



# Events in Multimedia Electronic Chronicles (E-Chronicles)

Utz Westermann, University of California, Irvine, USA

Ramesh Jain, University of California, Irvine, USA

---

## ABSTRACT

*The nature of information has changed significantly in the last two decades. Now information is multimedia, sensitive to its spatio-temporal roots, live, and dynamic. Current database and search technology is very limited in addressing organization, management, and access of emerging information systems. In this paper, we address some of the fundamental issues that must be addressed in general multimedia information management systems. We present multimedia electronic chronicles (e-chronicles) as an example of emerging systems that need such technology. We believe that to deal with dynamic information, events should be used as the fundamental basis in organizing and accessing information. Event models used in our approach capture the semantics involved in supporting e-chronicles. Our ideas are demonstrated in the context of an e-chronicle that is being implemented for a reconnaissance application using multiple disparate sensors.*

*Keywords:* content-based retrieval; data model; metadata; multimedia component architecture; multimedia database; person-machine interaction

---

## INTRODUCTION

Current information systems evolved from the data characteristics and application requirements of more than three decades ago. Information management has been in reactive mode for the last decade, trying to keep up with rapidly evolving novel applications, data characteristics, and data volumes. Multimedia, including text, has become ubiquitous; structured data is now a very small fraction of useful data in emerging applications; spatio-temporal and live data streams are now the norm; explora-

tion, not directed querying, is now the predominant mode of interaction. Great strides in technology areas like processing, storage, and communication were followed by new revolutions in software development and applications making novel applications emerge rapidly. We are seeing emergence of data rich applications in all aspects of our life ranging from business to healthcare, and education to homeland security. Two things are distinct in these applications in comparison to databases of yesteryears: they contain a vast amount of multime-

dia (both live and archived) data, and attention is moving away from examining isolated silos of data toward accommodating or creating more holistic pictures of situations from disparate data sources, in or near real time.

Most applications of databases and search engines have been evolving and transforming rapidly, yet the technology has remained tied to the approaches that were developed to solve mostly problems of structured and semi-structured alphanumeric persistent data a long time ago. Databases adopted and continue to extend the relational model to deal with problems in data management. Search engines derived their initial inspiration from information retrieval and developed everything around keywords. Both databases and search engines help us find information, but they use very different approaches to organize, access, and present information. One thing remains common between databases and search engines, however. Both use objects or entities as the basic concept to organize information. Databases model everything in terms of entities and limited implicit relationships among them. Search engines analyze each Web page to find keywords by throwing away everything but nouns and stemming them to find basic keywords and creating inverted file structures around these keywords. Text analysis tools, including SVM and data mining techniques, are applied to text to improve the ranking of relevant documents. These approaches only capture implicit statistical semantics in text documents.

Emerging applications increasingly deal with disparate multifarious data sources capturing measurements (including video) taken at different locations at different times to represent a holistic situation. We believe that next generation information systems pose unique challenges that require new thinking, which builds on techniques that have evolved in databases, XML, stream processing, statistical data mining, multimedia analysis, and search engines, but adopts a different perspective on all this collection of structured and unstructured data. When dealing with multimodal, live,

spatio-temporal data, it becomes obvious that *events* are at least as important, if not more, as *entities*. Current information systems are primarily organized around entities and cannot accommodate thinking about events.

In this paper, first we discuss an emerging class of application that we called Multimedia Electronic Chronicles, or e-chronicles (Jain, 2003b), and their requirements. In the last few years, e-chronicles are receiving attention in many different applications (Gong, Singh, & Jain, 2004; Jain, Kim, & Li, 2003; Kim, Podlasek, & Pingali, 2004; Pack, Singh, Brennan, & Jain, 2004; Pingali, Opalach, Jean, & Carlbom, 2002), but the focus has mostly remained on the application rather than the basic technology behind all these applications. We then outline our motivation for using events as the basic unit for organizing data and for user interaction. Next we present general architecture for e-chronicles and discuss how event-based approaches benefit in different modules of the system. We draw examples from several implemented systems in this section. We realized that a general theory and implementation of event models is essential to make progress in implementation of e-chronicles in particular and Multimodal Information Management Systems (MIMS) in general. This motivated us to develop a general model of events that we illustrate in the context of a specific e-chronicle application that we are building. Several emerging applications based on multifarious data sources capture dynamic situations. In these systems, events play a central role. Efforts to model those using traditional entity-relationships can only result in very limited dead-end solutions. Dynamic situations must be modeled using event-centered approaches.

## MULTIMEDIA ELECTRONIC CHRONICLES (E-CHRONICLES)

Humans always had fascination with history or experiences associated with past events. From early days of civilization such accounts were captured and stored in different forms such as cave paintings. Chronicles evolved as a primary mechanism of recording such experiences

once writing became a powerful communication and storage mechanism for human experiences. Historical accounts of important events in human society, minutes of a meeting, proceedings of a conference, wedding videos, surveillance videos, logs of visitors, and warehouses of all sales activities by Wal-Mart are different forms of chronicles. Advances in capture and storage technology in the last few decades are resulting in an increasing trend to record data related to important, and even not-so-important, events for later analysis. This is resulting in many novel applications with varying requirements.

Until a few years ago, it was natural to keep a record or log using alphanumeric techniques both in paper as well as in electronic form. Growth of On-Line Analytical Processing (OLAP) tools in the context of data warehouses was a result of applications that needed to understand characteristics of such large data collections. The popularity of OLAP tools is due to their ability to pull out data segments (commonly called data cubes) and process them to extract some important characteristics to understand a situation represented by the data. Data mining and related machine intelligence tools were then applied to these volumes of data, or logs, to extract more meaningful situational characteristics from this data. Collection of all these techniques has been commonly called business intelligence systems. Clearly the motivation behind OLAP and business intelligence approaches is situational understanding by analyzing logs of data.

Until electronic recording became possible, one could only keep a record of historical events by writing about it or taking pictures of it. Historical accounts written by people were subjective. Hence, it was common that the history was rewritten whenever kings changed. What we really need is a more objective recording mechanism. Photographs provide that, but only to a limited extent. Now, many corporations and legal systems require a video record of a meeting to be kept. We assume that the video is more objective than the record written by a person. And, an unedited video definitely

provides an objective recording of the proceedings in a meeting or events at a place. By placing different types of sensors at multiple locations in the environment, it is possible to record objective data related to a situation and then analyze it to understand the situation. There is an increasing trend resulting in increasing research in Continuous Archiving and Processing (Aizawa, Tacharoen, Kawasaki, & Yamasaki, 2004; Gemmell, Lueder, & Bell, 2003; Kim et al., 2004) of data research.

### **From Logs to E-Chronicles**

This trend in technology is resulting in interesting novel applications. In the last few years, there has been a lot of talk about LifeLog (Aizawa et al., 2004; Gemmell et al., 2003). LifeLog is a record of everything that a person does using video, audio, and other sensors. The goal is to record all this data, organize it so that it becomes easy to analyze the activities of this person, review the person's activities as needed, detect patterns, and organize the information. Moreover, LifeLogs of different people could be correlated to do routine tasks like coordinating meetings by considering not only calendars but also individual preferences inferred by these systems and thus provide "cognitive assistants" to them for organizing their activities. Some people (Gemmell et al., 2003) are interested in recording everything about a person's life to achieve a kind of digital immortality. As envisioned by them, digital immortality is a personal log that uses video and other recording mechanisms.

We believe that many applications of information systems belong to a new emerging field: multimedia electronic chronicles, or e-chronicles. An e-chronicle records events using multiple sensors and provides access to this data at multiple levels of granularity and abstractions, using appropriate access mechanisms in representations and terminology familiar to application users. However, if an e-chronicle contained just data streams using multiple sources of continuous high-fidelity multifarious and multimodal data streams for the life of a person, it will not be very useful. A

volume of raw logs, even high-fidelity multimedia logs, will become mostly write-only logs; they will be written but never accessed. To make them useful, e-chronicles should have continuous processing of raw incoming data into meaningful data events in each stream. Such events must be aggregated into episodes using domain models including ontologies and other useful concepts to organize and store this data and make it easily accessible to people as well as to other programs as required.

