

MMM2: Mobile Media Metadata for Media Sharing

Marc Davis, Nancy Van House, Jeffrey Towle, Simon King, Shane Ahern, Carrie Burgener, Dan Perkel, Megan Finn, Vijay Viswanathan, and Matthew Rothenberg

University of California at Berkeley School of Information Management and Systems

102 South Hall, Berkeley, CA 94720-4600

{marc, vanhouse, jtowle, simonpk, sahern, carrie, dperkel, megfinn, vijay, mroth}@sims.berkeley.edu

ABSTRACT

Cameraphones are rapidly becoming a global platform for everyday digital imaging especially for networked sharing of media from mobile devices. However, their constrained user interfaces and the current network and application infrastructure encumber the basic tasks of transferring, finding, and sharing captured media. We have deployed a prototype context-aware cameraphone application for mobile media sharing (MMM2) that aims to overcome these difficulties. MMM2 leverages the point of capture and of sharing to gather metadata, and uses metadata to support sharing. Based on the early results of the first 6 weeks of a six-month trial involving 60 users, indications are that with MMM2 users are actively capturing and sharing photos. The ability to automatically upload photos from a cameraphone to a web-based photo management application and to automatically suggest sharing recipients at the time of capture based on Bluetooth-sensed co-presence and sharing frequency promise to reduce the current difficulty of mobile media sharing.

Author Keywords

Cameraphones, mobile media metadata, photo sharing, context-aware media computing, social software, Bluetooth.

ACM Classification Keywords

H5.1. Information interfaces and presentation (e.g., HCI): Multimedia Information Systems; H.4.3. Information systems and applications: Communications Applications.

INTRODUCTION

In the last few years a new platform for digital imaging, the cameraphone, has rapidly grown in popularity around the world. In the first half of 2003, more cameraphones were sold worldwide than digital cameras, and the trend continues. Cameraphones combine several key features which make them a potentially transformative platform for digital imaging: media capture (photo, audio, video); programmable processing using open standard operating systems, languages, and APIs; wireless networking; personal information management functions; new user interface affordances (especially speech); automatically gathered contextual

metadata (when, where, and by and with whom media is captured and shared); and, as cell phones, are always-at-hand. Nonetheless, current media capture and sharing on cameraphones has been hampered by three key problems: getting media off device; finding and managing media assets; and sharing media. The affordances of the platform, however, contain within them the seeds of a solution: to enable users to automatically upload captured media, and to use contextual metadata to support the management and sharing of captured media. In 2003, we built, deployed, and tested our Mobile Media Metadata (MMM1) prototype running on the Nokia 3650 cameraphone with 60 users over 4 months to use contextual metadata and annotation at the point of capture to infer media content, to answer in effect the question: “What did I just take a photo of?” [2, 7] Based on findings from MMM1, we learned that in addition to capture, *sharing* is a key point in the mobile media lifecycle—one at which metadata is often already produced—and that sharing and metadata could be used in a mutually reinforcing way. In our Mobile Media Metadata 2 (MMM2) system we use sharing to gather metadata and metadata to support sharing to answer the question of “Whom do I want to share this photo with?”

We developed MMM2 to run on the Nokia 7610 cameraphone and have deployed it in a relatively large scale six-month trial to 60 users (40 students, 20 researchers). For the purposes of this paper we have removed the researcher usage data from the dataset reducing the total user population from 60 to 40 (if the researcher usage data were added, the results would be even more striking). Six weeks into the six month trial, we gathered preliminary, but significant, findings: on average per user per day far more pictures have been taken in MMM2 than MMM1 (more than 13 times more—see Table 1); and our use of metadata to support sharing (by automatically suggesting share recipients based on sharing frequency and Bluetooth-sensed co-presence) appears both to support a high degree of sharing (75% of personal photos taken were shared after the introduction of the automatic suggesting of share recipients) and to have increased the rate of sharing by 207% (36% of photos were shared before the introduction of the share guesser). These findings, supported as well by qualitative research with our users [6], indicate the value of the automatic upload of cameraphone photos to a web server as well as the automatic suggestion of sharing recipients.

STATS	MMM1	MMM2	DIFF
Users	38	40	5%
Days	63	39	-38%
Raw totals			
Personal photos uploaded	155	1478	854%
Total photos uploaded	535	1678	214%
Photos not uploaded	108	52	-52%
Average per user per day			
Personal photos uploaded	0.06	0.95	1363%
Total photos uploaded	0.22	1.08	381%
Photos not uploaded	0.05	0.03	-26%
Upload failure rate	16.8%	3.0%	-82%

Table 1. MMM1 vs. MMM2 Upload Statistics

THE PROBLEM: MOBILE MEDIA ASSET MANAGEMENT

MMM2 is designed to address three central problems in cameraphone use: getting photos off the phone; finding and managing photos; and sharing photos.

