Solving Simple Planning Problems with More Inference and No Search

Vincent Vidal[1] and Héctor Geffner[2]

[1] CRIL & Université d'Artois
rue de l'université - SP16, 62307 Lens Cedex, France
vidal@cril.univ-artois.fr
[2] ICREA & Universitat Pompeu Fabra
Paseo de Circunvalacion 8, 08003 Barcelona, Spain
hector.geffner@upf.edu

Abstract. Many benchmark domains in AI planning including Blocks, Logistics, Gripper, Satellite, and others lack the interactions that characterize puzzles and can be solved non-optimally in low polynomial time. They are indeed easy problems for people, although as with many other problems in AI, not always easy for machines. In this paper, we address the question of whether simple problems such as these can be solved in a simple way, i.e., without search, by means of a domain-independent planner. We address this question *empirically* by extending the constraint-based planner CPT with additional domain-independent inference mechanisms. We show then for the first time that these and several other benchmark domains can be solved with no backtracks while performing only polynomial node operations. This is a remarkable finding in our view that suggests that the classes of problems that are solvable without search may be actually much broader than the classes that have been identified so far by work in Tractable Planning.

1 Introduction

Simple problems can be hard for a general problem solver when the structure of the problems is not exploited. Domains like Blocks, Logistics, Satellite and others, for example, have all a low polynomial complexity (once the optimality requirement is dropped) and yet they all have been challenging for domain-independent planners until quite recently. Current planners solve these and other problems by exploiting structure in the form of heuristic functions that are extracted automatically and guide the search for plans [1, 2].

In this paper, we address the question of whether these and other simple, tractable domains can be solved by a domain-independent planner *with no search at all* by performing polynomial node operations only. The work in Tractable Planning addresses a related question by studying fragments of general planning languages such as Strips over which planning is polynomial [3–6]. Unfortunately, the fragments that have been identified so far as tractable remain somehow narrow and do not account for the tractability of the standard benchmarks.

In this work we approach this problem in a different way – *empirically* – by developing a general planning algorithm and showing that it solves these and other domains backtrack free, suggesting thus that the classes of problems that can be solved with no search by means of a domain-independent planner may be much broader than the ones that have been identified *theoretically* so far. Closing the gap between the empirical results and the theoretical accounts arises then as a key challenge that we hope to approach elsewhere.

The planning algorithm that we use for solving the various domains backtrack-free is an extension of the constraint-based planner CPT, an optimal temporal planner that combines a POCL branching scheme (for Partial Order Causal Link Planning [7]) with strong pruning mechanisms [8]. The extension is modular and takes the form of additional domain-independent constraints and inference mechanisms. We will refer to domain-independent planners that aim to solve simple problems in a backtrack-free manner by performing low polynomial operations in every node, as *easy planners.* The development of easy planners, we believe, is a crisp and meaningful goal, which may yield insights that an exclusive focus on performance may not, like the identification of broader tractable planning fragments, and the process by which people actually plan. Humans indeed are quite good at solving these simple problems, and while it is often assumed that this ability is the result of domain-dependent strategies, our results suggest that they may also result from simple but general inference mechanisms.

By itself, CPT like other SAT and constraint-based optimal planners [9–11], does not make for a good suboptimal planner and much less for an easy planner. Indeed, while SAT and constraint-based planners can be used with large, non-optimal planning horizons (which are upper bounds on the makespan of the plan), they face two problems: 1) SAT and CSP encodings based on one variable per time point, as normally used, become too large to handle for large planning horizons; and 2) the constraint that requires the goals to be true at the planning horizon becomes ineffective when the horizon is set too high.

In CPT the first is not a problem because, being a temporal planner, CPT uses temporal rather than boolean encodings; i.e. for each action in the domain, a single variable represents the starting time of the action in the plan. Thus, the use of a large bound on the admissible makespan of plans has a direct effect on the *upper bounds* of the temporal variables but not in their *number.*

CPT, on the other hand, does not escape from Problem 2: with a large bound on the makespan, the search becomes less constrained and focused, and even problems that are solved backtrack free with the optimal bound are not solved at all after thousands of backtracks when a larger bound is used instead. In this work, we tackle this problem by extending the inferential capabilities of CPT so that it relies less on inferences drawn from the bounding constraint and more on domain-independent inferences not captured by CPT. The new version of CPT, that we call eCPT, does simple but more extensive reasoning, making use and adapting techniques like landmarks [13, 14] and distances [15] among others.

The paper is organized as follows. We first review the CPT planner, discuss its strength as an optimal planner and its weakness as a suboptimal planner, and introduce extensions of the inferential machinery of CPT that render the search backtrack-free over a wide range of domains. We then evaluate the resulting planner, eCPT, and discuss implications and open ends.