A well-designed e-chronicle will maintain all detailed records, while providing summaries of important events as well as access to events at the required level of granularity. In addition to serving as an accessible automatically created multimedia log, it could serve as a powerful mechanism to identify abnormalities, inefficiencies, and unusual behaviors as well as conditions for appropriate actions and feedback. These systems could be as proactive as the application dictates to help users in their activities. Finally, with increasing progress in privacy and security issues, it will be possible to provide different levels of access privileges to different users, thus allaying privacy concerns of users.

Many variants of e-chronicles have already appeared and are being used. Data warehouses, video surveillance systems, meeting recording systems, sensor networks, and even blogs are early forms of e-chronicles. To facilitate development of these systems, we must identify the basic research and implementation issues in these systems and then develop appropriate approaches and systems for e-chronicles as a class of systems.

One of the most important and fundamental characteristics of e-chronicles is that they contain data that represents a dynamic evolving situation. Time is the most important parameter in a chronicle and many interactions with these systems are to find situations during a particular time period at a particular location involving particular objects. Traditional data organization approaches based on relational models or keywords were not designed considering these requirements. We believe that

we need to develop some fundamental concepts, tools, and techniques to address these new requirements.

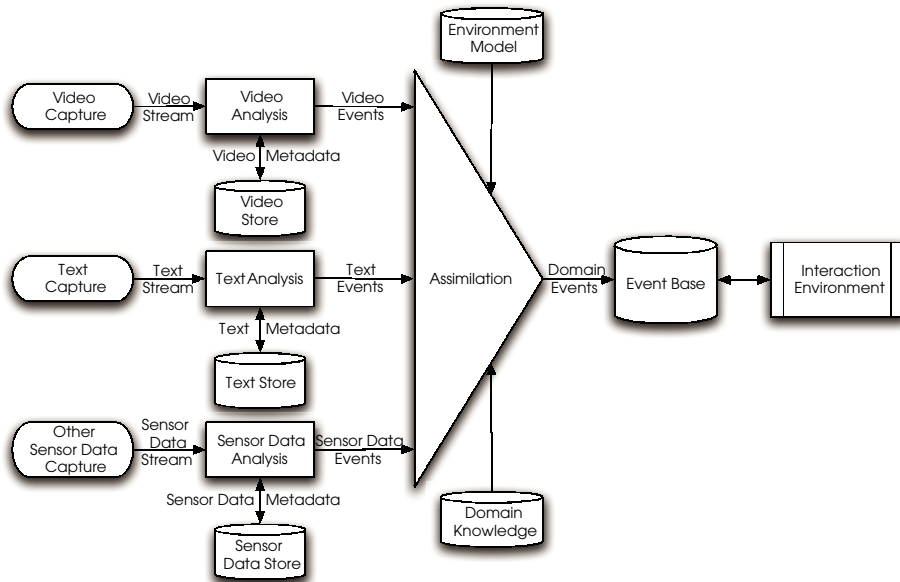
### **Events are Natural**

It is widely recognized in most areas of established sciences that static situations are different from dynamic situations and concepts, representations, and analysis tools, and techniques used for static and dynamic situations are different. For example, when Newtonian mechanics was developed, statics and dynamics were studied as independent, but obviously related, fields. Data and information management is now faced with this situation. Our current tools are very good in dealing with entities, objects, and keywords. To address the needs of Information Management (IM) in dynamic multimodal environments, we need new concepts and techniques for dynamic data in MIMS.

We believe that 'events' offer a concept to develop dynamic IM, much as 'objects' did for the current IM. Currently, 'event' is not a well-defined concept and has multiple interpretations depending on the context. This is very similar to the concept of 'object.' Objects can be defined differently based on the context. Objects become concrete when defined using implementation-oriented approaches

Events have been used in many areas of computer science, but mean different things in different areas. By developing concrete, well specified but general purpose, event models and trying to apply them to different applications, events can be made a basic concept for dynamic multimodal information management. There are strong and deep conceptual, engineering, and human centered design reasons to consider events as a primary source of organization structure for IM in those applications where one must deal with dynamic data in different modalities representing dynamic and often live situations. In a later section, we will make the concept of event concrete and show its implementation for a specific application in the context of ASSIST project, also discussed in more details later.

Figure 1. High-level architecture of an e-chronicle



## HIGH-LEVEL ARCHITECTURE OF E-CHRONICLES

An e-chronicle must capture data from disparate sources as it is coming from disparate sensors in real time. It is desirable to at least partially analyze the data and store it in properly indexed form. All the data must be organized and stored to provide interaction environment that allows finding information, insights, and experiences from a lifetime without spending a lifetime. Here we discuss major building blocks of an e-chronicle and relations among them. We also discuss some efforts to implement these blocks in various projects. Our goal here is to show the types of efforts that are going on and also to show how event-based representation and modeling will help in different modules of e-chronicles. The main components of an e-chronicle are shown in 0 and are briefly discussed here.

### Data Capture and Storage

An e-chronicle gets data from many disparate sources, from text to sensory input.

Much of this data is continuous and time sensitive. Time sensitive data should be captured and synchronized as it is collected so proper analysis and presentation can take place when required. Some data may be acquired separately and should also be synchronized with other data already captured. This problem of capturing data using multiple cameras, microphones, accelerometers, and other sensors in real time appears simple, but has been a difficult problem in real environments due to limited bandwidth available on the data buses via which sensors are connected to a computer, battery life, heat, imprecision of GPS positions, and real time processing of this data (e.g., Martin, Krösche, & Boll, 2006; Starner & Maguire, 1999). These problems appear simple, but are important hurdles in making e-chronicles practical for different applications.

Storage issues may also become significant in these applications. Large volumes of data related to the same event may be captured at different places. Moreover, since much of this data may be time sensitive, when presented



to users exact synchronization may be required. To avoid such issues careful consideration may be required in storage of data. Efforts to address these problems have started, but are clearly in very early infancy at this time (Kim, Gargi, & Jain, 2005).

### Analysis and Metadata

All data must be indexed for semantic access. This requires processing data and storing it along with all the metadata so later processes can utilize each data source. For example, an image could be interpreted in different ways depending on the thresholds for segmentation used in an early stage. The threshold is selected based on the assumptions related to the environment in which the image was taken. By storing the threshold value used in the segmentation or the environmental conditions assumed to select the threshold as metadata, the system will re-segment this and similar images based on this metadata. In fact, it is more important to use such metadata as semantics of an image in multimedia systems that are used in real world than storing the name of the camera. Analysis of data may also result in compression if appropriate events are detected and stored only in event form. Whether all original data must be recorded or one could only record detected events is an application dependent decision. The metadata recorded and events detected are all dependent on application level events and their manifestation as events in different data streams. Event detection in data streams is making good progress (Aizawa et al., 2004; Gemmell et al., 2003). Event related metadata is becoming increasingly used in simpler versions e-chronicle applications like image databases. In fact some very good work and experiments have shown the importance of events in organizing data that has relationship to time, such as photographs (Lim & Tian, 2003; Rodden & Wood, 2003). We believe that this is a very important observation that has not yet been used to its potential.

Another important research direction in this area is related to the type of metadata. Most

metadata is currently related to data. *In applications dealing with sensory data that is interpreted by systems, it is important to capture environmental parameters as metadata, but possibly more important is the fact that the parameters used by analysis programs are also captured as metadata.* This important direction was recently emphasized (Kim et al., 2005).

An e-chronicle like other MIMS must deal with events at different levels. Systems that focus on only one media like video start defining events based on media characteristics (Brown, Hampapur, Connell, et al., 2005; François, Nevatia, Hobbs, & Bolles, 2005; Nevatia, Zhao, & Hongeng, 2003) and are not effective in other applications and also to other media. To avoid media dependence of techniques and develop systems that will be scalable across media and applications, it is important to separate application events from media events and develop concepts that will allow combining cross media events and build application events for different applications and of different complexity. Some early ideas related to representing events of three types (data events, elemental events, and applications events) were presented in Jain (2003c).

### Assimilation

Designers of control and communication systems use a strong, domain-based information-assimilation approach to estimate system parameters that uses many disjoint and disparate information sources. The parameters of the mathematical system model are successively refined by observing the data as it becomes available. Each data source contributes to the model's refinement. The goal of the estimation process is to get the most precise model possible. At any time instant in this process, the model represents the current knowledge about the system based on evidence from all data sources. This approach is different from current information integration approaches used in many applications. In information integration systems, the system analyzes each data source and then combines results of data

sources with each other in a predefined configuration.

