Eventory – An Event Based Media Repository

Xiang-jun Wang Swathi Mamadgi Atit Thekdi Aisling Kelliher Hari Sundaram
Arts Media and Engineering Program, Arizona State University
e-mail: {xj, swathi.mamadgi, athekdi, aisling.kelliher, hari.sundaram}@asu.edu

Abstract

This paper focuses on the development of an event driven media sharing repository to facilitate community awareness. In this paper, an event refers to a real-world occurrence that unfolds over space and time. Our event model implementation supports creation of events using the standard facets of who, where, when and what. A key novelty in this research lies in the support of arbitrary event-event semantic relationships. We facilitate global as well as personalized event relationships. Each relationship can be unary or binary and can be at multiple granularities. The relationships can exist between events, between media, and between media and events. We have implemented a web based media archive system that allows people to create, explore and mange events. We have implemented an RSS based notification system that promotes awareness of actions. The initial user feedback has been positive and we are in the process of conducting a longitudinal study.

1. Introduction

This paper focuses on the development of an event driven media sharing repository. The key focus in this research is to create event driven mechanisms that facilitate community awareness. The long term goal is extract and track thematic communities that emerge through user interaction.

We are interested in community awareness. The problem of understanding community dynamics is an important one at the authors' institution (Arts Media and Engineering program). We are an interdisciplinary program (http://ame.asu.edu) drawing students and faculty from over ten disciplines. We collaborate via large teams on large scale multimedia systems. However, it has been very difficult to reflect on *how exactly* ideas / concepts from one domain influence the research output in other disciplines, or contribute to the project. While it is possible to speculate on links, or semantic similarities between sub-projects through manual intervention, we wish analyze the community behavior at a fine grained time-scale, as well as reveal these emergent linkages back to the users who are pursuing thematically similar ideas. Hence we proposed to develop a novel archive of human activity specific to members of the AME program (audio / video, publications, presentations, blog pages, sensor data) that additionally analyzed interactions of members of the community on that archive at fine time scale. We have termed this archive "eventory" – a media archive of events corresponding to meaningful activities at AME. The event based repository can reveal rich semantics with respect to everyday activities, and is readily generalizable to other social groups that interact with media.

Event based media organization is important in online narrative construction and sharing [13]. In the past, an amateur moviemaker was generally restricted to showing their personal narratives in the home environment. While the everyday publishing and sharing of personal media online is becoming more prevalent, as evidenced by sites such as YouTube [3] and Flickr [1], it is still a non-trivial exercise to create works for these sites that fulfill our narrative imagination and truly reflect the depth of our experience.

Media archiving for social networks seem to follow the dichotomy of media centric (e.g. Flickr / YouTube), and event centric (e.g. SEraja [2]) systems. The majority of commercial media archiving systems are media centric and focused on the collection of singular artifacts. Most systems remain specific to a particular type of media form, and few support multiple media types. For example, Flickr supports the sharing of images files only, while Youtube deals exclusively with video files. Nevertheless, these popular web systems have well developed infrastructure and intuitive, user-friendly interfaces. They provide multiple ways for exploring media and have good support for tagging and searching. To some extent, they also support the grouping and subscribing functions which are important in the social networks.

Event centric media organization is an emerging area of research. eChronicles [12] is an information system that provides rich documentation of real-world events. MediÆther [5] offers a distributed system for the propagation of and notification about multimedia events. SEraja is an event repository that supports the creation and search of events all around the world. However, these event based systems do not yet allow support for event-event relations, including temporal structures and arbitrary semantic relationships.

In our paper, an event refers to a real-world occurrence that unfolds over space and time. For example, a walk on

the beach, the birth of a baby, a performance by an AME student are all events in terms of this definition. Events are described using the facets of who, where, when, what and are re-mediated using images, sounds, videos and free text. These facets and the accompanying media support the understanding of events. We note that an event can have multiple media used to describe the event - the event "trip to china" may have hundreds of photographs associated with it. A single media clip can appear to belong to multiple events. This is complementary to the idea that an event can have multiple media to describe it. Our event model implementation supports the creation of events using the standard facets of who, where, when and what.

