

A Taxonomy and Representation of Sources of Uncertainty in Active Systems

Abstract. In recent years, there has been an increased need for the use of active systems - systems that include substantial processing which should be triggered by events. In many cases, however, there is an information gap between the actual occurrences of events to which such a system must respond, and the data generated by monitoring tools regarding these events. For example, some events, by their very nature, may not be signaled by any monitoring tools, or the inaccuracy of monitoring tools may incorrectly reflect the information associated with events. The result is that in many cases, there is uncertainty in the active system associated with event occurrence. In this paper, we provide a taxonomy of the sources of this uncertainty. Furthermore, we provide a formal way to represent this uncertainty, which is the first step towards addressing the aforementioned information gap.

1 Introduction and Motivation

In recent years, there has been a growing need for the use of active systems – that is, systems required to act automatically based on events. Examples of events include updates to a database, a room temperature reaching a certain threshold, and a bank account being overdrawn by more than a certain amount. The earliest active systems were active databases, *i.e.* databases in which automatic actions were required as a result of DB operations such as insertions, updates, and deletions. Currently, a major need for such active functionality is for systems in the realms of *Business Activity Monitoring (BAM)* and *Business Process Management (BPM)*. These are systems that monitor, streamline, and optimize the processes and activities which are at the heart of every business enterprise, and they comprise one of the largest of today’s emerging markets.

In order for an active system, in any domain, to react automatically to all events of interest, it must first be able to recognize these events. However, in many cases, there is a gap between the actual occurrences of events to which the system must respond, and the data generated by monitoring tools regarding these events. This gap results in **uncertainty**.

Example 1. A thermometer generates an event whenever the temperature rises above 37.5 °C. The thermometer is known to be accurate to within ± 0.2 °C. Therefore, when the temperature measured by the thermometer is 37.6 °C, there is some uncertainty regarding whether the event has actually occurred.

Another gap between the actual occurrence of events and the information available to the system is caused by the following: The information regarding the

occurrence of some events (termed *explicit events*) is signalled by *event sources* (e.g., monitoring tools such as the thermometer in Example 1). However, there are other events regarding whose occurrence explicit notification is never sent (events for which explicit notification is sent are termed *explicit events*, while others are termed *non-explicit events*). For an active system to respond to non-explicit events requires, in many cases, that the occurrences of these events be inferred based on the occurrence of other events (these events are termed *inferred events*). To facilitate such inference, several *event composition languages* have been defined (see Section 2) that make it possible to define a set of complex temporal predicates. However, uncertainty associated with explicit events also propagates to inferred events as well, as the following example shows:

Example 2. Consider a rule that states that event e_3 must be inferred whenever an event e_2 occurs after an event of type e_1 . Moreover, assume that it is known that both an event of type e_1 and an event of type e_2 have occurred, but there is uncertainty about the exact time of their occurrence, as follows: The occurrence of e_1 is known to be between time 2 and time 5, and the occurrence of e_2 is known to be between time 3 and time 7. Note that even though the rule is deterministic, the uncertainty regarding the occurrence time of e_1 and e_2 , and the fact that the occurrence of e_3 must be inferred based upon these occurrence times, generates uncertainty with regard to the occurrence of the event e_3 .

There are several possible ways to deal with uncertainty. The first (and currently the most prevalent) way is to ignore it. Ignoring uncertainty greatly simplifies the implementation of active systems. Handling uncertainty in this manner assumes (sometimes implicitly) that in the context in which the active system is used, uncertainty may be ignored without adversely affecting the outcome of the application. As an example, consider once more Example 1. A possible context for such an event is in a hospital ward, where the thermometer is monitoring the body temperature of a patient, and the purpose of the event is to alert medical staff to the possible need for medication. In this context, the uncertainty associated with the actual temperature can usually be safely ignored.

