Developing Domain-Independent Search Control for Europa2

Sara Bernardini

Istituto per la Ricerca Scientifica e Tecnologica (IRST)
Via Sommarive 18, 38055 Trento, Italy
bernardini@itc.it

David E. Smith

NASA Ames Research Center
Moffet Field, CA 94035–1000
desmith@arc.nasa.gov

Abstract

In the last few years, classical planners have achieved impressive results due to the development of problem relaxation techniques for computing distance estimates. In contrast, many large temporal planning systems used for practical applications have not benefitted from these techniques. Instead, these systems rely on careful engineering of the domain knowledge, together with carefully crafted domain-dependent control information. In this paper, we explain some of the characteristics of NASA's EUROPA2 planning system that make it difficult to directly apply the heuristic techniques developed for classical planning. However, we then borrow ideas from some of these methods to develop domain-independent heuristic techniques for EUROPA2. We show some promising initial results concerning their effectiveness.

Introduction

In the last decade, there have been significant improvements in the performance of automated planning systems. Key to this improvement has been the development of domain-independent heuristic techniques for estimating the distance between states and goals. Generally, these techniques rely on automatically generating a relaxed formulation of the planning problem and using a solution of this relaxed problem as a distance estimate. One popular method for doing this is to generate a plangraph (Blum & Furst 1997), extract a relaxed plan from it, and use the cost of this solution as the distance estimate (Hoffmann & Nebel 2001).

In contrast, many planners used for real-world applications, such as EUROPA (Frank & Jonsson 2003), ASPEN (Chien *et al.* 2000) and IxTeT (Ghallab & Laruelle 1994), have not benefitted significantly from these advances. Instead, these systems rely on careful engineering of the domain together with hand-crafted domain-dependent search control information. This process is generally quite painful, time consuming, and can lead to models that are not very robust to small changes in the domain or in the nature of the problems being solved. It would therefore be quite useful and desirable to have powerful domain-independent control techniques for these planning systems. Unfortunately, there are difficulties involved in doing this: the representation languages for these systems are quite different, allowing

much more complex temporal and metric constraints, and the search strategies employed by these systems cannot be characterized as either simple progression or regression. As a result, it is difficult to directly map the techniques from classical planning systems to these application systems.

In this paper, we develop novel domain-independent heuristic guidance techniques for the EUROPA2 planning system, currently being used for several NASA mission applications including MAPGEN, the ground-based daily activity planning system for the Mars Exploration Rover mission (MER) (Bresina *et al.* 2005). This planner has been shown to be extremely successful in solving complex real-world problems by providing the user with a powerful modeling language as well as a highly customizable solving engine. Nevertheless, EUROPA2 suffers from having little or no effective domain-independent heuristic guidance. Our technique borrows ideas from the work of Haslum and Geffner (2000), and Helmert (2006). In particular, we build transition graphs for the different state variables in EUROPA2, and use these graphs to compute distance estimates for choosing and resolving flaws within EUROPA2's plan refinement mechanism.

In order to explain our technique it is necessary to have some understanding of the EUROPA2 planning paradigm and search algorithm. We give a quick overview of the essentials in the next two sections. We then describe our technique for automatically deriving domain-independent heuristic estimates. We conclude by presenting some preliminary experimental results.

EUROPA2: Paradigm and modeling language

For EUROPA2, planning domains and problems are described using a declarative modeling language called NDDL (New Domain Definition Language). A **_planning domain_** $\mathcal{D}$ in NDDL is represented by the following elements:

- A set of **_timelines_**: $\mathrm{T} = \{\mathcal{T}_1, \mathcal{T}_2, \ldots, \mathcal{T}_n\}$, which are essentially variables capturing the evolution of a quantity or component over time
- A set of mutually exclusive **_activities_** associated with each timeline $\mathcal{T}_i$: $\mathcal{A}ct[\mathcal{T}_i] = \{a_1(\vec{x_1}, \delta_1), \ldots, a_n(\vec{x_n}, \delta_n)\}$ where $\vec{x}$ is the vector of the activity's parameters and $\delta = [\delta_{min}, \delta_{max}]$ is a mathematical interval in $\mathbb{N}$ representing the duration of the activity

- A conjunction of *temporal constraints* associated with each activity a_i: $\mathcal{C}[a_i] = c_1 \wedge \ldots \wedge c_n$, where a conjunct c_j can assume one of the following two forms:

  - $c_j = a_i$ temporal_relation a_k
    Such a conjunct is called a ***compatibility***. It is a qualitative (meets, met_by, etc.) or quantitative temporal constraint between the activity a_i and any other activity a_k belonging to the same timeline $\mathcal{T}$ or to another timeline $\mathcal{T}_h$. We talk about *internal* and *external* compatibilities and we indicate them as $\mathcal{C}^{\mathrm{I}}[a_i]$ and $\mathcal{C}^{\mathrm{E}}[a_i]$). The activity a_i is called *master* and the activity a_k is called a *slave*.

