Probabilistic Plan Evaluation
A Prerequisite for a Model of Information Value Dynamics

Markus Schaal*

Computergestützte Informationssysteme (CIS), Technische Universität Berlin,

Sekr. E-N 7, Einsteinufer 17, D-10587 Berlin, Germany

e-mail: schaal@cs.tu-berlin.de

November 30, 2000

Abstract

Time sensitive information supply is a common challenge in many information systems. The issue of determining the best time of information supply becomes even more important with the growing use of mobile communication devices. Therefore, the value of information for the information recipient needs to be studied with respect to time.

In this report, the impact of information on the information recipients plan and the utility of that plan are modeled in order to facilitate the study of time-dependent information value. Plan utility is modeled in terms of rewards for certain events and Bayesian networks are used for probabilistic evaluation of the expected reward.

The interaction of plan execution with the environment is formulated by event synchronization. Imprecise information about the expected time of external events is used here for the evaluation of synchronization with these events. The integration of such information stemming from distinct information sources is supported.

*This research was supported by the German Research Society, Berlin-Brandenburg Graduate School in Distributed Information Systems (DFG grant no. GRK 316).

1 Introduction

The evolution of the internet as global infrastructure for easy exchange of information gave rise to the development of information services that combine information from different sources. Currently, the advent and rapidly growing use of mobile connected devices such as laptops, mobile phones, and personal organizers is technologically improving the delivery of information on a just-in-time basis even though the quality of online connectivity varies widely between mobile and immobile devices.

This technological evolution is a challenge for the development and improvement of information systems, that fully exploit the emerging possibilities. Information logistics provides a framework for the development of intelligent services, that decide about information delivery with respect to content, presentation, location, mode, and the best time of information delivery.

The last aspect is rarely studied yet, even though numerous information systems do already inform the user in a timely manner, some of them are discussed in the sequel.

Query-based systems provide information as soon as possible after request. Time of request and technological limitations are solely responsible for the time of information supply.

Alerting Systems inform the user actively upon the occurrence of certain events that are specified in terms of profiles or activation triggers (cf. [HF99]). A reasoning mechanism for the optimal time of delivery could extend the ability of activation triggers.

Mixed-Initiative Assistance is a new paradigm rather than a class of information systems. Stemming from artificial intelligence (cf. Fergueson et al. [FAM96] and Horvitz [Hor99]), this term is used to denote an assistance systems ability to decide whether it is the users or the systems task to take initiative. Reasoning about the best time of taking the initiative seems to be an important issue in this area.

The application area considered here is individual trip planning assistance. This includes personal pre-trip and on-trip information services. Existing solutions in the public transport sector include query-based information systems for pre-trip planning such as HAFAS (`http://bahn.hafas.de`), EFA (`http://www.efa.de`), ASS, GEOFOX (`http://www.hbt-hamburg.de/cgi-bin/hvv/geofox`) and fahrinfo (`http://www.bvg.de/plan/fahrinfo.html`) and mobile solutions built on top of these such as TRIP (`http://www.trip.organon.net/`), HAFAS mobil (`http://www.hacon.de/hafas/mobil.html`) and EFApersonal. Other solutions exist for road networks (e.g. `http://www.routenplanung.de`).

These systems are two steps away from deciding about the best time of information supply. The first step is the processing of real-time data, the second step is the decision about the best time of information supply. Real-time data provides a dynamically changing knowledge base upon which an information system can judge the utility of the information already available to the information recipient as well as the utility of new information to be supplied.

Lets assume, that real-time data is available. The following scenario illustrates the problem of deciding about the best time of information supply.

> A business man plans the trip to a meeting. On the day before the trip, he asks a route planning system for the optimal connection and receives the information that it is best to take the train to his destination's main railway station and a bus from there to the final address.
>
> During the trip or, more precisely, during the train ride, the personal assistance system of the business man learns about a delay of the train, reducing the likelihood of catching the bus at the destination. A taxi might be the better choice.
>
> Shortly before arrival, the personal assistance system learns from the bus information system, that the bus will also run late, thus matching late arrival of the train.

The following conclusions can be drawn from the example:

1. Knowledge about the time of events and thus information to be given to the passenger is changing. Therefore, the trade-off between expected change and time-dependent information value needs to be considered for the decision about information transmission. Time-dependency of information value reflects the intuition of information to be useful at a certain time and to be useless at other times.

2. Knowledge about the time of events may stem from different information sources, such as the train and bus information systems respectively. In order to judge the combined effects of knowledge updates stemming from different sources, an appropriate reasoning mechanism has to be implemented for a personal assistance system.

The last issue, namely the judgement of information value with respect to knowledge updates stemming from different sources is the main issue of this report. We assume knowledge about the time of events to be given in form of probabilistic distributions over time and plan utility is evaluated with Bayesian networks in terms of expected rewards. Influence diagrams, which are an extension of Baysian networks, can then be used for reasoning about information supply (cf. Pearl [Pea91] for the use of influence diagrams for decision support).

Section 2 introduces the concepts that provide the basis of the probabilistic reasoning about event occurrence times. Section 3 contains a broader discussion and proper introduction to the example used throughout this report. Section 4 gives a formalization of the continuous-time Bayesian network for the reasoning about event occurrence times. Probabilistic plan evaluation is exemplified in Section 5. The results are summarized and future research is outlined in Section 6.

2 Basic Concepts

In this section, the basic concepts and terminology for the reasoning about actions, events and change are discussed in order to build a solid basis for further more detailed elaborations. Terms and concepts are mainly drawn from work in the field of artificial intelligence (cf. Pearl [Pea91] for probabilistic reasoning in general and Shafer et al. [SGS00] for the notion of events).

2.1 Events

Events and states are the basic entities of discrete change. These concepts have been settled in various disciplines of computer science independently. Therefore, the notion of events is ambiguous. In our understanding, events are characterized by the following properties.

Events are instantaneous. This reflects the intuitive assignment of duration to states, while events are instantaneous state changes.

Events are unique. In general, events are defined as state changes with time. In this case, uniqueness is implicit by specification of occurrence time. Here, events are introduced as abstract concepts, whose time is not known or not known precisely. Therefore, uniqueness has to be explicitly stated, distinguishing different events from each other. Reoccurrence of events is not possible, but replaced by a finite or infinite sequence of similar events.