Another way to deal with uncertainty, one that has been explored in the literature (again, see Section 2), is the following: Events are generated only when they occur with absolute certainty. Events regarding whose occurrence there is uncertainty (such as the one in Example 1) are recognized, but are dealt with using a mechanism external to the active system (e.g., a human operator). In this method, the event described in Example 1 would only be generated when the temperature read is above 37.7°C , i.e., in cases in which it is certain that the temperature is above 37.5°C . In addition, whenever the thermometer reads a temperature between 37.4°C and 37.6°C , the **possible** occurrence of such an event is signaled, and is dealt with outside the application.

A shortcoming of both methods is that uncertainty is not dealt with either explicitly or quantitatively within the active system. Explicit quantification is useful for carrying out automated actions based on mechanisms such as utility theory. Automation becomes mandatory for systems such as security systems

which must respond within a very short time span to events such as Denial of Service (DoS) attacks, and security breaches (it is worth noting that both DoS attacks and security breaches are examples of events which are neither signalled nor can be inferred with absolute certainty). BAM and BPM systems are additional examples of systems in which it is an asset to be able to respond automatically to uncertain information regarding events.

Therefore, quantitative and explicit handling of event uncertainty within active systems can greatly enhance these systems' abilities. However, this requires an inference mechanism regarding the occurrence of inferred events and event composition languages that can quantitatively handle uncertain information. A first step towards creating such languages is a formal definition of a quantitative representation of uncertainty. To ensure that such a representation is capable of capturing all the sources of uncertainty, the possible sources of uncertainty must first be listed and possibly classified. Therefore, in this article we provide a taxonomy of possible sources of uncertainties relating to event occurrence. We then describe a data structure that enables capturing this uncertainty quantitatively at the level of each individual event. To the best of our knowledge, ours is the first work that attempts to list and classify the sources of uncertainty associated with events in active systems, and that provides a formal framework for representing these uncertainties.

The rest of this article is organized as follows: In section 2 we present related work. Section 3 provides a taxonomy regarding event uncertainty. Section 4 formally defines a quantitative framework for representing the uncertainty associated with events. We conclude in Section 5.

2 Related Work

This section describes two types of relevant work. Works regarding uncertainty in the context of event composition are described in Section 2.1. Section 2.2 discusses quantitative formalisms for representing and reasoning about uncertainty. Such formalisms may serve as a possible basis for representing the event related uncertainty discussed in this work.

2.1 Existing Works on Event Composition Uncertainty

Very little work exists regarding uncertainty in the context of event composition. The overwhelming majority of works in this area either completely ignore the possibility of uncertainty in event composition (*e.g.* [7] and [1]), or address it in a very cursory fashion (*e.g.* [2]). The work in [2], for example, enables each event to have a certainty measure associated with it. However, the semantics of this certainty measure are not formally defined, and as a result, the measure is, in practice, never used. To the best of our knowledge, only the work in [8] begins to provide a more formal treatment of such uncertainty. However, this work does not describe the sources and types of uncertainties associated with events. An additional relevant work is [4]. This work addresses uncertainty resulting from

the limited accuracy of clock synchronization in distributed systems. In so doing, it highlights one type of uncertainty, that associated with the time at which the event occurred. In addition, this work proposes that such uncertainty be handled by the following process: Whenever the temporal uncertainty does not propagate to the inferred events, new events may be inferred by the event composition system. If, on the other hand, it is possible that this temporal uncertainty does propagate, a process external to the event composition system (*e.g.*, a human operator) is triggered to resolve this uncertainty. Therefore, even the single type of uncertainty (uncertainty regarding the occurrence time of events) which is addressed in this work is not handled, much less quantitatively represented, within the scope of an event composition system.

2.2 Formalisms for Quantitative Representation and Reasoning about Uncertainty

The most well known framework for quantitative representation and reasoning about uncertainty is probability. Other quantitative approaches include, among others, *Lower and Upper Probabilities*, *Dempster-Shafer Belief Functions*, and *Possibility Measures* (see [3]). In addition, *Fuzzy Sets* and *Fuzzy Logic* [9] are often considered for dealing with uncertainty.