  - $c_j = (\text{case } \gamma = 1 : \mathcal{C}_1^-); \ldots ; (\text{case } \gamma = m : \mathcal{C}_m^-)$
    A conjunct can also correspond to a choice between different conjunctions of compatibilities $\mathcal{C}_h^-$, where $\mathcal{C}_h^- = \{c_1 \wedge \ldots \wedge c_l\}$ and $c_i = a_i$ temporal_relation a_k. The choice between the conjunctions is regulated by the variable γ, which is called the ***guard*** of the case.

A rich set of temporal relationships is permitted in compatibilities, including: *equal*, *meets*, *contains*, *after*, *starts*, *overlaps* and all their inverse relations. These relations are similar to the thirteen temporal relations defined by Allen (Allen 1983)[1].

The above representation differs from a PDDL domain description in several respects: 1) it uses a variable/value representation (timelines/activities) rather than a propositional representation, and 2) there is no concept of state or action, only of activities and constraints between them. In this respect, models in NDDL look more like the schemas for SAT encodings of planning problems than PDDL models.

A ***planning problem*** $\mathcal{P}$ is represented by a pair $\mathcal{P} = \{\mathrm{H}, \mathcal{I}\}$ where:

- $\mathrm{H} \in \mathbb{N}$ is the end of the ***planning horizon***, meaning that we only care about the behavior of the system with respect to the temporal window $[0, \mathrm{H}]$.
- $\mathcal{I}$ is the ***initial configuration*** represented by a set of activities placed on their corresponding timelines. If we annotate an activity a by a time interval $\tau(a) = [st(a), et(a)]$ (indicating the temporal extent over which a holds), then, for each activity a_i in $\mathcal{I}$, it is possible either to specify the specific position of a_i on the timeline, that basically means fixing the start and end time of $\tau(a_i)$, or to leave a_i floating on the timeline between the origin and the horizon of the time axis.

The initial configuration $\mathcal{I}$ corresponds to both the initial state and the goal state as they are defined in classical planning. The activities in $\mathcal{I}$ that are placed at the beginning of the horizon correspond to the traditional initial state, while all the others generalize the classical notion of goal since

---

[1]It is worth noting that, although the temporal relationships defined by Allen are non-directional and can be inverted, compatibilities cannot be inverted. For example, the compatibility a_i meets a_k has a different semantics from a_k met_by a_i. The former means that if a_i exists on a timeline, a_k must also exist, while the latter means that if a_k exists on a timeline, a_i must exist.

they can be placed not only at the end of the horizon, but also in any other position. The initial configuration is also the initial partial plan that is turned into a final plan by the planning procedure. A final *plan* π is the configuration where all the timelines of the domain are fully covered by contiguous activities from the start to the end of the horizon. Some of the activities are those that appear in the initial configuration, the others are triggered by the applicable compatibilities associated with the initial activities and those that are incrementally added to the plan. In fact, a compatibility c that involves two activities a_i and a_k imposes that, once a_i has been chosen to be part of the plan and placed on its proper timeline $\mathcal{T}$, then the activity a_k must be necessarily introduced in the plan and placed over its timeline in such a way that the temporal contraint stated by c is satisfied. If an applicable temporal constraint contains a disjunction, that means that, once the master activity has been chosen to be a part of a plan, then at least one of its disjunctive slaves must be part of the plan. A plan is *consistent* when all the timelines are fully covered by activities, all the temporal contraints involving those activities are satisfied and all the variables are instantiated.

An example. As an illustration of a simple NDDL domain model, consider a rover equipped with a set of instruments to sample a geological site. We model the following subsystems as timelines: *Battery*, *Navigator*, *Controller*, *Instrument*$_1$, $\ldots$, *Instrument*$_n$. Each subsystem can only perform certain activities. An instrument, for example, can perform one of the following operations: *TakeSample(rock, 1)*, *Place(rock, 3)*, *Stow(2)*, *Unstow(2)* and *Stowed([1, +inf])*. The first activity consists in taking a sample of a rock at the site and lasts 1 time unit. The other specifications are similar. The constraints that regulate the behavior of an instrument are the following: in order to take a sample of a rock, the instrument must be first unstowed and then properly positioned in the vicinity of the rock. After taking the sample, the instrument can be placed in another position for performing another experiment or can be stowed. Those constraints are expressed by means of the internal compatibilities. We show just a few of them:

- *Unstow()* meets *Place(rock*$_i$*)*
- *Place(rock*$_i$*)* meets *TakeSample(rock*$_i$*)*
- case $\gamma = 0$: *TakeSample(rock*$_i$*)* meets *Stow()*;
  case $\gamma = 1$: *TakeSample(rock*$_i$*)* meets *Place(rock*$_j$*)*

Note that in this model, we have made no attempt to model or allow intermediate state between *Unstow*, *Place*, and *TakeSample* operations – something that is not possible in PDDL.