2 CPT

CPT is a domain-independent temporal planner that combines a branching scheme based on Partial Order Causal Link (POCL) Planning with powerful and sound pruning rules implemented as constraints [8]. The key novelty in CPT in relation to other formulations [7, 16, 17] is the ability to reason about supports, precedences, and causal links involving actions that are not in the plan. In this way, CPT can prune the start time and supports of actions that are not yet in the plan, rule out actions from the plan, detect failures early on, etc. The inferences in CPT are supported by a convenient representation of POCL plans in terms of variables and constraints. For example, for each action a in the domain there is a variable $T(a)$ that represents the starting time of a, and for each precondition p of a, a variable $S(p, a)$ that represents the supporter of precondition p of a. A causal link $a'[p]a$ is thus represented by the constraint $S(p, a) = a'$, while its negation is represented by the constraint $S(p, a) \neq a'$ which means that a' cannot produce p for a, i.e. the causal link $a'[p]a$ is forbidden. Unlike other POCL planners based on constraints however, [18–21], CPT represents and reasons with all these variables, whether or not the action a is part of the current plan.

CPT uses a simple extension of the Strips language that accommodates concurrent actions with integer durations. A temporal planning problem is a tuple $P = \langle A, I, O, G \rangle$ where A is a set of ground atoms, $I \subseteq A$ and $G \subseteq A$ represent the initial and goal situations, and O is the set of ground Strips operators, each with precondition, add, and delete list $pre(a)$, $add(a)$, and $del(a)$, and *duration* $dur(a)$. As is common in POCL planning, there are also the dummy actions *Start* and *End* with zero durations, the first with an empty precondition and effect I; the latter with precondition G and empty effects. As in GRAPHPLAN [22], two actions a and a' interfere when one deletes a precondition or positive effect of the other. CPT follows the simple model of time in [23] where interfering actions cannot overlap, and produces valid plans with minimum *makespan.*

The basic formulation of the CPT planner can be described in four parts: *preprocessing*, *variables*, *constraints*, and *branching*. After the preprocessing, the variables are created and the constraints are asserted and propagated. If an inconsistency is found, no valid plan for the problem exists. Otherwise, the constraint $T(End) = B$ for the bound B on the makespan, set to the earliest possible starting time of the action *End* (i.e.; $B = T_{min}(End)$ which is determined by preprocessing, see below), is asserted and propagated. The branching scheme then takes over and if no solution is found, the process restarts by retracting the constraint $T(End) = B$ and replacing it with $T(End) = B + 1$ (1 being the smallest time unit). The search is then restarted from scratch with the new bound, and this is repeated until a solution is found. For simplicity, we follow [8] and assume that no action in the domain can be done more than once in the plan. This restriction is removed in the last version of CPT, which is the one that we use, that introduces a distinction between action types and tokens. Such details, however, reported in [12], are not needed here and are omitted.

2.1 Preprocessing

In the preprocessing phase, CPT computes the heuristic values $h^2_T(a)$ and $h^2_T(\{p, q\})$ for each action $a \in O$ and each atom pair $\{p, q\}$ as in [24]. The values provide lower bounds on the times to achieve the preconditions of a and the pair of atoms p, q, from the initial situation I. The *(structural) mutexes* (pairs of atoms that cannot be true in a world situation) are then identified as the pairs of atoms p, q for which $h^2_T(\{p, q\}) = \infty$. An action a is said to *e-delete* an atom p when either a deletes p, a adds an atom q such that q and p are mutex, or a precondition r of a is mutex with p and a does not add p. In all cases, if a e-deletes p, p is false after doing a; see [25]. Finally, $T_{min}(End) = \max_{\{p,q\} \subseteq G} h^2_T(\{p, q\})$.

In addition, the simpler heuristic h^1_T is used for defining *distances* between actions [15]. For each action $a \in O$, the h^1_T heuristic is computed from an initial situation I_a that includes all facts *except those that are e-deleted by* a. The distances $dist(a, a')$ are then set to the resulting $h^1_T(a')$ values. These distances encode lower bounds on the *slack* that must be inserted between the completion of a and the start of a' in any legal plan in which a' follows a. They are not symmetric in general and their calculation, which remains polynomial, involves the computation of the h^1_T heuristic $|O|$ times.

2.2 Variables and Domains

The state of the planner is given by a collection of variables, domains, and constraints. As emphasized above, the variables are defined for each action $a \in O$ and not only for the actions in the current plan. Moreover, variables are created for each precondition p of each action a as indicated below. The domain of variable X is indicated by $D[X]$ or simply as $X :: [X_{min}, X_{max}]$ if X is a numerical variable. The variables, their initial domains, and their meanings are:

- $T(a) :: [0, \infty]$ encodes the starting time of each action a, with $T(Start) = 0$
- $S(p, a)$ encodes the support of precondition p of action a with initial domain $D[S(p, a)] = O(p)$ where $O(p)$ is the set of actions in O that add p
- $T(p, a) :: [0, \infty]$ encodes the starting time of $S(p, a)$
- $InPlan(a) :: [0, 1]$ indicates the presence of a in plan; $InPlan(Start) = InPlan(End) = 1$ (true)

In addition, the set of actions in the current plan is kept in the variable *Steps*; i.e., $Steps = \{a \mid InPlan(a) = 1\}$. Variables $T(a)$, $S(p, a)$, and $T(p, a)$ associated with actions a which are either in or out of the current plan (i.e., actions for which the $InPlan(a)$ variable is not set to either 0 or 1 yet) are *conditional* in the following sense: these variables and their domains are meaningful only under the assumption that they will be part of the plan. In order to ensure this interpretation, some care needs to be taken in the propagation of constraints as explained in [8].

