“First Have a Plan then Make Sure It Is a Good Plan” or Dealing with Underspecified Spatio-Temporal Relations in Unfamiliar Large-Scale Environments

Inessa Seifert
SFB/TR 8, University of Bremen
P.O. Box 330 440
28334 Bremen / Germany
seifert@informatik.uni-bremen.de

Abstract. This paper introduces an approach for dealing with underspecified spatial and temporal relations during early stage spatio-temporal planning in large-scale unfamiliar environments. For that purpose a partial spatio-temporal constraint satisfaction algorithm is introduced. The algorithm operates on a qualitative as well as metric representation of the given spatio-temporal relations. In order to reduce the computational complexity *structuring criteria* of the problem space are introduced. The paper outlines procedures for handling of overspecified spatio-temporal relations by removing of location assignments or fixed points in time, or shortening durations of events under consideration of the corresponding priorities. The algorithm concludes with ordering procedure for events with no temporal order but specified spatial assignments. The alternative spatial locations for events with underspecified spatial assignment are represented using a *qualitative distance scheme* or *region-based mapping* to the environment structure.

Introduction

Spatio-temporal planning is difficult and often annoying; especially in a large-scale unfamiliar environment which is experienced only seldom, e.g., an exhibition or a trade fair. The scientifically correct term for “difficult and annoying” is cognitively demanding. Imagine a situation where you have to plan a trade fair visit. Having a spatial layout of a trade fair and a list of events at hand try to answer the following questions:

• Do you have enough time to follow a hint of your boss to visit some particular exhibit?

• What can you do instead of listening to a scientific talk after you realized it’s been already full as you’ve arrived at the conference pavilion?

The illustrated problem domain comprises different events like presentations, tasks, meetings with people or inspections of trade fair booths. Such events are interrelated with each other through a set of spatio-temporal constraints: e.g. distances between the locations of booths and feasible temporal order of visits to the booths.

The main goal of a human spatio-temporal planning task is a compilation of a sequence of activities in a given environment. (cf. “…*human cognition is for action*”, Wilson, 2002). Ideally, such set of activities consists of a well-defined temporal order of events (e.g., e1 *before* e2, e2 *before* e3, etc.) and the corresponding locations. However, in the real world spatio-temporal events and relations between them are specified only to a certain degree, e.g., qualitatively, and go through a lot of changes especially during the early stage planning. Moreover, some of the constraints are even hard to verbalize due to user’s personal preferences or even to emotions (cf. Schlieder & Hagen, 2000). Consequently, the given spatio-temporal problem domain is ill-structured, i.e., there is no “perfect” plan (i.e., optimal solution) for there is no specific definition of such optimum and the important constraint satisfaction criteria (e.g., personal moods and emotions) cannot be mechanized (cf. Simon, 1973). Consequently, spatio-temporal problem solving cannot be totally outsourced to a computational constraint solver but rather requires collaboration with a human. Nevertheless, a typical spatio-temporal planning task can be outlined as follows:

First, events and relations between them which describe a spatio-temporal configuration should be specified. The second step is to identify possible contradiction resulting from overspecified relations. After that the contradictions should be relaxed in order to achieve a feasible order of activities. During the relaxation process some of the relations may become underspecified. The alternative arrangements caused by the underspecified relations should be compared to each other before a decision is made which of them can be taken as a result. The selected solution is a basis for accomplishing of activities in a given environment.

In order to figure out how collaborative assistance can be implemented, we have to consider the following questions:

- Which kind of spatio-temporal information is relevant for a spatio-temporal planning task (i.e., dimensions to be handled) and how can reasoning about spatio-temporal constraints be shared between a human and a collaborative assistance system?
- How can contradictions be resolved and constraints be relaxed by an assistance system in order to generate feasible solutions?
- How can appropriated subsets of feasible spatio-temporal solutions be selected or classified in order to reduce the cognitive load during human mental processing of the provided alternatives?