The external compatibilities for *Instrument*$_k$ govern its interactions with the activities on other timelines. For example, the instrument can take a sample of a rock only if the navigator has already reached that rock and persists in that position while the instrument is taking the sample:

Instrument$_k$.*TakeSample(rock*$_i$*)* contained_by
Navigator.*At(rock*$_i$*)*.

An initial configuration $\mathcal{I}$ for the rover domain can, for example, specify the level of the battery, the position of the

navigator and the status of the instruments at the start of the planning horizon and, furthermore, can establish that a sample of a particular rock should be taken within a certain time interval.

EUROPA2: search algorithm

The planning algorithm at the core of EUROPA2 can be thought of as an instance of *plan refinement search*; given a domain $\mathcal{D}$ and a problem $\mathcal{P}$, the algorithm starts from the initial configuration $\mathcal{I}$ and incrementally refines it by adding activities to the timelines, ordering those activities and binding variables until a final consistent configuration is found. This algorithm can also be seen as a search in the *space of partial plans* (McAllister & Rosenblitt 1991), where a ***partial plan*** Π consists of the following elements:

- For each timeline $\mathcal{T} \in \mathcal{D}$, a set of activities $\mathcal{A}ct_\Pi = \{t_1, t_2, \ldots, t_n\}$, which are not necessarily contiguous on time (actions in POCL)
- A ***temporal network*** $\mathcal{N}_\Pi$ representing all the start and end times of the activities in the plan and the constraints between them (ordering constraints in POCL).
- A set of ***flaws*** $\mathcal{F}_\Pi = \{f_1, f_2, \ldots, f_m\}$, where a flaw is an indication of a potential inconsistency in the partial plan. There are three types of flaws:
  - ***Open condition flaws***: They arise when applicable compatibilities are applied, triggering activities as slaves of masters that are already in the plan Π. Those slave activities are enforced to be part of the plan, but they are not yet associated with any timeline. We call them *free activities* (open preconditions in POCL).
  - ***Ordering flaws***: They arise anytime an activity is placed on a timeline and an ordering is required for the activity with respect to the other activities already on that timeline (threats in POCL).
  - ***Unbound variable flaws***: They arise when variables that have not yet been instantiated appear in the plan Π. Those variables are said to be *unbound*. There are two kinds of unbound variables: parameters of activities that are already in the plan and guards of applicable temporal constraints.

Refining a partial plan means to pick a flaw and resolve it. The process terminates when the set of flaws is empty. Each kind of flaw is solved in a different way.

- *Resolvers for open condition flaws*

  Flaws corresponding to free activities can be resolved in two ways:
  - ***Merging*** A free activity is merged with a matching activity already in the plan (similar to add-link in POCL planning). The operation of merging does not result in the addition of any new flaws to the current plan. An activity a is said to match an activity a' if a and a' unify and the temporal constraints involving a are satisfied by a'. Thus, a and a' can be considered the same activity and we do not need to introduce a in the plan. Consequently, the slaves of a are not fired, because they have been already triggered when a' was introduced in the plan.

  - ***Activation*** We introduce a new activity a in the current plan associating it with the proper timeline, but without choosing a specific time slot for it (similar to add-step in POCL planning). The compatibilities associated with a are applied and the slaves fired by those compatibilities are introduced as free activities. This results in both an ordering flaw, corresponding to the just activated activity, and a number of open condition flaws, corresponding to the added slaves.

- *Resolvers for ordering flaws*

  Once we have decided to place a new activity on a timeline, we need to choose where to put it with respect to the other activities already on that timeline. For this purpose, the temporal constraints involving the new activity are checked against the current temporal network. An ordering flaw is resolved by imposing ordering constraints among activities in such a way that the temporal network remains consistent and all the constraints are satisfied.

- *Resolver for unbound variable flaws*

  Unbound variable flaws are resolved by specifying a value in the domain of the variable. If the variable is a guard, the binding causes the introduction in the current plan of the slave activities associated with the chosen value.