2.3 Constraints

The constraints correspond basically to disjunctions, rules, and precedences, or their combination. Temporal constraints are propagated by bounds consistency [26]. The constraints apply to all actions $a \in O$ and all $p \in pre(a)$; we use $\delta(a, a')$ to stand for $dur(a) + dist(a, a')$.

- **Bounds:** For all $a \in O$, $T(Start)+\delta(Start, a) \leq T(a)$ and $T(a)+\delta(a, End) \leq T(End)$
- **Preconditions:** Supporter a' of precondition p of a must precede a by an amount that depends on $\delta(a', a)$:

$$T(a) \;\geq\; \min_{a' \in D[S(p,a)]} (T(a') + \delta(a', a))$$

$$T(a) \;\geq\; T(p, a) + \min_{a' \in D[S(p,a)]} \delta(a', a)$$

$$T(a') + \delta(a', a) > T(a) \rightarrow S(p, a) \neq a'$$

- **Causal Link Constraints:** For all $a \in O$, $p \in pre(a)$ and a' that e-deletes p, a' precedes $S(p, a)$ or follows a

$$T(a')+dur(a')+\min_{a'' \in D[S(p,a)]} dist(a', a'') \leq T(p, a) \;\vee\; T(a)+\delta(a, a') \leq T(a')$$

- **Mutex Constraints:** For effect-interfering a and a'[3]

$$T(a) + \delta(a, a') \leq T(a') \;\vee\; T(a') + \delta(a', a) \leq T(a)$$

- **Support Constraints:** $T(p, a)$ and $S(p, a)$ related by

$$S(p, a) = a' \rightarrow T(p, a) = T(a')$$

$$T(p, a) \neq T(a') \rightarrow S(p, a) \neq a'$$

$$\min_{a' \in D[S(p,a)]} T(a') \;\leq\; T(p, a) \;\leq\; \max_{a' \in D[S(p,a)]} T(a')$$

2.4 Branching

As in POCL planning, branching in CPT proceeds by iteratively selecting and fixing flaws in non-terminal states σ, backtracking upon inconsistencies. A state σ is given by the variables, their domains, and the constraints involving them. The initial state σ_0 contains the variables, domains, and constraints above, along with the *bounding constraint* $T(End) = B$ where B is the current bound on the makespan, which in the optimal setting is set to a lower bound and is then increased until a plan is found. A state is inconsistent when a non-conditional variable has an empty domain, while a consistent state σ with no flaws is a *goal state* from which a valid plan P with bound B can be extracted by scheduling the in-plan variables at their earliest starting times.

The definition of 'flaws' parallels the one in POCL planning expressed in terms of the temporal and support variables, with the addition of 'mutex threats':

[3] Two actions are effect-interfering in CPT when one deletes a positive effect of the other, and neither one *e-deletes* a precondition of the other.

– **Support Threats:** a' *threatens a support* $S(p, a)$ when both actions a and a' are in the current plan, a' e-deletes p, and neither $T_{min}(a') + dur(a') \leq T_{min}(p, a)$ nor $T_{min}(a) + dur(a) \leq T_{min}(a')$ hold,
– **Open Conditions:** $S(p, a)$ is an *open condition* when $|D[S(p, a)]| > 1$ holds for an action a in the plan,
– **Mutex Threats:** a and a' constitute a *mutex threat* when both actions are in the plan, they are effect-interfering, and neither $T_{min}(a)+dur(a) \leq T_{min}(a')$ nor $T_{min}(a') + dur(a') \leq T_{min}(a)$ hold.

Flaws are selected for repair in the following order: first Support Threats (ST's), then Open Conditions (OC's), and finally Mutex Threats (MT's). ST's and MT's are repaired by posting precedence constraints, while OC's are repaired by choosing a supporter, as usual in POCL planning.

3 eCPT

CPT is an optimal temporal planner with good performance which is competitive with the best SAT parallel planners when actions have uniform durations [8]. At the same time, for non-optimal planning, CPT has the advantage that the size of the encodings does not grow with the bound; indeed the bound in CPT enters only through the constraint $T(End) = B$, which affects the domain upper bounds of the variables but not their number. In spite of this, however, CPT does not make for a good suboptimal planner, because like SAT and CSP planners it still relies heavily on the bounding constraint which becomes ineffective for large values of B.

Figure 1 shows the performance of CPT for the Tower-n problem for several values of n and several horizons B. Tower-n is the problem of assembling a specific tower of n blocks which are initially on the table. This is a trivial problem for people, but as shown in [8], it is not trivial for most optimal planners. CPT, however, solves this problem optimally *backtrack free* for any value of n. As the figure shows, however, the times and the number of backtracks increase when the horizon B is increased *above* the optimal bound, and for large values of B, CPT cannot solve these problems after thousands of seconds and backtracks.

The figure also shows the performance of eCPT, the extension of CPT described in this paper. It can be seen that while the performance of CPT degrades with the increase of the bound B, the performance of eCPT remains stable, and actually *backtrack free* for the different values of B. eCPT exploits the flexibility afforded by the Constraint-Programming formulation underlying CPT, extending it with inferences that do not rely as much on the bound, and which produce a backtrack-free behavior across a wide range of domains. In this section we focus on such inferences.

