TuVista: Meeting the Multimedia Needs of Mobile Sports Fans

Frank Bentley and Michael Groble
Motorola Applied Research and Technology Center
1295 E. Algonquin Rd.
Schaumburg, IL 60196 USA
{f.bentley, michael.groble}@motorola.com

ABSTRACT

We describe the TuVista system, a service for viewing near-live sports content consisting of a multimedia editing/bundling station, a cloud-hosted metadata server, and a set of mobile clients. We begin by introducing TuVista I, a proof of concept experience prototype implemented quickly as a probe to understand multimedia needs at a live sporting event. After discussing the results of an initial field trial at Estadio Azteca in Mexico City, we describe the improvements in TuVista II to address the issues identified. These include rapid editing of multiple live video streams, push notifications of new content over XMPP, and an optimized metadata workflow for the content producer that reduced content publication time from fifteen minutes to less than 30 seconds. We conclude with a discussion of rapid prototyping and field deployments as a way to quickly identify user needs.

Categories and Subject Descriptors

H5.1 Information interfaces and presentation: (e.g., HCI): Multimedia Information Systems
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms: Design, Experimentation, Human Factors

Keywords

Mobile, multimedia, video editing, video annotation, metadata, XMPP, sports.

1. INTRODUCTION

While watching a television broadcast of a sporting event, viewers are often given a wide variety of information including player biographies, statistics, and multi-angle video instant replays. However, while at a live sporting event or mobile, sports fans often lack even the most basic information about the important plays in the game.

We have created the TuVista system to help address the needs of sports fans in these situations. When users are mobile, it is often not practical to watch an entire sports cast via mobile TV or a similar broadcast mechanism. Mobile users often exhibit small moments of focused use with mobile data applications [2, 10, 13]. Therefore, a content "snacking" [9] experience is more appropriate for mobile users. Sports fans in a stadium have bursts of time between plays or when leaving their seats to visit the concessions. During these times, they are missing the ability to relive the latest plays from other angles and view the latest stats for players, content that would be available to them in a television

MM'09, October 19-24, 2009, Beijing, China.

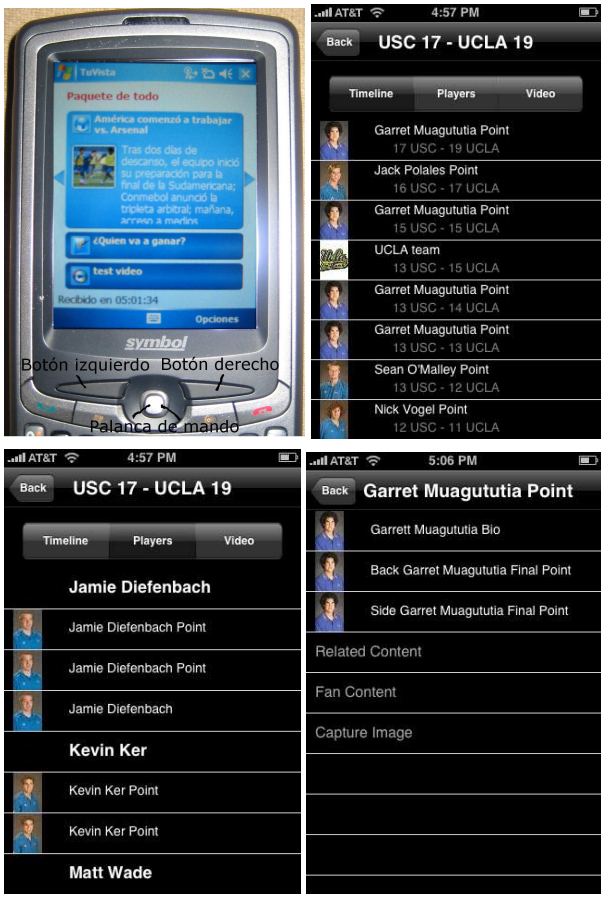

Figure 1. The TuVista Phase I (upper left) and II interfaces. In Phase 1, users could serially scroll through bundles of grouped content. Phase II provided various ways to access bundles based on metadata attributes.

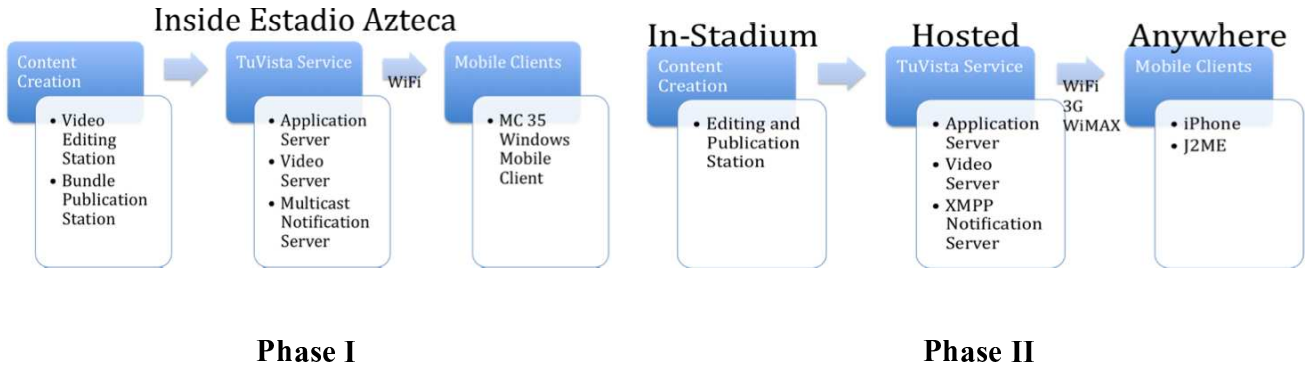

Figure 2: TuVista architectures for Phase I and Phase II. In Phase I, the goal was a rapid development cycle, and all components were installed in a sports stadium with access provided over wifi. Phase II introduces a hosted service and access via cellular/WiMAX.

broadcast. Likewise, fans on the go often have small bits of time that they can use to catch up on a game and having a list of key plays would help them to choose the content most desirable to watch in the short amount of time that they have available. Ethnographic work with sports fans has also shown a desire for replays of key events and access to statistics during a game as important for fans to feel engaged with the experience. [11]

To address these needs, we created the TuVista system. The system includes a backend editing tool, a web service, and mobile applications for several popular mobile phone platforms. The editing tool is the heart of the system and supports creating multi-angle, edited video clips, associating them with related metadata, and pushing them to mobile clients.