3 Taxonomy of Event Uncertainty

This section provides a taxonomy of event uncertainty as follows: Section 3.1 lays some groundwork by formally describing an event and its related data. Section 3.2 defines two dimensions that can be used to classify the uncertainties as relating to events, and Section 3.3 describes the causes of event uncertainties for the defined dimensions. ✓

3.1 Event Definition

A formal description of the concept of an event in active systems appears below in Definition 1:

Definition 1. An *event* is an **actual** occurrence or happening that is **significant** (falls within a domain of interest to the system), **instantaneous** (takes place at a specific point in time), and **atomic** (it either occurs or not) .

Examples of events include notifications of changes in a stock price, failures of IT components, and a person entering or leaving a certain geographical area.

A variety of data can be associated with the occurrence of an event. Two examples of such data are: The point in time at which an event occurred, and the new price of a stock in an event describing a change in the price of a specific stock. Some data are common to all events (*e.g.* their time of occurrence), while others are specific only to some events (*e.g.* data describing stock prices are

relevant only for stock-related events). The data items associated with an event are termed *attributes*.

Due to the information gap described in Section 1, it is necessary to differentiate between the actual event, and the information the system has regarding the event. In order to enable this differentiation, we define the notion of an *Event Instance Data (EID)*, which is the data structure that contains the information the system has regarding each event.

3.2 Dimensions of Event Uncertainty

When describing sources of event uncertainty, it is useful to classify the uncertainty according to two orthogonal dimensions: *Element Uncertainty* and *Origin Uncertainty*. The first dimension, Element Uncertainty, refers to the fact that event-related uncertainty may involve one of two elements:

1. *Uncertainty regarding event occurrence*: Such uncertainty is associated with the fact that although the actual event occurrence is atomic, i.e. the event either did or did not occur (see Definition 1), the active system does not know whether or not this event has in fact occurred. One example of such an event is the thermometer reading event from Example 1. Another example is money laundering, i.e. at any point in time, money laundering either was, or was not, carried out by some customer. However, an active system can probably never be certain whether or not money laundering actually took place.
2. *Uncertainty regarding event attributes*: Even in cases in which the event is known to have occurred, there may be uncertainty associated with its attributes. For example, while it may be known that an event has occurred, its time of occurrence may not be precisely known.

The second dimension, *Origin Uncertainty*, pertains to the fact that in an active system, there may be two types of events: Explicit events, signalled by event sources, and inferred events, which are inferred based on other events (see Section 1). In the following, events which serve as the basis for the inference of other events will be termed *evidence events*. Therefore, there are two possible origins for uncertainty:

1. *Uncertainty originating at the event source*: When an event originates at an event source, there may be uncertainty associated either with the event occurrence itself, or the event's attributes, due to a feature of the event source. Example 1, in which uncertainty regarding an event occurrence is caused by the limited measuring accuracy of a thermometer, illustrates such a case.
2. *Uncertainty resulting from event inference*: Due to some events being inferred based on other events, uncertainty can propagate to the inferred events. This is demonstrated by Example 2, in which there is uncertainty regarding the occurrence of event e_3 due to the uncertainty regarding the occurrence time of events e_1 and e_2 .

Example 1 demonstrates an event whose uncertainty originates at the event source, and concerns the actual occurrence of an event. Similarly, Example 2 describes an event whose uncertainty results from inference, and, again, involves the actual occurrence of an event. Below, two additional examples are provided: Examples 3 and 4. In Example 3, the uncertainty of an event originates from the source, but is limited to its attributes, rather than to its occurrence. Example 4 shows uncertainty regarding an event’s attributes resulting from event inference.