Events are local. Two instantaneous unique events may occur in either sequence if there is no total dependence between them. Therefore, event specifications need to be local to guarantee an unrestricted occurrence of other uncorrelated events. Causality has to be modeled explicitly.

Thus events are an abstract concept for change, that embodies the concepts of instantaneous occurrence, uniqueness, and locality as described above.

Upon this concept of events, physical movement can be defined in terms of events for *Arrival* and *Departure* of objects at certain geographic locations. On the informational level, information transmission can be defined by events for *Send* and *Receive* of information objects at certain physical objects.

2.2 Action, Goal, Plans, and Information

In a distributed world, autonomy plays an important role for the understanding of action, interaction, and change. Therefore, agents are modeled as autonomous objects. Agents may send or receive information. Additionally, they do directly interact with their environment through observation and action. The selection of actions is directed towards a goal. In this picture, the passenger from the introductory scenario is an agent. The recommendation to take the bus or taxi respectively is an information, and his experience of actually arriving at a location is an observation. The passenger's goal is to be in time at his meeting.

Actions are modeled as pairs of events, namely one event starting and one event finishing each action.

In general, agents act upon their observation of the world state, i.e. actions are selected according to the observed world state with respect to a knowledge base. Here, the knowledge base is reduced to a plan, namely a preselected sequence of events to be performed. In this case, agents simply act by selecting the next event in the plan and to perform it. An event may fail to occur in which case the plan gets stuck, e.g. if a passenger has planned to depart (by train) from some location but the train has already left.

Information is transferable knowledge. Here, the information received by an agent is either an initial plan or a plan update. A plan update does only affect parts of the plan that have not yet occurred from the agents point of view.

2.3 Probabilistic Projection

Decisions about information transmission are based upon expectations concerning the difference between the case of information transmission as opposed to the case of no such transmission. Since the effect of information transmission is assumed to help the information recipient in future situations, a model for the expected execution of plans is needed. The concept to be used here is probabilistic projection of the plan into the future in order to account for imprecise knowledge about the occurrence time of future events.

Numerous analytical and simulative models for stochastic processes, i.e. processes with probabilistic choice of alternatives and/or stochastic durations, have been developed. Stochastic Petri-nets are one exemplary model which combines the formalism of Petri-nets with stochastic time and selection of transitions. Analytically, stochastic Petri-nets comply to a certain class of continuous-time semi-Markov chains (cf. Ciardo et. al. [CGL94]).

In order to combine temporal knowledge from different sources, an absolute time line is needed for synchronization. Continuous time has been chosen here as a natural concept for this reference system. Absolute time will thus be represented by $\mathbb{R}_0^+$ thoughout this report.

Bayesian networks have been chosen here in order to reason about expected futures. The explicit modeling of event occurrence times with respect to an absolute time line requires an appropriate representation of time in Bayesian networks. The representation of time in causal networks (both Bayesian and non-Bayesian) has been studied before. A short overview is given in the next section.

2.4 Time in Causal Networks

Berzuini [Ber90] introduced a network of dates in order to reason about the probabilistic nature of event occurrence times for medical applications. Temporal random variables and continuous time is used in this work.

Other approaches employ Bayesian networks for reasoning about time. Dean and Kanazawa [DKK92] propose random variables for duration as a means to represent semi-Markov processes in probabilistic networks. Tawfik and Neufeld [TN94] employ Temporal Bayesian Networks (TBN) for the representation of probabilities as functions of time. Arroyo-Figueroa and Sucar [AFS99] model event occurrence times as nodes with respect to time intervals as being developed by Allen [All83]. This last approach is actually very similar to Berzuini's approach but it is restricted to a finite number of intervals for the occurrence time of events.

2.5 Reward, Utility and Information Value

Since information is modeled as plan update, comparison between the quality of two different plans, namely the original and the updated plan, is an indicator for the value of information transmission. The quality of a plan is the expected reward of the plan. Rewards are assigned to events of the plan occurring at specific times. Probabilistic plan evaluation gives the expected reward. The expected reward can be interpreted as utility of the plan with respect to the goal that is represented by the rewards.

3 An Example

Here, the information recipient is a human user. He tries to follow a given plan (including the use of different vehicles, both public and private traffic). The superior information of the expert system can be achieved by feeding it from real-time operating systems of public transport providers, floating car data, etc.

The example is settled in the application domain of individual trip planning assistance and is illustrated in Figure 1. Nodes depict events of interest and arrows show the temporal order of events. Two arrows going out from a node depict alternatives which exclude each other. Events are grouped together by boxes with labels. Each box corresponds to an object.

The passenger is in the center of the illustration. Events of different objects may be required to by synchronous in order to occur. This is symbolized by dotted lines between nodes.

For the trip to a meeting, two alternative routes are to be considered after arrival at the railway station by train. One involves the use of a bus, the other involves the use of a taxi for the trip to the final address. The illustration starts with Arrive(pas,A), where pas denotes the agent *passenger*, A denotes the location of the arrival station, and Arrive(pas,A) is a shorthand for *The event of* pas *arriving at location* A. From there, the passenger may either go to the bus stop (location B) or to the taxi stand (location C) and continue to his final address (location D). Arrivals and departures can only occur at the same time as the respective arrivals and departures of the transporting vehicles, namely bus, train, and taxi.

The alternatives of the passenger represent two potential plans, namely (Arrive(pas,A), Depart(pas,B), Arrive(pas,D)) and (Arrive(pas,A), Depart(pas,C), Arrive(pas,D)).