In the following, I'm going to refer to the results of the exploration studies described in (Seifert, 2005) which address spatio-temporal planning and reasoning problems and outline the common strategies to deal with underspecified relations. Subsequently, the requirements on the representation of the dimensions and the corresponding operations will be derived from exploration studies as well as from the literature concerned with human processing of spatial and temporal information. To obtain feasible solutions for a given problem the demonstrated dimensions are represented in form of a *constraint network*. Subsequently, a procedure to gain partial solutions will be outlined. A partial constraint satisfaction algorithm which operates on qualitative as well as metric spatio-temporal representation of events and relations between them will be demonstrated. Finally, it will be described how one of the possible selection criteria, i.e., region-based fine-to-coarse planning heuristic introduced by (Wiener & Mallot, 2003) can be applied to the generated partial solutions. The paper ends with a discussion of further selection criteria, outlook, open questions, and plans for future work.

What, where and when?

The exploration studies described in (Seifert, 2005) aimed at finding common human reasoning behavior during spatio-temporal planning tasks and strategies for dealing with overspecified and underspecified spatio-temporal relations.

The subjects have been given a list of events with overspecified as well as underspecified spatio-temporal events with a task to find an optimal and consistent temporal order of events. The description of binary constraints consisted of events and temporal orders which should hold between them: e.g. *before*, *after*, *meets*, *met-by*. The spatial constraints were expressed by assignment of certain locations and according distances between the events. The setting of events consisted of a topic, starting point in time, location, duration, and priority. The events in the given spatio-temporal configurations had at least a specified topic or location, and an optional indication of a point in time or location.

The protocols of the exploration studies have shown that subjects preferred to operate on qualitative rather than quantitative relations between the events during spatio-temporal planning. That corresponds to common findings which can be found in the literature concerned with cognitive aspects of spatial (e.g., Freksa, 1992) and temporal reasoning (e.g., Allen, 1983; Knauff et. al, 1995).

Violations of temporal constraints were caused by fixed points in time and corresponding spatial constraints, i.e., too large distances between some of the events. In order to achieve a feasible order of events the subjects preferred to shorten durations of events, or to change locations, and finally to relax points in time and temporal orders. Moreover, reasoning about temporal order of events was much affected by the corresponding priorities: e.g., "*low*", "*middle*" and "*high*", so it can be considered as *hierarchical* (cf., Qu & Beale, 1999). First, events with high priorities were considered, than with middle and after that those of the lowers priorities. Distances were often handled as temporal intervals to get from one location to another: e.g. a "*near*" distance can be reached within five minutes.

In the following, the properties of an atomic spatio-temporal event and consequently the dimension of the spatio-temporal problem domain are described by *topic*, *point in time*, *location*, *duration,* and *priority*. An event instance is identified by its topic or location and requires a minimal duration.

Requirements on the representation of the spatio-temporal problem domain

The requirements on the representation of the spatio-temporal problem domain for an artificial constraint solver can be summarized as follows:

- •It should allow for qualitative temporal reasoning,
- •Handle qualitative measures of distances and durations,
- •Allow for mapping of distance into temporal intervals,
- •And deal with hierarchical priorities.

Furthermore, it should provide precise information about spatial structure of a given environment, in order to generate feasible solutions calculating distances and handle precisely specified assignments for location or point in time. In the following, I'm going to provide a brief overview about existing techniques for dealing with such types of problems.

Existing approaches for spatio-temporal reasoning and partial constraint satisfaction problems

Since several decades, temporal and spatial reasoning has been in the focus of intensive research. Typically, two representation schemes, i.e., *interval based* (cf. Allen, 1983) and *point based* (cf. Vilain, Kautz, & van Beek, 1989) and various combinations (cf. Dechter, Meiri, & Pearl, 1991) are utilized for reasoning about temporal relations. Although, there are a set of temporal constraint satisfaction (TCSP) approaches, their main disadvantage is that they stop searching for a feasible solution as soon as no further valid assignments can be found. In order to solve overconstrained temporal reasoning problems another technique, called Partial Constraint Satisfaction (cf. Freuder & Wallace, 1992) stands in the focus of the investigations introduced by (Beaumont, et. al., 2001). However, the (PCS) algorithms proposed by this group operate primarily on interval-based representation of temporal constraints, produce partial solutions by weakening of temporal constraints, don't consider distances between the temporal intervals, and don't allow for handling of priorities.

In the following, reasoning about spatial constraints can be considered as one-dimensional and eventually mapped to temporal intervals. However, the operations handling distance intervals should be considered differently.