A key novelty in this research lies in the support of arbitrary event-event semantic relationships. This is important in extracting semantics out of community interaction. We facilitate global as well as personalized relationships between events. Global relationship types allow users of the system to establish a common understanding of the relationships between events. Personalized relationship can be defined by people's own words, languages and expressions. Each relationship in the eventory needs several properties to be defined. It must have a unique name, and the relationship type can be unary or binary. A binary relationship of the form x R y needs to define the attribute types of the entities x and y respectively as well as the relationship type. It also needs to define properties of the relation R. The relationships can be at multiple granularities – they can exist between events, between media, and between media and events.

We have implemented a web based media archive system that people can create, explore and mange events. Users can associate multimedia files with these events by uploading and sharing their media collections including images, videos, audios, documents, etc. The system supports complex event structures by allowing users to specify arbitrary relationships between media as well as between events. Through a notification mechanism, the system allows users to become aware of other users acting on their uploaded media, as well as events. We are currently in the process of conducting a longitudinal study to assess the impact of the system on community growth – the initial user feedback has been positive.

The rest of paper is organized as follows. The next section discusses the media archival problem. In Section 3 introduces the idea of events in terms of facets. Section 4 presents the event link semantics, while Section 5 discusses our current implementation. In Section 6, we discuss the current work and then summarize and present conclusions.

2. The Media Archival problem

In this section, we describe the current media archival problem, compare media sharing with experience sharing, explain why we need to improve current media archiving system and define overall challenges.

2.1 Media sharing vs. Experience sharing

Before examining the idea of media sharing in detail, let us first examine the photograph depicted in Figure 1. This photo was discovered randomly on Flickr and was not annotated with any title, tags or description. We can infer that the man show in the photograph is athletic and possibly engaged in some form of sports activity, although we cannot be sure. Without any supporting information or description, it is difficult to understand the significance or meaning of this image. In this case, we are experiencing the sharing of media, but have not the sharing of an understanding of the experience itself.

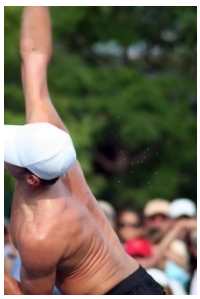

Figure 1: Untitled from Flickr

This is a familiar occurrence with online media sharing systems. The context for the media representation is missing as many people upload content without bothering with the hassle of adding additional metadata. As the image-taker, they are already familiar with the people, location and events shown in the image itself. They are not necessarily thinking of the needs of an audience entirely unfamiliar with the events depicted. However, we can possibly employ other strategies to help decipher the meaning of this mystery image. For example, we discovered several other photos [Figure 3 and Figure 2] uploaded to Flickr around the same time that seem to document the same or at least a similar event

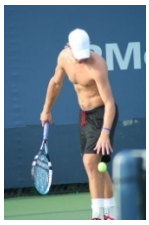

Figure 2: Untitled

In considering the additional information shown both in the photographs themselves and in the accompanying metadata, we can now understand the photos to be of tennis player Andy Roddick participating in a match in the US Open Tennis Championships. If we are curious to find out more about this match, we can use the descriptors from the photograph title ("Roddick") and content ("US Open") to search for video of the event on YouTube for example.

Figure 3: Flickr photo titled "Roddick"

A more efficient and readily understandable method for organizing and accessing such media content revolves around the idea of an event. An event called "US Open Men's Singles Final" is created that describes the location, time and players involved. All subsequently relevant uploaded images, audio and video content become associated with this event. By integrating descriptions of the event with multiple forms of media content, we are now no longer just sharing individual discrete media artifacts, but rather growing and exploring a shared experience.

3. Event Characterization

An event refers to a real-world occurrence that unfolds over space and time. For example, a walk on the beach, the hurricane of 2005, a trip to China, are all events in terms of this definition.

Typically, events are described in multimedia research using the facets of who, where, when, what and are re-mediated using images, sounds, videos and free text. These facets and the accompanying media support the understanding of events. We note that an event can have multiple media used to describe the event - the event "trip to china" may have hundreds of photographs associated with it. This event definition draws upon recent work by Westermann and Jain [13]. We note in the passing that events can comprise other events (e.g. a trip to china may include a walk on the beach) - the nature of the hierarchy is typically application and context dependent.

