TELLING STORIES WITH MYLIFEBITS

Jim Gemmell, Aleks Aris** and Roger Lueder**

*Microsoft Research, **University of Maryland

ABSTRACT

User authored stories will always be the best stories, and authoring tools will continue to be developed. However, a digital lifetime capture permits storytelling via a lightweight markup structure, combined with location, sensor and usage data. In this paper we describe support in the MyLifeBits system for such an approach, along with some simple authoring tools.

1. INTRODUCTION

MyLifeBits is a project dedicated to digital personal lifetime storage [2,5,6]. It is designed to be a personal lifetime store for all digital media and data including email, calendar events, contacts, documents, audio and video. Entities in MyLifeBits have type-specific attributes. For instance, a photo has a date-taken and a location-taken (latitude/longitude coordinates), while email has a sender attribute. Entities in the MyLifeBits store may also have typed links between them. For instance, there may be a link of type "person in photo" between a contact and a photo. Or, there may be an "annotation" link from a text document to a video, indicating that the text comments on the video. MyLifeBits uses a SQL Server database to store metadata.

The most obvious features to design have been those which help the user organize their life bits, and to find things again. However, there is also an element of passing the bits on for posterity, and it is questionable whether one's grandchildren will be able to make any sense of such a mountain of data. Likewise, it is not always easy to share experiences with others. Jain has suggested that we need to move from an overwhelming log of raw data to a useful *eChronicle* [8]. In this paper we describe some ways of extracting stories from one's life bits, starting with some simple authoring tools and moving on to more lightweight and even automatic methods.

2. STORY AUTHORING

The conventional way of telling stories is to author them. Early in our development, we made an effort to make authoring easy with Interactive Story By Query (ISBQ). ISBQ lets users make queries, and then drag and drop selections from the query result into a story. Two story types were supported: (1) A slide show and (2) a "time sheet".

The slide show authoring tool allows images in the query results pane to be dragged and dropped into a sequence. A caption may be added to each image, and an audio clip to be spoken (or dragged and dropped) to provide narration. Other audio clips may be added for background music. The resulting slide show thus has music and commentary. It is constructed using HTML+Time for easy sharing with others. Links are created from the story to all the media that they include. When played in the original MyLifeBits system, any image can be clicked to open a query window showing all items linked to that image.

A time sheet is a composition of multiple timelines. The user chooses the number of timelines desired and gives each a name. Items may be dragged and dropped into a timeline, where they are placed automatically according to their time. The timelines may have their scrolling locked together to allow comparison between the timelines.

3. LEVERAGING LOCATION & TIME

Location and time are powerful keys to unlocking memories. We believe that location awareness will become ubiquitous in media capture devices such as cameras, cell-phone cameras, and video cameras. As these devices already include clocks, it means we will be able to search for media based on place and time. In addition to search, location and time aid in storytelling. A list of locations covered in time order is a simple telling of the story of the trip, and by combining media with maps and animations we can create visualizations to tell the story in a much more compelling fashion.

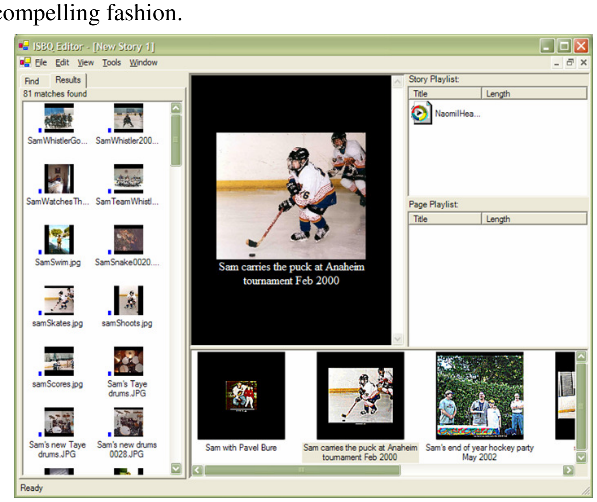

Figure 1: ISBQ slide show authoring interface.

