size of c [Dudziński and Walukiewicz, 1987].

3 Oversubscription Planning

In what follows we present a tractability analysis of the optimal oversubscription planning for fragments characterized by variable domain size and causal graph structure.

3.1 Additive State Value Function

Focusing on the additive state value functions, we start with a discouraging result, even for a severely structurally restricted fragment.

Theorem 1 *Optimal oversubscription planning for tasks with additive state value function, binary variable domains, and a causal graph with no edges, is weakly NP-complete.*

Proof: Membership in NP: cheapest plans are never required to change the value of any variable more than once, since all variables are independent of each other. The proof of hardness is by reduction from the 0-1 knapsack problem. Given items N , with a *profit* p_i and a *weight* w_i , for $i \in N$, and a *capacity* c , the oversubscription planning problem $\Pi = \langle V, O; s_0, \mathcal{C}, u, b \rangle$ is defined as follows. The variables $V = \{v_i \mid i \in N\}$, each with domain $\{0, 1\}$ represent

the items N . The initial state consists of all variables having the value 0. The value function u is defined to be p_i for the value 1 of variable v_i , and 0 otherwise, and the bound b is c . The operators $O = \{o_i \mid i \in N\}$ with $\text{pre}(o_i) = \{v_i = 0\}$, $\text{eff}(o_i) = \{v_i = 1\}$, and $\mathcal{C}(o_i) = w_i$, changing each variable from its initial value, with the cost equal to the weight of the item. A solution to the optimal oversubscription planning problem corresponds to the solution for the 0-1 knapsack problem and vice versa. ■

While the result sounds extremely discouraging, recall that the knapsack problem can be solved through dynamic programming in time polynomial in the problem description size and in the unary representation size of c , or a *pseudo-polynomial* time. Thus, in what follows, we search for a boundary of weak-NP-completeness. The first fragment of oversubscription planning we consider is a *fork-structured* causal graph with binary root domain and an additive state value function.

Theorem 2 *Given an oversubscription planning task $\Pi = \langle V, O; s_0, \mathcal{C}, u, b \rangle$ with an additive state value function and a fork causal graph rooted at $r \in V$, if $|\text{dom}(r)| = 2$, then an optimal plan for Π can be found in time polynomial in $\|\Pi\|$ and the unary budget representation.*

Proof: Let $\text{succ} = V \setminus \{r\}$ denote the leaf variables of Π . Observe that the fork structure of the causal graph $\text{CG}(\Pi)$ implies that all the operators in Π are unary-effect, and each leaf variable $v \in \text{succ}$ preconditions only the operators affecting v itself.

First, for $|\text{dom}(r)| = 2$, the algorithm below is based on the following three properties satisfied by the cheapest plans π for Π among those that achieve the same value. In what follows, we restrict our attention to cheapest plans only.

- (i) For any leaf variable $v \in \text{succ}$, the path $\pi \downarrow_v$ from $s_0[v]$ induced by π in the domain transition graph of v , $\text{DTG}(v)$, is either cycle-free or contains only zero-cost cycles. This is the case because otherwise all the nonzero-cost cycles can be eliminated from $\pi \downarrow_v$ while preserving its validity, violating the assumption that π is cheapest among those that achieve the same value. Without loss of generality, in what follows we assume that this path $\pi \downarrow_v$ in $\text{DTG}(v)$ is cycle-free; in the case of fork causal graphs, we can always select a cost-minimal π that satisfies this requirement for all $v \in \text{succ}$. Thus, we have $|\pi \downarrow_v| \leq |\text{dom}(v)| - 1$.
- (ii) Having fixed a sequence of value changes of r , the fork's leaves become mutually independent; that is, our ability to change the value of one does not affect our ability to change the value of any of the others.
- (iii) Because r is binary-valued, if $v \in \text{succ}$ is the “most demanding” leaf variable in terms of the number of value changes required from r by the operator preconditions along $\pi \downarrow_v$, then these are the only value changes of r along π , except for, possibly, a final value change to obtain a different value of r . Thus, in particular, we have $|\pi \downarrow_r| \leq \max_{v \in \text{succ}} |\text{dom}(v)|$.