Multimedia research is influenced in its description of events by well known journalistic practice. A key maxim in journalism is to use the six interrogatives - *who*?, *when*?, *where*?, *what*?, *why*? and *how*? to develop a comprehensive reportage of the event. This has become known as the five W's and one H.

A single media clip can appear to belong to multiple events. This is complementary to the idea that an event can have multiple media to describe it. For instance, a short video clip may be described as "a walk on the beach in Santa Barbara", "a technical discussion with colleagues on the beach", "debates about the Wednesday keynote talk of ACM Multimedia 2006", or "2006 conference trip to California". We shall clarify why these different descriptions emerge and are of value.

The rest of this section is organized a follows. We first revisit real-world events and the process of capturing them into multi-modal streams, as this process often induce changes to the values of the attributes ("Five Ws") relevant for event description and understanding. We go on to discuss the different instantiations of the "Five W's" in multimedia events. We argue that *who*, *when*, *where*, *what* along with the media, are the essential aspects that describe any multimedia event, whereas *how* and *why* refer to semantics that can only be understood by examining relationships between events, such relationship is essentially the role of context.

3.1 Real-world events captured in media

Media capture is typically a very small subset of the information relating to the event. For example, the attendees of the technical conference all sample different aspects of the "ACM Multimedia 2006" event. What they choose to record (e.g. a walk on the beach, the conference keynote, dinner), and how they choose to capture is informed by their situational context. For example a user may decide to take photographs at the conference dinner, while recording the speech at the keynote presentation on her mp3 player, for review later. Others may choose to document conference presentations on their computer via.

Every user who captures the event implicitly leaves out most of the people and the sub-events in the conference – e.g. the conference talks not attended, lunch conversations un-archived among others. Importantly, while no single user has a complete understanding of the conference event, the semantics of the conference event may be derivable by aggregating the capture across users.

The capture may not be faithful to the original experience. For example the semantics of an event may be altered if the person decides to take only black and white photographs - it is no longer possible to describe the event capture in terms of the color of the clothes worn by the participants. Note that at different stages of media production, such as premeditation, capture or editing, will change or create new metadata about the Five Ws in the event being captured. For the purpose of this paper we do not distinguish these changes incurred by different operations.

3.2 The Five W's of multimedia events

In this sub-section we shall examine the characteristics of the six multimedia event facets - the five W's and one H. Each facet can be described in very different ways, and the values are created at different stages of the media production process. We also note that aggregating the values across a specific attribute over different events can lead to more abstract event descriptions. For example, the events "phone meeting with USC", "eating lunch", "teaching class" when aggregated over the time dimension can lead to a more abstract event - "a day in the life of a teacher."

3.2.1. When – Time. The time attribute is one of the key components of an event, yet the description of time can take many forms. For example consider the event of "my trip to China." It is possible to specify the time attribute in a variety of ways – between "jan 1st 2007, and jan 17th 2007" (exact) "after christmas" (relative), "happy" (affective), "every year" (periodic), "in 2007" (at a coarse granularity).

Importantly, there are two temporal coordinates of interest - real-world time and media time. The real-world time refers to the absolute, unambiguous time that an event takes place in the physical world. Media time refers to the relative value within the media stream - e.g. 10 minutes after the news began, third shot in the film etc. We note that we can observe media time. Real world time can either be known, e.g., captured in media metadata such as EXIF data, or it can be hidden or inaccurate, e.g., the interview shot at 10 minutes of the news shot may be taken sometime during the day that is unknown to the viewers. Reconciliation of real-world clips that refer to the same event can become a challenging problem in the absence of temporal metadata.

The relationship between real-world time and the media time can vary widely across different content domains. Event recognition systems typically do not generalize well across domains since the assumptions about the relationship is implicitly coded and rarely explicit in the metadata. Many event recognition systems do not yet address the issue of recovery between media time and real-world event time. This has not yet been a critical problem as the recognition systems have currently been focused on recognizing the objects / people in the given media clip. However, this issue will become increasingly important when researchers begin to address the consumer need of retrieving media related to events (e.g. "media related to my son's second birthday party").