3.1 Impossible Supports

Many supports can be eliminated at preprocessing avoiding some dead-ends during the search. For example, the action $a' = putdown(b1)$ can never support

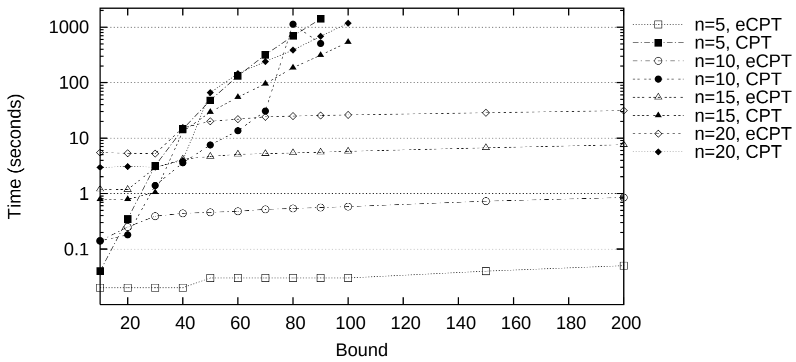

Fig. 1. Performance of CPT and eCPT over Tower-n for various numbers of blocks n and bounds B. Curves that diverge correspond to CPT; curves that remain stable to eCPT.

the precondition $p = handempty$ of an action like $a = unstack(b1, b3)$. This is because action a has another precondition $p' = on(b1, b3)$ which is e-deleted by a' (false after a') and which then would have to be reestablished by another action b before a. Yet it can be shown that in this domain, any such action b e-deletes p and is thus incompatible with the causal link $a'[p]a$.

More generally, let $dist(a', p, a)$ refer to a lower bound on the slack between actions a' and a in any valid plan in which a' *is a supporter of precondition* p *of* a. We show that for some cases, at preprocessing time, it can be shown that $dist(a', p, a) = \infty$, and hence, that a' can be safely removed from the domain of the variable $S(p, a)$ encoding the support of precondition p of a.

This actually happens when some precondition p' of a is not *reachable* from the initial situation that includes all the facts except those e-deleted by a' and where *the actions that either add or delete* p *are excluded.* The reason for this exclusion is that if a' supports the precondition p of a then it can be assumed that no action adding or deleting p can occur between a' and a (the first part is the systematicity requirement [7]). By a proposition being reachable we mean that it makes it into the so-called relaxed planning graph; the planning graph with the delete lists excluded [27].

This simple test prunes the action $putdown(b_1)$ as a possible support of the precondition $handempty$ of action $unstack(b_1, b_3)$, the action $stack(b_1, b_3)$ as a possible support of precondition $clear(b_1)$ of $pickup(b_1)$, etc.

3.2 Unique Supports

We say that an action *consumes* an atom p when it requires and deletes p. For example, the actions $unstack(b_3, b_1)$ and $pickup(b_2)$ both consume the atom *handempty*. In such cases, if the actions make it into the plan, it can be shown that their common precondition p must have different supports. Indeed, if an action a deletes a precondition of a', and a' deletes a precondition of a, a and a' are incompatible and cannot overlap in time according to the semantics. Then either a must precede a' or a' must precede a, and in either case, the precondition p needs to be established at least twice: one for the first action, and one for the second. The constraint $S(p, a) \neq S(p, a')$ for pairs of actions a and a' that consume p, ensures this, and when one of the support variables $S(p, a)$ or $S(p, a')$ is instantiated to a value b, b is immediately removed from the domain of the other variable.

3.3 Distance Boosting

The distances $dist(a, a')$ precomputed for all pairs of actions a and a' provide a lower bound on the slack between the end of a and the beginning of a'. In some cases, this lower bound can be easily improved, leading to stronger inferences. For example, the distance between the actions $putdown(b_1)$ and $pickup(b_1)$ is 0, as it is actually possible to do one action after the other. Yet the action $putdown(b_1)$ followed by $pickup(b_1)$ makes sense only if some other action using the effects of the first, occurs between these two, as when block b_1 is on block b_2 but needs to be moved on top of the block beneath b_2.

Let us say that an action a *cancels* an action a' when 1) every atom added by a' is e-deleted by a, and 2) every atom added by a is a precondition of a'. Thus, when a cancels a', the sequence a', a does not add anything that was not already true before a'. For example, $pickup(b_1)$ cancels the action $putdown(b_1)$.

When an action a cancels a', and there is a precondition p of a that is made true by a' (i.e., p is added by a' and is mutex with some precondition of a'), the distance $dist(a', p, a)$ introduced above becomes ∞ if all the actions that use an effect of a' e-delete p. In such case, as before, the action a' can be excluded from the domain of the $S(p, a)$ variable. Otherwise, the distance $dist(a', a)$ can be increased to $\min_b[dist(a', b) + dist(b, a)]$ with b ranging over the actions different than a and a' that either use an effect of a' but do not e-delete p or do not use necessarily an effect of a' but add p (because a' may be followed by an action c before a that e-deletes p but only if there is another action b between c and a that re-establishes p).