Our rapid video editing and transcoding system allows us to publish content within thirty seconds of the action occurring in the game, including video editing, metadata annotation, transcoding, and notification. This editing system required significant work in understanding the workflows of media creation and annotation during live, fast-action sporting events. From this knowledge, we were able to create a system that fit into the pace of a game to allow for quick editing and annotation of content.

Mobile phone multimedia systems provide extra challenges. The limited screen real estate and interaction paradigms make searching large collections of multimedia content difficult. Likewise, the fragmentation of third-party application support across handset manufacturers as well as differing video-codec support creates additional technical challenges for implementing mobile multimedia applications. This work focuses on the user interaction of the mobile application on multiple platforms and the system components that are needed in order to support this user experience.

We realized that building a complete system in this space would be a considerable effort. Before building a commercially viable system, we created a functional rapid prototype to test in a live game environment. By conducting this trial, we were able to quickly learn from real use in order to rapidly iterate on our initial designs based on the real needs of both content publishers and sports fans.

2. RELATED WORK

Over the past several years, multiple projects have been developed to deliver video of sporting events to mobile devices. The first of these was eStadium at Purdue. [1] eStadium provides a mobile web interface displaying a list of plays in the game. For each play, users can click on links in the webpage to launch the media player on the mobile device in order to view replay clips. While this works across many devices, often the lack of an integrated application can confuse users and leave them wondering how to get back to the play list. Also, eStadium relies on a single video feed for capturing content. We wanted to take the successful aspects of eStadium (on-demand clips, mobile access) and expand the content available as well as improve the overall user experience.

Another key area of research has been quickly associating metadata with video content. From Davis' Media Streams work back in 1993 [4], the concept of quickly applying a common set of metadata tags to video content has been a key component in managing large video repositories. Our system contains video over many games and seasons, so managing it and providing easy access to related content is extremely important.

Several researchers have addressed annotating sports content. Xu used live feeds from web-casting of actions in a game to associate actions to video clips taken at the same time in the game. [15] Zhu took this idea further by combining analysis of the video itself to help associate the feed data to the appropriate video clip. [18, 19] Unfortunately, these feeds are often quite behind the actual action in the game (30 seconds-1 minute). These delays are larger than our entire workflow time in the current system.

In the US, Major League Baseball (MLB) has launched an iPhone application that allows users to view the current status of a game and select video clips. The clips are single-feed from the live broadcast and during the 2008 season took between 10-20 minutes to appear on the handset after the action in the game occurred. To stay up to date, this application polls the MLB servers every 30 seconds to see if new content is available. We wanted to develop a better solution to push multi-angle clips with related content in a more real-time manner.

This work has set a strong foundation on which to create mobile multimedia systems. With this work in mind, we set out to create a system that would quickly deliver bundles of related multimedia content to fans interested in following a game.

2.1 Media Bundles

To help users receive content that is relevant to a particular point in the game, we use Media Bundles [14] which are a collection of related content defined in RDF or JSON notation. Media Bundles

allow the content creator to create bundles for plays in the game that contain video clips from multiple angles, bios and stats of the players involved, and links to related content such as previous goals scored by that player. Through some initial ethnographic research in stadiums with sports fans, we observed the desire for much of this ancillary information that is usually available to those watching on television. This information is currently lacking for fans in stadiums or following a game while mobile. The media bundle technology provided us with a platform to create these rich media experiences.

3. PHASE I TRIAL SYSTEM

Before committing to developing a large system in this space, we wanted to prove out the concept of delivering media bundles to sports fans within a live event. To this end, we created a trial system that could handle a few hundred users and utilized mostly off-the-shelf components for video editing and playback. Through this initial work, we sought to identify the key bottlenecks in the real-time production of this type of content in addition to understanding the experience from the point of view of both the in-venue content creator and the fans using the solution on their devices. This knowledge would then be used to design a real system that could be deployed for more than one game and for a wider audience.

3.1 Architecture

Our Phase I system included two workstations that we installed in the broadcast booth of Estadio Azteca in Mexico City. One of these machines was a Mac Pro with a video capture card and Final Cut Pro software for editing an incoming live video stream. In front of the capture card, we placed a 4-video monitor and a video switcher so that the operator could choose in real time the live camera input that fed into the capture card. Figure 3 illustrates this setup.

We were only able to capture one video stream at a time, however by using the jumbotron input as one of the cameras we could at times switch to receive two angles of a given play. The jumbotron feed was a "dirty" feed that included overlays for the score and game clock and switched between stadium camera feeds approximating video that might be broadcast. This included instant replays controlled by an operator in the booth with us. The other three camera inputs to our system were raw feeds from microwave cameras on the field or fixed cameras mounted around the stadium.

When key actions occurred in the game, the video editor in front of the editing workstation would create clips in Final Cut Pro and upload them to a server in a backroom of the stadium using a web interface that we developed. This process included setting in and out points in the video as well as any slow motion effects, and transcoding the video into WMV format at a size appropriate for mobile delivery.

The second computer, the Bundle Authoring Station, provided a web interface to create media bundles. The Bundle Author created new bundles when actions in the game occurred and attached the relevant stats, bio, and videos (uploaded from the video editor) to create a finished bundle. This interface included a simple list of content, and the author could check off a set of items to include them in a particular bundle. Through this simple interface, we hoped to learn what types of associations would be important to help automate in later versions of the system.