The basic algorithm in EUROPA2 is a *depth-first search* characterized by ***flaw selection***, ***flaw resolution*** and ***constraint propagation*** steps. Flaw selection identifies which flaw to resolve next. This is not a backtracking point, but, like variable ordering in constraint satisfaction, has a significant impact on the amount of search and backtracking required to find a solution. Flaw resolution deals with resolving a flaw by subsequently trying all the resolution options (activation and merging for open condition flaws, various activity orderings for ordering flaws and possible variable bindings for unbound variable flaws). This is a backtracking point because if a resolution option does not work, the algorithm tries another option until all options are exhausted. Operations of plan refinement are interleaved with *constraint propagation* on the constraint network underlying the current partial plan. Contraint propagation is mainly used to test partial plans for consistency, and discovers dead ends, which are either inconsistent partial plans or partial plans with flaws that cannot be solved. However, it also plays another major role: it provides the algorithm with a *look-ahead capability* that allows it to filter away infeasible flaw resolvers before the algorithm actually commits to them.

The search control problem

Through a combination of careful domain engineering and the crafting of domain-dependent search control information, a user can customize and control search, flaw selection and flaw resolution in EUROPA2. However, this process is painful, time consuming, and often leads to models that are not robust to further enhancements or changes. If EUROPA2 is run in the absence of domain-dependent heuristics, it inevitably experiences serious control problems. Plans are not found within a reasonable amount of time even for problems that are trivial for IPC planners.

As mentioned in the introduction, there has been considerable work in the classical planning community on devising domain-independent heuristics. Generally, these techniques involve solving some relaxed form of the planning problem in order to obtain heuristic distance estimates, which are then used to guide search. Simple but effective ways to obtain relaxed problems are, for instance, ignoring PDDL operator delete lists or decomposing the goal set of atoms into smaller subsets (Bonet & Geffner 2001), (Haslum & Geffner 2000). The computation of some of those heuristics rely on the explicit construction of a *reachability graph* (Blum & Furst 1997), (Hoffmann & Nebel 2001), while other methods perform *shortest path* calculations on a suitably defined atom space (Haslum & Geffner 2000). In addition to "distance-based" heuristics, other techniques have been proposed that work on multi-valued representations of planning problems. Fast Downward (Helmert 2006) extracts a heuristic function by constructing a *causal graph* of the domain and a *domain transition graph* for each state variable in the domain. The fist graph represents the critical interactions between state-variables, while the second graph describes the dependancies between the values of a single state variable.

Although plangraph distance estimates have been used effectively to guide POCL planners like RePop (Nguyen & Kambhampati 2001) and VHPOP (Younes & Simmons 2003), to date, EUROPA2 has not benefited from any of these techniques. There are several reasons for this, including: the variable/value representation, the lack of distinction between state and action, the lack of distinction between fact and goal, the lack of causality in the compatibilities, the large number of exogenous events and time constraints in many practical problems, and the bidirectional nature of the search strategy (which appears essential for domains involving many time constraints and exogenous events). All these factors make it difficult to directly map existing domain-independent search control strategies to EUROPA2. (Similar issues exist for other temporal planners like ASPEN (Chien *et al.* 2000) and IxTeT (Ghallab & Laruelle 1994).)

In the next section, we develop a domain-independent control strategy for EUROPA2 that builds on the ideas of distance-based estimations presented in (Haslum & Geffner 2000) and the construction of transition graphs described in (Helmert 2006).

A search control strategy for EUROPA2

In order to effectively guide search in EUROPA2, we need a method of assessing the impact of each possible flaw resolution on the cost of completing a partial plan. To do this, we build a set of transition graphs and use these graphs during planning to do distance estimation. More specifically, we construct a graph for each timeline in the domain, describing the possible *transitions* between the activities on that timeline. The nodes in the graph represent activities and the transitions are induced by the information available in the compatibilities for the activities participating in the transition. Constructing a useful transition diagram for a timeline essentially requires that we reconstruct the causality hidden

in the compatibilities for the different activities. This turns out to be a non-trivial task. For the sake of brevity, we do not discuss the topic here. Furthermore, a cost is associated with a transition that identifies the *temporal distance* between the activities involved in the transition. We calculate the cheapest path from any activity to any other activity by running an *all-pairs shortest path algorithm* on each graph.