In this way, the distance between the actions $putdown(a)$ and $pickup(a)$ in Blocks is increased by 2, the distance between $sail(a, b)$ and $sail(b, a)$ in Ferry is increased by 1, etc. The net effect is similar to pruning cycles of size two in standard heuristic search. Pruning cycles of larger sizes, however, appears to be more difficult in the POCL setting, although similar ideas can potentially be used for pruning certain sequences of commutative actions.

3.4 Qualitative Precedences

Unlike traditional POCL planners, CPT reasons with *temporal precedences* of the form $T(a) + \delta(a, a') \leq T(a')$ rather than with *qualitative precedences.* CPT is a CP-based temporal planner and such a choice arises naturally from the representation used. Yet, the constraint propagation mechanism, bounds consistency, is incomplete, and in a planning context, it is often too weak. In particular, bounds consistency does not capture *transitivity:* namely from the constraints $A < B$ and $B < C$, it does not entail $A < C$. Indeed if the initial domains of the variables A, B, and C is $[1, \ldots, 100]$, bounds consistency reduces the domains to $[1, \ldots, 98]$, $[2, \ldots, 99]$, and $[3, \ldots, 100]$ respectively, which do not make $A < C$ true for all value combinations. Transitivities, however, are important in planning, and thus eCPT incorporates, in addition to temporal precedences, qualitative precedences of the form $a \prec a'$ *not limited to the actions* a *and* a' *in the plan.* Such qualitative precedences are obtained every time a temporal precedence is asserted or entailed, and are kept closed under transitivity.[4] When a new qualitative precedence $a \prec a'$ is found, the transitive closure is computed as follows: if a belongs to the current partial plan, then for all a'' such that $a'' \prec a$, $a'' \prec a'$ is recorded; and if a' belongs to the plan, then for all a'' such that $a' \prec a''$, $a \prec a''$ is recorded. The same updates are incrementally performed for an existing relation $a \prec a'$ with a or a' not in the plan, as soon as a or a' make it into the plan.

Then two inference rules make use of these qualitative precedences for pruning further the domains of the support variables:

- for an action a' in the plan that adds a precondition p of an action a: if $a \prec a'$ then $S(p, a) \neq a'$
- for an action a' that adds a precondition p of an action a and an action b in the plan that e-deletes p: if $a' \prec b$ and $b \prec a$, then $S(p, a) \neq a'$

3.5 Action Landmarks

Like all POCL planners, CPT starts with a partial plan with two actions only: *Start* and *End*. In many cases, however, it is possible to infer easily that certain other actions must be in the plan as well. For example, if a block b_1 must be moved but is beneath two blocks b_3 and b_2 in that order, then the actions $unstack(b_3, b_2)$ and $unstack(b_2, b_1)$ will have to be taken at some point, and moreover, the first must precede the second. In eCPT we identify such necessary actions and a partial order on them in a preprocessing step, following the idea of *landmarks* introduced in [13], in the form presented in [14]. An action a is a *landmark* if the action *End* is not *reachable* when the action a is excluded from the domain (as mentioned above, an action a is reachable when it makes it into the relaxed planning graph). Also, a landmark action a *precedes* a landmark action b, when b is not reachable when the action a is excluded. Action landmarks and the partial order on them are computed in the preprocessing step and are included in the initial state of the planner along with the actions *Start* and *End*. This involves the calculation of $|O|$ relaxed planning graphs, one for each action in the domain.

[4] Temporal precedences are asserted as a result of the branching decisions corresponding to support and mutex threats, and are inferred when either supports are asserted or inferred, or when one of the disjuncts in a causal link or mutex constraint becomes false.

3.6 Branching and Heuristics

eCPT retains the same branching scheme as CPT and the same ordering: it first branches on support threats (ST's), then on open conditions (OC's), and finally on mutex threats (MT's). The heuristic for selecting the support threats and open conditions however, is slightly different.

Support threats $\langle a', S(p, a)\rangle$ are selected in eCPT minimizing $T_{min}(a)$, breaking ties by first minimizing $T_{max}(p, a)$, and then with the slack based criterion used in CPT. Open supports $S(p, a)$ are selected minimizing $T_{max}(p, a)$, breaking ties minimizing $slack(a', a)$ where a' is the producer of a in $D[S(p, a)]$ with min $T_{min}(a')$. Also the constraint posted in the second case is $S(p, a) = a'$, and if that fails, $S(p, a) \neq a'$.

4 Experimental Results

While our motivation behind the development of eCPT is to study empirically the possibility of solving a wide variety of planning benchmarks with no search, we report also results that are illustrative for assessing eCPT as either an optimal or suboptimal planner in relation with state-of-the-art systems such as FF or SATPLAN04. FF [27] is a suboptimal, sequential planner winner of the 2nd Int. Planning Competition, while SATPLAN04 [9] is an optimal parallel planner winner of the Optimal Track of the 4th and last Int. Planning Competition that relies on the Siege SAT solver [28]. The instances and domains are all from the 2nd and 3rd Int. Planning Competitions [29, 30], and the results have been obtained using a Pentium IV machine running at 2.8Ghz, with 1Gb of RAM, under Linux. The time limit for each problem is 30 minutes, and all times include preprocessing. Since FF and SATPLAN04 cannot handle temporal domains, we consider only the formulation in which all actions have unit duration. The bound B on the makespan for suboptimal eCPT is then set to 200 which is well above the optimal makespan in these benchmarks.