Figure 3. The Bundle Authoring and Video Editing workstations were installed in a corner of the broadcast booth. On top of the desk is a 4-screen video preview so that the Video Editor could choose the appropriate feed.

When the bundle author completed this process and chose to publish a bundle of content, a multicast notification packet was sent out over the in-stadium wifi network containing an RDF description of the bundle and its content.

Users in the stadium had access to Motorola Symbol MC35 mobile devices with the mobile TuVista client installed. The client displayed all published media bundles and allowed users to scroll through bundles for the current game. The screen displayed one bundle at a time and users could see short summaries of the content in that bundle (see upper left image in Figure 1). Users could click on any item in a bundle to view that particular content item. Content items included videos, text, and photos. The client also supported polls whereby fans could vote on questions determined by the Bundle Author.

3.2 Phase I Trial

In order to gauge the usefulness of such a mobile sports system, we conducted a live field trial during the semi-finals of the South American Cup in April 2008. We provided 60 MC35 devices to 40 fans in a controlled-access restaurant area of the stadium overlooking the field and 20 devices to fans in various skyboxes. Participants were selected from fans who happened to be attending the game that day and were not recruited in advance. All fans in the restaurant were approached as they entered and asked if they would like to try a new mobile sports system during the game, however participants from the skyboxes were selected in advance by stadium personnel.

Participants were given a one-page introduction to the system with a basic overview of the features and navigation and were informed that there would be a short interview after the game. Members of the research team were stationed in the restaurant and regularly stopped by the skyboxes to observe the use of the system throughout the event. Following the game, researchers conducted short semi-structured interviews about the use of the system and asked for specific feedback regarding particular screens and features. Participants also completed a questionnaire consisting of questions about their usage of the system, photo/video capture during the game on other devices, and demographic information. For their effort, participants received 250 pesos (roughly $US25).

Two members of the research team served as the Bundle Author and Content Editor and spent the game in the broadcast booth. These users made a record of issues that they faced while running the system during a live event.

All interaction with the system was also logged and we were able to view content requests per user as well as interactions from the Bundle Author and Content Editor as they created content during the game.

3.3 TuVista in Use

The use of the system during the trial focuses on two main user populations, the editor/author creating the content and the fans receiving the content throughout the game.

The beginning of the game started slowly in terms of action on the field. Pre-game bundles were published including line-ups for each team, video highlights from the quarterfinal game the week before, and historical information about the team including the lyrics to the fight song (which plays repeatedly throughout the game over the PA system).

As a few goals were scored, the Video Editor was able to keep up, however the complexity of using Final Cut with live content meant that it took minutes to properly edit and transcode the content. Scanning linearly through the video proved to be a difficult task to accomplish quickly in order to set in/out and slow motion points. Table 1 shows the publication times for each bundle from the time of the action in the game until content was visible on the mobile devices. As you can see, the second half proved quite difficult as several goals were scored within minutes of each other forcing the Video Editor to scroll through the linear history of the video to find the key moments. Because of this, bundle publication took up to 36 minutes as the editor got several plays behind.

For the fans, the content was seen as a useful extension to the game. Particularly during down time in the game (injuries, setting up corner kicks, halftime, post-game) participants were seen huddled around the nearest device watching the key plays again and again. A graph showing the times in the game when videos were viewed appears in Figure 4. Users appropriated the device into their own celebrations of the game. After the game, a group of fans in the restaurant was observed holding the device in the air replaying the game-winning goal while dancing together and singing the team song. At other times participants were observed sharing the device with fans at nearby tables. TuVista became a social experience that enhanced the live game for those using it. This social use of mobile video was also seen by O'Hara et al in their study of mobile video use [8] and we believe is an important use case to consider when producing videos and designing interfaces for browsing multimedia on a mobile device.

Time of Action	Editing Time (mins)
21:11	8
21:16	8
21:33	7
21:37	9
21:39	17
22:16	8
22:18	14
22:20	17
22:26	23
22:34	18
22:55	36

Table 1. Editing time for content in the Phase I system. As you can see, many events occurred closely together in the second half of the game, dramatically increasing editing time.

Overall, the bundle concept was understood by nearly all participants. Users appreciated the grouping of related content into bundles and the relative simplicity of the interface. Our participants also liked that the most recent bundle would pop to the top as soon as it arrived so that they could glance down and see if there was new content available. However, there was nearly universal agreement that it took far too long for content to arrive after the event occurred in the game. After using our system, fans generally agreed that about 30 seconds would be a reasonable time to wait between an action in the game and the videos appearing on their devices.

While we had expected that the most recent bundle would be what fans were most interested in at any given time, we saw fans regularly going back to older bundles to relive key aspects of the game with their friends. As in the post-game dancing example, these participants wanted to see every goal that led their team to victory. This led us to think more about the browsing interface and metadata structure to help users more easily navigate to older plays in the game such as viewing content by player or type of action in the game (e.g. all goals).

In addition to observing users during the game, we published a bundle just before the end of the first half containing a coupon for a free drink. We only left it published for fifteen minutes in order to see how many users were paying attention to the mobile device. Thirty-two out of the 40 participants in the restaurant redeemed the coupon showing that fans maintained a high level of interaction with the device during the game. Our server logs of video access also corroborate this.

These challenges of rapidly editing video to reduce bundle publication time, quickly associating metadata, and allowing for flexible viewing of content based on metadata would be some of our main focus areas for the next iteration of TuVista.