Getting Photos off the Phones

In our qualitative research both with our users and other cameraphone and digital camera users as well [7, 6], the primary complaint is the difficulty of getting photos off the phone and onto the web so they can be effectively managed and shared. The problem is that there are many ways, but by and large no simple and easy way, to transfer photos from the cameraphone to the web. The plethora of transfer techniques—email, Bluetooth, Infrared, USB cable, MMS—as well as the lack of a uniform destination for these various means of transfer lead to user confusion. As a result, with many cameraphone users, their photos get “stuck” on the device.

Finding and Managing Photos

The popularity of cameraphone photos is also the root of their problems: as users take more and more photos on their cameraphones the affordances of the cameraphone user interface—small screen, slow scrolling, and difficult text entry of annotations and queries—make it exceedingly hard to find photos once the collection is larger than can be easily browsed. The multiple destinations for cameraphone photos also contribute to the challenges of finding and managing them: if users can get the photos off the phone, they are often distributed across a variety of disconnected locations: on the phone, in email attachments, on the web, on the PC, etc.).

Sharing Photos

Sharing photos on cameraphones, like uploading them, suffers from the problem of there being many ways (MMS, email, http upload, Bluetooth and infrared locally), but no easy consistent way to share photos from the cameraphone. On most cameraphones and wireless carriers sharing faces the following additional problems: phone hardware and software incompatibilities (can the recipient receive an MMS, etc.); multi-step, time-consuming process (often

email to web account, and then share from web); uncertainty about the cost model (who pays how much for what?); managing lists and groups of sharing recipients on the limited phone UI; and finding or typing in email addresses. MMM2 works to address these problems by: enabling automatic background uploading of cameraphone photos to a web-based photo management and sharing application; acquiring system and user created metadata to support photo management and sharing; and automatically suggesting sharing recipients on the phone and in the web application based on sharing frequency and Bluetooth-sensed co-presence.

RELATED WORK

Our prior work in mobile media computing has addressed using contextual information at the point of capture to infer media content [2] and several recent systems have attempted to aid cameraphone photo sharing through simplified or automatic uploading of photos and offering lists of sharing recipients on the phone and on the web [1, 4, 5]. MMM2 also enables automatic background photo uploading, but overcoming some of the limitations of prior work, leverages metadata about prior frequent recipients and Bluetooth-sensed [3] co-present potential recipients to automatically rank and suggest sharing recipients for a cameraphone photo.

THE SOLUTION: MOBILE MEDIA METADATA

By gathering metadata at the point of capture and through the sharing process itself, MMM2 produces and uses metadata in a mutually reinforcing way to automatically suggest likely sharing recipients. At capture time, MMM2 automatically records: *spatial context* (through CellID, Bluetooth-enabled GPS devices, and fixed location Bluetooth objects such as PCs); *temporal context* (through network time servers); and *social context* (who is co-present determined by the association of sensed mobile Bluetooth devices and user names). Additional metadata may be entered by the user at capture time (an optional caption—see Figure 1 right); at share time on the phone (individual and lists of sharing recipients, albums, and additional textual messages for email or blogging—see Figure 2); and on the web (the same metadata as at share time on the phone plus comments and faceted metadata from a variety of RDF ontologies).

System Description

The MMM2 architecture consists of three major components: the context logger, the metadata backend, and the photo management UI. The context logger (developed by and modified in cooperation with the University of Helsinki Department of Computer Science Context Project: <http://www.cs.helsinki.fi/group/context/>) is a Series 60 application that both captures and uploads newly created media (photos, videos, and sound clips) and related contextual metadata. This information is uploaded via http to a servlet, parsed, and stored in our metadata backend.

The metadata backend consists of two major parts, both

written in Java: the faceted metadata engine and the photo management engine. The photo management engine is backed by a relational database and stores MMM2-specific information, such as captions, comments, and sharing data as well as all information uploaded by the context logger. The faceted metadata engine is also backed by a relational database, but this database is structured to store graph information expressed in RDF triples, using the Jena toolkit, and is designed to be accessible by other media metadata applications.