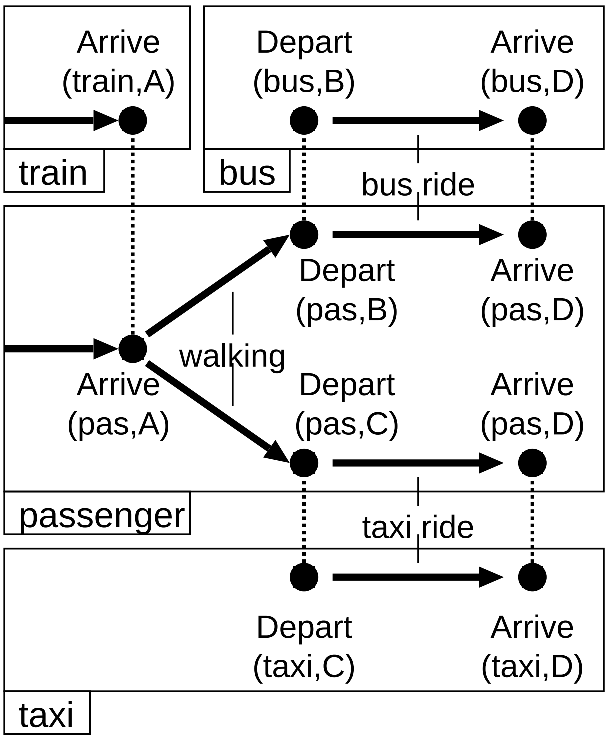

Figure 1: Trip-Example

Not depicted in the illustration, but equally important is the goal of the passenger. We assume a deadline problem here, with deadline $t_{deadline}$ being the latest time to be at location D, i.e. the meeting.

Whether the deadline will be met depends on various relationships between the occurrence times of the events, some of them are listed here.

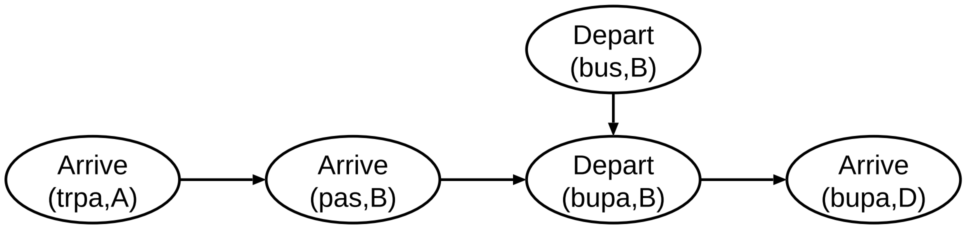

Figure 2: Bayesian Network (BN) for probabilistic projection (take bus)

- Late occurrence of Arrive(pas,A) may result in non-occurrence of Depart(pas,B), i.e. the passenger may miss the bus if the train is running late.
- Late occurrence of both Arrive(pas,A) and Depart(bus,B) may result in late occurrence of Depart(pas,B), i.e. if both bus and train are delayed, the passenger may catch the bus, but at a later time.
- Late occurrence of Depart(bus,B) may result in late occurrence of Arrive(bus,D), i.e. if the bus is delayed at departure time, he might be as well at arrival time.

In order to evaluate the utility of a plan alternative (e.g. for taking the taxi) with respect to the goal (here: the deadline), probabilistic plan evaluation will be performed with Bayesian networks.

This is illustrated in Figure 2 for the alternative of taking the bus. The nodes in the Bayesian network are the random variables for the occurrence times of events denoted by the respective label. The events considered here are not the same as the ones depicted in Figure 1. Arrive(trpa,A), Depart(bupa,B) and Arrive(bupa,D) are the respective joint events of both passenger and train or bus respectively. Arrive(pas,B) has been added in order to account for the fact, that the passenger will not arrive at the same time as he departs by bus.

The qualitative model of conditional dependence that is depicted in Figure 2 is justified as follows:

- Arrive(pas,B) depends on Arrive(trpa,A) by the duration for walking from A to B. Immediate start for walking is assumed here.
- Depart(bupa,B) depends on Arrive(pas,B) and Depart(bus,B) since the passenger needs to be at B before the departure of the bus. Depart(bus,B) is assumed being independent of Arrive(pas,B) as its major goal is to keep the bus schedule.
- Arrive(bupa,D) depends on Depart(bupa,B) and is determined by the travel time of the bus.

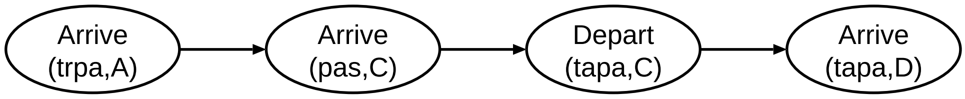

Figure 3: Bayesian Network (BN) for probabilistic projection (take taxi)

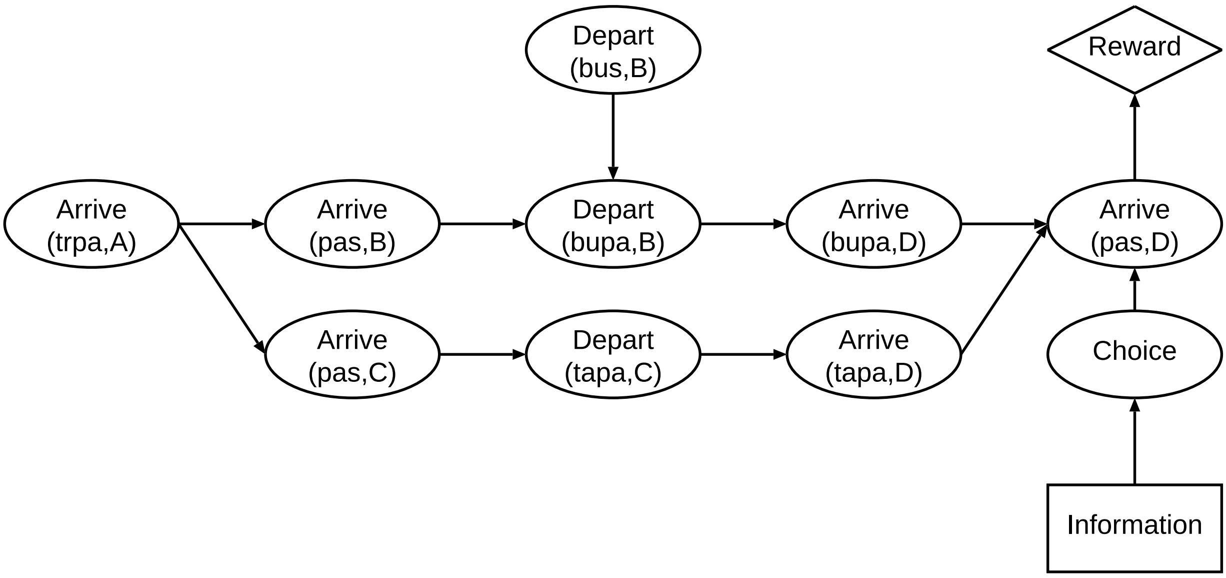

Figure 4: Influence diagram for the trip-example

It is important to note, that Depart(bupa,B) may fail to occur, since there is the possibility of Arrive(pas,B) occurring after Depart(bus,B). Therefore, the random variable Depart(bupa,B) needs a special value for non-occurrence (plan failure) which must be propagated along with the plan.