An assimilation approach can efficiently deal with real-time data by keeping only what is important for the goal of the system. Most applications collect real-time data, and it is very important to know that all data is not equal and the importance is context dependent. Information systems must learn to ignore irrelevant data.

This approach also allows a very smooth and effective introduction of semantics in the process. In our approach this semantics is represented in the event models used in the system. This perspective requires that strong event models should be available not only for application-level events, but also how these events should be transformed into elemental and then data events. This requires development of event translation algorithms. There are some information integration oriented approaches that use similar techniques to combine information from multiple sources into application events (Atrey, Kankanhalli, & Jain, 2005; Ihler, Kirshner, Ghil, Robertson, & Smyth, 2005). This approach may be useful in some applications, but approaches that will focus more on assimilating information from multiple sources are required.

### Unified Indexing

Current indexing techniques either focus on low level data and organizational issues or result in data silos based on the type of data. An approach to data silo breakdown is to build a unifying indexing system that introduces a layer on top of the metadata layer of each data silo, or disparate data source. The layer uses an event-based domain model and metadata to construct a new index that is independent of data type.

An *event model* parses the data as it is coming and assimilates data to build a situation model that reflects knowledge about the event on the basis of information collected so far. The system essentially creates a list of spatio-temporal events as they take place. This becomes the database that describes domain semantics and links these events to individual

data streams. We call this collection an *event base* that stores the name and nature of the event and all other relevant information. The relevant information may not be available at the time the event is created. If so, when it becomes available, the system attaches it to the event. Thus, the event base is an organic database that keeps growing as a result of many different processes running. The event base also stores links to original data sources, which means the system can present the appropriate media in the context of a particular event.

A user interacts with the event base directly, and the event base uses original data sources as required. This has several important implications: The system can pre-process important information related to events and objects according to its domain knowledge. It can present information using domain-based visualization, and it can provide unified access to all information related to an event, independent of the time the data became available or was entered in an associated database. Because of these characteristics, unified indexing is the backbone of an e-chronicle environment.

### Interaction Environment

In addition to informational, experiential, and relational components, an event has three important dimensions: what it is (its name and class), where it took place, and when it took place. A user navigating through an event base is interested in finding all events of a class that occurred at a particular location and time. Location can be specified using some kind of map — geographic, schematic, or conceptual. Time could be organized as a timeline or using some other mechanism.

In our earlier work (Gong et al., 2004), we presented an interaction environment in which users may select one or two event classes or navigate through class hierarchies. The depth of the hierarchy depends on the model used in the application and the data available for these events. Different selection tools could be used for different areas. Event can be selected using selection, location by using panning and zoom-

ing in a map, and time selection in time line could also be based on zooming and panning.

In the multidimensional search environment mentioned previously, it is important to introduce What-You-See-Is-What-You-Get (WYSIWYG) characteristics. Once a user selects event classes, part of the map, and time, all events and their selected attributes are presented by the system. By displaying events on a map as well as on a timeline, the WYSIWYG approach maintains event context in a multidimensional environment. A user action in one part automatically updates results in other windows. As users change the search criteria, they get immediate results with minimal latency. The system displays results continuously, making it easier to hypothesize about data relationships. It may be possible to test a hypothesis by linking such a system to data-mining tools.

By selecting an event in any of the three windows (what, when, or where), detailed information, experiences or relationships could be visualized or explored further. The system provides all relevant data sources (audio, video, or text) and any other event characteristics stored in the event base. Interaction environment should also use context and personalization to present appropriate information to each user.

## THE EC-ASSIST E-CHRONICLE

We have implemented e-chronicles for several different application domains, such as research, meetings, and work processes (for example, Gong et al., 2004; Jain et al., 2003; Kim et al., 2004). Our focus has been mostly on applying the event-centric approach to information management and exploration in those domains and demonstrating the benefits of this approach. The basic e-chronicle functionalities as discussed in the section on “High-Level Architecture of E-Chronicles” — support for the analysis and handling of metadata, data assimilation, unified indexing, and an interaction environment — were treated mainly in an ad-hoc fashion.

We are currently applying the experience from those early ad-hoc e-chronicle prototypes

in the DARPA project “Electronic Chronicling and Group Wear for Advanced Soldier Information Systems and Technology” (EC-ASSIST). While the project goal is to develop an e-chronicle that helps soldiers equipped with wearable sensors — cameras, GPS receivers, accelerometers, microphones — to analyze the course of reconnaissance missions and prepare after-mission reports, we want to develop more generic and broadly applicable solutions for the basic e-chronicle functionalities. While our project partners work on solutions for sensor data analysis and metadata extraction (IBM T. J. Watson Research Center, GeorgiaTech), information assimilation (GeorgiaTech, MIT), and interaction environments (IBM T. J. Watson Research Center), our focus is on unified event modeling and storage (see also Terpstra, Behnel, Fiege, et al., 2003).

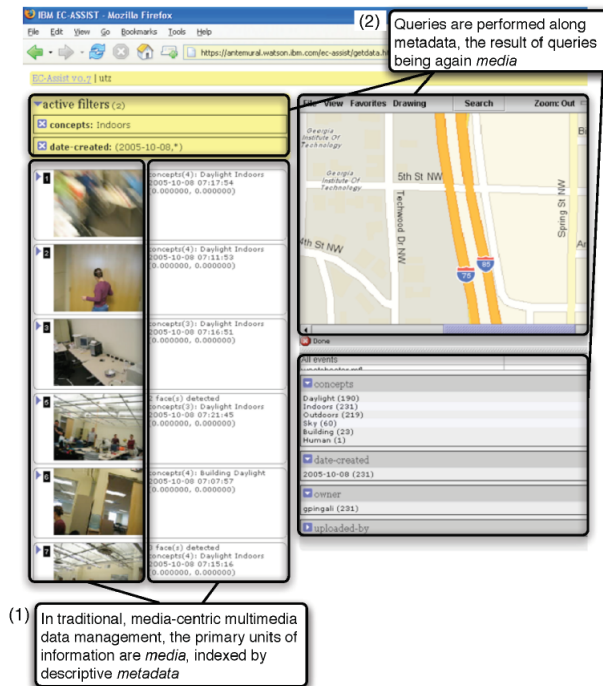
In the following, we illustrate many of the high-level rationales behind e-chronicles presented in the previous sections from the concrete perspective of the EC-ASSIST project. In particular, we motivate the need for the event-centric approach of e-chronicles to media management in EC-ASSIST instead of a traditional, media-centric approach as found in today’s multimedia databases and multimedia retrieval systems. We then outline challenges and requirements one faces with regard to the modeling of events in e-chronicles. We describe the overall design E, a generic E event model for e-chronicle applications addressing these challenges. Finally, we give an architectural overview of the event database component based on E that is used in EC-ASSIST.

## Event-Centric Media Management

Current infrastructure components for multimedia data management such as multimedia databases and multimedia retrieval systems are media-centric. They focus on media and their descriptive metadata, often even limiting themselves to a single media type such as text, image, or video (King, Popitsch, & Westermann, in press). The events documented by media — if at all considered — form just one part of media metadata. This classic media-centric view-



Figure 2. Media-centric media management illustration



point does not suffice the needs of users in EC-ASSIST; instead an event-centric e-chronicle application is needed. Also, conventional media-centric multimedia data management components do not provide an adequate underpinning for the implementation of e-chronicles.