The photo management UI is a suite of JSP pages that access data through either the photo management engine or the faceted metadata engine. These pages run on the same server as the metadata backend and access the data through our API, as both run in the same JVM. The UI component provides services for both our PC-based browser component and the Series 60 Opera browser component. The Opera browser component uses style sheets specially designed to work on a mobile handset.

Cameraphone Usage Scenarios

In MMM2 there are two main usage scenarios on the cameraphone, with additional functionality on the web application. Our focus here is on the capabilities of MMM2 to support automatic uploading and sharing of photos from the cameraphone.

Capturing ► (Captioning) ► Background Uploading

The goal of our design of the image capture interaction was to keep it as simple as possible. We use the built-in camera software that ships with the phone, using a background pro-

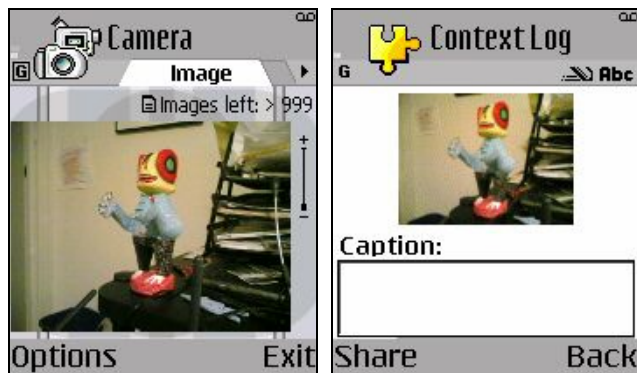


Figure 1. Nokia Camera App Screen and MMM2-Modified Context App Upload Screen

cess to monitor this application’s activity. Immediately after a new photo is captured (see Figure 1 left), the user is prompted to caption and/or share the image. The user may optionally enter text in the caption field, which makes use of the t9 dictionary functionality. The user has the choice of returning to the camera application (by pressing ‘Back’) or sharing the image (by pressing ‘Share’). Meanwhile, the photo begins uploading in the background. If a data network is not available, the photo is queued until the phone re-enters the coverage area.

Capturing ► (Captioning) ► Sharing

The sharing process begins in the same way as the background upload process. After media capture, the user may optionally caption the image, then decide whether to share immediately or simply let the image upload in the background for later use.

If ‘Share’ is clicked, the Opera web browser is launched and directed to a user-specific sharing page. From this page, the user may share with any number of suggested recipients, share with recipients from his or her address book, share with a pre-determined list of people (like a mailing list), post to a blog, and/or add the photo to an album.

The Suggested Recipients list is currently based on two sets of information: Bluetooth-sensed co-presence and a user’s sharing history. The co-presence list is built from the list of nearby devices and the stored association of person names and Bluetooth IDs entered by MMM2 users. Any person known to the MMM2 system may be suggested to a user who captures media in his or her presence, which encourages serendipitous media sharing.




Figure 2. MMM2 Share-a-photo Screen (Joined Screenshots)

The current, very simple Share Guesser, which assembles the Suggested Recipients list, first selects users who are present when the shared photo was captured. Then the list of users with whom the sharer has most frequently shared in the past is appended to the Bluetooth co-presence list. Finally, users with whom the shared photo has already been shared are removed from this list. The resulting list is ordered with Bluetooth co-present users at the top of the list, followed by frequent share recipients in descending order of frequency.

INITIAL RESULTS AND FINDINGS

We tracked many quantitative aspects of the uploading and sharing of photos in MMM2 as well as conducted qualitative user research (7 person focus group, 9 person usability study, and 12 interviews described in [6]). Our findings show both that a large number of pictures were uploaded and shared and offer explanations of the factors contributing to these phenomena.

Photo Capture and Uploading

Compared to last year's MMM1 Project [7, 2], which used the Nokia 3650, with MMM2 and the Nokia 7610 students on average uploaded far more photos with MMM2 (See Table 1): 1.08 photos per user per day vs. 0.22 photos per user per day (an increase of 381%). But when we remove from these figures the photos that students were assigned to take and upload (5 photos per student in MMM2, and 10 photos per student in MMM1), and examine the photos they took and uploaded for personal or other non-assigned reasons (listed as "Personal photos" in Table 1), the difference is even more striking: 0.95 photos per user per day vs. 0.06 photos per user per day (an increase of 1363%). This remarkable increase in cameraphone use between MMM1 and MMM2 appears to have a variety of contributing factors we gathered from our qualitative user research: better image quality (VGA vs. 1 megapixel image resolution, "night mode" for low light, and digital zoom); familiarity of the user population with cameraphones (12 prior camera-phone users this year vs. 1 last year); the availability of only one, rather than 2 camera applications in MMM2 vs. MMM1; automatic background upload of photos to the web photo management application; and automatic support for sharing on the cameraphone and on the web.