From (iii), it is sufficient to go over the polynomial number of value changing sequences of the root variable r , finding for each such sequence the best plan in which the root variable performs exactly that sequence. Thus, in what follows, assume that σ is the sequence of value changes performed by the root variable r . From (ii), all other variables become mutually independent. Thus, for $v \in succ$, and for each domain value $\vartheta \in dom(v)$, let $\pi_v^\sigma(\vartheta)$ denote the cheapest sequence of operators that corresponds to a path in $DTG(v)$ from the initial value $s_0[v]$ to ϑ that can be achieved under the support of σ . In other words, $\pi_v^\sigma(\vartheta)$ can be extended by σ to a plan for the task obtained by projecting away all other successors of r . Since finding such a sequence corresponds to solving oversubscription planning optimally for a planning task with 2 variables, all such sequences for all domain values can be obtained in polynomial time.

We now construct a multiple choice knapsack $MK(\sigma)$ problem as follows:

- For each state variable $v \in succ$ we have a class $N_v = \{\vartheta \in dom(v) \mid u[v](\vartheta) > 0\} \cup \{s_0[v]\}$ of items to pack.
- The weight of each item $\vartheta \in N_v$ is the cost of $\pi_v^\sigma(\vartheta)$.
- The profit of each item $\vartheta \in N_v$ is the value $u[v](\vartheta)$.
- The overall capacity c is set to be $b - \mathcal{C}(\sigma)$, where $\mathcal{C}(\sigma)$ is the cost of a cheapest sequence of operators performing the sequence of changes σ of the variable r .

Solving the multiple choice knapsack $MK(\sigma)$ problem optimally results in one value $\vartheta_v \in N_v \subseteq dom(v)$ being selected for each v , maximizing the summed value such that the sum of the costs of the cost-minimal paths from $s_0[v]$ to ϑ_v is under the capacity $b - \mathcal{C}(\sigma)$. A corresponding plan π_σ is then constructed by appending the cost-minimal paths $\pi_v^\sigma(\vartheta_v)$ between the r -changing operators, constructing a cost-minimal path from $s_0[v]$ to ϑ_v for each v . Note that the plan π_σ achieves the maximal value among all plans in which the value changes of r are exactly σ . The cost of the resulted plan π_σ is thus $\mathcal{C}(\pi_\sigma) = \mathcal{C}(\sigma) + \sum_{v \in succ} \mathcal{C}(\pi_v^\sigma(\vartheta_v)) \leq b$. Going over the polynomial number of possible sequences σ , we can now find the plan that achieves the maximal value among all the plans π_σ .

It is straightforward to verify that the complexity of the above procedure equals to the complexity of the *multiple choice knapsack problem*, which can be solved in time polynomial in the problem description size and in the unary representation size of c . Since $c \leq b$, the complexity of the above procedure is polynomial in $\|\Pi\|$ and the unary representation size of b . To prove correctness, we show that the plan π' returned by the procedure for the task Π satisfies $\hat{u}(\pi') \geq \hat{u}(\pi)$ for some optimal plan π for Π . Given an oversubscription planning task Π , let π be an optimal plan for Π with $\mathcal{C}(\pi) \leq b$ and all $\pi_{\downarrow v}$ for the leaf variables v being cycle-free. Let σ be the corresponding sequence of value changes of r . It is sufficient to show that for the plan π_σ found by our algorithm we have $\hat{u}(\pi_\sigma) \geq \hat{u}(\pi)$. For each $v \in succ$, we have $\pi_{\downarrow v}$ corresponding to a path in $DTG(v)$ from the initial value $s_0[v]$ to $\vartheta_v = s_0[\llbracket \pi \rrbracket][v]$ that can be achieved under the support of σ . Thus, $\mathcal{C}(\pi_{\downarrow v}) \geq \mathcal{C}(\pi_v^\sigma(\vartheta_v))$, the cheapest such path. Therefore, the collection $\{\vartheta_v \mid v \in succ\}$ is a (not necessarily

optimal) solution to the multiple choice knapsack problem. Thus, $\sum_{v \in succ} u(s_0[\llbracket \pi \rrbracket][v]) \leq \sum_{v \in succ} u(s_0[\llbracket \pi_\sigma \rrbracket][v])$. Since the root value for the resulting state for π is the same as for π_σ , we thus have $\hat{u}(\pi) \leq \hat{u}(\pi_\sigma)$, as desired. ■

Since additive state value functions may be used for encoding classical planning goals, hardness results from classical planning can be translated into corresponding results for oversubscription planning. In particular, relaxing the restriction on the root variable domain size causes the cost-optimal classical planning fragment become (strongly) NP-complete [Katz and Keyder, 2012], giving us the corresponding result for oversubscription planning.