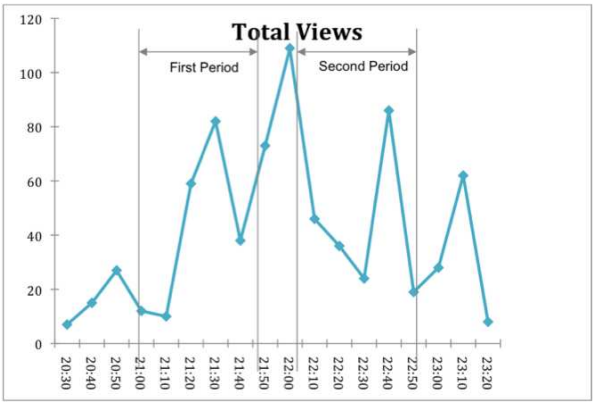

Figure 4. Usage of the TuVista system during the Phase I trial. Content views for each 10 minutes are displayed over the course of the game.

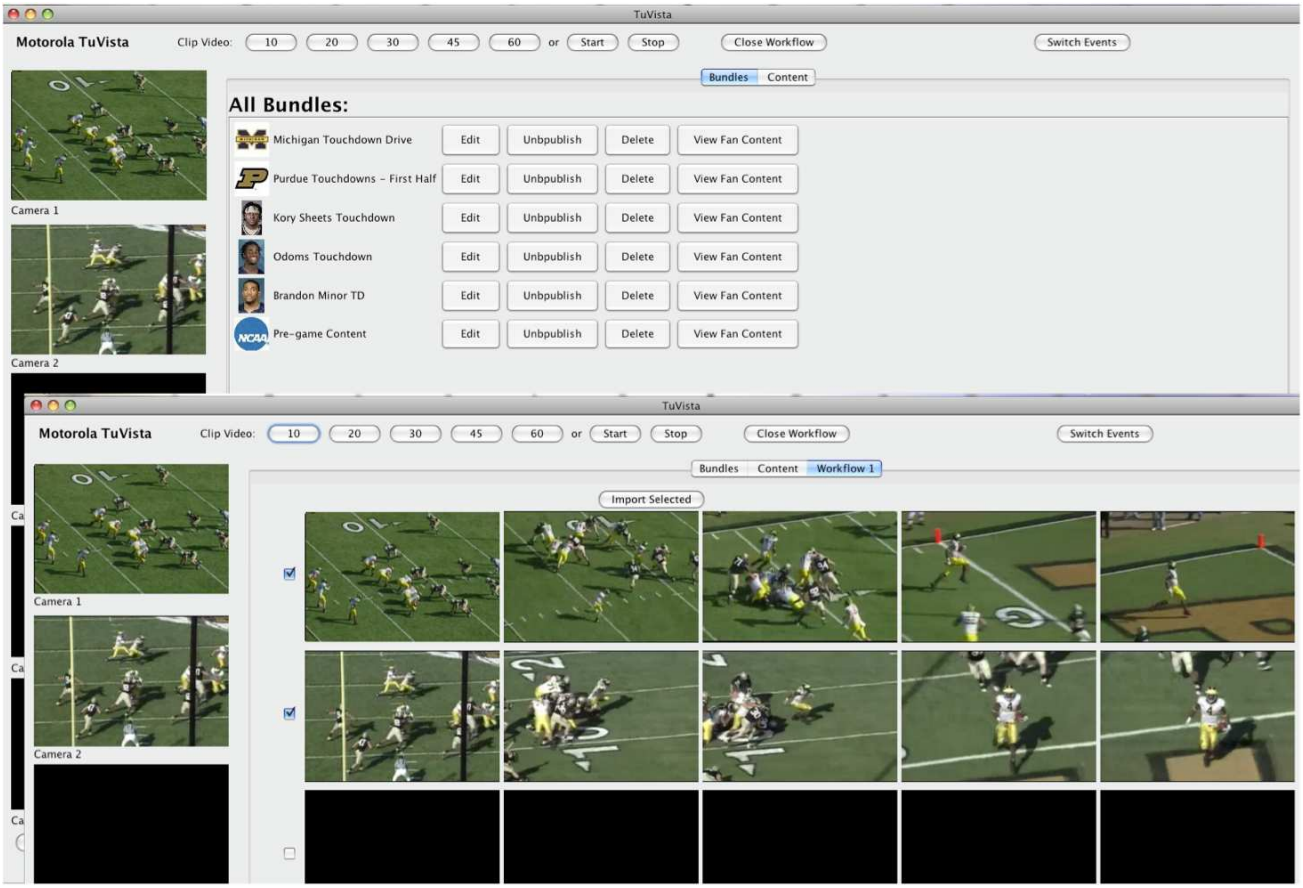

Figure 5. The backend publication tool. During the game, the Editor sees a view like the one on the top. Live video previews for each of the incoming cameras appear on the left so that the Editor can monitor the current state of the game. Buttons on the top allow for clipping the last n seconds of video or for starting and stopping capture for a given play. Once video is captured, it is displayed in an array allowing for simple one-click editing of in and out points and addition slow motion effects.

4. PHASE II SYSTEM

Using our learnings from the Phase I trial, we set out to create a system that made it easy to edit multiple video angles and associate related content. We set a goal to reduce the editing, annotation, bundling, transcoding, and delivery cycle to less than 30 seconds as this met with fans' expectations of when they would expect to see content after a play occurred.

Phase II of TuVista is intended to serve multiple leagues and games simultaneously and provide access to content for fans both inside and outside of a stadium over wifi and cellular data networks. To support this, several major architectural changes were required, including changing from an in-stadium server to a solution hosted on the Internet. In addition to the architectural changes to the system, we put our effort into six areas:

- Video Editing and Bundle Authoring
- Workflow Management
- Related Content
- Device Reach and Notifications
- Faceted Browsing
- Content Transcoding

4.1 Video Editing and Bundle Authoring

In the trial of the Phase I system, it became obvious that requiring two people to run the system would be unrealistic in practice inside most stadiums. Often, space is quite limited and the time of professionals in the stadium is at a premium. To reduce this need to one person, we focused on improving the video editing and bundle authoring workflow. In the Phase II system, these tasks are accomplished on one workstation using a custom Java application. This application can be seen in Figure 5.