Example 3. Consider a case in which an event must be generated whenever the temperature read by a thermometer changes. Assume that the thermometer under discussion has the same accuracy as the one defined in Example 1. Furthermore, assume that the new temperature is an attribute of this event. In this case, there is no uncertainty regarding the actual occurrence. There is only uncertainty regarding this new temperature attribute.

Example 4. Consider a case similar to the one in Example 2, in which e_3 should be inferred whenever an event e_2 occurs after an event of type e_1 . Assume that the inferred event e_3 has an attribute a_1^3 whose value is the sum of the value of the attribute a_1^1 of event e_1 and the value of the attribute a_1^2 of event e_2 . Now assume that both e_1 and e_2 are known to have occurred with certainty, but there is uncertainty regarding the value of attribute a_1^1 . In this case there is uncertainty only regarding the value of attribute a_1^3 , and this uncertainty results from event inference.

The above examples demonstrate that the two dimensions are indeed orthogonal. Therefore, uncertainty associated with events could be mapped into one of four quadrants, as shown in Figure 1. In addition, due to the orthogonality of these dimensions, we define four types of event uncertainty: *Uncertainty regarding event occurrence originating at an event source*, *uncertainty regarding event occurrence resulting from inference*, *uncertainty regarding event attributes originating at an event source*, and *uncertainty regarding event attributes resulting from event inference*.

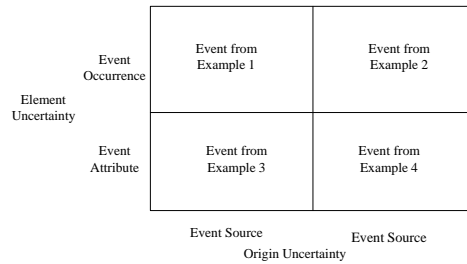


Fig. 1. Display of quadrants defined by uncertainty dimensions, as well as the mapping of the events in Examples 1-4 to these quadrants

3.3 Causes of Event Uncertainty

This section describes, at a high level, the various causes of event uncertainty, according to the dimensions defined in Section 3.2.

Causes of Uncertainty Originating at the Source Uncertainty regarding event *occurrence* originating at an event source is caused by one of the following:

1. **An unreliable source:** An event source may malfunction and indicate that an event has occurred even if it has not. Similarly, the event source may fail to signal the occurrence of an event which has, in fact, occurred.
2. **An imprecise event source:** An event source may operate correctly, but still fail to signal the occurrence of events due to limited precision (or may signal events that did not occur). This is illustrated by Example 1.
3. **Problematic communication medium between the event source and the active system:** Even if the event source has full precision, and operates correctly 100% of the time, the communication medium between the source and the active system may drop indications of an event's occurrence, or generate indications of events that did not occur.
4. **Uncertainty due to estimates:** In some cases, the event itself may be a result of a statistical estimate. For example, it may be beneficial to generate an event whenever a network Denial of Service (DoS) event occurs, where the occurrence of such a DoS event is generated based on some mathematical model. However, as the mathematical model may produce erroneous results, this event also has uncertainty associated with it.

Uncertainty regarding the *attributes* originating at the event source can also be caused by any of the above reasons. An unreliable or imprecise source may be unreliable or imprecise regarding just the attribute values. Similarly, the communication medium may garble just the values of attributes, rather than the fact that the event did or did not occur. Finally, estimates may also result in uncertainty regarding event attributes.

In distributed systems, there exists an additional cause for uncertainty regarding the special attribute capturing the occurrence time of the event. This is due to the fact that in distributed systems, the clocks of various nodes are usually only guaranteed to be synchronized to within some interval of a global system clock (see [4]). Therefore, there is uncertainty regarding the occurrence time of events as measured according to this global system clock.

Note that in both of the above cases, uncertainty regarding a specific event may be caused by a combination of factors. For example, it is possible that both the event source itself and the communication medium simultaneously corrupt information sent regarding the same event.