Theorem 3 *Optimal oversubscription planning for tasks with an additive value function and a fork causal graph rooted at a ternary domain variable is strongly NP-complete.*

Proof: Membership in NP: cheapest plans are never required to change the value of any non-root variable v more than $|dom(v)|$ times, and r , the only variable that can support such changes, changes its value at most $|dom(r)| = 3$ times to support each change. The proof of hardness is by reduction from bounded classical planning. Given a classical planning task $\Pi_G = \langle V, O; s_0, G, \mathcal{C} \rangle$ with fork structured causal graph rooted at $r \in V$ with $|dom(r)| = 3$ and a bound b , the additive state value function u is defined by mapping the goal values $G[v]$ for $v \in \mathcal{V}(G)$ to 1 and others to 0. Thus, the value of a goal state will be exactly $|\mathcal{V}(G)|$, and strictly smaller for all non-goal states. Thus, optimal oversubscription planning plans that achieve the value $|\mathcal{V}(G)|$ correspond to bounded planning plans and vice versa. ■

We now move to the second fragment, an *inverted fork-structured* causal graph with constant-bounded sink domain.

Theorem 4 *Given an oversubscription planning task $\Pi = \langle V, O; s_0, \mathcal{C}, u, b \rangle$ with an inverted fork causal graph with sink $r \in V$, if $|dom(r)| = O(1)$, then an optimal plan for Π can be found in time polynomial in $\|\Pi\|$ and the unary budget representation.*

Proof: Let $pred = V \setminus \{r\}$ denote the predecessors of r and let $|dom(r)| = d$. Observe that the inverted-fork structure of the causal graph $CG(\Pi)$ implies that all the operators in Π are unary-effect, and that the sink r preconditions only the operators affecting r itself. Hence, in what follows we assume that $u[r](\vartheta) > 0$ for some $\vartheta \in dom(r)$; otherwise Π breaks down to an oversubscription planning problem over a set of independent variables, which can be solved by a multiple choice knapsack. Likewise, from the above properties of Π it follows that, if π is a cheapest plan for Π among those that achieve the same value, then the path $\pi_{\downarrow r}$ from $s_0[r]$ to $s_0[\llbracket \pi \rrbracket][r]$ induced by π in $DTG(r)$ is either cycle-free or contains only zero-cost cycles. The latter can be safely eliminated from π , and thus we can assume that $\pi_{\downarrow r}$ is cycle-free. The algorithm below is based on the following three properties satisfied by the cheapest plans π for Π among those that achieve the same value. In what follows, we restrict our attention to cheapest plans only.

- (i) There are $\Theta(|O|^d)$ cycle-free paths in the domain transition graph $DTG(r)$ from $s_0[r]$ to all $\vartheta_r \in \text{dom}(r)$.
- (ii) Having fixed a sequence σ of value changing operators of r , the inverted fork's parents change their values to support the preconditions required along σ independently of each other.
- (iii) For each variable $v \in \text{pred}$, and each pair of v 's values $x, y \in \text{dom}(v)$, the cost-minimal path $\pi_v(x, y)$ from x to y in $DTG(v)$ can be computed in poly-time. The whole set of such cost-minimal paths can be computed using $|V| - 1$ applications of the Floyd-Warshall algorithm on the domain transition graphs of the sink's parents pred , in $O(\sum_{v \in \text{pred}} |\text{dom}(v)|^3)$.

From (i), it is sufficient to go over the polynomial number of value changing sequences of the sink variable r , finding for each such sequence the best plan in which the sink variable performs exactly that sequence. Thus, in what follows, assume that σ is the sequence of value changing operators of the sink variable r . Given σ , for each $v \in \text{pred}$, let π_v be some cheapest sequence of v -changing operators achieving the values of v in the order they appear in the preconditions along the sequence σ . Such sequences can be constructed by a simple concatenation of the sequences computed in (iii), with the last achieved value for each parent v is the precondition value that appears last in the sequence σ , with π_v being empty if no precondition is required from v along σ . Let ρ_σ be the applicable sequence of operators obtained by merging σ with all π_v for $v \in \text{pred}$ and let $s_\sigma = s_0 \llbracket \rho_\sigma \rrbracket$ be the state resulting from applying ρ_σ in the initial state.