The technology used to obtain location may be GPS, detection of RFID tags at some known location, or one of many other methods. Regardless of which approaches become most widespread, we can now anticipate ubiquitous location-enabled devices and begin work on the software that will let us leverage them. At present, we carry a pocket-size GPS independent of our media capture devices. Time correlation between the GPS record and the media is then used to infer a location for the media. GPS records are also interesting apart from any media;

0-7803-9332-5/05/$20.00 ©2005 IEEE

they show the "trail" you have traveled. To date, we have only experimented with photos, although much of our discussion would extend to video or audio in a straightforward way. In addition to off-the-shelf cameras, we have also experimented with SenseCams (described below).

GPS location readings, consisting of date/time, latitude, longitude, elevation, and precision, are loaded into a table in the MyLifeBits database by an import program. Date/time correlation between photos and GPS readings may then be used to set the location in the photos whenever possible. However, the user may elect to not set the location in photos based on time, because they may obtain some photos from a third party, or may have loaned the GPS to someone else. Furthermore, some photos may already have their location set by a third party. So, while inferring location based on time correlation with a GPS is a feature that we support, it cannot be used in all instances, and location must be stored as an attribute for each photo. Note that many GPS "track points" will not have any corresponding photo, but are still very interesting and worth saving as a record of the user's location.

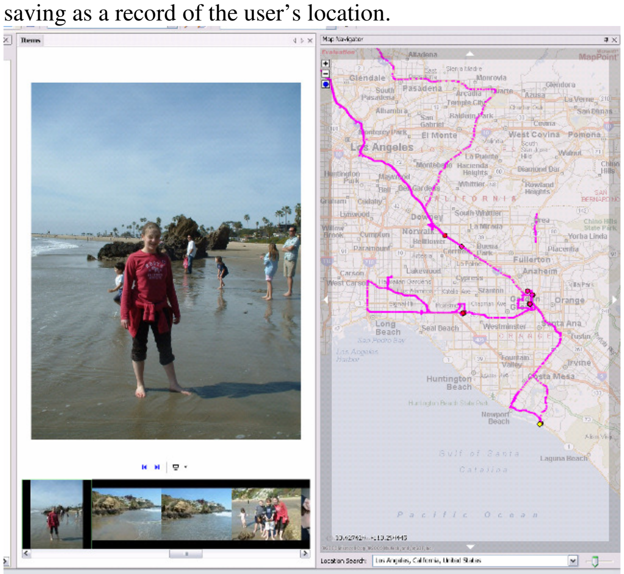

Figure 2 – Map UI map on right shows large red dots where photos are taken and small pink dots for GPS track points. The pane on left shows photos taken in the area shown on the map, in film strip view with a large preview.

Photos have their meta-data stored in the MyLifeBits database, while their image content is a file in an NTFS folder tree that MyLifeBits is configured to monitor. In the case of commercial cameras, it is simple enough to load photos into an NTFS folder. A little more is involved in the case of a SenseCam. SenseCam is a device that combines a camera with a number of sensors, including a 2-channel accelerometer, a digital light sensor, a temperature sensor and a passive infra red sensor for detecting living beings. Photos are triggered by sensor data and/or time interval. It is worn on a pendant hung from one's neck. The sensors are used to automatically take pictures at "good" times [6]. Sensor information is also recorded, and is uploaded along with the photos. An import program uploads SenseCam photos and sensor data into MyLifeBits. The photos are in JPEG format and are stored just like any other JPEG photos in MyLifeBits, with attributes that include date/time taken, location, and camera make. The sensor data is stored in tables in the MyLifeBits database. All sensor values include the date/time of the sensor reading. With SenseCam, photo-taking is passive, freeing the user up to just experience rather than worry about taking photos [6].

A map UI element in the MyLifeBits interface marks photo locations by a large red dot and track point locations with a smaller pink dot (Figure 2). Moving the mouse over the corresponding dot for a photo pops up a thumbnail of the image in a direction that avoids being partially drawn (Figure 3). A location (address, zip code, or city/state) may be entered into the "Location Search" text box to set the map location. The +, -, o on the left of the map support zoom in, zoom out, and zoom out to the entire earth, respectively. A region of the map may also be selected to zoom into using the mouse, while a clickable arrow appears on mouse over near the map edges to support panning. Zooming and panning on the map issues a new query for photos from the region of the visible map to be shown in the corresponding pane. Photos that have no location (e.g., old photos) can be dragged and dropped onto the map to set their location. The basic mapping interface was derived from the WWMX, and more details on its appearance and operation can be found in a paper describing WWMX[12].