3.2.2. Where – Location. Space is a key multimedia variable to index and interpret events. Prior work [10] reveals that time and locations are the most important attributes for people to recall real-world events. Similar to time, the description of location can also take many forms e.g. "500 W. 120th St., New York City" (absolute and exact); "five miles northeast" (relative); "besides a lake" (approximate). It can also be used at different granularities, "seventeenth floor auditorium in the GE Building" or "around New York City".

Space, like time is used in two coordinate systems - the absolute spatial location where an event occurs and the display space where creators can reorganize elements to communicate a specific affect / meaning. The relationship between absolute real-world event locations and their corresponding media locations is not a straightforward mapping as in some instances for time. Changes in the real-world event location are usually not reflected in changes to any object / person's location in the video. Geo-spatial visualization of media (e.g. [4]) is a possible way in event location changes can correspond to media location changes. We note that in creative domains, there is very little relationships between real-world event locations and how they are manifest on screen. Film-makers routinely alter our perception of space (as well as time) though clever event capture, and event editing [6,7].

3.2.3. Who – Subject. The *who* field has typically referred to the subject in the media clip – *who is in the photo*?. However, this can quickly get complex given the entire media processing chain. For example, one could ask *who took the photo*?, *who edited the photo*?, *who posted the photo online*?, *who has seen the photo*? While the first three questions are directly connected to the event itself (event participation and event capture), the last three questions are about operations on the media clip that re-mediates the original event (media editing, communication and viewing). These different attribute values are useful in different contexts. For example in the case of the identity of the person who edits the media clip, this may be important in the context of a media production house, where the "event" refers to an edit of the raw capture, not the original event itself. It may be important to retrieve the media clips for the editing event based on the media clip editor.

3.2.4. What – Actions, activities and their aggregates. The *what* field describes the action taking place in the media clip. It is answers the question *what is happening in this clip*?. For example, the clip containing the stroll by the beach the *what* field would be described as "walking," "stroll." The answer also depends on the granularity (or levels of abstraction). For example, this could be alternatively answered as – *John walking on Venice beach* (highly specific), to *a person walking* (abstract).

The answer also depends on the event to which the media belongs. For example, in the capture part of the processing chain, this can be asked as *what is the equipment on which this media clip is being recorded*? For editing – *what are the processing / layering effects in this clip*?. For transmission, *what is the data representation format*?, and for viewing *what program do I need to view this media file*?. The utility of the questions depend on the event, and the specific user context.

3.2.5. Why – Event Context. The *"why"* field answers the question *why did this event happen*? For example, "why did John's party take place?" This question cannot be answered by examining a single event in isolation. In this case, the reason why the party took place could depend on another event - "john received a raise." The set of events that are needed to understand the semantics of a specific event, form the context of the event. Note that this is different from the event description context. The event description context are the set of conditions that affect the *values* of the attributes of the event (i.e. the specific values taken by who, where, what etc.).

3.2.6. How – Event Dynamics. The *"how"* answers a subtly different question from *why* or *what*. It is the answer to the question *how did this event come about*?. In the preceding example, the answer to john's question would be related to events that describe shopping, determining the guest lists for the vents, cooking etc. The "how" helps understand the event dynamics, and like the *why* field, cannot be answered by examining a single event. We note that since the *why* and the *how* attributes of an event depend on other events current multimedia research does not try to answer these questions from a single media instance.

3.2.7. Why are the semantics variable? A media clip can belong to many different events - this grouping can be due to a single user, or can happen due to people who interact with it as part of the production process. Consider a segment that appears on broadcast news. This segment can exist in multiple events based on the specific context of the event as well as the people to participate in its production and consumption. The reporter could annotate the clip as "my first story as a reporter." The editor could store it as "file edited before computer crashed," while the user could annotate it as "the clip that I was watching last." These are all different events - each one is a valid event, in the

context of the person who is annotating the event, and goals of such event descriptions are different - the report is using it as memory, the editor as contextual information, and the user as a mechanism to go back to the clip that she was watching last.