We now construct a multiple choice knapsack $MK(\sigma)$ problem as follows:

- For each state variable $v \in \text{pred}$ we have a class $N_v = \{\vartheta \in \text{dom}(v) \mid u[v](\vartheta) > 0\} \cup \{s_\sigma[v]\}$ of items to pack.
- The weight of each item $\vartheta \in N_v$ is set to be the cost of $\pi_v(s_\sigma[v], \vartheta)$.
- The profit of each item $\vartheta \in N_v$ is the value $u[v](\vartheta)$.
- The overall capacity c is set to be $b - \mathcal{C}(\rho_\sigma)$.

Solving the multiple choice knapsack $MK(\sigma)$ problem optimally results in one value $\vartheta_v \in N_v \subseteq \text{dom}(v)$ being selected for each v , maximizing the summed value such that the sum of the costs of the cost-minimal paths from $s_\sigma[v]$ to ϑ_v is bounded by $b - \mathcal{C}(\rho_\sigma)$. A corresponding plan π_σ is then constructed by appending the cost-minimal paths $\pi_v(s_\sigma[v], \vartheta_v)$ to the end of ρ_σ . Note that π_σ achieves maximal value among all plans in which the value changes of r are exactly σ . The cost of π_σ is thus $\mathcal{C}(\pi_\sigma) = \mathcal{C}(\rho_\sigma) + \sum_{v \in \text{pred}} \mathcal{C}(\pi_v(s_\sigma[v], \vartheta_v)) \leq b$. Going over the polynomial number of possible sequences σ , we can now find the plan that achieves the maximal value among all the plans π_σ .

The complexity of the above procedure equals to the complexity of the *multiple choice knapsack problem*, which can be solved in time polynomial in the problem description size and in the unary representation size of c . Since $c \leq b$, the complexity of the above procedure is polynomial in $\|\Pi\|$ and the unary representation size of b . To prove correctness, we show that the plan π' found by running the procedure on task Π satisfies $\hat{u}(\pi') \geq \hat{u}(\pi)$ for some optimal

plan π for Π . Given an oversubscription planning task Π , let π be an optimal plan for Π with $\mathcal{C}(\pi) \leq b$ and $\pi \downarrow_r$ for the sink variable r being cycle-free. Let σ be the corresponding sequence of value changes of r . It is sufficient to show that for the plan π_σ found by our algorithm we have $\hat{u}(\pi_\sigma) \geq \hat{u}(\pi)$. For each $v \in \text{pred}$, we have $\pi \downarrow_v$ corresponding to a path in $DTG(v)$ from the initial value $s_0[v]$ to $\vartheta_v = s_0 \llbracket \pi \rrbracket [v]$, passing through the precondition values required along σ . Let s_v be the last such value, if exists, otherwise $s_v = s_0[v]$. Note that $s_v = s_\sigma[v]$ for all parents $v \in \text{pred}$. The value s_v separates the sequence of operators $\pi \downarrow_v$ into two sequences $\pi \downarrow_v (s_0[v], s_v)$ from the initial value to s_v and $\pi \downarrow_v (s_v, \vartheta_v)$ from s_v to the final value ϑ_v . Thus, $\mathcal{C}(\pi \downarrow_v (s_v, \vartheta_v)) \geq \mathcal{C}(\pi_v(s_\sigma[v], \vartheta_v))$, the cheapest such path. Therefore, the collection $\{\vartheta_v \mid v \in \text{pred}\}$ is a (not necessarily optimal) solution to the multiple choice knapsack problem. Thus, $\sum_{v \in \text{pred}} u(s_0 \llbracket \pi \rrbracket [v]) \leq \sum_{v \in \text{pred}} u(s_0 \llbracket \pi_\sigma \rrbracket [v])$. Since the sink value for the resulting state for π is the same as for π_σ , we thus have $\hat{u}(\pi) \leq \hat{u}(\pi_\sigma)$, as desired. ■

Here as well, relaxing the restriction on sink variable domain size make optimal oversubscription planning (strongly) NP-complete, even if the domain transition graphs of all the state variables are strongly connected.