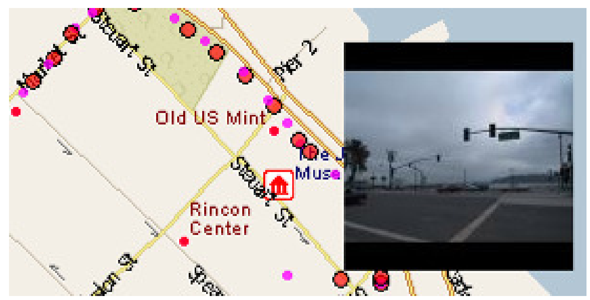

Figure 3 pop-up image when mouse hovers over the corresponding dot on the map

After some time of using a GPS and taking photos, a selected map region may become saturated with points. One way to filter out undesired clutter on the map is to restrict photos and track points by time. To this end, we divide the photos and track points into clusters and consider each a "trip." The clustering algorithm simply looks for gaps in space or time above a given threshold (e.g. time gap > 90 min, or location gap > 1 mile) in order to divide the data. The trips are displayed in a list box below the map. The top line allows the selection of all track points and photos by clicking on "All Trips." The number in parentheses to the left of each line, e.g., (1154) to the left of "All Trips", indicates how many photos/points that trip includes. When a trip is selected from the list, only its photos and track points appear on the map (Figure 4). An alternative to filtering by trip is to view all trips, but connect the dots belonging to the same trip with a line to help distinguish between trips. The user can select this with the "show connections" checkbox .

The interface described so far supports searching and browsing based on location and time. However, a static map of track points and photo locations is still hiding a lot of useful information. One does not know the direction traveled, or the time spent in different places. We have created a trip "replay" visualization that animates the trip to show this information to the user. During the replay, a curve grows from the starting point of the trip towards the ending point, with the time spent at each point proportional to real time of the trip. As the curve hits locations where an image has been taken, a preview of the image is displayed below the map (Figure 5). During the replay, the images are retrieved from the database on demand and cached to improve the performance. The same cache is used to store and retrieve the pop-up images (as seen on Figure 3).

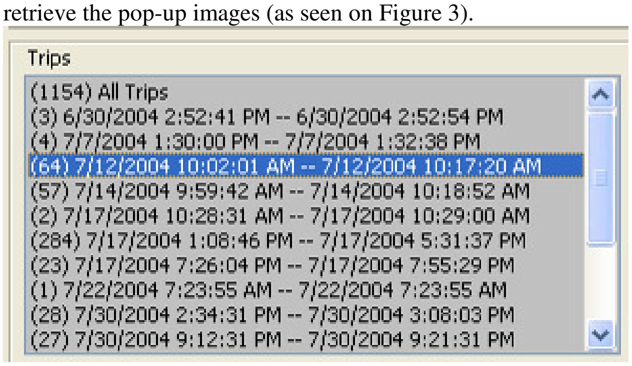

Figure 4 Selecting a trip

Figure 5 Replay of a trip

To start the replay one presses the play button. During the replay, the "Replay Progress" slider shows the progress of the replay (in terms of time) advancing from left to right. The user can press pause button to suspend the replay. While paused, the progress slider may be dragged to a desired point of the replay, and replay can be resumed from that point. During playback, one can also directly drag the progress slider to any desired point of the replay and the replay will continue from that point when the user releases the slider. While the user drags the slider, the curve dynamically grows and shrinks. Simultaneously, the picture box updates to show the current image.

The user can select the duration for the trip replay (default: 5 seconds). In trips with highly variable speed (e.g. walking down the street taking a few pictures, stopping for a minute to take several photos of an interesting feature, and then continuing to walk with a few snap shots) some photos may be displayed for a long time, while others will flip by so fast as to be unrecognizable. To compensate, the user may switch replay from "proportional" to "sequential", which cause the replay to show each photo for the same amount of time and disregard the actual rate of the trip.