The plurality of descriptions for the *same media clip* is clearly necessary, as the intended use of each description is different and is driven by the specific user context. We would like to emphasize that the media clip-event relationship is a dynamic construct, and is driven by user need - indeed this can change over the lifetime of media use for a single person. For example, soon after my trip to china, I would annotate each event in detail - however a few years later, user may just want to aggregate all of these events and would like to recover them as "travel before graduation." Clearly, such aggregation could not have been anticipated with high probability until the user *actually graduated* and had an opportunity to reflect on past events.

Affect can have a significant role to play as well. For example, consider a couple on vacation in India. They label the different events, associate media clips of each event. They may label some of these events as "happy." A few years later, if the couple separates, then this usually will have a significant effect on the semantics of the events. One person may now label all of these events as "sad" and may be unwilling to retrieve any clip associated with such events. In conclusion, event semantics are highly variable across people, and may change over time for a single person, due to contextual dependencies.

3.3 Event Model

In this section we shall discuss event relationships – as they address the *why* and the *how* aspects of events. We first discuss the spatial and temporal relationships and then generalize to other relationships. We note that each event has the standard attributes of *who, where*, *when*, *what* and *media* associated with it. As these are standard attributes, we shall not discuss them in detail.

There are two types of events in the Eventory: physical event and system event. Physical event is the real world event that users can create and publish through the web interface of Eventory. System event is essentially the records of user interaction. For example, user Bob logged into Eventory and uploaded a video. The system event is important to collect the data of user interaction, which make it possible to analyze how they interact with system and other users.

In this our model, event, media and user are inter-related by their IDs. Each media contains the event ID and user ID it is associated with. Each of event, media and user can be connected by using the definition in types of links table and recorded in relationship table.

Event structure in this model defines event recurrence as none repeating, daily, weekly, monthly and yearly. It supports attributes such as interval of every two repeating events (e.g., four years for Olympics), whether event is repeating by day or date (monthly event can recur every first day or first Monday of the month), and on which date event is repeating (for weekly event, Monday, Friday, etc.).

Events may have temporal and spatial relations. Some events repeat, evolve and develop. For example, the event of Olympics repeats every four years. Additionally, there are also changes in location, people, program, etc. Spatial relationships amongst events include – a meeting that is held alternatively in two different locations each week.

Event relationships can be generalized beyond temporal and spatial relations, to arbitrary semantic relationships between events. For example two events may be connected using a 'is-part-of' relationship. For example consider tennis matches that are part of the Olympic games. Each is an event, but part of a larger Olympic games event. Events may have a 'cause and effect' relationship. For example, because of the leg injury, Beckham had to quit the game. Events can be related to each other through evolutionary aspects. For example, the Soccer World League Cup that evolves from Italian League Cup to European League Cup to worldwide. We can regard the later events as developed instances of the original event.

Current event models seldom deal with the problem of relationship because it is hard to list all the relationships for event. In prior work [11] the authors defined several specific relationships such as Parent-Child Relationship, Category Relationship and Domain Relationship.

In our approach, we have generalized this idea to an arbitrary semantic relationship. Rather than defining the semantics of the relationship, we allow the users to specify arbitrary semantics to the relation – the semantics will be clear to the user – the application need not be aware of the relationship. However we do impose constraints *on the semantics of the specification of the relationship.*

4. Event Link Semantics

The key idea in this paper is the introduction of arbitrary semantic relationships between two events (ref. Figure 4). In this section we first introduce the idea of personalized vs. global semantics, and then discuss restrictions on the specification of the relationship. Finally we discuss semantic granularities.

4.1 Personal vs. Global

In our approach, we facilitate global as well as personalized relationships between events. Global relationship types allow users of the system to establish a common understanding of the relationships between events. These are words whose semantics are unambiguous. For example, we have defined "similar" and "related to" as global relations between events. Personalized relationship can be defined by people's own words, languages and expressions. So it might be only meaningful to the person

who creates it, or a group of people who understand the context of definition. For example, a user might create a relationship "meta" to connect two events – the meaning of which may only be known to her. The important issue is if she uses this relationship consistently, then she will be able to browse the events in the eventory in a consistent manner.