Theorem 5 *Optimal oversubscription planning for tasks with an additive value function and an inverted fork causal graph with strongly connected domain transition graphs of all the state variables is strongly NP-complete.*

Proof: Membership in NP: cheapest plans are never required to change the value of the sink variable r more than $|\text{dom}(r)|$ times, and for an operator o changing the sink variable, each parent v can achieve o 's precondition with at most $|\text{dom}(v)|$ value changes. The proof of hardness is by reduction from bounded classical planning. Given a classical planning task $\Pi_G = \langle V, O; s_0, G, \mathcal{C} \rangle$ with inverted fork structured causal graph with sink $r \in V$ and a bound b , the additive state value function u is defined by mapping the goal values $G[v]$ for $v \in \mathcal{V}(G)$ to 1 and other to 0. Thus, the value of a goal state will be exactly $|\mathcal{V}(G)|$, and strictly smaller for all non-goal states. Thus, optimal oversubscription planning plans that achieve the value $|\mathcal{V}(G)|$ correspond to bounded planning plans and vice versa, and the hardness stems from the corresponding result for cost optimal planning [Helmert, 2004]. ■

3.2 0-Binary Value Function

The results of Theorem 1 leave virtually no hope for deriving tractable fragments of oversubscription planning for the additive state value functions, and thus here we turn our attention to other families of value functions. One such family that was shown useful for deriving heuristic values is 0-binary value function [Mirkis and Domshlak, 2013]. In case of a 0-binary value function compactly described by a set of partial states with a non-zero value, let us denote such a set by S_f for a value function f . The following theorem establishes a connection between tractable fragments of cost-optimal classical planning and oversubscription planning with the aforementioned restriction on the value function.

Theorem 6 *Given an oversubscription planning task $\Pi = \langle V, O; s_0, \mathcal{C}, u, b \rangle$ with a 0-binary value function u described by a set S_u of partial states with a non-zero value, if the classical planning task $\Pi_G = \langle V, O; s_0, G, \mathcal{C} \rangle$ is optimally solvable in time polynomial in $\|\Pi_G\|$ for any $G \in S_u$, then optimal oversubscription planning for Π is also solvable in time polynomial in $\|\Pi\|$ and $|S_u|$.*

Proof: For each partial state $G \in S_u$, solve the classical planning task $\Pi_G = \langle V, O, s_0, G, \mathcal{C} \rangle$ optimally. Let π_G be some cost-optimal plan for the classical planning task Π_G . If for some $G \in S_u$ we have $\mathcal{C}(\pi_G) \leq b$, then π_G is an optimal plan for Π . Otherwise, the empty plan is an optimal plan for Π . ■

Note that even with such 0-binary value functions optimal oversubscription planning is at least as hard as cost bounded classical planning for tasks of the same structure, for the same reasons as for additive value functions.

Following Theorem 6, we can obtain tractable fragments of oversubscription planning from the tractable fragments of cost-optimal planning whose tractability is not dependent on a specific goal state. In particular, the following tractable fragments fit the setting of Theorem 6: forks with binary root variable domain and inverted forks with constantly bounded sink variable domain [Katz and Domshlak, 2008b]; polytrees with either 1-dependent operators or $O(1)$ -bounded in-degree [Katz and Domshlak, 2008a]; hourglasses with binary center variable domain and operators with at most one precondition or semiforks with binary center variable and $O(1)$ -bounded size hat [Katz and Keyder, 2012].