In this new system, we wanted to provide multiple angles of video clips to our users. Much of our work on the publication tool has focused on the simple five-second editing of multi-angle video. While the game is in progress, the left hand side of the screen contains video previews from each of the cameras so that the operator can follow the action of the game. Buttons are placed at the top of the interface to grab pre-determined lengths of video (e.g. the last 10, 30, or 60 seconds of action) from all cameras. There is also a button to start/stop recording for plays with a defined beginning and end (e.g. for American Football). This interface removes the need for a linear video editing solution like Final Cut Pro and reduced the rough cutting of multiple angle video to a single click.

Figure 6. The content publication flow in the Phase II system with typical times for each step. Total editing and publication time was reduced to an average of approximately 30 seconds.

Upon requesting video for a given time period, the interface provides a screen displaying thumbnails from even intervals in each video from that period (e.g. a thumbnail for each 5 seconds of the last 20 seconds). In and out points as well as slow motion in and out points can be set simply by clicking on the appropriate frames of each video, turning multi-stream video editing into less than a five second process. These large targets allow for easy navigation and reduce the time needed to move the mouse. The operator can click between frames to reveal a mid-point frame for more accuracy in clipping content. This screen provides savings of minutes over traditional linear editing for each of the video feeds.

Following the setting of editing points, all videos are transcoded from raw to a high quality MPEG4 format and uploaded to the server where they are further transcoded into formats suitable for all of the supported mobile devices that we support. Before uploading content, the operator can choose to associate an action in the game with the clips (e.g. a Goal by Cabañas). These actions can be received from a live event feed (similar to [1, 15]) or can be manually entered by using pulldowns of players and valid actions for the given sport.

Once the new content items are uploaded, the author can view related content that she may wish to bundle with the video that was just captured. This content is retrieved based on the metadata associated with the clips. For example, the player's biography or statistics can be displayed on this screen, or a related advertisement for a given player or action in a game could be listed. The operator chooses the content for the bundle by clicking on the items and the bundle is then published and pushed out to users.

This new interface brings average bundle creation time down to about 30 seconds for one person from fifteen minutes for two people in the Phase I system and makes it much easier to include clips of the action from multiple live cameras. Figure 6 illustrates a typical breakdown of times for each step in this workflow.

4.2 Tabs for Workflows

While running the Phase I trial, we saw that often actions in the game occur back-to-back forcing the editor to work on multiple publication flows at once. Although we designed the new interface such that video editing can be completed in a few seconds, we realized that interesting plays in sports events often occur quite close in time to each other. To help the editor with these situations, we designed the editing interface around workflow tabs. Each time a new video is clipped by pressing the timer buttons, a new tab is opened for editing, annotating, and bundling that particular event in the game. Multiple tabs can be open at once, and the editor can freely switch between them, editing video in one tab while uploading or adding suggested content in another.

These tabs allow for editing content of even the fastest moving sporting events, something that would be very important for our Phase II trials at volleyball games where plays are often just a few seconds apart from each other. This functionality would have also helped us in Phase I when there were many goals in a short period of time.

4.3 Related Content

When creating bundles, the editor needs quick access to related content to include in the bundle. For example, a bundle for a goal in a soccer match might include player biographies, stats, or videos of that player from previous matches. The editor may also wish to include a related advertisement for products that particular players endorse. In Phase I, the Bundle Author spent time visually searching for the appropriate content to bundle together with a video in a long list of content items. As Zaletelj points out, sports are very dynamic events and it is not known what supporting content will be relevant until after the event occurs. [17] For Phase II, we wanted to make sure our publishers could get access to this content quickly through automatic recommendation.

To make this task as simple as possible, we created a screen where editors can see a list of related content and add it to the bundle that they are currently composing. We use the metadata associated with the current video clips as a seed and search our database via a REST request for the player(s) and action(s) involved in the current play. Lists of related content to the player and action are shown as well as any relevant advertisements or user-generated content taken at the time of the action that have already been uploaded. In a typical usage scenario, this screen would show player bios for the players involved in the seed content items as well as previous videos of those players in action. If one of the players is sponsored by a particular brand, an advertisement for that brand would also be shown. Each item is displayed with a check box and the editor can choose which to include with the captured content to form a bundle. The Related Content screen forms the final screen in the bundle production workflow. Our goal was to be able to bundle content together in five seconds from the ingestion of content items, and this simple check-box style interface provided a quick means to accomplish this.

We also use this same mechanism to allow fans to view additional related content from inside of a bundle on their phones. Fans can view all content in the system for a player or action by choosing "Related Content" from any bundle or content item. They are then

provided with a list of metadata from the content items in that bundle, and upon selecting a piece of metadata (e.g. a player's name), they are taken to a content list of all content in the system that matches that metadata value. This provides an easy way to leverage large databases of existing content and helps fans get to the content they are interested in without the need for cumbersome text or voice searches. This functionality was influenced by the faceted browsing work from Berkeley [16] and metadata browsing interfaces developed at Motorola for searching personal multimedia content [3].

4.4 Device Reach and Notifications

While Phase I was focused on a particular mobile device, we wanted our Phase II system to be used by anyone carrying a smart phone into the stadium. The hassles of renting out devices and handling repairs would add too much complication and additional cost to a service such as this. We also saw a need to support users outside the stadium in keeping up with key plays and reliving events in the game. Therefore, we created several versions of the application for popular devices like the iPhone and Windows Mobile as well as a J2ME application for Motorola and Nokia devices. Because users of the system may now be on wifi or cellular networks, the multicast updates used in Phase I can no longer work to keep clients up-to-date. Often subscriber devices are behind NATs and most Internet routers do not forward multicast packets. To overcome this while still delivering push updates of new content, we opted to use XMPP for real-time notification. [12]

Upon entering the mobile application, an XMPP connection is opened with the server. When a user enters the bundle list for an event, the JSON description of that event returned from the server contains information on the XMPP channel associated with bundle updates from that event. The client joins this channel and keeps the interface updated whenever a new message arrives signaling that new content has been published, modified, or recalled. The updates arrive as XMPP messages containing the description of the Media Bundle in JSON format. This connection adds minimal overhead and bundles are usually received within a second of publication time. Because it does not rely on multicast, this mechanism can function on any network and users are free to access the system over 3G, wifi, WiMAX, or any other Internet connection.