4.2 Relationship properties

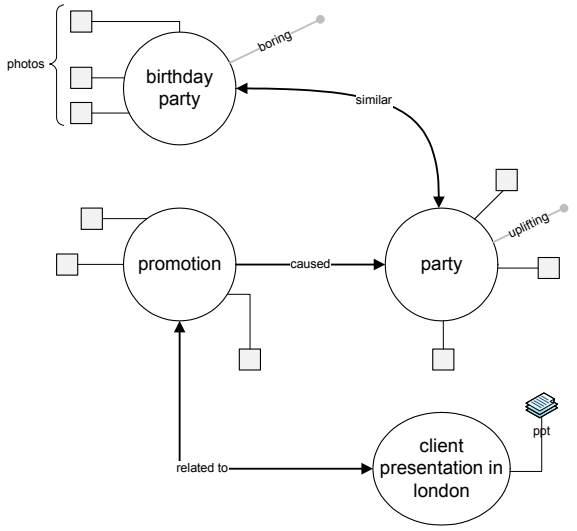

Figure 4: Event relationships. The figure shows both binary relationships between events ("caused") as well as unary relationships ("uplifting").

Each relationship in the eventory needs several properties to be defined. It must have a unique name, and the relationship type can be unary or binary.

The relationship type allows us to specify properties of a single event as well as how two events may be connected. For example, an event A: Italy wins the 2006 World Soccer Cup, event B: Italians celebrate on the street. A user may specify the unary relationship "rocks" to event A. For this user, event A "rocks." A binary relationship can link both events. For example a binary relationship such as "leads to" can connect the two events. After linking event A to B, we have "event A leads to event B".

A binary relationship of the form *x* R *y* needs to define the attribute types of the entities *x* and *y* respectively. It also needs to define properties of the relation R.

- *Attribute types*: Both *x* and *y* can be people, event, media and location. These specifications enable us to filter the events based on attribute semantics. If a specific user creates a relationship "similar" where attributes semantics of *x* and *y* are both media, then the system will not allow the user to specific the following relationship: Bob "similar" Phoenix, as people cannot be similar to places (as specified by the semantics of the relation "similar"). Note that for a unary relationship, we need to specify one attribute type.
- *Relationship types*: A binary relation *x* R *y* can be reflexive, symmetric, and transitive. *Reflexivity:* If *x* R *x*, then R is reflexive. For example relationship "equal" has reflexivity because A is equal to A. *Symmetry:* If *x* R *y* implies *y* R *x*, then R is symmetric. For example, relationship "related" has symmetry because A is related to B implies B is related to A. *Transitivity:* If *x* R *y* and *y* R *z* implies *x* R *z*, then R is transitive. For example the relationship "family" if John is related to Mary via the relation "family" and Mary is related to Joe via "family" then John and Joe are related via the same relation.

4.3 Semantic Granularities

The relationships can be at multiple granularities – they can exist between events, between media, and between media and events. For example, we could have a "promotion" event that *causes* a "party" event. The party may have many photographs / videos associated with it – these may be *related* to media captured at the "birthday party" event.

5. System Implementation

We now discuss the implementation details of the Eventory. Eventory is a web media archive system that people can create, explore and mange events. Users can associate multimedia files with these events by uploading and sharing their media collections including images, videos, audios, documents, etc. The system supports complex event structures by allowing users to specify arbitrary relationships between media as well as between events. Through a notification mechanism, the system allows users to become aware of other users acting on their uploaded media, as well as events.

The Eventory has been implemented using Ruby on Rails Web Application framework. The core of the Eventory is made up of a relational database (MySQL 5.0), hosting the database schema proposed in data model section. On top of this relational schema, we provide an object-oriented implementation using the lightweight object relational mapping framework offered by Rails. With the support of AJAX, Eventory supports dynamic and instant interaction. In the next three sections we discuss event creation, media upload and event notification.