Representation of the spatio-temporal events and relations

The main goal of a common spatio-temporal problem solving task is to obtain an optimal temporal order of events which are interrelated through spatial and temporal constraints. Current Temporal Constraint Satisfaction algorithms aim at finding complete solution or determining one or more consistent scenarios (cf. Schwalb & Dechter, 1997). However, what we need is a set of partial solutions, where the spatial or temporal constraints can be widened or removed and allow the solution to be found (cf. Freuder & Wallace, 1992).

In order to obtain a feasible solution for a given problem the constraints are represented in a *constraint network (CN)* which consist of a set of variables $X = \{X_1, …., X_n\}$. Each of the variables is associated with a *domain* of discrete values: $D_1, …., D_n$ and a set of constraints, $\{C_1, …., C_n\}$ which is expressed as a relation defined on some subset of variables (cf. Dechter, 1992).

In the given problem domain the variables are spatio-temporal events and the binary constraints which guarantee a feasible temporal order of events can be expressed: *{before, after, meets, met-by}*. Spatial constraints are expressed by corresponding distances between the events, which can be specified as temporal intervals. The *domain* of the variables is defined by a set of the introduced dimensions *{topic, time, location, duration, priority}*. For the corresponding domain values the following conditions hold:

- Each event is valid if a topic or location is specified, i.e., (topic ≠ *{Ø}* or location ≠ *{Ø}*) and duration is not zero, i.e., (duration ≠ *{Ø}*).
- **Topic** is specified by a string containing a name of the event or its description or *{Ø}*.
- **Time** is defined by starting point in time of the event in the format: (*hours, minutes*) or *{Ø}*.
- **Location** is specified by metric *(x, y)* coordinates on a given floor plan represented as a two dimensional matrix consisting of raster points, or *{Ø}*.
- **Duration** is specified either by precise temporal measures, e.g., in minutes and qualitatively: *{short, middle, long}*.
- **Priority** is defined as *{low, middle, high}*.

In contrast to classical temporal reasoning problems each binary relation between pairs of events has to be obtained not only from temporal information, but also from the time required to overcome the distance between the corresponding locations of events. In the following the binary relation including distance between the corresponding locations of events is termed as *spatio-temporal relation.*

Spatio-temporal relations between events can be defined by the thirteen qualitative relations described in (Allen, 1983) in his interval algebra, denoted by the labels before (*b*), after (*bi*), during (*d*), contains (*di*), overlaps (*o*), overlapped-by (*oi*), meets (*m*), met-by (*mi*), starts (*s*), started-by (*si*), finishes (*f*), finished-by (*fi*), and equals (*e*). In the following, I'm going to refer to a subset of the introduced relations which are relevant for identification of constraint violations consisting of 9 relations, where the 4 relations: starts (*s*), started by (*si*), finishes (*f*), and finished by (*fi*) are combined into a single during (*d*) relation.

The following table comprises a mapping from points in time into the basic qualitative relations including distance between a pair of events denoted as *t* and *s*, where *tStart* is a starting point and *tEnd = tStart + duration* the ending point of the event *t*, and accordingly *sStart* is a starting point and *sEnd = sStart + duration* an ending point of the event *s*:

Table 1. Qualitative spatio-temporal relations including distance between the events

Relation between t and s	**point-based expression**
before (*b*)	*tEnd + distance < sStart*
after (*bi*)	*tStart > sEnd + distance*
equals (*e*)	*(tStart = sStart) and (tEnd = sEnd or tEnd+distance=sEnd or sEnd+distance=tEnd)*
overlaps (*o*)	*(tStart < sStart) and (tEnd + distance > sStart) and (tEnd + distance < sEnd)*
overlapped-by (*oi*)	*(tStart > sStart) and (tStart < sEnd + distance) and (tEnd > sEnd + distance)*
meets (*m*)	*(tEnd + distance = sStart)*
met-by (*mi*)	*(sEnd + distance = tStart)*
during (*d*)	*((tStart > sStart) and ((tEnd = sEnd + distance) or (tEnd > sEnd + distance)) or ((tStart > sStart or tStart=sStart) and (tEnd < sEnd + distance)))*
contains (*di*)	*((tStart < sStart) and ((tEnd+ distance = sEnd)*

	or (tEnd + distance > sEnd))) *or ((tStart < sStart or tStart=sStart)* *and (tEnd + distance > sEnd)*