4 Net-Benefit Planning

Following the suggestion of Aghighi and Jonsson (2014), we now turn our attention to the net-benefit planning. Focusing on the additive state value functions, observe that the reformulation suggested by Keyder and Geffner (2009) can be adapted to the multi-valued SAS⁺ formalism as follows. Let $\Pi = \langle V, O; s_0, \mathcal{C}, u \rangle$ be a net-benefit planning task. The *cost-optimal planning encoding* $X(\Pi) = \langle V, O'; s_0, G, \mathcal{C}' \rangle$ of Π is constructed by adding an auxiliary value ϵ_v to each variable v with non-zero utility values. The goal G is the collection of these values. Let $m_v = \max\{u(\vartheta) \mid \vartheta \in \text{dom}(v)\}$ be the maximal utility value for each variable. First, the operator set O' includes all the operators O from the net-benefit planning task. Further, for each variable v and each value $\vartheta \in \text{dom}(v)$, O' consists of an operator $o(v, \vartheta)$ with $\text{pre}(o(v, \vartheta)) = \{v = \vartheta\}$, $\text{eff}(o(v, \vartheta)) = \{v = \epsilon_v\}$, and $\mathcal{C}'(o(v, \vartheta)) = m_v - u(\vartheta)$. Each such operator changes a single variable, and thus does not contribute any edges to the causal graph. In words, the reformulation allows to achieve the goal from any state by paying the difference between the maximal possible utility and the obtained one. The suggested encoding turns out to be a useful tool, allowing us to derive tractability of optimal net-benefit planning from the results for cost-optimal planning. We start with the fork fragment.

Theorem 7 *Given a net-benefit planning task $\Pi = \langle V, O; s_0, \mathcal{C}, u \rangle$ with an additive state value function and a*

fork causal graph rooted at $r \in V$, if $|\text{dom}(r)| = 2$, optimal net-benefit planning for Π is polynomial in $\|\Pi\|$.

Proof: Let $X(\Pi)$ be the cost-optimal planning task obtained from Π by the adapted reformulation of Keyder and Geffner (2009). Thus, $X(\Pi)$ has a fork structure with an at most ternary root variable domain. In the case of a ternary root domain, the goal value of the root is a terminal value, with no outgoing edges, and thus the proof of Theorem 4 in Katz and Domshlak (2010) can trivially be adapted to $X(\Pi)$. ■

Note that the simplest fragment from Theorem 1 for oversubscription planning, namely binary variable domains and causal graph with no edges is a subcase of the fork fragment above, and thus is also polynomial, in contrast to the results of Theorem 1. While the fork fragment requires a minor adaptation of the original proof for the cost-optimal planning fragment, the tractability of the inverted fork fragment stems directly from the cost-optimal planning result.

Theorem 8 *Given a net-benefit planning task $\Pi = \langle V, O; s_0, \mathcal{C}, u \rangle$ with an inverted fork causal graph with sink $r \in V$, if $|\text{dom}(r)| = O(1)$, optimal net-benefit planning for Π is polynomial in $\|\Pi\|$.*

Proof: Let $X(\Pi)$ be the cost-optimal planning task obtained from Π by the adapted reformulation of Keyder and Geffner (2009). $X(\Pi)$ has an inverted fork structure with constant valued root variable domain, and thus is tractable due to Theorem 5 of Katz and Domshlak (2010). ■

The hardness results also hold for fragments that correspond to Theorems 3 and 5. The proofs follow the same path as the proofs for these theorems and thus are omitted here.

5 Summary and Future Work

We presented new results for the complexity of optimal oversubscription and net-benefit planning for additive state value functions, drawing a boundary between tractable and NP-complete cases for net-benefit planning and between weakly NP-complete and strongly NP-complete cases for oversubscription planning. We show several fragments that are polynomial for optimal net-benefit planning but weakly-NP-complete for optimal oversubscription planning. As even severely structurally restricted fragments of optimal oversubscription planning are already weakly NP-complete, the feasibility of exploiting these fragments for deriving informative utility estimates remains an open question. Further, an investigation is needed into which additional restrictions should be imposed to make optimal oversubscription planning tractable. Theorem 1 hints that the only possible direction is to restrict the value functions.

Following that direction, for 0-binary value functions, we present a general result relating tractable fragments of oversubscription and classical planning, for value functions that are compactly represented by a set of partial states for non-zero utility values. Another option would be to compactly represent a 0-binary value function by the set of states with zero utility values. One interesting question is whether similar results can be devised for such value functions.

Acknowledgments

We thank Emil Keyder for discussions and comments that greatly improved the paper.