The formal model of uncertain events as nodes in Bayesian networks is given in section 4.4 and considers the concept of non-occurrence.

Figure 3 illustrates the other plan alternative, namely taking the taxi. Here, the departure time of the taxi is not modeled as it depends heavily on the arrival time of the passenger. Instead, Depart(tapa,C) is directly and only dependent on Arrive(pas,C). All other dependencies are in analogy to the previously discussed alternative of taking the bus.

In order to decide about the best information for the passenger, the Bayesian networks will be extended with decision nodes representing alternative actions and value nodes, representing cost or benefit. The resulting network is an influence diagram (cf. Shachter [Sha86] [Sha87] and Pearl [Pea91]). Efficient solutions for the computation of joint distributions of Bayesian networks and Influence Diagrams are known (cf. Lauritzen and Spiegelhalter [LS88]).

The influence diagram in Figure 4 contains the probabilistic projections for both plan alternatives with common node Arrive(trpa,A) (to the left). Since no information about exclusion is modeled in this network, the respective branches for the plan alternatives are to be interpreted as potential plans. The additional nodes are Arrive(pas,D), Choice, Information, and Reward. They are explained in the sequel:

- Choice is the decision of the passenger and has bus and taxi as possible values.
- Arrive(pas,D) depends on Choice, because the arrival time depends on the plan alternative. The distribution in time is to be copied from Arrive(bupa,D) for Choice=bus and from Arrive(tapa,D) for Choice=taxi.
- Reward is a value node in an influence diagram. Unlike chance nodes, the name for conventional nodes in Baysian networks, value nodes are not specified in terms of conditional probabilities. Instead, they assign a certain reward to every outcome of the chance nodes being parents to the value node. Reward assigns one for Arrive(pas,D) less or equal to $t_{deadline}$ and zero for Arrive(pas,D) greater than $t_{deadline}$ or non-occurring.
- Information is a decision node in an influence diagram. Decision nodes represent alternative actions. Arrows going from decision to chance nodes depict conditional dependence. Information has possible values take bus, take taxi and no info.
- Choice depends on Information. We assume even, that the passenger follows the advice given to him, i.e. in the example, the conditional probability of Choice=bus is one for Information=Take bus.

For given marginal distributions and known conditional distributions, the optimal decision can be calculated based upon the joint distribution of all variables. In the example, marginal nodes would be Arrive (trpa,A) and Depart (bus,B). Concrete calculations are given in Section 5.

Since marginal distributions may change in the course of time, a typical scenario may look like this:

One day before the trip: Trains and buses are expected to depart and arrive in time.
Decision: take bus.

During train ride: Train is running late, likelihood of missing the bus raises.
Decision: take taxi.

Short before arrival in A: Bus is also delayed. The delay raises the likelihood of catching the bus. The bus is still in time with greater probability than provided by the taxi.
Decision: take bus.

So what is the best information policy with respect to the passenger? One day before the trip, the passenger may have inquired for the optimal route to take, and taking the bus was, according to the scenario, the optimal decision. Is it neccessary to inform the passenger about the changed decision during the train ride? If not, what are the arguments?

Even though the consideration of information policies is not within the scope of this report, two observations should be kept in mind for future work:

Optimal information may change due to changes in the knowledge base. As a consequence, early information is neighter better nor worse than late information.

Information timing is important for the assumed direct influence on the choice of the passenger. Information need to reach the passenger before the alternative has to be chosen.

According to the second observation information should be given to the passenger before Arrive(trpa,A).

Intuitively, it seems to be correct to wait as long as possible with information supply, since late information is assumed to be better than early information. But the first observation suggests a slight variation in the example. If the information system learns about the delay of the bus only after Arrive(trpa,A), it would update information according to the delay of the train only.

4 Formal Model

The formalization of the concepts being introduced informally in Section 2 is built upon the example from the last section.

4.1 Events

For reasoning about events, the following types are to be considered.

- *(abstract) events* are the pure unique events as being described in Section 2.1. Specifically, they are not associated with any notion of occurrence or occurrence time. Abstract events are denoted by $e_1, e_2, \ldots$ and the set of abstract events is denoted by E. Causal knowledge about events (including temporal relationships) will always be specified with respect to abstract events. Abstract events may also be referred to by dropping the attribute *abstract* and simply write event.

- *concrete events* model concrete potential occurrences of abstract events at a specific point in time. Concrete events are denoted by $\dot{e}_1, \dot{e}_2, \ldots$ and the set of concrete events is denoted by $\dot{E}$. A function $t : \dot{E} \mapsto \mathbb{R}_0^+$ maps concrete events into $\mathbb{R}_0^+$, the continous space of time points. The semantics of causal dependencies between events are specified with respect to concrete events.

- *uncertain events* model uncertainty about the potential occurrence and occurrence time of abstract events. Uncertain events are denoted by $\tilde{e}_1, \tilde{e}_2, \ldots$ and the set of uncertain events is denoted by $\tilde{E}$. In Section 4.4, uncertain events are modeled by a continous probability measure for the possibilities of occurrence at some point in time and of non-occurrence. Uncertain events are used for the probabilistic evaluation of plans by projection into the future.

In order to correlate different types of events, that represent the same abstract event, an equivalence relation $=_e$ is defined for the disjoint union of E, $\dot{E}$ and $\tilde{E}$. For an abstract event e, $[e]$ denotes the equivalence class that contains the abstract event itself and all concrete and uncertain events representing e.

Example: $e =$Arrive(pas,A) is an abstract event. $\dot{e} \in [e]$ with $t(\dot{e}) = t$ is the concrete event of passengers arrival in A at time t.