Given a domain $\mathcal{D}$ and a timeline $\mathcal{T} \in \mathcal{D}$, the **Activity Transition Graph** for $\mathcal{T}$ is a *directed weighted graph* $\mathcal{G}^A[\mathcal{T}] = \{V, E, \mathcal{L}_E\}$, where V is the set of vertexes, E the set of edges and $\mathcal{L}_E$ is a weight function that assigns a numeric weight to each edge in the graph. The graph is developed as follows:

- We create a node $v \in V$ for each grounded activity a that can appear on $\mathcal{T}$.
- For each activity a that belongs to $\mathcal{T}$, we examine the internal and external compatibilities for a, respectively $\mathcal{C}^I[a]$ and $\mathcal{C}^E[a]$, with the purpose of defining the transitions in the graph. In particular, the set $\mathcal{C}^I[a]$ will specify the edges appearing in the transition diagram, while the set $\mathcal{C}^E[a]$ will dictate additional *conditions* on those edges.
- Consider the set $\mathcal{C}^I[a]$ of internal compatibilities for a. All of these compatibilities must be either generalized `meets` or `met_by`, because activities on the same timeline cannot overlap. We define the transitions into a and out of a as follows:

  - The possible transitions out of a are described by `meets` compatibilities: for each $c \in \mathcal{C}^I[a]$ such that $c = a$ `meets` a', we add a directed edge $e \in E$ between the node corresponding to a and the node corresponding to a'. The edge e is labelled with the lower bound of the duration δ of the activity a.
  - The possible transitions into a are described by `met_by` compatibilities: for each $c \in \mathcal{C}^I[a]$ such that $c = a$ `met_by` a', we add a directed edge $e \in E$ between the node corresponding to a' and the node corresponding to a. The edge e is labelled with the lower bound of the duration δ of the activity a'.

  Note that there may be more than one edge into or out of a because of the presence of unbound guards in the specification of the temporal constraints involving a.

- We now consider the external compatibilities $\mathcal{C}^E[a]$ for a and divide them into two further categories:

  - `meets`, `starts` and `contains` compatibilities, which specify that the activity a must start at a particular time at or before the start of another activity a'. We will assume that these compatibilities are describing "side effects" of the activity a and we will ignore the compatibilities in this category. (These side effects might cause interference with the behavior of other timelines, but we neglect this point here.)
  - `met_by`, `ends` and `contained_by` compatibilities, which specify that the activity a must start after the start of another activity a', or that only specify that a must start after some particular time. We will assume that these compatibilities describe "requirements" for

the activity a.[2] For such compatibilities, we do not add any edge in the graph, but we keep track of them by associating a set of *conditions* with the appropriate incoming edge for the node representing a (we call the set $\mathcal{C}ond(a)$).

For the rover example, Fig. 1 shows the activity transition graph for the *Instrument$_i$* timeline, assuming there are only two rocks in the domain.
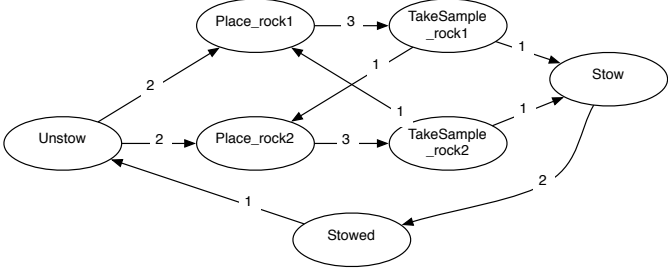

Figure 1: Activity Transition Graph for *Instrument$_i$*

Given an activity transition graph $\mathcal{G}^{\mathrm{A}}[\mathcal{T}]$ for a timeline $\mathcal{T}$, we define $\mathcal{C}ost_{\mathrm{SP}}(a_1, a_2)$ to be the cost of the shortest path between a_1 and a_2 in the graph. Using an ***all pairs shortest-path*** algorithm we can precompute and store this information for each timeline prior to beginning planning.

We now consider how to make use of this information to do flaw resolution. Consider a partial plan Π with an open condition or ordering flaw f, and suppose that f has a possible resolution r. We define the cost of the resolution, $\mathcal{C}_{res}(r)$, as follows for merging and placement:

- *Merging* the activity a with some existing activity a' on the timeline $\mathcal{T}$:
$$\mathcal{C}_{res}(r) = 0$$

- *Placing* the activity a in an empty slot s on the timeline $\mathcal{T}$. The activity a can be compatible with more than one empty slot on $\mathcal{T}$. Given one of those empty slots s, the activity a_i preceding the slot s, and the activity a_{i+1} following the slot s (see Fig. 2):
$$\mathcal{C}_{res}(r) =$$
$$\mathcal{C}ost_{\mathrm{SP}}(a_i, a) + \mathcal{C}ost_{\mathrm{SP}}(a, a_{i+1}) - \mathcal{C}ost_{\mathrm{SP}}(a_i, a_{i+1})$$

The first definition corresponds to the intuition that the operation of merging has little cost, since it does not modify the partial plan except for adding new temporal constraints. Moreover, it narrows the current set of flaws while not adding any new flaws. The second definition estimates how well the activity a fits in the empty slot s on $\mathcal{T}$. Without a, there is a cost $\mathcal{C}ost_{\mathrm{SP}}(a_i, a_{i+1})$ of going from the activity a_i preceding s to the activity a_{i+1} following s. By inserting a in the slot s, we instead incur the cost $\mathcal{C}ost_{\mathrm{SP}}(a_i, a)$ of getting from a_i to a, plus the cost $\mathcal{C}ost_{\mathrm{SP}}(a, a_{i+1})$ of getting