Causes of Inferred Uncertainty Uncertainty regarding event *occurrence* resulting from inference has the two possible causes:

1. **Propagation of Uncertainty:** An inferred event can be a result of a deterministic rule. However, when there is uncertainty regarding the events based on which the new event is inferred, there is uncertainty regarding the inferred event. This is depicted in Example 2.
2. **Uncertain Rules:** The rule itself may be defined in an uncertain manner, whenever an event cannot be inferred with absolute certainty based on other events. An example of this is money laundering, where events denoting suspicious transactions only serve to indicate the possible occurrence of a money laundering event. Usually, a money laundering event cannot be inferred with certainty based on such suspicious transactions.

Note that these two causes may be combined. That is, it may happen that not only is the inference of an event based on an uncertain rule, but also uncertainty exists regarding the occurrence (or attribute values) of one of the events which serve as evidence for this inference.

Regarding uncertainty of inferred event *attributes*, the possible causes depend on how these attributes are calculated from the attributes of the evidence events. The most intuitive way to calculate such inferred attributes is using deterministic functions defined over the attributes of the evidence events. In such a case, the only cause of uncertainty in the inferred attributes is propagation of uncertainty from the attributes of the evidence events. This is because the uncertainty regarding the event attributes is defined as the uncertainty regarding the attribute values given that the event occurred. Therefore, the rule cannot induce uncertainty regarding the attribute values of the inferred events. However, if the inference system makes it possible to define the attribute values of the inferred events in a non-deterministic manner, uncertain rules may also be a cause of inferred attribute uncertainty.

Figure 2 summarizes the causes of uncertainty

	Origin: Event Source	Origin: Event Inference
Uncertainty Regarding Event Occurrence	<ul style="list-style-type: none"> • Unreliable Source. • Imprecise Source. • Problematic Communication Medium • Estimates 	<ul style="list-style-type: none"> • Propagation of Uncertainty • Uncertain Rules
Uncertainty Regarding Event Attributes	<ul style="list-style-type: none"> • Unreliable Source. • Imprecise Source. • Problematic Communication Medium • Estimates • Time Synchronization in Distributed Systems 	<ul style="list-style-type: none"> • Propagation of Uncertainty

Fig. 2. Causes of uncertainty

4 Representing Uncertainty

To represent the uncertainty associated with events we have chosen to use probability theory. We made this decision for the following reasons:

- Probability theory is a well-understood and powerful tool.
- Under certain assumptions, probability is the only "rational" way to represent uncertainty (again, see [3]).
- There are well-known and accepted methodologies for carrying out inferences based on probability theory, involving structures known as Bayesian Networks (see [5]).
- Probability theory can be used together with utility theory (see [6]) for automatic decision making. Such automatic decision making would facilitate the implementation of automatic actions by active systems.

Once the decision to use probability theory has been made, it must be shown how probability theory can be used to represent the uncertainty associated with events. This representation is described in section 4.1. In addition, a quantitative approach based on probability theory requires the quantification of various probabilities. Section 4.2 discusses how using the taxonomy defined in this work can help in quantifying these probabilities.

4.1 Representing Uncertainty of Individual Events

As defined in Section 3.1, for each event, the information held by the event composition system is contained in a data structure called an EID. This EID must be capable of incorporating all relevant data about an event, including its type, attributes such as times of occurrence, and uncertainty quantification. In active systems with no uncertainty, each event can be represented by a single tuple of values $Val = \langle val_1, \dots, val_n \rangle$ (one value for each type of data associated with the event).