If a trip overlaps itself (e.g., in a round trip, where one returns by the same path) the animation will not be visible where overlap occurs. In order to avoid this problem, we use a different color for the leading part of the animated line from that of the previously rendered part. Thus, the leading portion of the animation is always visible, even when it overlaps previous sections of the trip. We are experimenting with how much of the line to have in a different color and currently expose the setting in the UI. "Full" selects just a single color, while "Aged" allows the user to select the amount in a different color as a percentage of total trip time. Using a percentage of total trip time makes the amount in a different color vary according to the speed traveled, giving the user a way to compare speed during the trip.

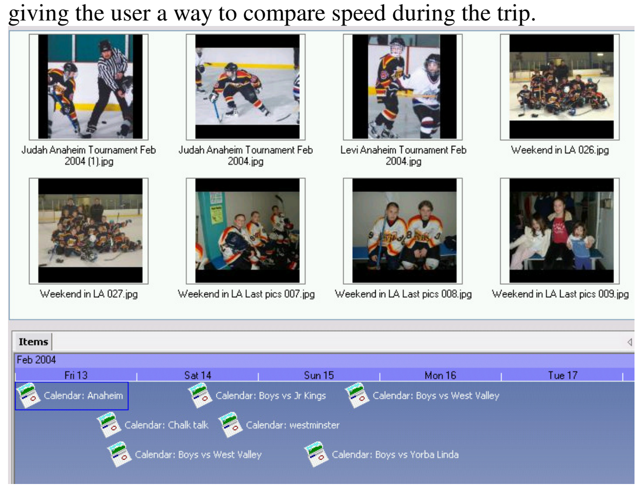

Figure 6: Selecting a calendar event from the timeline (below) shows all photos of the event in the related items pane.

The use of time for search and browsing is pervasive in MyLifeBits. Time attributes can be used in search and refinement of searches. For example, the user may search for all files created in a given date range. Or, suppose a search for emails results in a large number of hits. The user may select to refine by date received, and MyLifeBits will cluster the emails by date on the fly. The user can then pick a cluster to narrow in on. MyLifeBits logs most user activity and it is a simple right-click operation to see the history of that item. For example, the history of a photo could show when it was taken, and every time the JPEG file has been opened. MyLifeBits also supports "pivots" on time corresponding to the user remembering that something happened at the same time. For instance, a user may click on an old appointment in their calendar and ask to see everything that overlaps in time with it, and thus find photos taken at that time.

However, just as location data may not be relevant to a photo taken at the same time, these time overlaps may or may not be significant. For example, photos received from others may overlap with an event on your calendar but have nothing to do with it. For that reason, links ought to be created between media and associated events, for example, a link from a photo to an event indicating it is a photo of that event. Our photo import wizard allows the option of automatically using the time overlap to create such a link; it need not be used in the case that the photos came from someone else.

4. THUMBS UP/PUBLISH

Marking up media with any kind of meta-data is usually considered a chore; something that users are not usually willing to do. However, we have observed that there are some opportune moments for obtaining information from the user, so long as it is available at the moment the notion strikes them and that the mode they desire is supported. Thus, it is our goal to have a number of lightweight markup methods available ubiquitously, including: voice/text annotations, thumbs up/thumbs down ranking, and an indication of whether the item can be made public or should be kept private. These tools must be readily available in the camera cell-phone that is taking pictures, in the photo import wizard, in the searching/browsing tools, and even in the screen saver – and we have implemented all of these in MyLifeBits.

Putting this lightweight markup together with comprehensive usage logging allows us to create a new environment for story-telling. The user simply captures and uses their material as usual. Throughout all phases of the media creation and usage chain there are opportunities where some media can be flagged with a thumbs up, or simply observed from usage to be popular. Events in the calendar can also be flagged as popular.

We can then query the database to show popular events or events linked to popular media on a timeline. The user can click the event to see this popular media; if they are interested they can broaden the scope to see less popular media that was also captured. The map view can be used to see where media was captured and animation can be applied as described above. Entries that are marked both as popular and as public can be exported to a blog or other sharing point and formatted to be attractive.