These access networks introduce complexity in terms of access control. Often, content rights owners only have rights for particular geographical areas. Stadium owners often can only distribute content within the stadium and broadcasters sometimes have blackout markets. In our Phase II system, we utilize IP filtering to help limit content to particular networks and have the option to require users to enter a code when accessing content over cellular networks. This code may be something unique printed on a ticket or pass distributed within the venue.

4.5 Faceted Browsing

During our Phase I trial, users could only navigate sequentially through the bundles published during an event. It was difficult for users to have a good idea of the number of bundles published or the content of these bundles without scrolling through them all. To address this problem and help users find the content that they are interested in, we wanted to provide multiple ways to view the content for an event. We solved this by creating a tabbed interface on the client that provides views for a Timeline, a Player List, and All Videos as seen in Figure 1. The timeline shows each bundle in the order that it was published. The player list organizes the bundles by the players that feature in them, allowing users to follow the actions of their favorite player throughout the game, or quickly get to a given play by remembering the player involved in that action. Finally the video list provides a sequential list of all video content items published in the event, allowing for the user to keep up with the very latest highlights if they are not interested in all of the related content in the bundles. This is in addition to the related content links in each bundle which provide access to all content for the players or actions featured in a given bundle.

The editor also has multiple ways to view content. The content list can be filtered by player, team, or action to find the appropriate content for inclusion in a given bundle. The publisher can also view client usage statistics for each content item in this interface.

4.6 Content Transcoding

As we now have multiple different types of mobile devices accessing content, it is necessary to transcode and appropriately size content for each device. When new content is uploaded to the system, multiple versions are generated automatically in order to support the most common phones.

We also faced the problem of identifying the video capabilities of particular phones and what type of phone was accessing the service. In J2ME, most phones do not put the standard User Agent header in HTTP requests making it difficult for us to know what type of device is making a particular call. J2ME also does not have any APIs for identifying the phone model. To get around these concerns, we place the contents of the J2ME system property for supported video encodings into the HTTP Accept header as well as the screen size in the UA-windowpixels header.

As we identify new devices, we maintain a mapping between received values and formats that are known to work. When a client requests the metadata for a particular content item, we include the URL for the most appropriate version for that device.

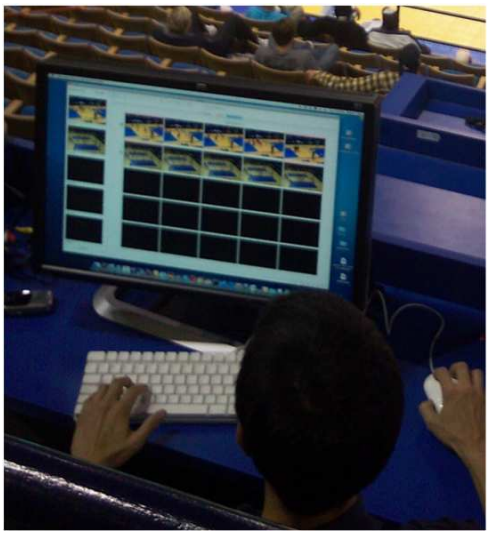

Figure 7. Improved video editing flow in the Phase II trial. Live video previews appear on the left with buttons along the top to grab set lengths of video from the cameras. Using this interface, we reduced multi-angle editing from an average of 15 minutes to an average under 30 seconds.

4.7 User Generated Content

During the Phase I trial, we observed many participants capturing photos on their cell phones or with digital cameras during the game. After the game we asked them about the photos they captured and their plans for them. Most participants intended to share the photos with others soon after the game. They also showed the photos to friends at the game directly from the camera screen.

As capturing media is a core part of the fan experience at a sporting event, we wanted to make that a part of our Phase II system. We wanted to support both personal content that users wished to share with friends and family as well as content of the game that all fans could benefit from. We designed the user-generated content feature such that photos and videos uploaded to the system through the application would automatically be associated with the content bundle closest to the time of capture. Since most cell phones keep their time updated automatically, the time from the phone can be considered accurate enough to use for finding the event in the game that a media item corresponds to.

Upon uploading content, users can decide to keep it private and just forward it to select friends, or they can keep it open for anyone to view. For each bundle, a "Fan Content" link appears in the mobile application that allows access to the related fan content for that action in the game.

5. PHASE II TRIAL

For our Phase II trial, we moved to UCLA and tested the system at both men's and women's volleyball games. Volleyball is a much faster sport than soccer and we believed it would be a much more difficult test for our redesigned editing interface. In volleyball, there is a score almost every minute, with many of these being "interesting" plays for fans (e.g. spikes, aces, dynamic saves, etc.). This dramatically increased the responsibilities of the bundle author from a soccer game with less than a dozen interesting events over 90 minutes of play.

The first UCLA trial consisted of six users, all affiliated with the research team. This trial was mainly to functionally test the system and to focus on the video editing and bundle creation components of the system. A UCLA student ran the system and it was the first time we were able to get feedback from an operator not on our research team.

The second trial was open to UCLA students attending the game. We had a total of 16 users for this event and it was open to both wifi and cellular access, allowing us to prove out our notification infrastructure as well as our transcoding system to support multiple phone types and bitrates. Notices were posted in several classes the week before the game and students could email the researchers if they were interested in participating.