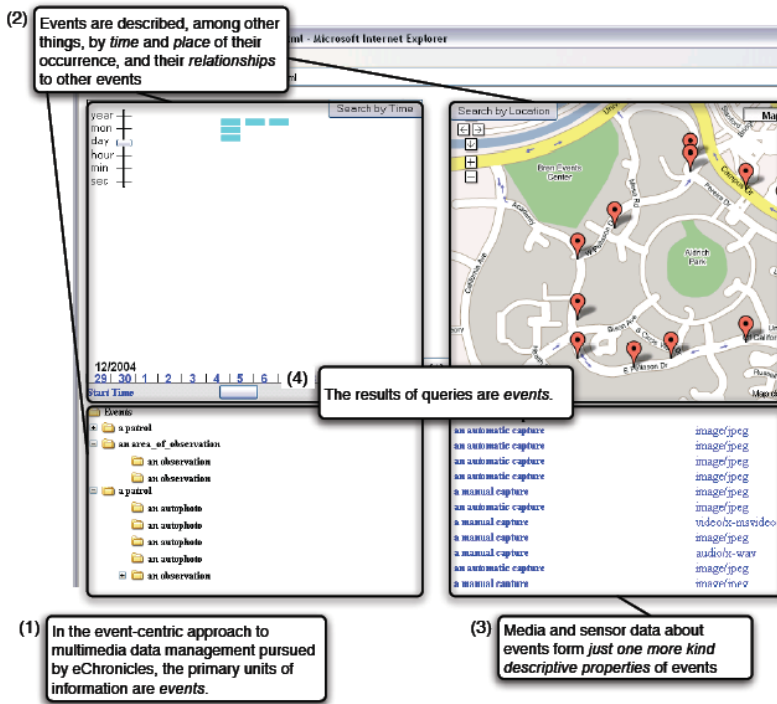
To support these claims, Figure 2 and Figure 3 illustrate the differences between traditional media-centric media management and event-centric media management in e-chronicles using an early media-centric user interface prototype for EC-ASSIST that runs on top of a classic media-centric multimedia database (Figure 2) and an experimental alternative event-centric e-chronicle user interface we have implemented on top of an event base (Figure 3). Both user interfaces directly reflect the respective approach to multimedia data management of their backend.

In the classic media-centric multimedia database for EC-ASSIST, the various media

(videos, photos, and audio recordings) created both manually and automatically during missions are indexed by metadata including the context of capture (time, GPS position of capture, soldier state), content categories determined by content analysis (indoors, gunshot, car, etc.), and further descriptive data such as the results of speech transcription and face and number plate detection. Querying is performed along these metadata, and the results of queries are media.

In the event base underlying the EC-ASSIST e-chronicle, the events that occurred during missions form the primary information unit. They are described by time and location of their occurrence and relationships to other events, especially composition relationships; all media and sensor data that document events are also considered descriptive data about events. The event base supports the querying of events along these descriptive

Figure 3. Event-centric media management illustration



properties, with the result of queries being again events.

A major problem of applying media-centric multimedia data management for EC-AS-SIST is the clear impedance mismatch between the primary objects of interest to data management and the primary objects of interest to users. Soldiers want to know of important events that occurred on their missions such as area observations, encounters, or insurgents arrested and the characteristics of those events. Media components are just sources of information among others from which insights about the events can be obtained. But media-centric data management facilitates querying of media along their metadata. This encourages the creation of media type-specific information silos. This also hinders the implementation of e-chronicles with media-centric multimedia data management components, as these typically do

not provide immediate support for the querying of events.

In e-chronicles, events also serve as a unified index onto any media and sensor data that document a given event regardless of modality. Unified indexing via events thus provides a multimodal abstraction from raw media and sensor data, which promotes the application of multimodal content and sensor analysis for event detection and inference. In EC-AS-SIST, a recent field trial revealed that audio recordings during missions are heavily distorted. This leaves the detection of events like soldiers greeting people by means of speech transcription alone very unreliable. However, the situation can be improved by assimilating speech with accelerometer data for detecting handshakes. In contrast, media-centric multimedia data management and infrastructure components typically keep different types of media

or sensor data isolated from each other, providing no direct links for supporting multimodal event detection.

Unified indexing via events has a further positive side effect on scalability. When analyzing data from hundreds of missions and selecting an area on a map view, media-centric multimedia data management will usually interpret the query such that any scrap of media captured during a mission in that area is to be returned. But the soldier more likely wants to be informed about the important events that occurred in the area, abstracting from the potentially many media that document them. E-chronicles and their event-centric media management share this interpretation.

Moreover, events are a modeling construct that can be seamlessly applied to many different abstraction levels. The event-centric approach of e-chronicles to multimedia data management in EC-ASSIST can capture low-level technical media creation and content analysis events such as “manual photo” or “gunshot.” But it can also aggregate them (e.g., via spatiotemporal clustering) to domain events like “observation” or “encounter” closer to the user’s thinking. Although media and their metadata could be regarded as media creation and content analysis events in media-centric multimedia data management, complex events cannot be represented without providing further data modeling constructs such as media folders or complex media objects (King et al., in press).

By treating events as first-class citizens independent of media, e-chronicles can handle multiple views onto the same media and sensor data gracefully. In EC-ASSIST, two soldiers may very well have different opinions about the course of a mission and the reasons for this course. In e-chronicles, different events can refer to the same media and sensor data thereby capturing different views and interpretations of these contents. By storing interpretations of content in media metadata, in contrast, media-centric multimedia data management tends to superimpose global views onto content.

In analyzing missions in EC-ASSIST, not only events are of user interest but also the associations between events, like structural, causal, temporal, and spatial relationships. Due to their focus on media, however, media-centric multimedia data management typically considers only associations between media. This makes the representation of inter-event associations with traditional media-centric multimedia data management components difficult.

Finally, events are also central concept in distributed systems and data stream management (Atrey et al., 2005; Carzaniga, Rosenblum, & Wolf, 2001; Demers, Gehrke, Hong, et al., 2005; Pietzuch, Shand, & Bacon, 2003). The event-centricity of e-chronicle applications suits the distributed nature of media and sensor data production in EC-ASSIST. It also opens up the system to future extensions such as the analysis of live events from ongoing missions or the provision of mobile notification systems that inform soldiers on patrol about important events in close by.

### **A General Event Model for E-Chronicles: Design Considerations**

In previous implementations of e-chronicles, the data models we applied for the representation and storage of events were ad-hoc, directly tailored to the particular kinds of events encountered in an application domain (Appan & Sundaram, 2004; Gong et al., 2004; Jain et al, 2003; Kim et al., 2004; Pingali et al., 2002). Our ambitions in the EC-ASSIST project are to not only come up with an event model that is only suitable for capturing events occurring during reconnaissance missions but to come up with an event model that is generally applicable to many different domains. Such a general event model would provide the foundation for reusable and interoperable event management components:


- Generic event bases for storing events
- Event query languages and event query processors based on them

- Event rule languages and event inference engines based on them
- Interoperable sensor data analysis and assimilation tools
- Distributed event notification systems
- Reusable interaction environments and user interface components for event exploration

A common multimedia event model may even constitute the germ of an Event Web (Jain, 2003a): a worldwide infrastructure of interlinked events and related media and sensor data as opposed to the network of interlinked hypermedia documents provided by the World Wide Web. In order to meet these expectations, several requirements need to be considered when designing a general event model:

- **Formal definition:** A formally defined event model will provide a foundation for the development of event query and rule languages with sound operator semantics.
- **Globally unique identification:** Events should be uniformly and uniquely identifiable, allowing one to address and interrelate events regardless of their origin.
- **Expressiveness:** An expressive event model will be more broadly applicable. It should provide rich primitives for the representation of elementary aspects of event description: the temporal aspect — the time when an event takes place, the spatial aspect — the place where an event occurred, the informational aspect — the incident signified by an event and the involved entities, the experiential aspect — media that can be consumed by the user providing documentation on the course of event, the structural aspect — the sub-events that occurred during the course of a more complex composite event, and the causal aspect — events that were the cause of a given event.
- **Media-awareness:** In order to be capable of serving as unified indexes onto media and sensor data, the event model should provide primitives for referencing media, sensor data, and parts of media or sensor data for a broad range of modalities. It should also be able to capture basic metadata about the referenced media to aid selection.
- **Media-independence:** Despite a desirable degree of high media awareness, an event model for e-chronicles should nevertheless maintain media independence. Since events may be documented by many different modalities, they do not form a descriptive property of a single medium alone and cannot reasonably be made part of its metadata; events exist independent of documenting media.
- **Knowledge-awareness:** One should be able to relate events not only to media or sensor data that document them but also to the entities of an application domain that are involved in an event. It should thus be possible to address to domain knowledge in various external information sources that may range from domains ontologies or databases to address books or calendars.
- **Relationship support:** To a user working with an e-chronicle, just as important as events are the relationships between them, such as temporal, structural, or causal ones. Representing relationships explicitly and not implicitly by references or pointers between events is essential, as it allows one to establish different, even conflicting, views of different users on the structure and cause of the same set of events.
- **Uncertainty:** Current media and sensor data analysis methods for the automatic detection of events, event properties, and associations between events are unreliable, imprecise, and possibly ambiguous. This holds especially true for very uncontrolled environments such as outdoors reconnaissance patrols in EC-ASSIST. During a recent field trial, for example, automatic detection of cars on photos as well as automatic recognition of their number plates has proven to be very unreliable. Thus, the validity of a car detection event will be highly uncertain just as the validity of a descriptive property carrying the number of the car attached to that event. A suitable multimedia event model, and indirectly, any query and rule languages

Figure 4. A typical event model



on top of it, must be able to handle such uncertainty.