5.1 Event creation and exploration

A simple web based interface allows the user to create events (ref. Figure 5). The figure shows an example of event creation, where event attributes of name, participants, location, start date, start time, end date, end time, tags and descriptions can be provided. Also, event structure such as recurrence can be defined. In the figure we can see that the

author can define the event to repeat monthly. Options to allow for other frequencies are available (week / day).

Figure 5: Event Creation page

Event name redundancy can occur due to the fact that participants of the same event can give different names to describe the same event. A participant may not be aware that the event has already been created. For example, event "US Open 2007" can be named "US Open tennis 07", "US Open Championship 2007" etc., while they are referring to the same event. This creates confusion while associating media to events. To address this problem, we provide keyword hints at the time of user input. With the support of AJAX, keyword will be captured and searched in event database while the user is typing the name of the event. If there are events with the similar name, these names will be revealed to notify user that this event probably already exists, as shown in the yellow region of Figure 5. We provide a calendar interface for efficient management of chronological data and to display the existing events. Thus users can directly view and manage events through calendar.

5.2 Media Sharing

The Eventory supports the sharing of all types of media files including videos, photos as well as text documents. A media upload must have a name, description, set of tags, author. The system records the time that the media was uploaded. Once the media has been uploaded, the user can locate an existing event to which the media can be added. Other functions such as display media and delete media are also available.

We allow media to be uploaded independent of whether it is to be associated with an event. The main reason for this flexibility is that people, who may be interested in sharing media, may not be interested in creating events – indeed one of the key aspects of Web 2.0 sites such as Flickr lies in their ability to allow users to quickly upload and tag media. This flexibility creates the problem of resolving the semantics of the media, as they are not attached to an event.

5.3 Community awareness

We have designed notification and subscription systems to provide passive and active forms of awareness. These notification systems are important as the allow people to become aware of each other activities on the eventory. For example, on Flickr, this is implemented as "new comments" whenever someone comments on the current users photos. Notification is automatically generated by system when there are related updates about the user. This is termed passive, as the user the system provides this notification automatically. The user can also actively subscribe to content in which she is interested, these could be events, users, as well as the result of any valid query in the system.

5.3.1. Notification. The analysis of system level events (clicks / tagging / creation of a link etc.) allows us to easily develop a notification system that is tailored to each user. We only notify users of new links relevant to them due to a concern of overloading the users with too many system level notifications. Notification can be accessed from "my links" in the control panel page. Whenever there are new links, the relevant links will appear as "you have new stuff!" Users can click the hyperlink to view detailed information.

5.3.2. Subscription. In this section we describe ability of subscribing to other users of the Eventory. This is done in two parts: syndication (creating the RSS feed) and aggregation (combining the feeds for each user).

In our system a user can subscribe to events (say a "party" event), people (e.g. all events corresponding to Mary). In the future we are also planning to allow for subscription to arbitrary queries – e.g. all events that that took place at school / events similar to Anne's party etc. In the Eventory, events or media uploaded by a particular user are pushed to other users, thereby updating them on the relevant changes to the Eventory (e.g. they may be subscribing to another users profile).

The RSS syndication standard requires three compulsory channel elements which are title: The name of the channel, link: The URL to the corresponding HTML website and description: Phrase or sentence describing the channel. We include pubDate (posting/publication date) in our summary, which is one of the optional elements in RSS 2.0 specification. Aggregation is the process of combining multiple syndication feeds. In the Eventory when a current user subscribes to other users, their syndicating profile pages are registered to an aggregator built in current user's profile page (i.e. all profile pages have syndicating and aggregating capabilities). This is useful because the current

user can just view a summary view of the feeds on her profile page, rather than visit each individual profile. In the future, we are planning to develop user-developed filters that allow a fraction of the subscribed content to be visible.

6. Current work

The eventory is a work in progress – we are in the process of longitudinal evaluation within the Arts Media and Engineering Community at ASU. The community is of useful size – roughly 70~80 people and we believe this feedback would be very useful in scaling the system as well as in enhancing the capabilities of the system. The preliminary feedback on the eventory has been positive, and we are investigating the following aspects.