Users loaded the TuVista application onto their own phones before the game. Our users were able to access the system by downloading an ad-hoc distributed application onto their iPhones/iPod Touch devices or a J2ME application onto compatible Motorola or Nokia phones. During the game, they could use the application as much or as little as they wanted. All participants completed a survey after the game and received a $US25 gift card.

As the game progressed, our bundle author created video content items of interesting plays and created bundles consisting of multi-angle clips of a given play along with the bio of the player that was primarily involved in the action. Before the event, we created pre-game bundles with starting lineups of each team, the team's fight song, and a poll asking which team will win the game.

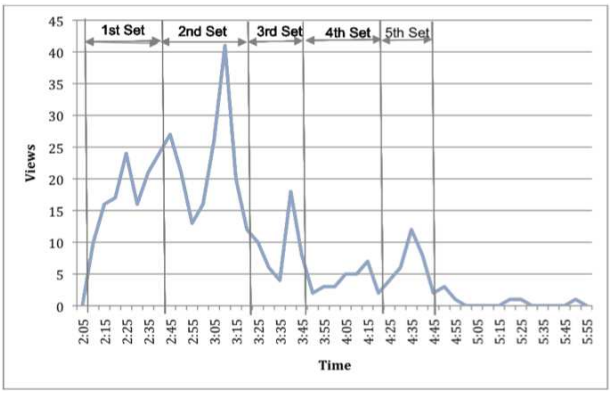

Figure 8. Usage of the TuVista system during the Phase II trial. Content views for each 5 minutes are displayed over the course of the game. Note that several users had to leave after the second set.

As in the first trials, we observed TuVista to be a social experience. Participants were seen sharing the device with their friends and watching videos with those around them. Fewer fans than we expected utilized the user-generated content feature, but we think this may be due to the fact that they had to create an account on the system before posting content and the interface did not make this clear.

Our main focus in these trials was proving our workflow in the challenging situations of a fast moving sport. We were able to reduce the average bundle publication time to 29.8 seconds from action in the game to the time that content appeared on the users' mobile devices. For these games, we produced many more bundles than in the soccer trials, up to 67 per game. With this number of bundles, having multiple ways for users to browse content became important and our users reported using the player and video facets to get to the content they were interested in.

6. DISCUSSION

Our work on the TuVista system has demonstrated the usefulness of mobile sports clips delivered in near-real time to fans following the game. When users are mobile or at a game, clips are important, as users are not often in a position to watch an entire live broadcast. Whether viewing the application to catch up on a game from afar or to relive a play at the stadium, having an appropriate number of clips from key events quickly available has been a critical component of our system and one most appreciated by users. Being able to search for content also became important as users wanted to relive past events together with their friends at the game.

Because the entire solution of TuVista involves many components and many stakeholders, following a traditional user-centered-design process from the start and observing and interviewing each group proved too costly. Instead, by taking some educated guesses and building a rapid functional prototype, we were able to dive into the situations ourselves and learn in the context of a live event. These learnings proved invaluable in validating the concept with multiple stakeholders and for informing the design of the next iteration of our system.

We see rapid testing of prototypes in the field as a means of providing a quick "inspiration for design." Dourish discusses how ethnographic-style field study can be used to help inspire new concepts [5] and we see similar benefits from introducing new technology into the hands of users and observing how they bring these experiences into their lives. These quick studies can often illuminate more about potential bottlenecks and areas that need more focus than much longer studies observing key stakeholders perform tasks in systems that already exist. When working in a new area of technology development, a probe like this quick application can help researchers observe issues more similar to those that would be faced in a final deployment.

Because participants in research studies often cannot accurately predict how they would interact with a new technology or even an old technology in a new setting, we believe that taking new concepts into the field early on in development is the best way to see how a concept works or fails in the hands of real users.

Through developing this system, we faced many challenges that each would delay the total publication time if not adequately solved. Editing multiple video streams was the largest challenge, and we believe that the interface that we've designed does a respectable job at allowing our back end publishers to rapidly identify segments of video that are of interest and quickly edit the boundaries. Taking the system to UCLA several times helped us to get quick feedback from our Content Editor there and helped us make several quick iterations of the backend workflow.

The real-time nature of sports video meant that our system had to be as fast as possible. Through creating our wizard-like multi-tabbed workflow for editing and annotating video, we've sped up the process while allowing for tasks like transcoding to occur asynchronously. Getting new bundle notifications pushed to fans within seconds was always our number one focus.

Annotation proved to be a second large challenge. Assigning the appropriate players and actions to a set of video clips was critical to the operation of our system. Often leagues and other content owners have large repositories of existing content from previous games and seasons. However, this content often goes unused and is underutilized by fans. With appropriate annotations assigned, we've shown how it is possible to use new content as a jumping-off point to explore related content based on its metadata. This expands TuVista from being a mobile experience for an event to truly being about experiencing the entire sport/league it is used for.

The rich browsing experiences that we have developed for the client provide interesting new ways to navigate content with richly structured metadata in order to discover additional content items that are relevant to the user. As databases of sports clips continue to grow, allowing users easy access to these archived clips becomes more important and novel ways to browse through the content to support users' interests become increasingly important.

Still, there are challenges that remain. Ideally, we'd like to import real-time play information from existing services so that clips from the game can be automatically annotated, saving a step for the editor. We would also like a mobile application that could be run in the background and notify fans when interesting events occur in the game. Currently, the backgrounding of mobile applications is not possible on the iPhone and most J2ME-capable devices. Future work on Android may help us to learn about differences in use between a foreground and background application.

7. FUTURE WORK

We hope to continue testing the TuVista Phase II system in larger events in order to better understand usage behavior of fans during an event. We're also interested in testing the editing tool on different types of sports in order to discover different needs for more distributed sports such as golf where the editor likely cannot see all of the action taking place in front of her.

We're particularly interested in usage over time and differences in end user behavior when following a game inside a stadium or in front of a live broadcast versus using the system while mobile as the only way to keep up with a live game. The social aspects of the system are also interesting and we'd like to conduct a larger ethnographic-style study to help understand how the application is shared with co-present others.