Table 1 shows for each domain, the total number of instances, the number of instances solved by eCPT, the number of instances solved backtrack free (and in parenthesis, the max number of backtracks over problems solved with backtracks), and the max number of nodes generated (in POCL planning, this number is different than the number of actions in the plan). For illustration purposes, the number of instances solved and the corresponding max number of nodes generated are reported also for FF [27]. *As it can be seen, eCPT solves 339 out of 350 instances, 336 of them backtrack free, including all the instances of Blocks, Ferry, Logistics, Gripper, Miconic, Rovers and Satellite* (the 11 unsolved instances are actually all caused by memory limitations in the Claire language rather than time). This is quite remarkable; these are instances that were challenging until very recently. eCPT solves actually 3 instances more than FF over this set of problems, eCPT having best relative coverage in Blocks and DriverLog, and FF in Depots and Zeno. In the last domain from IPC-3, Freecell, FF solves more instances than eCPT, which no longer exhibits a backtrack-free behavior.

Table 1. eCPT vs. FF: Coverage over various simple domains, showing # problems solved, backtrack free (max # backtracks), and max # of nodes generated.

	#pbs	eCPT			FF	
		solved	b.-free (max b.)	max nd	solved	max nd
blocks	50	50	50 (0)	275	42	146624
depots	20	18	16 (4)	285	19	166141
driver	20	17	16 (5)	176	15	4657
ferry	50	50	50 (0)	1176	50	201
gripper	50	50	50 (0)	201	50	200
logistics	50	50	50 (0)	273	50	2088
miconic	50	50	50 (0)	131	50	76
rovers	20	20	20 (0)	207	20	3072
satellite	20	20	20 (0)	249	20	5889
zeno	20	14	14 (0)	70	20	933

This domain, however, causes difficulties to FF as well due to the presence of dead-ends [31].

Not all the new inference rules are needed to generate the backtrack free behavior in every domain; yet Impossible Supports appears to be critical for Depots, Distance Boosting for Depots, DriverLog and Ferry, Qualitative Precedences for all domains except Blocks, and Action Landmarks for Blocks. In addition, often some disjunctions of rule sets are critical as well. For example, while either Qualitative Precedences or Unique Supports can be removed in Blocks without generating additional backtracks, the removal of both sets of rules does cause backtracks. Also, the modified heuristics are crucial for all domains.

Information about the runtime of eCPT over the various domains can be seen in Figure 2, with Table 2 providing additional details for selected instances in comparison with FF. As it can be seen, the runtimes for eCPT tend to scale well although they do not compete with the runtimes of FF (except for a few Depots instances): FF generates many more nodes but does so faster. Plan quality measured in the number of actions in the plan is better for FF in domains like Logistics or DriverLog, which may have to do with the fact that eCPT computes concurrent plans.

The scatter plots in Figures 3 and 4 compare eCPT respectively with CPT and SATPLAN04 as *optimal* planners; dots above (below) the diagonal indicate instances where eCPT is faster (resp. slower), while dots on the right (top) border are unsolved by eCPT (resp. the other planner). eCPT solves 207 out of the 350 instances, while CPT and SATPLAN04 solve 179 and 180 instances respectively. eCPT generates many fewer nodes than CPT, often running orders-of-magnitude faster (although not always so as the additional overhead does not always pay off).

As *suboptimal planners,* eCPT solves 339 out of the 350 instances, while CPT solves only 66 instances, with runtimes often above 1000 seconds, resulting usually in poor plans. This is because, as mentioned above, CPT behavior degrades

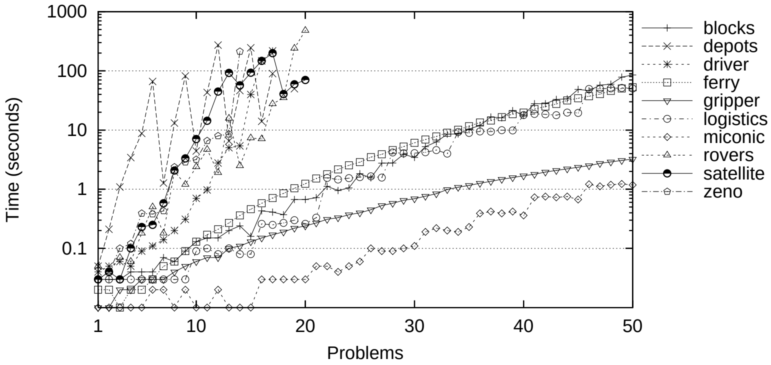

Fig. 2. eCPT running times on all domains.

Table 2. eCPT vs. FF: further details on a few instances.