4.2 Plans and Constraints

Plans are defined on top of abstract events, because they do not refer to the actual occurrence of any event.

Definition 4.1 *A plan is a finite sequence of distinct events* $(e_1, \ldots, e_n)$*, i.e.* $e_i \neq e_j$ *for* $1 \leq i < j \leq n$.

Example: With e_1=Arrive(pas,A), e_2=Arrive(pas,B), e_3=Depart(pas,B) and e_4=Arrive(pas,C), the first plan alternative depicted in Figure 1 is represented by plan (e_1, e_2, e_3, e_4).

For projection of plan execution, all influences on the temporal distance between and non-occurrence of events in the plan have to be considered. Temporal distance between events is caused by inherent duration or waiting for external events. Non-occurrence is caused by failure to synchronize with external events. Inherent duration and synchronization with external events are thus the concepts underlying these influences.

Inherent duration is the minimal temporal distance between subsequent events given by some kind of implicit knowledge about natural laws governing the occurrence of these events.

Definition 4.2 *Let p be a plan, $E(p)$ the set of events in p. The duration constraint $duration(e_1, e_2, d)$ with $e_1, e_2 \in E(p)$ and $d \in \mathbb{R}_0^+$ is met by concrete events $\dot{e}_1 \in [e_1]$ and $\dot{e}_2 \in [e_2] \Leftrightarrow t(\dot{e}_1) \leq t(\dot{e}_2) + d$.*

Example: e_2 occurs not earlier than 5 minutes after e_1, accounting for the time needed to walk from location A to location B.

Definition 4.3 *Let p be a plan, $E(p)$ the set of events in p. e_x an external event. The synchronization constraint $sync(e_i, e_x)$ with $e_i \in E(p)$ is met by concrete events $\dot{e}_i \in [e_i]$ and $\dot{e}_x \in [e_x] \Leftrightarrow t(\dot{e}_i) = t(\dot{e}_x)$.*

Example: Depart(pas,B) happens synchronously with Depart(bus,B), accounting for the fact of simultaneous departure of passenger and bus.

The next event in a plan does occur, if it can be synchronized. The time depends on the time of synchronization or, alternatively, on the minimum temporal distance from the previous event.

4.3 Rewards

Rewards are given as functions over time per event e.

Definition 4.4 *Let p be a plan, $E(p)$ the set of events in p. The reward for execution of p is given by $R : E(p) \times \mathbb{R}_0^+ \mapsto \mathbb{R}_0^+$ with $R(e, t)$ being the reward for event occurrence $\dot{e}$ with $\dot{e} =_e e$ and $t(\dot{e}) = t$.*

Example: For the example (cf. Section 3), the deadline for e=Arrive(pas,D) can be modeled by reward

$$R(e,t) = \begin{cases} 1 & t \leq t_{deadline} \\ 0 & t > t_{deadline} \end{cases} \tag{1}$$

All other events receive no rewards.

4.4 A probability measure for uncertain events

Imprecise information about the timing of external events can be modeled by different models and on different levels of granularity. A continuous probability measure for the occurrence time of events is used here.

Generally, a continuous probability space $(\Omega, \mathcal{A}, P)$ is defined by $\Omega = \mathbb{R}$ or $\mathbb{R}_0^+$ with $\mathcal{A} \subseteq 2^\Omega$ being a σ-algebra containing at least all intervals in Ω. For any $A \in \mathcal{A}$ a probability is assigned by probability measure P with $P(A) = \int_{t_a}^{t_b} f(t)\,dt$, where f is a non-negative function. Additionally, $\int_0^\infty f(t)\,dt = 1$ since $P(\Omega) = 1$.

Here, the possibility of non-occurrence shall be modeled by an extra discrete atomic *event.* The term *event* here refers to an element of $\mathcal{A}$ and atomic means, that it consists of a single element of Ω.

It is a wide-spread practice to define Ω as a compact interval $[0, \infty]$ and to assign non-occurrence to the value ∞. We will follow this approach but we will represent the probability measure by a pseudo-density function for occurrence and a separate dirac impulse for non-occurrence.

Definition 4.5 $(\Omega, \mathcal{A}, P(e))$ *is a probability space for event e with*

- $\Omega = [0, \infty$ *for the continuum of time* $(\mathbb{R}_0^+)$ *and the possibility of non-occurrence* (∞)*,*
- $\mathcal{A}$ *is a* σ*-algebra containing all intervals of type* $[a, b)$ *and* ∞*,*
- $P(e)$ *is a probability measure with* $P(A)$ *for* $A \in \mathcal{A}$ *given by*

$$P(A) = \begin{cases} \int_A f(t)\,dt & \infty \notin A \\ \int_A f(t)\,dt + p & \infty \in A \end{cases} \tag{2}$$

with density function $g(t) = f(t) + p\delta_\infty(t)$ *and* $p = 1 - \int_0^\infty f(t)\,dt$. $f(t)$ *is a non-negative, bounded, piecewise continuous function and* $p\delta_\infty(t)$ *is a dirac impulse at* ∞ *of size p.*

Notation: For unindexed events e, P_e and f_e are used instead of $P(e)$ and f, for indexed events e_i, P_i and f_i are used.

We call $f_e(t)$ the pseudo-density function representing the uncertain event $\tilde{e}$.

The definition allows both discrete, interval-based approaches to probabilistic reasoning about the occurrence time of events and continuous reasoning with arbitrary intervals.

4.5 Quantification of Conditional Dependence

Probabilistic projection requires two basic operations. These operations are encoded into the conditional dependencies between the nodes. The first (Duration) is required, if the next event of the plan occurs only after a certain duration, the second (Wait or miss) is required, if the next event of the plan can only occur at the same time as a specific external event.

Mixed cases, i.e., both *Duration* and *Wait or miss* for the same event are not considered here, they can be separated by adding potential events into a plan.

Duration Let e_1, e_2 be subsequent events of a plan and $duration(e_1, e_2, d)$ is the constraint to be obeyed. d is a continuous random variable on $\mathbb{R}_0^+$ represented by $f_d(t)$. For given pseudo-density function $f_1(t)$,

$$f_2(t) = \int_0^t f_1(t-\tau) \cdot f_d(\tau)\, d\tau. \tag{3}$$

a kind of a pseudo-convolution.