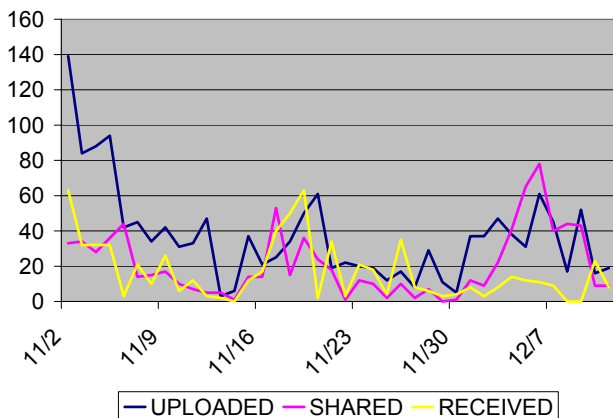


Figure 3. Photos Captured/Uploaded, Shared, and Received.

Photo Sharing

Removing the 200 photos that were required to be taken and shared in a class assignment from the totals, 1478 photos were uploaded between the introduction of MMM2 on 11/02/2004 and 12/11/2004, of which 840 were shared (57%). If we look just at the sharing data for shares after the share guesser was introduced (11/15/2004), the results are even more striking: 790 photos were uploaded, of which 591 were shared (75%). The MMM2 share guesser listed the names users wanted to share within the top 5 entries in the Suggested Recipients list 60% of the time and in the top 10 entries 77% of the time. The use of Bluetooth-sensed co-presence data increased the share guesser accuracy over frequency of sharing data alone by 15%—this contribution will likely increase as we obtain more Bluetooth co-presence data when we are able to override Bluetooth being automatically powered off periodically on the Nokia 7610.

Users found the share guesser to be accurate and helpful in most situations: "The Suggested Recipients are usually right on." As illustrated in Figure 3, the increase in photo sharing after the introduction of the sharing guesser subsided briefly over the Thanksgiving vacation as students had local, not national data plans, and could not use their cameraphones outside the Western United States. After returning from break, photo sharing continued to increase and we look forward to observing the trends over the next 4½ months of data collection.

CONCLUSION

MMM2 addresses fundamental challenges in mobile media application design by enabling the automatic upload of photos and providing an easy sharing interface on the camera-phone that automatically suggests sharing recipients based on Bluetooth-sensed co-present people and frequency of prior sharing. Future work will improve the share guesser by using additional metadata as well as machine learning techniques to better predict with whom the user wants to share cameraphone photos (both locally and remotely) based on the spatial, temporal, and social context of capture and the user's sharing history.

ACKNOWLEDGMENTS

We wish to thank our students and our sponsors—British Telecom, Hewlett-Packard, Nokia, France Telecom, Ricoh, TEKES, Opera, HIIT, and the UC Discovery Grant for Digital Media—for their generous support of this work.

REFERENCES

1. Counts, S., and Fellheimer, E. Supporting Social Presence through Lightweight Photo Sharing On and Off the Desktop. In *Proc. CHI 2004*, ACM Press (2004), 599-606.
2. Davis, M. King, S., Good, N., and Sarvas, R. From Context to Content: Leveraging Context to Infer Media Metadata. In *Proc. MM 2004*. ACM Press (2004), 188-195.
3. Eagle, N., and Pentland, A. Social Serendipity: Proximity Sensing and Cueing. MIT Technical Report (2004): <http://vismod.media.mit.edu/tech-reports/TR-580.pdf>.
4. Gossweiler, R., and Tyler, J. Plog: Easily Create Digital Picture Stories through Cell Phone Cameras. In *Proc. IWUC 2004*, INSTICC Press (2004), 94-103.
5. Sarvas, R., Viikari, Mikko, Pesonen, J., and Nevanlinna, H. *MobShare*: Controlled and Immediate Sharing of Mobile Images. In *Proc. MM 2004*, ACM Press (2004), 724-731.
6. Van House, N., Davis, M., et al. The Uses of Personal Networked Digital Imaging: An Empirical Study of Cameraphone Photos and Sharing. In *Ext. Abstracts CHI 2005*, ACM Press (2005).
7. Wilhelm, A. Takhteyev, Y. Sarvas, R., Van House, N., and Davis, M. Photo Annotation on a Camera Phone. In *Proc. CHI 2004*. ACM Press (2004), 1403-1406.