	CPU time (sec.)		Actions		Nodes	
	eCPT	FF	eCPT	FF	eCPT (bkts)	FF
bw-ipc48	59.51	-	74	-	281 (0)	-
bw-ipc49	78.37	-	80	-	282 (0)	-
bw-ipc50	85.09	0.02	88	86	235 (0)	195
log-ipc48	50.56	0.20	164	142	261 (0)	515
log-ipc49	51.54	0.50	176	171	273 (0)	1252
log-ipc50	50.39	0.43	161	154	245 (0)	1147
depots06	66.23	-	68	-	160 (0)	-
depots07	1.27	0.01	28	25	68 (0)	142
depots08	13.13	579.89	75	43	206 (0)	172478
driver14	5.40	0.09	48	45	75 (0)	1209
driver15	39.91	0.03	69	44	130 (0)	161
driver16	147.15	-	107	-	163 (5)	-

quickly as the bound on the makespan is pushed above the optimal value leaving the problems unconstrained and the search unfocused.

5 Discussion

The task of solving simple planning problems in a *domain-independent* way with no search, by performing low polynomial operations in every node, is a crisp and meaningful goal, which may yield insights that an exclusive focus on performance may not, like the identification of broader tractable planning fragments, and the process by which people actually plan. In this work we have shown that a suitable extension of the temporal planner CPT achieves this behavior over a wide range of benchmark domains. The new constraints and inference mechanisms have been obtained from observing the behavior of CPT over various domains.

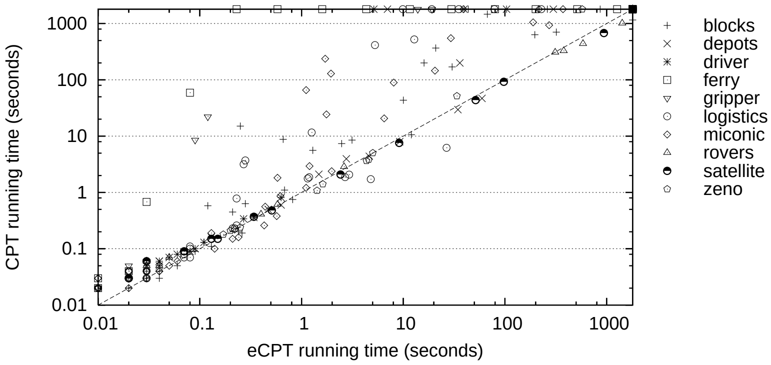

Fig. 3. eCPT vs. CPT for optimal planning.

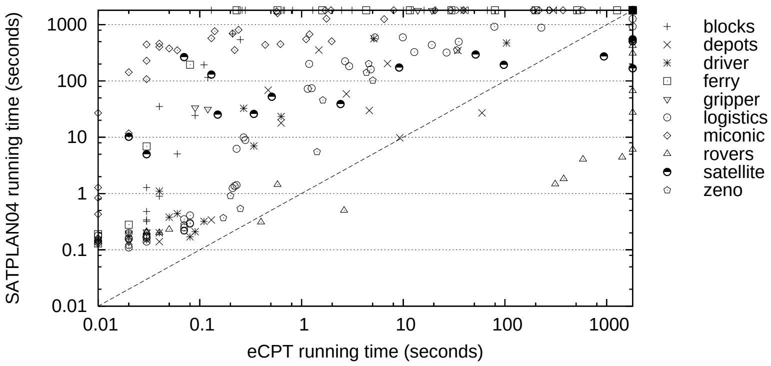

Fig. 4. eCPT vs. SATPLAN04 for optimal planning.

The fact that this fine grain analysis is possible, and that the results can be easily incorporated into the planner, is a clear benefit of the CP formulation, which thus provides a way for making use of (human) *domain-specific* analysis for improving the performance of a *domain-independent* planner. We have also empirically evaluated the resulting planner eCPT, as a suboptimal and optimal planner, and have shown significant gains over CPT.

The finding that a few *inference rules* is all that it takes to render the search backtrack free in domains which until recently were considered challenging for planners, bears some similarity with the *empirical observation* in [2] that a simple domain-independent *heuristic function* can effectively guide the search for plans in many domains, an idea exploited in many current planners. The two devices for taming the search, however, are different: heuristic estimators provide *numeric* information to weight alternatives, the inference rules provide *structural* information to discard alternatives. We believe that it should be possible to *prove* some domains backtrack free for eCPT, and in this way identify new abstract classes of tractable problems. Current classes, as defined in [3–6], remain somewhat narrow, and do not account for the tractability of existing benchmarks [32]. In the future, we want to investigate the gap between the empirical results and current theoretical tractability accounts, and explore the possibility of obtaining the behavior of eCPT from a *general inference engine* and not a customized implementation.

Last but not least, we have recently studied in detail the traces for some of the problems considered, and noticed that in several cases, some of the decisions taken are the result of ties. Since we are interested in making the planning process transparent such ties pose a problem for justifying the decisions that are taken. We decided then to explore *all* the branches that are tied, rather than selecting the first branch only, hoping that all tied branches would lead to backtrack-free solutions. We discovered however that in some problems, this was not the case; namely, in some cases the way in which the code was breaking ties mattered, and in such cases, this had to do, for example, with the order of the actions in the PDDL file. Our next goal is thus not only to extend the range of domains that are solved backtrack-free but to do so in a *robust* way, meaning that decisions that are found to be equally good according to the criteria that are *explicit* in eCPT, should all lead equally well to the solution (namely backtrack free). This requires refining the rules and the selection criterion in eCPT still further. The most recent results that we have obtained, suggest that this is possible too.