- **Extensibility and adaptability:** A broadly applicable event model for e-chronicle applications needs to be able to adapt to different needs for the representation of events in different application domains. Events in a research e-chronicle certainly differ from events in the EC-ASSIST e-chronicle.

Figure 4 provides a sketch of a class of event models commonly encountered in current e-chronicle applications (Appan & Sundaram, 2004; Gong et al., 2004; Jain et al., 2003; Kim et al., 2004; Pingali et al., 2002), event querying, inference, and storage mechanisms (François et al., 2005; Kim et al., 2005; Kosch, Böszörményi, Bachlechner, et al., 2001; Liu, Gupta, & Jain, 2005; Pack et al., 2004), and distributed event notification systems and data stream management systems (Atrey et al., 2005; Boll & Westermann, 2003; Carzaniga et al., 2001; Demers et al., 2005; Luckham, 2002; Pietzuch et al., 2003; Olsson & Nilsson, 2001). This class of models uses more or less structured tuples for the representation of events. As indicated in the sketch, these tuples usually capture an event's informational aspect by elements providing an event id, the event type, and possibly a set of additional, generic attribute-value pairs to further describe the "what?" of an event. The temporal aspect is usually specified by either a time stamp or a time interval in a given unit of measure, the spatial aspect by a geographic point or region. For capturing the structural and causal aspects of an event, a tuple may hold sets of ids of subevents or causing events; for capturing the experiential as-

pect, a tuple may carry a set of URLs pointing to documenting media and sensor data.

While looking reasonable at first glance, such an event model does not adequately account for the previous design considerations:

- Uncertainty is not considered.
- Predefining the representation of an event's temporal and spatial aspect in the described way may not be adequate for every application. Time stamps and geographic points are good representations of the time and place of discrete events, such as a gunshot in the EC-ASSIST scenario. Time and place of a patrol event with a temporal or spatial extension are better represented using time intervals and geographic regions. It can also be desirable to express time and location of events in terms of temporal or spatial concepts. For example, precise geographic coordinates for capturing individual buildings of a village in terms of geographic regions may be unavailable. It can thus make sense to express the spatial aspect of events that occurred in the village by referring to spatial concepts that represent those buildings. Finally, there may be events whose time or location of occurrence is not known exactly, but only in relation to other events. For such events, it is appropriate to express the temporal and spatial aspects in form of temporal and spatial relationships to other events.
- The use of references to capture subevents and causal events imposes a single structure and causality onto an event as a global fact. But as already explained: soldiers in the ECASSIST e-chronicle application, for instance, may have different views onto the



course of a patrol and the reasons for that course.

- The experiential aspect of an event is not captured sufficiently in the sketched model, since media metadata is not considered.

### The E Generic Event Model for E-Chronicle Applications

Considering these difficulties of conventional event models to serve a broader variety of e-chronicle applications, we have chosen a different approach toward event modeling in EC-ASSIST. For the development of generic E event model, we have applied the following design principles (Terpstra et al., 2003):

- **Cover the primary objects of interest:** This includes occurrences like events and activities but also the entities — concepts in external knowledge sources and media and sensor data — to which these occurrences are related. This provides E's media- and knowledge-awareness.
- **Address the essential aspects of event description:** E explicitly takes account of the temporal, spatial, informational, experiential, structural, and causal aspects of event description. Application-specific aspects can be introduced if feasible.
- **Provide rich event descriptors:** Events can be described by attribute values of arbitrary domains, simple tags, references to concepts, media, and other sensor data, and may participate in n-ary relationships with other events. Thus, E is highly expressive.
- **Permit the use of any kind of descriptor for any description aspect:** E does not prescribe which of the descriptors to use for the representation of which aspect in an application domain. As demanded by our requirements, E is thus highly adaptable to individual application needs.
- **Offer generality:** E is a generic model that does not predefine any types of events, concepts, media, or descriptors. This makes the model highly extensible to application needs. For the definition of the applicable types in a domain, E relies on the availability of event

schema or ontology languages. Such a language is still a work in progress.

- **Support uncertainty:** Any primary object and any descriptors attached to them can be assigned confidence values, allowing one to express uncertainty not only at the level of events but also at the level of event properties.

Figure 5 gives an overview of the E event model using object-oriented modeling. The reason for using a semi-formal approach for the specification of the model was to be able to prototype and apply the model quickly. We are currently in the process of formalizing the model.

As one can see, E distinguishes three kinds of model constituents — primary objects of interest, descriptors applied to describe those objects, and the description aspects addressed by a descriptor. Each constituent is globally identified by a universally unique identifier (UUID), which is easy to produce in distributed settings without synchronization overhead. The type of a constituent, similar to the Resource Description Framework (RDF), and the schema, from which the type is taken, are identified by uniform resource identifiers (URIs) that can be freely chosen for an application. This similarity allows us to consider tools from the RDF community for the definition of event schemas. Moreover, it is possible to specify the creator of each constituent (for example, a user or software component), the time of its technical creation, and the creator's confidence in its validity. This introduces uncertainty to E and may be useful for the assimilation and synchronization of data from different sources. The confidence of a model constituent is a real number in the range from 0 to 1.0, with 1.0 meaning that the constituent's creator considers the validity of the of component as factual and 0 meaning that the constituent is considered invalid. Description aspects always have a confidence of one.

Figure 6 shows the primary objects in the E event model. These are categorized into two groups: event-like occurrences — which sub-

Figure 5. E model overview (UML class diagram)

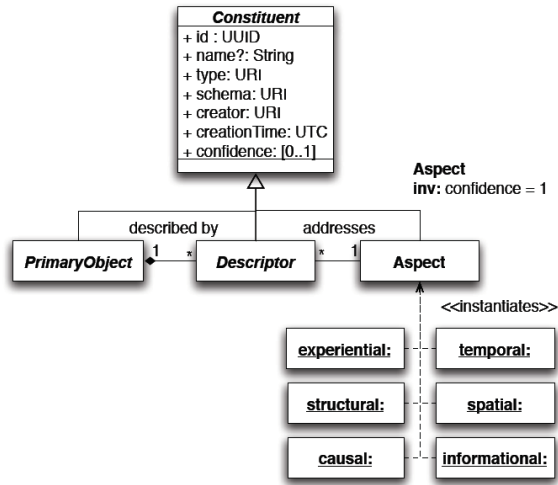
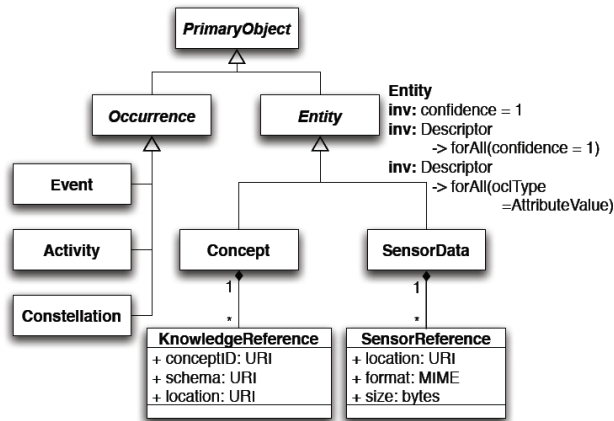