5. RELATED WORK

Other projects for personal storage in a similar vein to MyLifeBits include *Placeless docs* [3] and *Haystack* [7]. Systems that emphasize metadata and annotations for photos in particular include *PhotoFinder* [9] and Shoebox [11]. Kim et al [10] have prototyped a set of personal chronicling tools to support knowledge workers on the job with similar features to what we have discussed, including ubiquitous media tagging, linking items related to an event, and sharing. Davis advocates guiding the user in the capture phase with feedback to aid in later story creation [3]. Appen et al have developed an event-based storytelling framework [1]

6. CONCLUSION AND FUTURE WORK

User authored stories will always be the best stories, and authoring tools will continue to be developed. However, a digital lifetime capture permits storytelling via a lightweight markup structure, combined with location, sensor and usage data. In this paper we have described our support for such an approach, along with some simple authoring tools, in the MyLifeBits system. In the future, we plan to develop software to automatically blog entries marked public, and to perform user studies involving our software and the SenseCam.

7. ACKNOWLEDGEMENTS

We would like to acknowledge Ron Logan and Kentaro Toyama for their contributions to the mapping software, and Lyndsay Williams and Ken Wood for including us in the SenseCam project. Gordon Bell, Jim Gray, and Ben Shneiderman have provided guidance to our work.

8. REFERENCES

[1] Appan, Preetha, Sundaram, Hari and Birchfield, David, Communicating everyday experiences, *Proc. First ACM Workshop on Story Representation, Mechanism & Context,*, New York, New York, Oct. 2004, pp. 17-24.

[2] Aris, Aleks, Gemmell, Jim and Lueder, Roger, Exploiting Location and Time for Photo Search and Storytelling in MyLifeBits, *Microsoft Research Technical Report* MSR-TR-2004-102, October 2004.

[3] Davis, Marc, "Active Capture: Integrating Human-Computer Interaction and Computer Vision/Audition to Automate Media Capture." *Proc. IEEE International Conference on Multimedia and Expo (ICME 2003)*, Baltimore, Maryland, IEEE Computer Society Press, Vol. II, 185-188, 2003.

[4] Dourish, Paul, Edwards, Keith, LaMarca, Anthony, Lamping, John, Petersen, Karin, Salisbury, Michael, Terry Douglas and Thornton, Jim, Extending Document Management Systems with User-Specific Active Properties, *ACM TOIS*, 18(2), 2000, pp. 140-170.

[5] Gemmell, Jim, Lueder, Roger, and Bell, Gordon, The MyLifeBits Lifetime Store, *ACM SIGMM 2003 Workshop on Experiential Telepresence (ETP 2003)*, November 7, 2003, Berkeley, CA.

[6] Gemmell, Jim, Williams, Lyndsay, Wood, Ken, Lueder, Roger, and Bell, Gordon, Passive Capture and Ensuing Issues for a Personal Lifetime Store, Proc. *ACM Workshop on Continuous Archival and Retrieval of Personal Experiences (CARPE '04)*, Oct. 15, 2004, New York, NY, USA, pp. 48-55.

[7] Huynh, David, Karger, David, and Quan, Dennis, Haystack: A Platform for Creating, Organizing and Visualizing Information Using RDF, *Semantic Web Workshop*, 2002.

[8] Jain, Ramesh, Multimedia Electronic Chronicles, *IEEE Multimedia*, July 2003.

[9] Kang, Hyunmo and Schneiderman, Ben, Visualization Methods for Personal Photo Collections: Browsing and Searching in the PhotoFinder, *Proc. IEEE International Conference on Multimedia and Expo*, 2000, pp. 1539-1542.

[10] Pilho Kim, Mark Podlaseck, and Gopal Pingali, Personal Chronicling Tools for Enhancing Information Archival and Collaboration in Enterprise , *Proc. ACM Workshop on Continuous Archival and Retrieval of Personal Experiences (CARPE '04)*, Oct. 15, 2004, New York, NY, USA, pp. 56-65.

[11] Rodden, Kerry and Wood, Kenneth R., How do people manage their digital photographs? *Conference on Human Factors in Computing Systems (CHI 2003)*, April 2003, pp 409-416.

[12] Toyama, Kentaro, Logan, Ron, Roseway, Asta, and Anandan, P., Geographic location tags on digital images, *Proceedings of ACM Multimedia '03*, November 2003, Berkeley, CA, USA, pp. 156–166.