Wait or Miss Let e_1, e_2 be subsequent events of a plan, e_x an external event and $e_2 =_t e_x$ is the constraint to be obeyed. For given pseudo-density functions $f_1(t)$ and $f_x(t)$,

$$f_2(t) = f_x(t) \int_0^t f_1(\tau)\, d\tau \tag{4}$$

It should be noted, that $\int_0^\infty f_2(t)\, dt$ is not equal to one even if $\int_0^\infty f_x(t)\, dt = 1$ and $\int_0^\infty f_1(t)\, dt = 1$. This is intended, since the case of missing the bus needs to be accounted for the probability of the non-occurrence.

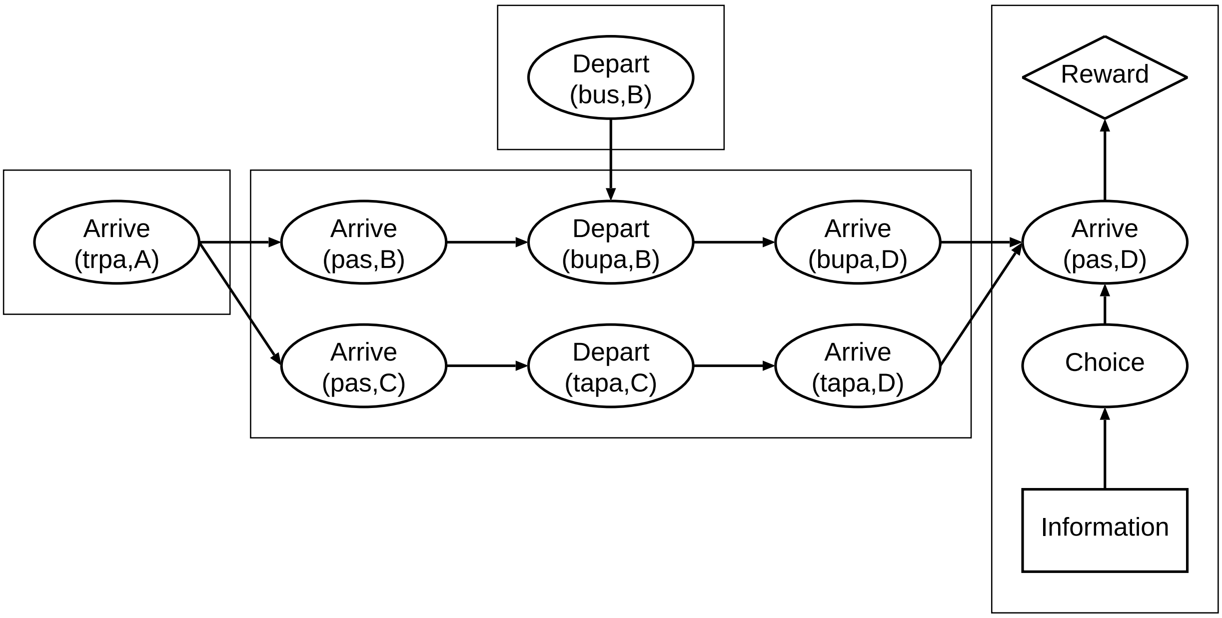

Figure 5: Influence diagram with indexed events

5 Probabilistic Evaluation by Example

The influence diagram from Figure 4 is shown in Figure 5 again. Nodes are grouped here with respect to their intended use for decision about information supply.

The marginal distributions of Arrive(trpa,A)) and Depart(bus,B) represent knowledge about the time and occurrence of external events. This knowledge can be updated by direct replacement of the marginal distributions. This reflects the fact that evidence influencing these nodes will be obtained outside the model, e.g. by the railway information system or the bus information system, and transferred into the model by update of marginal distributions.

The nodes at the center of the illustration represent the potential events of two plan alternatives. Conditional dependence is modeled by application of the concepts being introduced in Section 4.5. It is important to observe, that the events here are indeed potential events, since no decision about the actual plan of the passenger is modeled in this group of nodes.

The nodes to the right are merely for the decision about information supply. Arrive(pas,D) is an abstract event for the combined arrival at D. It is important to observe, that Arrive(pas,D) is not the joint event of Arrive(bupa,D) and Arrive(tapa,D), since they are exclusive to each other. Instead, the pseudo-density function of Arrive(pas,D) is the weighted sum of the pseudo-density functions for Arrive(bupa,D) and Arrive(tapa,D) where the weights are given by the decision of the passenger in node Choice.

In order to solve the influence diagram, the HUGIN-toolset (`http://www.hugin.com/`), a leading software solution in the field of Bayesian network analysis is applied.

In HUGIN, discrete random variables have to be used instead of continuous random variables. The continuous random variables are thus represented by intervals of time. Marginal distributions for Arrive(trpa,A) and Depart(bus,B) are shown in Figures 6(a) and 6(b).

Applying the results of Section 4.5 to the example from Section 3, we get the following conditional dependencies:

- (Walking from A to B) Arrive(trpa,A), Arrive(pas,B) have temporal distance due to duration $d_{walking}$, the time neccessary for walking from A to B. With deterministic duration $d_{walking} = 5$, the conditional dependency table for the example is depicted in Figure 6(c) following the concept of *Duration*. For a deterministic (discrete) duration, convolution (cf. Section 4.5) reduces to a shift on the time line, being represented by the functional dependence given in the table.
- (Wait for or miss the bus departure) Arrive(pas,B), Depart(bupa,B) have temporal distance due to synchronization between Depart(bus,B) and Depart(pas,B). Depart(bupa,B) represents the joint occurrence of both events. Since the bus does not wait, its departure time is independent of the arrival time of the passenger at B. However, Depart(bupa,B) can fail to occur due to synchronization failure. The conditional dependency table for the example is depicted in Figure 6(e) following the equation for *Wait or miss* (cf. Section 4.5).
- (Bus ride from B to D) Depart(bupa,B), Arrive(bupa,D) have temporal distance due to duration d_{bus}, the time for the bus ride. $d_{bus} = 25$ is a deterministic (discrete) duration and is modeled the same way as walking from A to B.
- (Walking from A to C) Arrive(trpa,A), Arrive(pas,C) have temporal distance due to duration $d_{walking}$, the same time as for walking from A to B.
- (Waiting for taxi departure) Arrive(pas,C), Depart(tapa,C) have temporal distance due to duration $d_{waittaxi}$, the time needed for a taxi to show up and start. $d_{waittaxi}$ is 5, 10 or 15 with probabilities 0.7, 0.2 and 0.1 respectively. The conditional dependency table is depicted in Figure 6(d).
- (Taxi ride from C to D) Depart(tapa,C), Arrive(tapa,D) have temporal distance due to duration d_{taxi}, the time for the taxi ride. $d_{taxi} = 25$ is a deterministic (discrete) duration and is modeled the same way as walking from A to B.