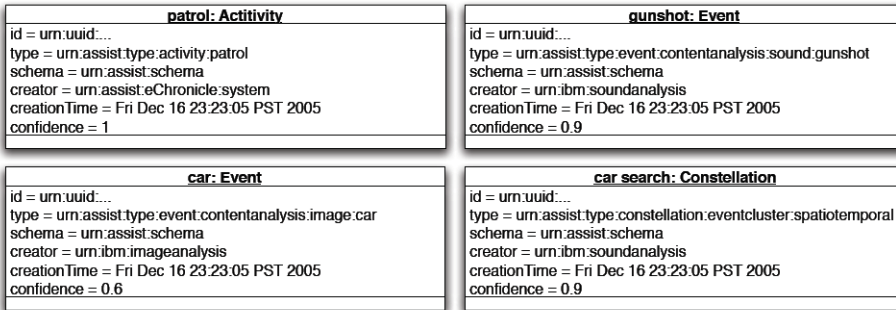
Figure 6. Primary objects in E (UML class diagram)



sume telic events, atelic activities, and constellations — and entities — which subsume abstract concepts and sensor data. Constellations constitute n-ary relationships between occurrences, including other constellations, introducing explicit support for relationships as first-class object to E. E treats constellations as oc-

currences, because the fact that events are in a certain relationship to each other may signify an important incident in an application. Concepts represent important notions of an application domain involved in occurrences. They may be augmented with references to external knowledge sources — databases, ontologies,

Figure 7. Occurrences example (UML object diagram)



calendars, and so forth — carrying further information about them. Sensor data represent arbitrary pieces of media or sensor data documenting occurrences. They may be augmented with pointers to physical storage and network access locations of those data. Concepts and sensor data establish the media- and knowledge-awareness of the model.

Figure 7 gives some examples of E primary objects as they occur in the EC-ASSIST e-chronicle, showing types, creators, creation time, and confidence values. The activity “patrol” represents a reconnaissance mission. The events “gunshot” and “car” represent events that have been automatically detected by sound and image analysis from the audio captured with one of the microphones attached to a soldier and a photo taken with a camera attached to a soldier’s helmet.

E permits occurrences to have a confidence of less than 1. In the example, this feature has been used by the content analysis components that detected the events to express the limited reliability of their applied event detection methods.

The depicted constellation represents a cluster of events closely related in time and space that may indicate a larger incident of importance in a mission: a car search in this case. More details on the idea behind the use of spatio-temporal clusters in EC-ASSIST are provided in the “E Event Base” section.

Figure 8 shows some examples of entities as they appear in EC-ASSIST. The sensor data “microphone recording” and “helmet photo” capture the audio and image data from which the events “gunshot” and “car” of the previous figure have been derived. The sensor data provide a logical abstraction; the raw data, which have been uploaded to a Web server after the mission, are referenced by corresponding sensor sources. The concept “Aberdeen training grounds” constitutes a spatial concept that can be used for the description of the patrol locations. It points to a GIS database that offers further data about that location using a knowledge source. E regards concepts and pieces of sensor data as facts. Unlike occurrences, they always have a confidence of 1.

Figure 9 shows the different kinds of descriptors E offers. Descriptors are attached to primary objects and may not exist on their own, thus forming second-class objects in E. Any primary objects can be described by arbitrary attribute values. Furthermore, document descriptors allow occurrences to refer to sensor data that document them. The kind of documentation provided is given by the document descriptor’s type. E offers selectors to address only parts of sensor data with document descriptors. Occurrences may also refer to concepts of an application domain using reference descriptors. The role the concept plays for an

Figure 8. Entities example (UML object diagram)

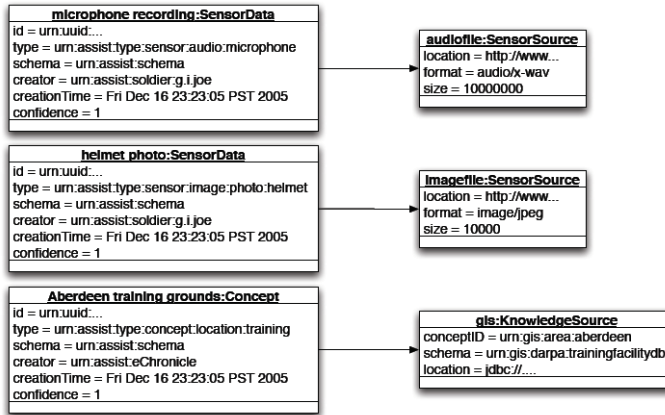
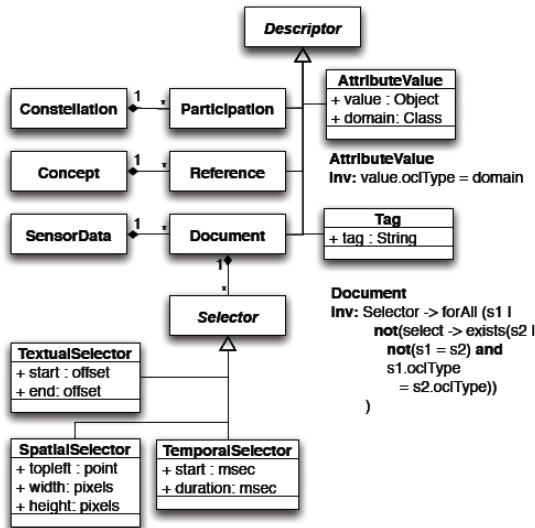


Figure 9. Descriptors in E (UML class diagram)



occurrence is again given by the descriptors type. Quite similar in purpose to reference descriptors are tag descriptors. But instead of linking to a formal, predefined concept, the concept referred to by a tag is specified by a free

text string. This enables “folksonomy” applications where events are being relaxedly annotated with free text tags in a collaborative effort. Finally, an occurrence may participate in a constellation via participation descriptors. The

role in which the occurrence participates is again given by the descriptor type.

It is our intention with E to focus on the representation events and other occurrences in different application domains. While we want to relate occurrences to the concepts of an application domain involved, it is not our goal to also describe the concepts of the domain and their interrelationships and thereby re-implement what is already available with established knowledge representation models such as RDF or Topic. Therefore, we deliberately constrained E not to allow descriptors to be attached to entities, with the sole exception of attribute values. The latter exception turned out to be useful for pragmatic reasons, since one may like frequently accessed basic metadata about sensor data and properties of concepts to be directly represented in E.

Figure 10 gives examples of descriptors that could be attached to the occurrences and entities of Figure 7 and Figure 8 in the EC-ASSIST scenario. For reasons of space, the figure does not show the various description aspects the descriptors address as the independent model constituents that they are but indicates them using the UML stereotype notation within the descriptors. Attribute value descriptors are used to represent the temporal and spatial aspects of the “gunshot” and “car” events by means of time points and GPS positions. These can be obtained by correlating time and GPS sensor data with content analysis. Using document descriptors, both events refer in their experiential aspect to the “microphone recording” and “helmet photo” sensor data from which they have been detected via content analysis. Selectors are used to demark the occurrence of the events in the sensor data more precisely. The “gunshot” and “car” events further participate in the “car search” constellation via participation descriptors. This indicates that they are members of the spatio-temporal cluster of events represented by the constellation. A tag descriptor has been manually attached to the constellation by a soldier during after-mission analysis to indicate that the larger incident denoted by the event cluster is a car search.

The temporal aspect of the “patrol” activity is represented by an attribute value carrying a time interval spanning the duration of the mission. A reference descriptor pointing to the “Aberdeen training grounds” concept is used to capture the spatial aspect of the activity.

Descriptors of occurrences may have a confidence value of less than 1. This permits E to take account of the uncertainty behind descriptors that have been automatically extracted out of media or sensor data with content and sensor data analysis methods of limited reliability. In the example, this is used for the attribute value descriptor attached to the “car” event that holds the license plate number of the detected car extracted from the “helmet photo” sensor via unreliable OCR methods. Attribute value descriptors that are attached to entities always have a confidence value 1. In the example, the “helmet photo” sensor data features an attribute value descriptor capturing the color histogram of the picture.