---

[2]The causality for contains and contained-by compatibilities in NDDL is not always clear. The contained interval could be a temporary *effect* of the containing activity, or it could be a *condition* that must hold in order for the containing activity to function as desired. For present purposes we will assume that the contained interval is an effect rather than a condition.

from a to a_{i+1}. The difference of these costs is an indication of the penalty incurred by placing a in the slot s. It represents the difference between the shortest path to go from a_i to a_{i+1} going through a and the direct shortest path from a_i to a_{i+1}. Clearly, if a is part of the direct shortest path, the measure is zero.


Figure 2: Placing the activity a on the timeline $\mathcal{T}$

If $\mathcal{R}[f] = \{r_1, \ldots, r_k\}$ is the set of possible resolutions for a flaw f, we define the ***Cheapest Local Resolution*** as:
$$\mathrm{CLR}(f) = min_{r_i \in \mathcal{R}[f]} \mathcal{C}_{res}(r_i)$$

By using the $\mathrm{CLR}(f)$ for a placement flaw f, we prefer to place the activity in a slot where it causes the smallest increase in the net cost for the timeline $\mathcal{T}$. This provides an initial good estimate of cost since it generally prefers merging (cost zero) to other possibilities, prefers slots with low cost paths to higher cost paths, and avoids slots where no transition is possible (infinite cost).

The above scheme is fairly simple because it does not consider the interactions between an activity a and activities on other timelines. In particular it neglects the requirements that must be satisfied on other timelines when placing a in a slot s. It also does not consider the *side-effects* that might result on other timelines by placing a in slot s. Omitting the side-effects is similar to "ignoring delete lists" used in many current planning systems, and we do not consider it further here. However, if we want to compute a better estimation of the cost of placing an activity on a timeline, we should consider the costs of the *conditions* that must be satisfied on the other timelines in order to make the placement possible. The information regarding conditions is available in the transition graph for the activity, since each edge is annotated with a set of conditions involving activities on other timelines (and hence appearing in other transition graphs). There are a number of possibilites for estimating the costs of satisfying these conditions. All those options basically try to estimate the cost of achieving a condition a' on a timeline $\mathcal{T}'$ by analyzing the transition graph for $\mathcal{T}'$ and calculating the $\mathrm{CLR}(f')$, where the flaw f' corresponds to the placement of a' on $\mathcal{T}'$. Two issues must be addressed when following conditions back to their transition graphs:

- *Duplication:* Conditions may be repeated for several edges along a shortest path, so we must avoid including the cost of a single condition more than once.

- *Recursion*: We could continue chasing back the conditions along the shortest path for each condition c_i, trying to get a better estimate of the cost of obtaining it. This process might never end, because conditions for achieving c_i might belong to the original timeline.

We have developed an algorithm for calculating costs of conditions that gets around those problems by first recursively collecting all the conditions into a set, and then adding up the

CLRs of the conditions. This approach avoids double counting and recursion because each condition can appear at most once in the set. Space limitations prevent us from presenting the details of the algorithm so we give only a sketch here. Given a flaw f for placing an activity a on a timeline $\mathcal{T}$, the process aims at collecting the set of all the conditions on all the timelines that should be satisfied in order to perform that placement. The final cost of the placement is then taken as the sum of the CLR of the flaw f (as before) plus the costs for this set of conditions. The set is developed by recursively going backward to the graphs of the conditions for a, finding the paths to achieve them, and unioning their conditions to the set, while taking particular care that no duplicates are added. Since there are a finite number of nodes and edges in the transition graphs of the domain, this process will terminate. By doing this, we are in essence collecting the entire set of steps (over all timelines) that are necessary in order to place a on $\mathcal{T}$. This set can be seen as a relaxed plan for a.

So far, we have presented the algorithm used by the flaw resolution procedure when it has to estimate the cost of resolvers for open condition and ordering flaws. Due to space limitations we do not give the details of the algorithm to treat *unbound variable flaws*, which is based on the same concepts and mechanisms that we have just described. In fact, choosing a resolver for an unbound variable flaw means choosing a value for a guard variable, which in turn corresponds to enforcing one set of compatibilities instead of another. In order to rank the different possible choices for a guard variable, we need to evaluate how difficult it is to achieve the compatibilities associated with that choice. Each compatibility will raise an open condition flaw or an ordering flaw and we have shown how to estimate the cost of resolving these kinds of flaws. Once we have the cost of each flaw triggered by binding the unbound variable with a certain value, we pick the value associated with the lowest cost and assign it to the guarded variable. As far as the *flaw selection* procedure is concerned, we can repeat the same argument. We have also considered other more traditional heuristics for flaw selection, such as choosing the variable with "Minimum Domain Size", but these heuristics resulted in very poor performance.