The other nodes and conditional dependencies are added as follows (cf. Section 3):

- Decision node Information has values Take bus, Take taxi, and No info.
- Chance node Choice has values bus and taxi. The conditional dependency table for the example is depicted in Figure 6(f). Passengers choice follows the information given to him and in the case of no information, he will take either plan alternative with equal probability.

time	probability
0-5	0.90
5-10	0.10

(a) Arrive(trpa,A)

time	probability
10-15	0.80
15-20	0.20

(b) Depart(bus,B)

Arr(trpa,A)	0-5	5-10
5-10	1.00	0.00
10-15	0.00	1.00

(c) Arrive(pas,B)

Arr(pas,C)	5-10	10-15
10-15	0.70	0.00
15-20	0.20	0.70
20-25	0.10	0.20
25-30	0.00	0.10

(d) Depart(tapa,C)

Dep(bus,B)	10-15		15-20	
Arr(pas,B)	5-10	10-15	5-10	10-15
10-15	1.00	0.50	0.00	0.00
15-20	0.00	0.00	1.00	1.00
n-occ	0.00	0.50	0.00	0.00

(e) Depart(bupa,B)

Information	Take bus	Take taxi	No info
bus	1.00	0.00	0.50
taxi	0.00	1.00	0.50

(f) Choice

Choice	bus					...	taxi			
Arr(tapa,D)	35-40			40-45		...	45-50	50-55		
Arr(bupa,D)	35-40	40-45	n-occ	35-40	40-45	...	n-occ	35-40	40-45	n-occ
35-40	1.00	0.00	0.00	1.00	0.00	...	0.00	0.00	0.00	0.00
40-45	0.00	1.00	0.00	0.00	1.00	...	0.00	0.00	0.00	0.00
45-50	0.00	0.00	0.00	0.00	0.00	...	1.00	0.00	0.00	0.00
50-55	0.00	0.00	0.00	0.00	0.00	...	0.00	1.00	1.00	1.00
n-occ	0.00	0.00	1.00	0.00	0.00	...	0.00	0.00	0.00	0.00

(g) Arrive(pas,D)

Arr(pas,D)	35-40	40-45	45-50	50-55	n-occ
reward	1.00	1.00	0.00	0.00	0.00

(h) Reward

Figure 6: Marginal distributions and conditional dependencies

- Chance node Arrive(pas,D) is a copy of the value of Arrive(bupa,D) for Choice=bus and a copy of the value of Arrive(tapa,D) for Choice=bus. The conditional dependence table is to big to be printed in full, but illustrative parts of it are depicted in Figure 6(g).
- Value node Reward=1 for Arrive(pas,D)≤ 45 and Reward=0 for Arrive(pas,D)> 45 or non-occurrence. The table is shown in Figure 6(h).

Computations have been performed for three different scenarios. The scenarios are inspired by the introductory example and are named *pre-trip*, *on-trip I (short after departure)*, and *on-trip II (short before arrival)* respectively. The three scenarios differ only in the marginal distributions for Arrive(trpa,A) and Depart(bus,B). Both taxi and bus need the same time to D from B or C respectively. Therefore, occurrence and time of Depart(bupa,B) and Depart(tapa,C) are the determining factors for meeting the deadline.

In Figure 7, the results are shown for all three scenarios. For ease of representation, several random variables are combined into a single table, even though not all possible values need to be in the domain of the respective random variable. Departure from B or C after time 20 (marked by a horizontal line) does not meet the deadline.

The marginal distributions Arrive(trpa,A) and Depart(bus,B) differ at the probabilities marked by $\triangleleft$ from the respective previous scenario. The resulting changes for Depart(bupa,B) and Depart(tapa,C) are marked by + and − behind the probabilities.

Separate tables are given for the decision about information supply with the expected reward for any possible outcome. In the *pre-trip* scenario, Take bus is the optimal information with expected reward 0.96. This changes in the *on-trip I* scenario, where Take taxi is the optimal information with expected reward 0.7 due to the risk of missing the bus. Short before arrival (*on-trip II*), Take bus is the optimal information again with expected reward 1.00 due to the delay of the bus. The expected reward is not a probability. Since the reward is one for Arrive(pas,D) occurring before the deadline and zero for Arrive(pas,D) occurring after the deadline or not at all, it may though be interpreted as the probability of meeting the deadline.

It seems to be counterintuitive, that the probability of being in time is higher if both the passenger and the bus are delayed and not in the *pre-trip* case. However, in the *pre-trip* case, there is still a possibility of the passenger being late and the bus being in time.