- *Expressiveness*: We are extending our work to include structural relationships in location, identity and action – this is a generalization of the current work in temporal structures. This is useful when meeting repeat with some common people, or when events can take place at the same location. We are also working to integrate sensor data as part of the event definition (pressure information from the couch / chairs etc.).
- *Community extraction*: We are builing upon recent work on community extraction [9] on blogs to extraction of thematic communities in the eventory.
- *Summarization / Visualization*: We are working on developing summarization of community activity based on recent work [8] and developing novel visual representations, to promote awareness.

7. Conclusions

This paper focused on the development of an event driven media sharing repository. The goal was to create event driven mechanisms that facilitate community awareness. We were motivated by a desire to reveal inter-disciplinary interdependencies in the authors' institutions.

In our paper, an event referred to a real-world occurrence that unfolded over space and time. Our event model implementation supports the creation of events using the standard facets of who, where, when and what. A key novelty in this research was in the support of arbitrary event-event semantic relationships. We facilitated global as well as personalized relationships between events. Each relationship in the eventory needed several properties to be defined – unary / binary, attribute type and relationship type. The relationships could be defined at multiple granularities.

We implemented a web based media archive system that people can create, explore and mange events. Users can associate multimedia files with these events by uploading and sharing their media collections including images, videos, audios, documents, etc. The system supported complex event structures by allowing users to specify arbitrary relationships between media as well as between events. Through a notification mechanism, the system allows users to become aware of other users acting on their uploaded media, as well as events. We are currently in the process of conducting a longitudinal study to assess the impact of the system on community growth – the initial user feedback has been positive.

We are currently working on extending the expressiveness of the system, as well as developing algorithms for community extraction and summarization.

8. References

[1] *Flickr* http://www.flickr.com.

[2] *SEraja* http://www.seraja.com/.

[3] *YouTube* http://www.youtube.com.

[4] P. Appan and H. Sundaram (2004). *Networked multimedia event exploration*, Proc. ACM Multimedia 2004, also AME TR-2004-10, pp. 40-47, Oct. 2004., New York, New York.

[5] S. Boll and U. Westermann (2003). *MediÆther: an event space for context-aware multimedia experiences*, Proceedings of the 2003 ACM SIGMM workshop on Experiential telepresence, ACM Press New York, NY, USA, 21-30, Nov. 2003, Berkeley, CA.

[6] D. Bordwell and K. Thompson (2001). Film art : an introduction. McGraw Hill 0072317256 (pbk. alk. paper) New York.

[7] M. Davis (2003). *Editing Out Video Editing*. IEEE Multimedia. **10:** pp. 54-64.,

[8] Y.-R. Lin and H. Sundaram (2007). *Blog Antenna: Summarization of Personal Blog Temporal Dynamics Based on Self-Similarity Factorization*, to appear in Proc. International Conference and Multimedia Expo (ICME) 2007, July 2007, Beijing, China.

[9] Y.-R. Lin, H. Sundaram, Y. Chi*, et al.* (2006). *Discovery of Blog Communities based on Mutual Awareness*, Third Annual Workshop on the Weblogging Ecosystem: Aggregation, Analysis and Dynamics, at the 15th Annual World Wide Web Conference - WWW 2006, also AME-TR-2006-03, May 2006, Edinburgh, Scotland.

[10] M. Naaman, Y. J. Song, A. Paepcke*, et al.* (2004). *Automatic organization for digital photographs with geographic coordinates*, Proceedings of the 4th ACM/IEEE-CS joint conference on Digital libraries, pp. 53-62, 2004, Tuscon, AZ, USA.

[11] R. Singh, Z. Li, P. Kim*, et al.* (2004). *Event-based modeling and processing of digital media*, Proceedings of the 1st international workshop on Computer vision meets databases, ACM Press, 19-26, Paris, France.

[12] U. Westermann and R. Jain (2006). *Events in multimedia electronic chronicles (e-chronicles)*. International Journal of Semantic and Web Information Systems **2**(2): 1-27.

[13] U. Westermann and R. Jain (2007). *Toward a Common Event Model for Multimedia Applications.* IEEE Multimedia **14**(1): 19-29.