In order to obtain feasible partial solutions classical backtracking algorithms and specialized variations are utilized (cf., Dechter, Schwalb & Dechter, Beaumont et. al.). Such algorithms provide an exhaustive search through the problem space consisting of all possible instantiations of the variables and corresponding relations. However, there are methods to reduce the search space: dividing the problem space into smaller sub-problems reduces its computational complexity. Furthermore, searching in a sorted problem space requires less backtracking (cf., Dechter, 1992) and achieves the largest, i.e., "best", solution more quickly. Therefore, the given spatio-temporal problem space can be classified and sorted due to the specified and underspecified dimensions of the spatio-temporal events: e.g., events with no defined location but point in time, or in opposite events with defined locations and various combinations.

In the following, I'm going to describe the procedure for generation of alternative partial solutions.

Generation of partial solutions

First, the set of all events, building a spatio-temporal configuration should be cut into the following subsets:

Spatio-temporal events: events with both fixed point in time and an assignment of location.
Temporal-only events: events with only a fixed point in time.
Temporal events: events with fixed points in time, in other words a union of a set of spatio-temporal events and *temporal-only* events.
Spatial-only events: events with an assignment of location.
Other events: events with no particular assignment of either time or location.
Each of the obtained subsets divides the spatio-temporal problem domain in sub spaces, which can be handled separately according to the specified dimensions of the corresponding events.

The set of temporal events is used to identify overconstrained situations and generate alternative temporal orders by weakening assignments of fixed points in time or locations. The subset of spatio-temporal events derived from the resulting feasible set of temporal events builds a set of *spatio-temporal reference intervals* (cf. Allen, 1983). The spatio-temporal reference intervals are handled separately to obtain a feasible order of *temporal-only* as well as *spatial-only* events lying between the two corresponding edge points and locations. Consequently, the main procedure for dealing with overspecified as well as underspecified sets of events can be described as follows:

1. Order the set of temporal events by time,
2. Generate partial solutions by relaxation of fixed point in time, shortening of duration to a tolerable degree or removing the assignment of location,
3. Select the largest solutions, with the least changes.
4. Initialize the reference intervals with corresponding *temporal* and *reachable-spatial-only* events,
5. Generate alternative orders of events for each of the spatio-temporal reference intervals,
6. Select the largest feasible solutions,

7a. Apply region-based structuring to the generated solutions,
7b. Represent alternative allocations of events by qualitative distance scheme.

In subsequent sections, I'm going to describe each step of the introduced procedure in detail.

Handling overspecified events

A partial solution consists of a set of all available events structured into the following subsets: a sorted list of *temporal events*, list of *reference intervals,* list of *spatial-only* events, list of *temporal-only* events and list of events which have duration, but neither temporal nor spatial assignment, termed *other-events* and a list with the *resulting order* of events. There are several methods that can be used to obtain a partial solution (Freuder & Wallace, 1992):

- Remove variables from the problem,
- Remove constraints from the problem,
- Weaken constraints in a problem,
- Widen the domains of variables to include extra values.

Removing a variable from the problem is a very drastic approach to obtain a partial solution. Conversely, if we remove the binary temporal constraint which holds between each pair of events, we won't be able to obtain a feasible order of events. Weakening the temporal constraints is equal to removing those constraints. However, weakening a spatial constraint, i.e., an assignment of some particular location to reduce the distance between the events can be taken into consideration. Subsequently, in order to keep the temporal order of events valid also the fourth method can be applied: e.g., relaxation of fixed points in time between overlapping events, and shortening duration of a corresponding event with lower priority. The procedure to deal with overspecified events can be described as follows:

1. Take an event from the list of temporal events, in the following, termed as the current event,
2. If the current event is the last one, do nothing and stop searching,
3. Take the next event from the list of temporal events,
4. Prove the spatio-temporal relation between the events,
5. If no contradiction occurs, repeat the procedure with next event.
6. If there is a conflict, try to shorten the duration of the first event in the pair to a tolerable degree, but only if its priority is lower than of the second event.
7. Repeat the procedure with the changed duration.
8. If the duration cannot be shortened, weaken the location or/and time assignment under consideration of respective priorities of events.
9. If the location of the first event in the pair is assigned, propagate the location constraint to the weakened event.
10. Move the corresponding weakened events into the appropriated *spatial-only, temporal-only* or *other-events* subsets.
11. Repeat the procedure with the changed values of events.