The model does provide different suggestions for information supply for different scenarios represented by different knowledge in terms of marginal distributions. However, the conclusion for the decision about information supply is not straightforward. If the best possible choice is changing all the time, it might be more secure for the passenger to take the taxi and avoid informational confusion.

time	A(trpa,A)	D(bus,B)	D(bupa,B)	D(tapa,C)
0-5	0.90	0.00	0.00	0.00
5-10	0.10	0.00	0.00	0.00
10-15	0.00	0.80	0.76	0.63
15-20	0.00	0.20	0.20	0.25
20-25	0.00	0.00	0.00	0.11
25-30	0.00	0.00	0.00	0.01
n-occ	0.00	0.00	0.04	0.00

(a) pre-trip

information	exp. reward
Take bus	0.96
Take taxi	0.88
No info	0.92

(b) exp. reward for pre-trip

time	A(trpa,A)	D(bus,B)	D(bupa,B)	D(tapa,C)
0-5	0.00◁	0.00	0.00	0.00
5-10	1.00◁	0.00	0.00	0.00
10-15	0.00	0.80	0.40−	0.00−
15-20	0.00	0.20	0.20±	0.70+
20-25	0.00	0.00	0.00	0.20+
25-30	0.00	0.00	0.00	0.10+
n-occ	0.00	0.00	0.40+	0.00

(c) on-trip I (short after departure)

information	exp. reward
Take bus	0.60
Take taxi	0.70
No info	0.65

(d) exp. reward for on-trip I

time	A(trpa,A)	D(bus,B)	D(bupa,B)	D(tapa,C)
0-5	0.00	0.00	0.00	0.00
5-10	1.00	0.00	0.00	0.00
10-15	0.00	0.00◁	0.00−	0.00
15-20	0.00	1.00◁	1.00+	0.70
20-25	0.00	0.00	0.00	0.20
25-30	0.00	0.00	0.00	0.10
n-occ	0.00	0.00	0.00−	0.00

(e) on-trip II (short before arrival)

information	exp. reward
Take bus	1.00
Take taxi	0.70
No info	0.85

(f) exp. reward for on-trip II

Figure 7: Three different scenarios

6 Summary

In this report, we presented a formal model for the probabilistic evaluation of plans with imprecise knowledge about the time of future events. Bayesian networks have been employed for the reasoning about possible plan executions with nodes representing event occurrence times. The chosen approach

- supports the combination of imprecise temporal knowledge from different sources,
- allows the easy integration of temporal knowledge with varying granularity due to its foundation on continuous time and
- facilitates the comparison of the utility of different plans for the decision about alternatives with respect to the combined knowledge.

For the example used in this report, a qualitative analysis of the relationships between event occurrence times and a quantitative analysis of these relationships based upon qualitative considerations has been demonstrated. The resulting Bayesian network has been extended with a simple decision model for information supply in the form of an influence diagram.

Probabilistic plan evaluation forms a solid basis for the study of information value dynamics. Future research is directed towards the construction of decision models for reasoning about the time of information supply and multiple information supply at different points in time.

Acknowledgement This research was supported by the German Research Society, Berlin-Brandenburg Graduate School in Distributed Information Systems (DFG grant no. GRK 316). I would like to thank Prof. Lenz for his valuable guidance towards related work. The Information Logistics Project, guided by Prof. Weber and being located at Fraunhofer ISST[1] and CIS[2], serves as an inspiring environment for my work, providing both scenarios and discussions with practioners. Thanks to Susanne Busse, Juliane Dehnert, Annika Hinze, Ralf Kutsche, Roland Müller, Felix Naumann, Torsten Schlieder, Srinath Srinivasa, Asuman Sünbül, Tjark Wahnfried, Karsten Weicker and Nicole Weicker for comments on earlier drafts of this work and valuable discussions.

[1]Fraunhofer-Institut für Software- und Systemtechnik (ISST), Berlin/Dortmund

[2]Computergestützte Informationssysteme (CIS), Technical University, Berlin

References

[AFS99] Gustavo Arroyo-Figueroa and L. Enrique Sucar. A Temporal Bayesian Network for Diagnosis and Prediction. In *Proceedings of the Fifteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI–99)*, pages 13–20, San Francisco, CA, 1999. Morgan Kaufmann Publishers.

[All83] James F. Allen. Maintaining Knowledge about Temporal Intervals. *Communications of the ACM*, 26:832–843, 1983.

[Ber90] C. Berzuini. Representing Time in Causal Probabilistic Networks. In M. Henrion, R. D. Shachter, L. N. Kanal, and J. F. Lemmer, editors, *Uncertainty in Artificial Intelligence 5*, volume 10 of *Machine Intelligence and Pattern Recognition*, pages 15–28. North-Holland, Amsterdam, 1990.

[CGL94] Gianfranco Ciardo, Reinhard German, and Christoph Lindemann. A Characterization of the Stochastic Process Underlying a Stochastic Petri Net. *IEEE Transactions on Software Engineering*, 20(7):506–515, July 1994.

[DKK92] Thomas Dean, Jak Kirman, and Keiji Kanazawa. Probabilistic Network Representations fo Continuous-Time Stochastic Processes for Applications in Planning and Control. In James Hendler, editor, *Artificial Intelligence Planning Systems: Proceedings of the First International Conference (AIPS 92)*, pages 273–274, College Park, Maryland, USA, June 1992. Morgan Kaufmann.

[FAM96] George Ferguson, James Allen, and Brad Miller. Towards a Mixed Initiative Planning Assistant. In B. Drabble, editor, *Proceedings of the 3rd International Conference on Artificial Intelligence Planning Systems (AIPS-96)*, pages 70–77. AAAI Press, 1996.

[HF99] Annika Hinze and Daniel Faensen. A Unified Model of Internet Scale Alerting Services. In *Proceedings of the International Computer Science Conference*, Hong Kong, 1999.

[Hor99] Eric Horvitz. Principles of Mixed-Initiative User Interfaces. In *Proceedings of ACM CHI 99 Conference on Human Factors in Computing Systems*, volume 1 of *Characters and Agents*, pages 159–166, 1999.

[LS88] S. Lauritzen and D. J. Spiegelhalter. Local computations with probabilities on graphical structures and their application to expert systems (with discussion). *Journal of the Royal Statistical Society series B*, 50:157–224, 1988.

[Pea91] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann, 1991. (Revised 2nd Edition).

[SGS00] Glenn Shafer, Peter R. Gillet, and Richard Scherl. The Logic of Events. *ANNALSMAI: Annals of Mathematics and Artificial Intelligence*, 2000. to appear.

[Sha86] Ross D. Shachter. Evaluating Influence Diagrams. *Operations Research*, 33(6), 1986.

[Sha87] Ross D. Shachter. Probabilistic Inference and Influence Diagrams. *Operations Research*, 36(4):589–604, 1987.

[TN94] Ahmed Y. Tawfik and Eric Neufeld. Temporal bayesian networks. In *Proceedings of First International Workshop on Temporal Repres entation and Reasoning (TIME)*, Pensacola, Florida, 1994. FLAIRS.