References

- [Aghighi and Bäckström, 2015] Meysam Aghighi and Christer Bäckström. Cost-optimal and net-benefit planning - A parameterised complexity view. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)*, pages 1487–1493, 2015.
- [Aghighi and Jonsson, 2014] Meysam Aghighi and Peter Jonsson. Oversubscription planning: Complexity and compilability. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI 2014)*, pages 2221–2227. AAAI Press, 2014.
- [Bäckström and Klein, 1991] Christer Bäckström and Inger Klein. Planning in polynomial time: the SAS-PUBS class. *Computational Intelligence*, 7(3):181–197, 1991.
- [Bäckström and Nebel, 1995] Christer Bäckström and Bernhard Nebel. Complexity results for SAS⁺ planning. *Computational Intelligence*, 11(4):625–655, 1995.
- [Brafman and Domshlak, 2003] Ronen I. Brafman and Carmel Domshlak. Structure and complexity in planning with unary operators. *Journal of Artificial Intelligence Research*, 18:315–349, 2003.
- [Briel *et al.*, 2004] Menkes Van Den Briel, Romeo Sanchez, Minh B. Do, and Subbarao Kambhampati. Effective approaches for partial satisfaction (over-subscription) planning. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI 2004)*, pages 562–569. AAAI Press, 2004.
- [Domshlak and Mirkis, 2015] Carmel Domshlak and Vitaly Mirkis. Deterministic oversubscription planning as heuristic search: Abstractions and reformulations. *Journal of Artificial Intelligence Research*, 52:97–169, 2015.
- [Dudziński and Walukiewicz, 1987] Krzysztof Dudziński and Stanisław Walukiewicz. Exact methods for the knapsack problem and its generalizations. *European Journal of Operational Research*, 28(1):3–21, 1987.
- [Edelkamp, 2001] Stefan Edelkamp. Planning with pattern databases. In Amedeo Cesta and Daniel Borrajo, editors, *Proceedings of the Sixth European Conference on Planning (ECP 2001)*, pages 84–90. AAAI Press, 2001.
- [Giménez and Jonsson, 2008] Omer Giménez and Anders Jonsson. The complexity of planning problems with simple causal graphs. *Journal of Artificial Intelligence Research*, 31:319–351, 2008.
- [Helmert, 2004] Malte Helmert. A planning heuristic based on causal graph analysis. In Shlomo Zilberstein, Jana Koehler, and Sven Koenig, editors, *Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS 2004)*, pages 161–170. AAAI Press, 2004.
- [Katz and Domshlak, 2008a] Michael Katz and Carmel Domshlak. New islands of tractability of cost-optimal planning. *Journal of Artificial Intelligence Research*, 32:203–288, 2008.
- [Katz and Domshlak, 2008b] Michael Katz and Carmel Domshlak. Structural patterns heuristics via fork decomposition. In Jussi Rintanen, Bernhard Nebel, J. Christopher Beck, and Eric Hansen, editors, *Proceedings of the Eighteenth International Conference on Automated Planning and Scheduling (ICAPS 2008)*, pages 182–189. AAAI Press, 2008.
- [Katz and Domshlak, 2010] Michael Katz and Carmel Domshlak. Implicit abstraction heuristics. *Journal of Artificial Intelligence Research*, 39:51–126, 2010.
- [Katz and Keyder, 2012] Michael Katz and Emil Keyder. Structural patterns beyond forks: Extending the complexity boundaries of classical planning. In Jörg Hoffmann and Bart Selman, editors, *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2012)*, pages 1779–1785. AAAI Press, 2012.
- [Keyder and Geffner, 2009] Emil Keyder and Héctor Geffner. Soft goals can be compiled away. *Journal of Artificial Intelligence Research*, 36:547–556, 2009.
- [Mirkis and Domshlak, 2013] Vitaly Mirkis and Carmel Domshlak. Abstractions for oversubscription planning. In Daniel Borrajo, Subbarao Kambhampati, Angelo Oddi, and Simone Fratini, editors, *Proceedings of the Twenty-Third International Conference on Automated Planning and Scheduling (ICAPS 2013)*, pages 153–161. AAAI Press, 2013.
- [Smith, 2004] David E Smith. Choosing objectives in over-subscription planning. In Shlomo Zilberstein, Jana Koehler, and Sven Koenig, editors, *Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS 2004)*, pages 393–401. AAAI Press, 2004.