Implementation and Experimental Results

We have some preliminary experimental results for the proposed heuristics within EUROPA2. The current implementation is in C++ and the results were obtained using a Pentium IV machine running at 1.8GHz with 1Gb of RAM. Our current implementation includes pre-processing to construct the activity transition graphs and shortest-path tables, and the simple versions of the flaw resolution and flaw selection procedures. We have not yet fully implemented the more complex heuristics that recursively chain back through transition graphs to account for the conditions on graph edges. The unavailability of a benchmark set of domains written in NDDL makes performing experimentations within EUROPA2 very laborious, since domains and problems have to be manually provided. In order to carry out a more comprehensive evaluation of the performance

of EUROPA2, we are developing an automatic translator from PDDL2.1 to NDDL, building on the translator from PDDL2.1 to SAS$^+$ tasks presented in (Edelkamp & Helmert 1999) and (Helmert 2004). The translator will provide us with the opportunity to use the benchmark sets developed for the International Planning Competition. Although these results are preliminary, we present them as an indication of the fact that it is possible to successfully export key techniques developed by the classical planning community into a very different framework such as EUROPA2. In particular, we aim at showing that, if we introduce automatically derived heuristics into EUROPA2, it can work reasonably well on domains that are not specifically tailored to fit its features, without the use of hand written control rules.

We discuss the tests of the proposed heuristic on two standard domains: TOWER-n and LOGISTICS. Both these domains are particularly difficult for the standard version of EUROPA2, because they involve many causal disjunctive constraints and just a few simple temporal constraints. In our translation, an activity can correspond to either an action or a proposition. The constraints involving actions describe the conditions and the effects of actions on the behavior of the other activities. The constraints involving atoms basically express the different kinds of axioms that occur in SAT encodings of planning problems. For each activity, they explain under which circumstances that activity can be started and terminated. Frame axioms are particularly critical for EUROPA2 for two reasons. First, they introduce many disjunctions in the domain specification. Second, since EUROPA2 works bi-directionally, it can happen that a constraint explaining how an activity can be terminated is prematurely applied. That results in an early action commitment that is completely unmotivated with respect to the achievement of the goal. The proposed heuristic overcomes the two problems by postponing those kinds of constraints and binding disjunctive guards in an effective way.

We compare EUROPA2 with a closely related planner, CPT (Vidal & Geffner 2006), which was awarded distinguished performance in optimal planning for temporal domains at IPC-2006. CPT is based on a simple extension of the STRIPS language where concurrent actions with an integer duration are allowed. A contraint programming formulation is extracted from the initial problem specification. The domain theory is hence expressed in terms of variables, their domains and constraints corresponding to disjunctions, rules and temporal restrictions. The inference machinery over this CP formulation provides a powerful pruning mechanism for discarding partial solutions generated by a classical POCL branching schema. The novelty of CPT is the ability to perform inference not only on the actions already in the partial plan, but on all the actions in the domain. Constraint propagation in EUROPA2 offers some look-ahead capability, but not the full reachability analysis provided in CPT.

The TOWER-n domain deals with the construction of a tower made of n blocks $b_1, \ldots, b_n$. Eventually, the block b_1 should be on top and b_n on the table. We consider two different initial configurations: (1) all the blocks are on the table; (2) all the blocks are on the table, except for b_n which

is on top of b_1. If the original planner is run on those problems without the use of hand written control rules, it does not manage to find a plan within a time bounds of hours, even for instances with only three blocks. On the other hand, when we introduce the heuristic estimators, EUROPA2 performs extremely well. If we consider the initial configuration (1), a solution is found by pure inference and no search. In Figure 3, we show the performance of EUROPA2 on this problem considering instances from two to fifty blocks. CPT, like EUROPA2, does not rely on search to solve this problem and outperforms other planners such as BLACKBOX (Kautz & Selman 1999). Considering the initial configuration (2), EUROPA2 finds a plan with only shallow backtracking and the performance is comparable with that shown in Figure 3. Although this problem appears trivial for classical planners, powerful systems such as FF (Hoffmann & Nebel 2001) cannot solve it.
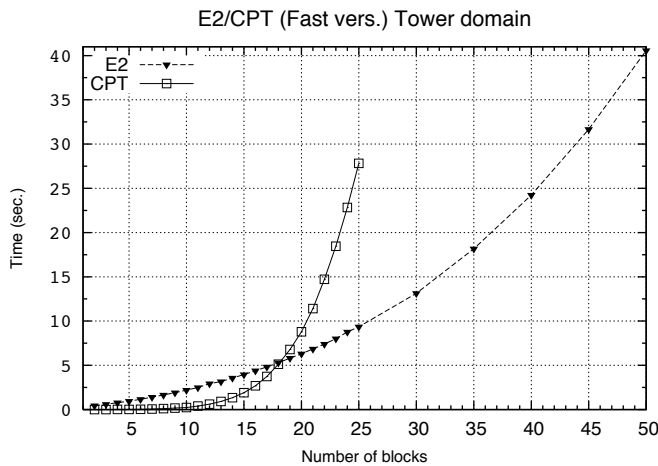