To capture the uncertainty associated with an event instance, we extend each EID to be a Random Variable (RV) such that the possible values of EID are taken from the domain $V = \{notOccurred\} \cup V'$, where V' is a set of tuples of the form $\langle val_1, \dots, val_n \rangle$. In addition, the semantics associated with the probabilities on this RV are the following: The probability that the value of E belongs to a subset $S \subseteq V \setminus \{notOccurred\}$ is the probability that event e has occurred, and that the value of its attributes is some tuple of the form $\langle val_1, \dots, val_n \rangle$, where $\{\langle val_1, \dots, val_n \rangle\} \in S$. Similarly, the probability associated with the value $\{notOccurred\}$ is the probability that the event did not occur. Such a RV is illustrated in Example 5

Example 5. In case of the thermometer related event appearing in Example 1, assume the following: whenever the thermometer reads 37.5°C , the probability that the temperature is 37.3°C is 0.1, the probability that the temperature is

37.4 °C is 0.2, the probability that the temperature is 37.5 °C is 0.5, the probability that the temperature is 37.6 °C is 0.2, and the probability that the temperature is 37.7 °C is 0.1. Therefore, the probability that the event did not occur is 0.3 and the probability that the event did occur is 0.7. Moreover, assume that whenever the event does occur, the temperature is an attribute of this event.

Assume that at time 5, the thermometer registers a reading of 37.5 °C. This information would be represented by an EID E with the following set of values: $\{notOccurred\}$ - indicating that the event did not occur, $\{5, 37.5\text{ °C}\}$ - indicating that the event occurred at time 5 with temperature 37.5 °C, $\{5, 37.6\text{ °C}\}$ - indicating that the event occurred at time 5 and the temperature was 37.6 °C, and $\{5, 37.7\text{ °C}\}$, indicating that the event occurred at time 5 with temperature 37.7 °C. Moreover, the probabilities defined over E are the following: $\Pr(E = \{notOccurred\}) = 0.3$, $\Pr(E = \{5, 37.5\text{ °C}\}) = 0.5$, $\Pr(E = \{5, 37.6\text{ °C}\}) = 0.2$ and $\Pr(E = \{5, 37.7\text{ °C}\}) = 0.1$.

Note that, by definition, the set of possible values of the EID RVs contain information regarding both the occurrence of the event and its attributes. Therefore, this representation is sufficient to represent the uncertainty regarding both occurrence and attributes.

4.2 Inferring the Probabilities

The probabilities representing the various uncertainties must be quantified for the above formal representation of event uncertainty. For uncertainties resulting from event inference, these probabilities depend on the definition and semantics of the rules. Therefore, their quantification will not be discussed in this work (for an example of such rules and semantics, see [8]).

For uncertainty originating at the event source, the uncertainty type (i.e. whether the actual occurrence of the event or an attribute value of the event) defines the probabilities that must be quantified, as follows:

1. **Uncertainty regarding event occurrence originating at the event source:** In this case, there is uncertainty because one of the following may occur: Either the actual event occurs but is not signaled, or a signal regarding the occurrence of an event is generated without the actual event having occurred. Therefore, once it is recognized that this is a cause of uncertainty, it is clear that the following probabilities must be quantified: The probability that the event occurred but was not signalled, and the probability that the event was signalled without such an event occurring. More detailed information regarding the exact source of uncertainty can be used to assist in the quantification of these probabilities. For example, if the cause is an unreliable event source, reliability information provided by the manufacturer of the event source may be used to quantify the probability. Likewise, if the source is a problematic communication medium, information regarding the network topology and the components in the network may provide useful information for this quantification.

2. Uncertainty regarding event attributes originating at an event source:

The uncertainty in this case arises from either of the following: Either an incorrect value of the attribute is provided, or no value for the attribute is provided at all. In the first case (i.e., incorrect value), a probability space that quantifies the probability of all possible values of the attribute given the value signalled by the event source must be provided. In the second case, a probability distribution of all possible values of the attribute is required. Again, more detailed information regarding the exact source of uncertainty may help both in defining the set of all possible values of the attributes, and providing the required probability space. To illustrate this, consider Example 5: Knowledge of the imprecision inherent in the thermometer makes it possible to define the value of possible ranges (e.g., $37.3^{\circ}\text{C} - 37.7^{\circ}\text{C}$ when the thermometer registered a temperature of 37.5°C). Similarly, in the case of time in distributed systems, the reading registered by a local clock on a system participating in a time synchronization algorithm is known to be within some interval of a global time (see [4]).