Acknowledgments

V. Vidal is partially supported by the "IUT de Lens", the CNRS and the "région Nord/Pas-de-Calais" under the COCOA program. H. Geffner is partially supported by Grant TIC2002-04470-C03-02, MCyT, Spain.

References

1. McDermott, D.: A heuristic estimator for means-ends analysis in planning. In: Proceedings of AIPS-96. (1996) 142–149
2. Bonet, B., Loerincs, G., Geffner, H.: A robust and fast action selection mechanism for planning. In: Proceedings of AAAI-97, MIT Press (1997) 714–719
3. Bylander, T.: The computational complexity of STRIPS planning. Artificial Intelligence **69** (1994) 165–204
4. Bäckström, C., Nebel, B.: Complexity results for SAS$^+$ planning. Computational Intelligence **11** (1995) 625–655
5. Jonsson, P., Bäckstrom, C.: Tractable planning with state variables by exploiting structural restrictions. In: Proceedings of AAAI-94. (1994) 998–1003
6. Brafman, R., Domshlak, C.: Structure and complexity of planning with unary operators. JAIR **18** (2003) 315–349
7. McAllester, D., Rosenblitt, D.: Systematic nonlinear planning. In: Proceedings of AAAI-91, Anaheim, CA, AAAI Press (1991) 634–639

8. Vidal, V., Geffner, H.: Branching and pruning: An optimal temporal POCL planner based on constraint programming. In: Proceedings of AAAI-2004. (2004) 570–577
9. Kautz, H., Selman, B.: Unifying SAT-based and Graph-based planning. In Dean, T., ed.: Proceedings of IJCAI-99, Morgan Kaufmann (1999) 318–327
10. Rintanen, J.: A planning algorithm not based on directional search. In: Proceedings of KR'98, Morgan Kaufmann (1998) 617–624
11. Do, M.B., Kambhampati, S.: Solving planning-graph by compiling it into CSP. In: Proceedings of AIPS-00. (2000) 82–91
12. Vidal, V., Geffner, H.: Branching and pruning: An optimal temporal POCL planner based on constraint programming (long version). Technical report (2004)
13. Porteous, J., Sebastia, L., Hoffmann, J.: On the extraction, ordering, and usage of landmarks in planning. In: Proceedings of ECP-01. (2001) 37–48
14. Zhu, L., Givan, R.: Heuristic planning via roadmap deduction. In: 4th. Int. Planning Competition Booklet (ICAPS-04). (2004)
15. Van Beek, P., Chen, X.: CPlan: a constraint programming approach to planning. In: Proceedings AAAI-99. (1999) 585–590
16. Kambhampati, S., Knoblock, C., Yang, Q.: Planning as refinement search: A unified framework for evaluating design tradeoffs in partial-order planning. Artificial Intelligence **76** (1995) 167–238
17. Weld, D.S.: An introduction to least commitment planning. AI Magazine **15** (1994) 27–61
18. Laborie, P., Ghallab, M.: Planning with sharable resources constraints. In Mellish, C., ed.: Proceedings of IJCAI-95, Morgan Kaufmann (1995) 1643–1649
19. Joslin, D., Pollack, M.E.: Is "early commitment" in plan generation ever a good idea? In: Proceedings of AAAI-96. (1996) 1188–1193
20. Penberthy, J.S., Weld, D.S.: Temporal planning with continous change. In: Proceedings of AAAI-94. (1994) 1010–1015
21. Jonsson, A., Morris, P., Muscettola, N., Rajan, K.: Planning in interplanetary space: Theory and practice. In: Proceedings of AIPS-2000. (2000) 177–186
22. Blum, A., Furst, M.: Fast planning through planning graph analysis. In: Proceedings of IJCAI-95, Morgan Kaufmann (1995) 1636–1642
23. Smith, D., Weld, D.S.: Temporal planning with mutual exclusion reasoning. In: Proceedings of IJCAI-99. (1999) 326–337
24. Haslum, P., Geffner, H.: Heuristic planning with time and resources. In: Proceedings of European Conference of Planning (ECP-01). (2001) 121–132
25. Nguyen, X.L., Kambhampati, S.: Reviving partial order planning. In: Proceedings of IJCAI-01. (2001) 459–466
26. Marriot, K., Stuckey, P.: Programming with Constraints. MIT Press (1999)
27. Hoffmann, J., Nebel, B.: The FF planning system: Fast plan generation through heuristic search. Journal of Artificial Intelligence Research **2001** (2001) 253–302
28. Ryan, L.: Efficient algorithms for clause-learning sat solvers. Master's thesis, Simon Fraser University (2003)
29. Bacchus, F.: The 2000 AI Planning Systems Competition. Artificial Intelligence Magazine **22** (2001) 47–56
30. Long, D., Fox, M.: The 3rd international planning competition: Results and analysis. Journal of Artificial Intelligence Research **20** (2003) 1–59
31. Hoffmann, J.: Local search topology in planning benchmarks: An empirical analysis. In: Proc. IJCAI-2001. (2001) 453–458
32. Helmert, M.: Complexity results for standard benchmark domains in planning. Artificial Intelligence **143** (2003) 219–262