Figure 3: Results for the Tower domain

LOGISTICS is the classical problem of moving packages between different cities using trucks and planes. This problem differs from TOWER since it presents many independent subgoals, while TOWER is characterized by many dependent subgoals. The performance for LOGISTICS follows the same trend we showed for TOWER. The original version of EUROPA2 fails to find a plan with time bounds of hours for trivial instances, such as those involving three packages and two cities, while it succeeds in solving big instances without search when it uses the proposed heuristics.

Conclusions

We have developed novel domain-independent search control techniques for the EUROPA2 planning system. These techniques construct transition graphs for each timeline in the domain model and use these graphs to estimate the cost of resolving flaws in different ways. This information is used to guide both flaw selection, and flaw resolution. Although our experimental results are preliminary, they suggest that EUROPA2 can get by with far less domain-dependent guidance, and can successfully function as a general purpose engine if it makes use of these powerful domain-independent heuristics.

References

Allen, J. 1983. Maintaining knowledge about temporal intervals. *Communications of the ACM* 26(11):832–843.

Blum, A., and Furst, M. 1997. Fast planning through planning graph analysis. *Artificial Intelligence* 90:281–300.

Bonet, B., and Geffner, H. 2001. Planning as heuristic search. *Artificial Intelligence* 129((1-2)). Special issue on Heuristic Search.

Bresina, J.; Jonsson, A.; Morris, P.; and Rajan, K. 2005. Activity planning for the Mars Exploration Rovers. In *Proc. of the Fifteenth International Conference on Automated Planning and Scheduling (ICAPS-05)*, 40–49.

Chien, S.; Rabideau, G.; Knight, R.; Sherwood, R.; Engelhardt, B.; Mutz, D.; Estlin, T.; B.Smith; Fisher, F.; Barret, T.; Stebbins, G.; and Tran, D. 2000. ASPEN - Automated planning and scheduling for space missions operations. In *International Conference on Space Operations (SpaceOps 2000)*.

Edelkamp, S., and Helmert, M. 1999. Exhibiting knowledge in planning problems to minimize state encoding length. In *Proc. of the Fifth European Conference on Planning (ECP'99)*, 135–147.

Frank, J., and Jonsson, A. 2003. Constraint based attribute and interval planning. *Journal of Constraints* 8(4):339–364. Special Issue on Planning.

Ghallab, M., and Laruelle, H. 1994. Representation and control in ixtet, a temporal planner. In *Proc. of the Second International Conference on Artificial Intelligence Planning Systems (AIPS-94)*, 61–67. AAAI Press.

Haslum, P., and Geffner, H. 2000. Admissible heuristics for optimal planning. In *Proc. of the Fifth International Conference on Artificial Intelligence Planning and Scheduling (AIPS-00*, 140–149.

Helmert, M. 2004. A planning heuristic based on causal graph analysis. In *Proc. of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS-04)*, 161–170.

Helmert, M. 2006. The Fast Downward planning system. *Journal of Artificial Intelligence Research* 26:191–246.

Hoffmann, J., and Nebel, B. 2001. The FF planning system: fast plan generation through heuristic search. *Journal of Artificial Intelligence Research* 14:253–302.

Kautz, H., and Selman, B. 1999. Unifying sat-based and graph-based planning. In *Proc. of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI-99)*.

McAllister, D., and Rosenblitt, D. 1991. Systematic non-linear planning. In *Proc. of the Ninth National Conference on Artificial Intelligence (AAAI-91)*, volume 2, 634–639. AAAI Press.

Nguyen, X., and Kambhampati, S. 2001. Reviving partial order planning. In *Proc. of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI-01)*, 459–466.

Vidal, V., and Geffner, H. 2006. Branching and pruning: An optimal temporal pocl planner based on constraint programming. *Artificial Intelligence* 170(3):298–335.

Younes, H., and Simmons, R. 2003. Vhpop: Versatile heuristic partial order planner. *Journal of Artificial Intelligence Research* 20:405–430.