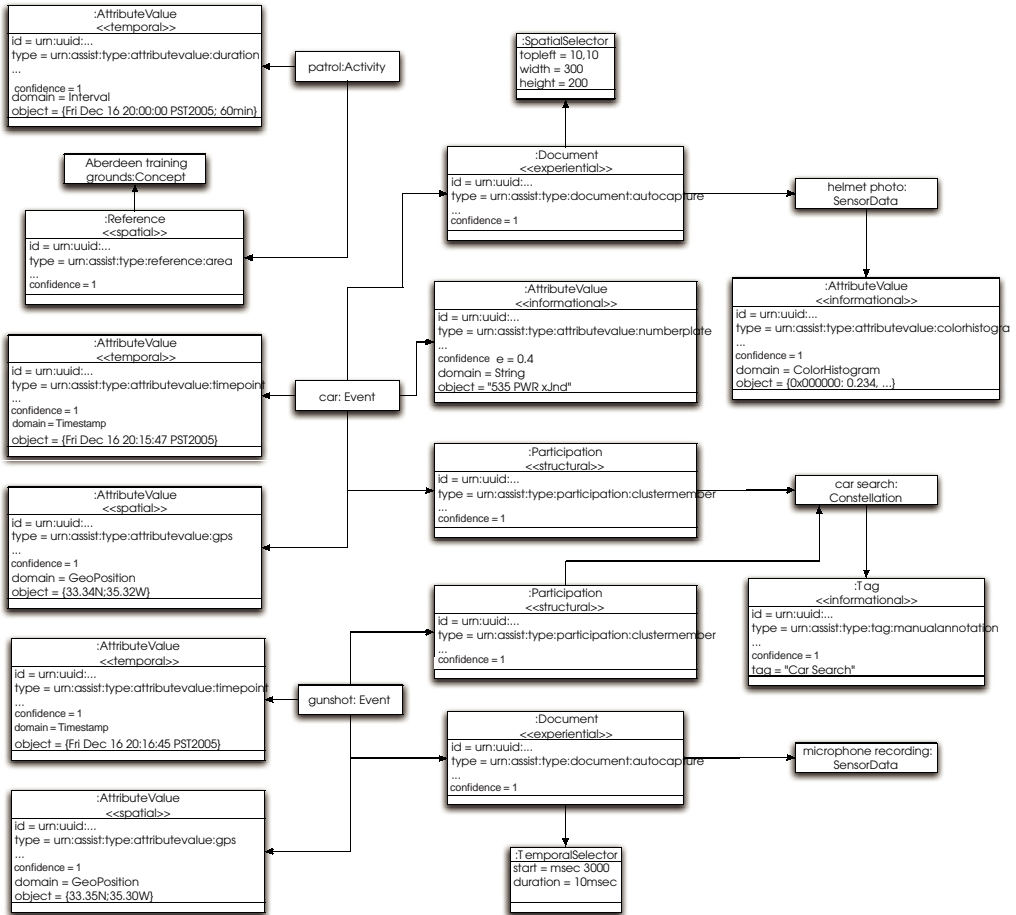
### The E Event Base

On the basis of E, we have implemented an initial event base prototype for EC-ASSIST. Given the genericity of E, however, this event base will be applicable to a broad spectrum of e-chronicle applications.

Figure 11 gives a high-level overview of the event base. For its implementation, we have made use of the Ruby on Rails Web application framework. The core of the event base is made up of a relational database server (currently MySQL 5.0) hosting a database schema that constitutes a straightforward translation of the object model for E outlined in the “E Generic Event Model for E-Chronicle Applications” section to the relational data model. Using the lightweight object-relational mapping framework offered by Rails, we then provide an object-oriented 1:1 implementation of the E object model on top of the relational schema. The implementation provides basic accessors to and constructors for E’s constituents as well as navigation methods for traversing event structures expressed in E.



Figure 10. Descriptors example (UML object diagram)

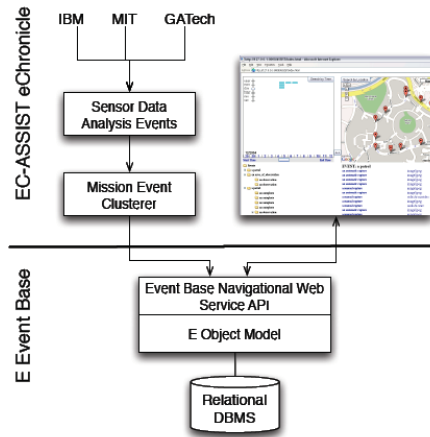


Based on the E object model, we have then implemented a Web service API with basic functionality for adding, modifying, deleting, and retrieving the model constituents stored in the event base for e-chronicle applications. As we have not developed an event retrieval language yet, the API offers primarily navigational access to the event base. Given that the predominant access mode to event data in e-chronicles is by interactive exploration, this is adequate. However, it is our goal to merge our past research efforts with regard to event query engines (Kim et al., 2005; Liu et al., 2005)

with the E event model to also obtain adequate declarative querying facilities for E.

The EC-ASSIST e-chronicle is the first application running on top of the E event base. The user interface, which we have already introduced in the “Event-Centric Media Management” section, makes use of the navigational Web service API for the browsing and filtering of events along time, location, and structure. Via that API, events that occurred during a patrol are also ingested into the event base. This is a two-step process. In a first step, the events detected by the sensor data and media analy-

Figure 11. E Event Base Architecture



sis methods of our project partners are gathered after the mission. These constitute low-level, sensor data events focused on one modality — such as the “car” and “gunshot” event in the examples of the “E Generic Event Model for E-Chronicle Applications” section — all of which are indexed by time and GPS position of their occurrence. There is the need to assimilate these data-level events from different modalities and to aggregate them to the real-world mission events they document such as “car search” that are closer to the thinking of soldiers. For this purpose, we are evaluating the use of spatio-temporal clustering techniques. The basic rationale behind this approach is that if something of interest is going on in a mission, there will be a spatio-temporal concentration of events — either because soldiers create more media to document the events going on or simply because they stay in the area of the event for a longer time. Moreover, sensor data analysis results in such an uncontrolled environment as a military reconnaissance patrol tend to be very unreliable, leaving time and GPS position the most reliable characteristics available for high-level event detection. First results obtained from the application of spatio-tempo-

ral clustering for mission event detection on real-world patrol data from a recent evaluation of EC-ASSIST with real soldiers are very promising.

The E event base currently lacks an event schema language for E. As there is no component in the event base that ensures the validity of the stored events with regard to the allowable event types and their permissible properties in a domain, applications such as the EC-ASSIST e-chronicle are thus burdened with extra code to ensure consistency of the event base contents themselves. We are evaluating the use of ontology languages such as RDF Schema and OWL from the Semantic Web domain with regard to their suitability for the specification of event schemas for E.

## CONCLUSION

We presented a class of emerging applications, Multimedia Electronic Chronicles, that are becoming increasingly popular in disparate application areas ranging from personal information management to business activity monitoring. All these applications must organize time sensitive multimodal data and provide access to this data in semantic terms. Current data-

base and search engine techniques are unable to address these applications effectively because both these areas were developed for static entities, and e-chronicles are fundamentally related to dynamic situations and must deal with spatio-temporal data.

We presented event-based organization of e-chronicles as a potential solution. The unified event model could be used in multiple applications. This model can deal with informational, experiential, and basic data related to events and can allow relational and causal links to explore different types of relationships among events. For addressing issues in analysis of experiential data and maintaining causal and relational information involving time, novel approaches to capturing semantics are required. These approaches may require development of powerful event ontologies for different application domains. Most current systems capture object information, but are not designed for event information. Similarly, current approaches to address metadata are powerful for representing semantics of data. However, in many applications the semantics depends on the parameters used in an interpretation process, and current systems and standards have not even started thinking about these cases.

Though, we are involved in multiple applications of e-chronicles, in this paper we discussed a system called ASSIST that is being developed under the sponsorship of DARPA in partnership of Georgia Tech, IBM, MIT, and UC Irvine. We briefly presented the project here, and then focused on presenting details of the event model and its implementation in e-chronicle-like applications.

The research presented here is in its early stage, but is addressing a very important and fundamental problem. The issue of organizing and semantically accessing multimodal situational data is becoming increasingly important in all aspects of the next generation of information systems. The semantic of experiential data such as images, video, and audio is more complex than the semantics in text. Such experiential data is becoming an increasingly larger percentage of our information systems,

including the Web. It is important to address issues that will allow capturing and representation of semantics in these systems so our systems can reason with them.

## ACKNOWLEDGMENTS

We are very thankful to Bo Gong and Srikanth Agaram for contributing to discussions and ongoing implementations of the ASSIST project. We also thank Susanne Boll, Rahul Singh, Sharad Mehrotra, and several other members of RESCUE team for discussions related to the role of event models in many emerging applications.