We believe that the TuVista editing interface represents a step forward in multi-angle video editing, annotation, and publication and hope to continue our deployments to larger venues. Also, we believe that the methods of early-stage rapid prototyping and evaluation will help us as we approach other future endeavors.

8. ACKNOWLEDGMENTS

We would like to thank Televisa and UCLA for providing access to sporting venues and accommodating our often-unusual requests (including building a plywood table in the broadcast booth of a 120,000 seat stadium). We also thank Larry Marturano, Rajiv Taneja, and Mat Hans for their Larry Marturano, Rajiv Taneja, for their management support of the project. We would also like to thank Siddhartha Mal for being our Editor for the trials at UCLA.

9. REFERENCES

[1] Ault, Aaron, Krogmeier, James V., Dunlop, Steven R., Coyle, Edward J. eStadium: The Mobile Wireless Football Experience, iciw,pp.644-649, 2008 Third International Conference on Internet and Web Applications and Services, 2008

[2] Bentley, F. and Metcalf, C. The Use of Mobile Social Presence, IEEE Pervasive Computing. 2009. (in press)

[3] Bentley, F. and Metcalf, C. Flexible Views: Annotating and finding context-tagged mobile content. Ubicomp 2006 workshop on Pervasive Image Capture and Sharing. September, 2006.

[4] Davis, M. Media Streams: An Iconic Visual Language for Video Annotation. Telektronikk 4.93(1993): 59-71.

[5] Dourish, P. 2006. Implications for design. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Montréal, Québec, Canada, April 22 - 27, 2006). R. Grinter, T. Rodden, P. Aoki, E. Cutrell, R. Jeffries, and G. Olson, Eds. CHI '06. ACM, New York, NY, 541-550.

[6] Gandhi, B., Martinez, A., Bentley, F. Intelligent multimedia content management on mobile devices. ICME 2004: 1703-1706.

[7] Hwang, A., Ahern, S., King, S., Naaman, M., Nair, R., and Yang, J. 2007. Zurfer: mobile multimedia access in spatial, social and topical context. In Proceedings of the 15th international Conference on Multimedia (Augsburg, Germany,

September 25 - 29, 2007). MULTIMEDIA '07. ACM, New York, NY, 557-560.

[8] O'Hara, K., Mitchell, A. S., and Vorbau, A. 2007. Consuming video on mobile devices. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (San Jose, California, USA, April 28 - May 03, 2007). CHI '07. ACM, New York, NY, 857-866.

[9] Miller, N. "Minifesto for a New Age." Wired Magazine. Issue 15.03. March 2007. Available online: http://wired-vig.wired.com/wired/archive/15.03/snackminifesto.html

[10] Naaman, M., Nair, R., and Kaplun, V. 2008. Photos on the go: a mobile application case study. In Proceeding of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems (Florence, Italy, April 05 - 10, 2008). CHI '08. ACM, New York, NY, 1739-1748. DOI= http://doi.acm.org/10.1145/1357054.1357326

[11] Ojala, T., Korhonen, J., Sutinen, T., Parhi, P., and Aalto, L. 2004. Mobile kärpät: a case study in wireless personal area networking. In Proceedings of the 3rd international Conference on Mobile and Ubiquitous Multimedia (College Park, Maryland, October 27 - 29, 2004). MUM '04, vol. 83. ACM, New York, NY, 149-156. DOI= http://doi.acm.org/10.1145/1052380.1052401

[12] Saint-Andre, P., Miller, J. Extensible Messaging and Presence Protocol (XMPP): Core, draft-ietf-xmpp-core-24, May 6, 2004.

[13] Sodergard, C. (ed) 2003. Mobile Television – Technology and User Experiences. Report on the Mobile TV Project. VTT Publications.

[14] Taskiran, C., Li, D., Wang, W., and Gandhi, B. Seamless Content Consumption across Media Delivery Channels and Devices Using Media Bundles, Proceedings of the Society of Cable Telecommunications Engineers Conference on Emerging Technologies, January 23–25 2007, Houston, TX.

[15] Xu, C., Wang, J., Wan, K., Li, Y., and Duan, L. 2006. Live sports event detection based on broadcast video and web-casting text. In Proceedings of the 14th Annual ACM international Conference on Multimedia (Santa Barbara, CA, USA, October 23 - 27, 2006). MULTIMEDIA '06. ACM, New York, NY, 221-230. DOI= http://doi.acm.org/10.1145/1180639.1180699

[16] Yee, K., Swearingen, K., Li, K., and Hearst, M. 2003. Faceted metadata for image search and browsing. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Ft. Lauderdale, Florida, USA, April 05 - 10, 2003). CHI '03. ACM, New York, NY, 401-408. DOI= http://doi.acm.org/10.1145/642611.642681

[17] Zaletelj, Janez. 2008. Recommender System for the Multi-Channel TV Production. In Proceedings of EuroITV 2008.

[18] Zhu, G., Xu, C., Huang, Q., Gao, W., and Xing, L. 2006. Player action recognition in broadcast tennis video with applications to semantic analysis of sports game. In Proceedings of the 14th Annual ACM international Conference on Multimedia (Santa Barbara, CA, USA, October 23 - 27, 2006). MULTIMEDIA '06. ACM, New York, NY, 431-440. DOI= http://doi.acm.org/10.1145/1180639.1180728

[19] Zhu, G., Huang, Q., Xu, C., Rui, Y., Jiang, S., Gao, W., and Yao, H. 2007. Trajectory based event tactics analysis in broadcast sports video. In Proceedings of the 15th international Conference on Multimedia (Augsburg, Germany, September 25 - 29, 2007). MULTIMEDIA '07. ACM, New York, NY, 58-67. DOI= http://doi.acm.org/10.1145/1291233.1291250