It is worth noting that in all of the above cases, available knowledge regarding the event source may be insufficient to uniquely define the required probabilities. In such cases, techniques such as the *Principle of Indifference* (see [3]) may be used for probabilistic quantification (the Principle of Indifference assumes that all outcomes are equally likely given no evidence to the contrary).

In many cases, the event source itself is not capable of producing the representation of an EID as defined in Section 4.1. Therefore, a component that transforms the information received by an event source to the EID format required by the active system would be useful. This is depicted in Figure 3. In this architecture, the *Event to EID Translator* component has the responsibility of transforming the information generated by the event source to the Random Variable EID format defined in Section 4.1.




Fig. 3. Probability Augmentation Architecture

To understand how this component would work, consider again the thermometer event defined in Example 1. In this case, the event source would generate an event whenever the temperature of the thermometer exceeds 37.5°C , with an attribute detailing the temperature registered. If, for instance, the registered temperature is 37.5°C , the Event to EID translator would create an EID as defined in Example 5.

The Event to EID translator would need to have knowledge of the specific event source for each event for the EID to be correctly created. In addition, assumptions such as the Principle of Indifference may also have to be incorporated.

5 Summary, Discussion and Future Work

In this paper, a taxonomy of the uncertainty relating to events in active systems was presented according to two orthogonal dimensions: Element Uncertainty, i.e., whether the uncertainty involves the actual occurrence of an event or an attribute of it; and Origin Uncertainty, which specifies whether the uncertainty originates from an event source or is a result of event inference. The value of this taxonomy is that it clearly maps the possible types of uncertainties associated with events in active systems. Such a formal definition is a prerequisite to ensuring that a quantitative representation sufficiently captures all the required information about such uncertainty. In addition, a formal representation based on probability theory was defined, and it was shown how the uncertainty taxonomy is useful in quantifying the required probability.

There are many avenues open for future research. For example, while the taxonomy provides some assistance in the quantification of probabilities, additional work is required to ensure that such quantification can always be obtained or estimated. In addition, future research should also be carried out to consider alternative representations of uncertainty, such as Dempster-Shafer belief functions. Such alternate representations may be more appropriate in some settings, and may be better suited to deal with a lack of quantitative information.

References

1. A. Adi, D. Botzer, and O. Etzion. The situation manager component of amit - active middleware technology. In *NGITS '02: Proceedings of the 5th International Workshop on Next Generation Information Technologies and Systems*, pages 158–168. Springer-Verlag, 2002.
2. A. Adi and O. Etzion. AMIT - the situation manager. *VLDB Journal*, 13(2):177–203, 2004.
3. J. Y. Halpern. *Reasoning About Uncertainty*. MIT Press, 2003.
4. C. Liebig, M. Cilia, and A. Buchman. Event composition in time-dependant distributed systems. In *Proceedings of the Fourth IFCIS International Conference on Cooperative Information Systems*, pages 70–78. IEEE Computer Society Press, 1999.
5. J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufman, 1988.
6. H. Raiffa. *Decision Analysis : Introductory Lectures on Choices under Uncertainty*. Addison-Wesley, 1968.
7. C. S. Snoop: An expressive event specification language for active databases. *Data and Knowledge Engineering*, 14(1):1–26, 1994.
8. S. Wasserkrug, A. Gal, and O. Etzion. A model for reasoning with uncertain rules in event composition systems. In *Proceedings of the 21st Annual Conference on Uncertainty in Artificial Intelligence (UAI-05)*, pages 599–606. AUAI Press, 2005.
9. L. A. Zadeh. Fuzzy sets. *Information Control*, 8:338–353, 1965.