The partial solution with the least modifications is the "best" one. The usage of priorities allows on the one hand for avoiding of unnecessary search through the problem space; however, equal priorities contribute to branching of the possible partial solutions, which must be taken into account. The resulting partial solutions with alternative *temporal event* lists can be handled now separately.

Handling underspecified events

After obtaining a feasible temporal order of events with specified points in time the herewith temporally structured problem space can be cut into *spatio-temporal reference intervals,* further called just reference intervals. A reference interval contains the following:

- Two temporally ordered events holding a valid fully specified spatio-temporal relation, so called *edge points* with fixed points in time and assigned locations, in the following termed as *start* and *end*,
- A set of *temporal-only* events with no spatial assignment during the given reference interval,
- A set of *reachable-spatial-only* events, which are obtained with help of the following heuristic, whereas *distance sum* is the distance from the given *spatial-only* event to *start*

edge point plus distance from the given *spatial-only* event to *end* edge point, plus *duration* of event: (see Figure 1):

duration of the reference interval - ∑ durations of the temporal events – *distance sum* > 0.

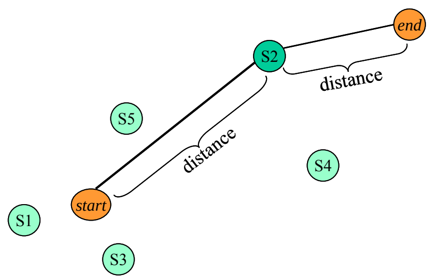

Figure 1: Example of a spatio-temporal reference interval with the corresponding *reachable spatial-only* events: *distance sum (*S2*)* = distance (start, S2) + distance (S2, end) + duration (S2).

After dividing the problem space into sub spaces consisting of reference intervals with corresponding *temporal-only* and *reachable-spatial-only* events the search for a feasible order of *reachable-spatial-only* events which each reference interval can be started. However, the introduced heuristic doesn't guarantee that every *reachable-spatial-only* event fits only in a single reference interval. From that follows, the *reachable-spatial-only* events should provide a list of reference intervals, during which they can be accessed.

Generating a feasible order of spatial only events

First, we have to obtain the order of the *reachable spatial-only* events included in the reference interval. For this purpose a classical branch and bound algorithm with a corresponding *path cost* is used. The *path cost* consists of summarized distances and durations of events in the preliminary resulting order of events including the *start* edge point, the duration of the current event, and the distance between the current event and the *end* edge point: (see Figure 2).

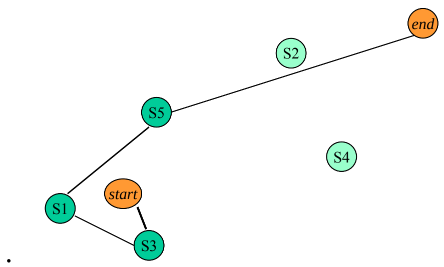

Figure 2: Example of a path cost for the branch and bound algorithm: path cost (S5) = distance (start, S3) + duration (S3) + distance (S3, S1) + duration (S1) + distance(S1, S5) + duration(S5) + distance(S5, end)

The algorithm takes a *reachable-spatial-only* list of events as input, sorted by the *distance sum* of the spatial events (see Figure 1). The resulting shortest path with the biggest number of events provides the best partial solution. In the best case the algorithm terminates as soon as the number of the available *reachable-spatial-only* events is exceeded. However, what should be done if too many events are hypothetically reachable but not all of them fit into the given interval due to their duration and the corresponding *path cost*? In the worst case the *reachable-spatial-only* events should be handled according to their priority or the possibility to be reached during another reference interval. The procedure to resolve such conflicts can be described as follows: if the number of the available *reachable-spatial-only* events has not been yet exceeded but the *path cost* is greater then the duration capacity of the reference interval:

- Find in the list of the preliminary order an event with the longest *distance sum* and the lowest priority.

- If the next event from the set of the *reachable-spatial-only* events has a larger *distance sum* and a lower priority compared to the event found in the resulting set, skip the event form the list of available events. Repeat the procedure with the next event of the list of available *reachable-spatial-only* events.
- If the found event has the same or a lower priority and a larger *distance sum* remove this event from the preliminary resulting list and repeat the procedure with consequently reduced *path cost*.