## REFERENCES

- Aizawa, K., Tacharoen, K., Kawasaki, S., & Yamasaki, T. (2004). Efficient retrieval of life log based on context and content. In *Proceedings of the 1st ACM Workshop on Continuous Archival and Retrieval of Personal Experiences (CARPE'04)*, New York.
- Appan, P., & Sundaram, H. (2004, October). Networked multimedia event exploration. In *Proceedings of the 12th ACM International Conference on Multimedia (MULTIMEDIA'04)*, New York.
- Atrey, P. K., Kankanhalli, M., & Jain, R. (2005, October). Time-line based information assimilation in multimedia surveillance and monitoring systems. In *Proceedings of the Video Surveillance and Sensor Networks, Singapore*.
- Bacon, J. (2001). Generic support for distributed applications. *IEEE Computer*, 33(3).
- Boll, S., & Westermann, U. (2003, November). MediÆther — An event space for context-aware multimedia experiences. In *Proceedings of the 2003 ACM SIGMM Workshop on Experiential Telepresence (ETP'03)*, Berkeley, CA.
- Brown, L., Hampapur, A., Connell, J., et al. (2005, October). IBM SmartSurveillance System (S3): An open and extensible architecture for smart video surveillance. In *Proceedings of the 10th IEEE International Conference on Computer Vision (ICCV 2005)*, Beijing, China.

- Carzaniga, A., Rosenblum, D. S., & Wolf, A. L. (2001). Design and evaluation of a wide-area event notification service. *ACM Transactions on Computer Systems, 19*(3).
- Demers, A., Gehrke, J., Hong, M., et al. (2005). *A general algebra and implementation for monitoring event streams* (Tech. Rep. No. TR2005-1997). Cornell University Library.
- François, A. R. J., Nevatia, R., Hobbs, J., & Bolles, R. C. (2005). VERL: An ontology framework for representing and annotating video events. *IEEE MultiMedia, 12*(4).
- Gemmell, J., Lueder, R., & Bell, G. (2003, November). The MyLifeBits lifetime store. In *Proceedings of the 2003 ACM SIGMM Workshop on Experiential Telepresence (ETP'03)*, Berkeley, CA.
- Gong, B., Singh, R., & Jain, R. (2004, October). ResearchExplorer: Gaining insights through exploration in multimedia scientific data. In *Proceedings of the 6th ACM SIGMM Workshop on Multimedia Information Retrieval (MIR'04)*, New York.
- Ihler, A., Kirshner, S., Ghil, M., Robertson, A., & Smyth, P. (2005). *Graphical models for statistical inference and data assimilation*. Unpublished manuscript.
- ISO/IEC JTC1/SC29. (2003). *Information technology — Multimedia content description interface — Part 5: Multimedia description schemes* (International Standard 15938-5:2003). International Organization for Standardization/International Electrotechnical Commission (ISO/IEC).
- Jain, R. (2003a). Experiential computing. *Communications of the ACM, 46*(7).
- Jain, R. (2003b). Multimedia electronic chronicles. *IEEE MultiMedia, 10*(3).
- Jain, R. (2003c, March). Out-of-the-box data engineering: Events in heterogeneous environments. In *Proceedings of the 19th International Conference on Data Engineering (ICDE 2003)*, Bangalore, India.
- Jain, R., Kim, P., & Li, Z. (2003, November). Experiential meeting system. In *Proceedings of the 2003 ACM SIGMM Workshop on Experiential Telepresence (ETP'03)*, Berkeley, CA.
- Kim, P., Gargi, U., & Jain, R. (2005, November). Event-based multimedia chronicling system. In *Proceedings of the 2nd ACM Workshop on Continuous Archival and Retrieval of Personal Experiences (CARPE'05)*, Singapore.
- Kim, P., Podlaseck, M., & Pingali, G. S. (2004, October). Personal chronicling tools for enhancing information archival and collaboration in enterprises. In *Proceedings of the 1st ACM Workshop on Continuous Archival and Collaboration in Enterprises (CARPE'04)*, New York.
- King, R., Popitsch, N., & Westermann, U. (2006). METIS: A flexible foundation for the unified management of multimedia assets. *Multimedia Tools and Applications*.
- Kosch, H., Böszörményi, L., Bachlechner, A., et al. (2001, September). SMOOTH — A distributed multimedia database system. In *Proceedings of the 27th International Conferences on Very Large Databases (VLDB 2001)*, Rome, Italy.
- Lim, J. H., & Tian, Q. (2003). Home photo content modeling for personalized event-based retrieval. *IEEE MultiMedia, 10*(4).
- Liu, B., Gupta, A., & Jain, R. (2005, November). MedSMan: A streaming data management system over life multimedia. In *Proceedings of the 13th ACM International Conference on Multimedia (MULTIMEDIA'05)*, Singapore.
- Luckham, D. (2002). *The power of events*. Boston: Addison-Wesley.
- Martin, J. D., Krösche, J., & Boll, S. (2006, March). Dynamic GPS-position correction for mobile pedestrian navigation and orientation. In *Proceedings of the 3rd Workshop on Positioning, Navigation and Communication 2006 (WPNC'06)*, Hanover, Germany.
- Nevatia, R., Zhao, T., & Hongeng, S. (2003). Hierarchical language-based representation of events in video streams. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW'03)*, Madison, WI.
- Olsson, D., & Nilsson, A. (2001). MEP — A Media Event Platform. In *Proceedings of the*

- 34th Hawaii International Conference on System Sciences (HICSS-34)*, Maui, HI.
- Pack, D., Singh, R., Brennan, S., & Jain, R. (2004). An event model and its implementation for multimedia representation and retrieval. In *Proceedings of the 2004 IEEE International Conference on Multimedia and Expo (ICME'04)*, Taipei, Taiwan.
- Pietzuch, P. R., Shand, B., & Bacon, J. (2003, June). A framework for event composition in distributed systems. In *Proceedings of the 4th ACM/IFIP/USENIX International Conference on Middleware (Middleware '03)*, Rio de Janeiro, Brazil.
- Pingali, G. S., Opalach, A., Jean Y. D., & Carlbom, I. B. (2002). Instantly indexed multimedia databases of real-world events. *IEEE Transactions on Multimedia*, 4(2).
- Rodden, K., & Wood, K. R. (2003, April). How do people manage their digital photographs? In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'03)*, Ft. Lauderdale, FL.
- Seraja Pvt. Ltd. (2006). *The EventWeb*. Retrieved March 6, 2006, from <http://www.seraja.com>
- Starner, T., & Maguire, Y. (1999). Heat dissipation in wearable computing aided by thermal coupling with the user. *ACM Journal on Mobile Networks and Applications*, 4(1).
- Snoek, C. G. M., & Worring, M. (2005, August). Multimedia event-based video indexing using time intervals. *IEEE Transactions on Multimedia*, 7(4).
- Terpstra, W. W., Behnel, S., Fiege, S., et al. (2003, June). A peer-to-peer approach to content-based publish/subscribe. In *Proceedings of the 2nd International Workshop on Distributed Event-Based Systems (DEBS'03)*, San Diego, CA.
- Westermann, U., & Jain, R. (2006, April). E — A generic event model for event-centric multimedia data management in eChronicle applications. In *Proceedings of the 2006 IEEE International Workshop on Electronic Chronicles (eChronicle 2006)*, Atlanta, GA.

*Utz Westermann is a postdoctoral scholar at the University of California, Irvine. He is currently working on event-centric multimedia data management. Before coming to California, he was an ERCIM Fellow at VTT Electronics Oulu, Finland. Dr. Westermann obtained his doctoral degree from the Technical University of Vienna, Austria, and his diploma degree from the University of Ulm, Germany.*

*Ramesh Jain is currently the Bren professor in Information & Computer Sciences at University of California, Irvine. He is a Fellow of ACM, IEEE, AAAI, IAPR, and SPIE. He recently co-founded Seraja to address needs of emerging EventWeb. Earlier Ramesh co-founded three companies Imageware, Virage, and PRAJA. He was the founding EIC of IEEE Multimedia and is chairman of ACM SIG MULTIMEDIA.*