However, if both events have the same priority and *distance sum*, the possibility to reach those events during other reference intervals has to be considered. Finally, the partial solution with the largest number of events that can be reached during the given interval provides the "best" solution.

Determining the resulting order of events

Now we have a set of reference intervals with ordered *spatial-only* events and corresponding *temporal-only* events. In order to obtain the resulting order of events the procedure can be described as follows (see Figure 3):

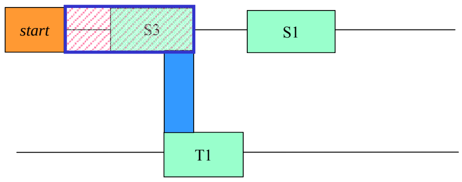

Figure 3: Determining the order of temporal only and spatial events represented on the two parallel temporal axes.

- Take the a pair of events from the resulting order of the spatio-temporal events, included the first edge point event *start*, (for example *start* and *S3*), and a temporal event (*T1*) from the list of *temporal-only* events,
- Determine a preliminary ending point in time of the corresponding spatial event *S3*.
- Build a new temporal interval *I* which begins with the end of the *start* and ends with the preliminary ending point of *S3* (see Figure 3).
- Determine a qualitative temporal relation between *I* and the temporal event *T1*.
- If the obtained temporal relation is *during* or *overlaps* put *T1* before the *spatial-only* event *S3*.
- Calculate the preliminary starting point for *S3* taking into account the duration and point in time of *T1*, in case it has been placed before *S3*.
- Repeat the procedure with the next triple of events (*S3*, *S1* and *T2*).

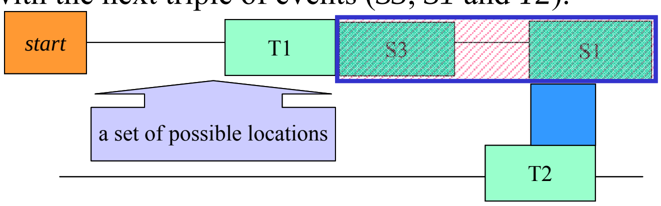

Figure 4: Repeating the procedure for the next triple of events represented on the two parallel temporal axes.

In opposite to the *spatial-only* events ordering of *temporal-only* events can be determined in a single spatio-temporal reference interval. That means that some of the reference intervals may contain redundant assignments of *spatial-only* events of the resulting spatio-temporal orders. There are several possibilities to handle this problem. The first one is to merge the reference intervals before handling the *temporal-only events* and obtain the shortest path considering all possible resulting orders of spatial events. The second one is to apply the *region-based selection criterion* described in the following section and hope that the redundancies can be eliminated. Another option is to let the user decide which of the spatio-temporal configuration she prefers. Therefore, a user-friendly interaction model adapted to the human spatio-temporal planning and reasoning strategies is needed. However, such interaction model is a subject of further research.

Representing alternative locations using qualitative distance scheme

Due to points in time of temporal events as well as time available during each spatio-temporal reference interval, many possible locations for the corresponding *temporal-only* events can be taken into consideration.

Since distances between the spatio-temporal events can be expressed qualitatively, the alternative assignment of locations for *temporal-only* and the remaining events (i.e., *other-events*) can be represented trough a qualitative distance scheme: "*near*", "*not-so-far*", and "*far-away*" regarding the corresponding *spatial-only* and events with high priority or the *edge points* of the reference intervals.

Region-based representation of alternative locations

Series of psychological experiments done by (Wiener & Mallot, 2003) concerned with processing of spatial environment structured in geographic regions have shown that humans avoid changing of regions during spatial problem solving tasks like navigation and route-planning. Furthermore, human processing of spatial knowledge is considered hierarchical (cf. Hirtle, 1998; Tversky, 1993). Consequently, the alternative sets of spatial assignments can be combined into spatial regions, which can be also structured hierarchically. However, such regions should be mapped to the spatial structure of the given environment. At this juncture some of the possible spatial assignments and redundancies in the allocation of the *spatial-only* events can be omitted.

Outlook and future work

In the scope of the paper I've introduced procedures for dealing with overspecified as well as underspecified spatio-temporal relations under consideration of temporal as well as spatial constraints. Algorithms and corresponding structuring criteria of the problem space for generation of alternative spatio-temporal orders of events were discussed. However, in order to achieve an adequate or "optimal" solution, collaboration between a human and an artificial constraint solver is required. User-friendly collaboration can be achieved by introducing the possible alternatives in form of qualitative distance schemes, as well as hierarchical or region-based representation of alternative spatial allocations of events.

However, a suitable interaction model for collaborative spatio-temporal planning and reasoning remains a subject of further research.

Acknowledgements

I gratefully acknowledge financial support by the Deutsche Forschungsgemeinschaft (DFG) through the Transregional Collaborative Research Center SFB/TR 8 Spatial Cognition: Reasoning, Action, Interaction (project R1-[ImageSpace]) and valuable comments on the paper and discussion about my work with my colleagues Holger Schultheis, Thomas Barkowsky and Sven Bertel.

References

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

Beaumont, M., Sattar, A., Maher, M., and Thornton, J. (2001) *Solving Overconstrained Temporal Reasoning Problems*. In M. Brooks, D. Corbett, and M. Stumptner (Eds.): AI 2001, LNAI 2256, pp. 37−49, 2001. Springer-Verlag Berlin Heidelberg 2001.

Dechter, R., (1992) "*Constraint Networks (Survey).*" In Encyclopedia of Artificial Intelligence, 2nd edition, 1992, John Wiley & Sons, Inc., pp. 276-285.

Dechter, R., Meiri, I., & Pearl, J.. (1991). *Temporal constraint networks*. Artificial Intelligence 49: 61–95.

Freksa, C., (1992). *Using orientation information for qualitative spatial reasoning*. In A. U. Frank, I. Campari, U. Formentini (Eds.), Theories and methods of spatio-temporal reasoning in geographic space, pp. 162–178. Springer, Berlin.

Freuder, E., & Wallace, R. (1992) *Partial constraint satisfaction*. Artificial Intelligence, 58(1):21–70, 1992.

Hirtle, S. C. (1998). *The cognitive atlas: using GIS as a metaphor for memory*. In M. Egenhofer & R. Golledge (Eds.), Spatial and temporal reasoning in geographic information systems (pp. 267-276). Oxford University Press.

Knauff, M., Rauh, R., & Schlieder, C. (1995). P*referred mental models in qualitative spatial reasoning: A cognitive assessment of Allen's calculus*. Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society (pp. 200-205). Mahwah, NJ: Lawrence Erlbaum.

Schlieder, C., & Hagen, C. (2000). I*nteractive layout generation with a diagrammatic constraint language*. In C. Freksa, C. Habel, & K. F. Wender (Eds.), Spatial cognition II - Integrating abstract theories, empirical studies, formal methods, and practical applications (pp. 198-211). Berlin: Springer.

Seifert, I. (2005). *Towards cognitively adequate interaction for mental model-based spatio-temporal assistance*. Proc. of 2005 AAAI Spring Symposium on "Reasoning with Mental and External Diagrams: Computational Modeling and Spatial Assistance".

Simon, H.A. (1973). *The Structure of Ill Structured Problems*. Artificial Intelligence 4 (1973), North-Holland Publishing Company.

Schwalb, E., and Dechter, R., (1997) "*Processing Disjunctions in Temporal Constraint Networks*." In "Artificial Intelligence", volume 93, pp. 29-61, 1997.

Tversky, B. (1993). Cognitive maps, cognitive collages, and spatial mental models. In A. Frank & I. Campari (Eds.), Spatial information theory (pp. 14-24). Berlin: Springer.

Vilain, M., Kautz, H., & van Beek, P. (1989). *Constraint propagation algorithms for temporal reasoning*: A revised report. In Daniel S. Weld and Johan de Kleer (eds.), Readings on Qualitative Reasoning about Physical Systems. Morgan Kauffman, pages 373–381.

Wiener, J.M., Mallot, H.A. (2003) *'Fine-to-Coarse' Route Planning and Navigation in Regionalized Environments. Spatial Cognition and Computation 3(4)*, pp. 331-358

Wilson, M. (2002). *Six views of embodied cognition*. Psychonomic Bulletin & Review, 9: 625-636

Qu, Y. & S. Beale A. (1999) *Constraint-Based Model for Cooperative Response Generation in Information Dialogues*, (abs). Copyright AAAI. In Proceedings of AAAI-99, Orlando, FL, July 1999.