Increasing contribution in online communities using alternative displays of community activity levels

Patrick Barry, Uri Dekel, Neema Moraveji and Justin Weisz
School of Computer Science
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213
{pbarry,udekel,nmoravej,jweisz}@andrew.cmu.edu

Abstract

Our goal is to increase contributions to an online community by making information about the level of activity in the community salient to its members. We specifically focus on the LiveJournal web log community, where a user's community is defined by their personal list of friends, and contributions are made by either posting new content to one's own journal, or by commenting on existing journal entries.

Our theory-based approach to increasing contributions is discussed in depth, along with results from an experiment designed to measure the effectiveness of various activity level displays.

1 Introduction

Online communities, like their physical-world counterparts, live or die by the participation of their members. Experience shows that it is not enough to have a community that is centered on a certain topic or goal. There has to be a critical mass of people keeping the community alive, contributing time and resources to the group [9]. Since the social and economic capital of online communities is produced by these contributions [3], finding means for increasing contributions is of utmost importance, and is the primary goal of this work.

Our goal is to increase the number of contributions made in an online community, with a focus on online journals (also known as ”web logs”, or just ”blogs”), by making interactions with the community more synchronous. We achieve this goal by presenting users with a display of the *activity level* of their community, designed to inform users how active their community is. We believe that by increasing a user's synchronous awareness of community interaction, we can induce into users a much greater desire to contribute and interact within the community.

Outline: The rest of this paper is organized as follows: Section 2 describes our approach to increasing contributions from the social science literature, as well as LiveJournal, the specific community we focus on. Section 3 gives background information on ambient devices in general, as well as information regarding the specific devices we use to display activity level information. Our experiment and hypotheses are detailed in Section 4. Results from our experiment are presented and discussed in Section 5. Finally, we present our conclusions and suggest avenues for further research in Section 6. Appendix A presents the architecture of the software we built for use in our experiment.

2 Theoretical background

In order to satisfy our goal of increasing participation in an online community, we first built a model describing factors which increase contribution, by consolidating theories of social presence, computer mediated communication, and communication context cues. This section describes our model, which is depicted in Figure 1.

2.1 Causal model of contribution

Social presence is the ability of a medium to provide members the affirmation of peer presence; i.e. the degree to which interaction in a medium subjectively compares to face-to-face interaction. Viegas and Donath [14] have found that “[social] presence has a strong impact on the style of discourse, for participants often feel compelled to constantly post messages so that they will not be forgotten by the others”, and thus it is a strong generator of the desire to contribute. The need and desire for social presence has been observed in many groups. For example, teenagers often feel compelled to appear available to others in an attempt to maintain their status among their peers [5, 6]. Furthermore, Short et al. [12] argue that an increase in the social presence of group members will increase one's efforts in maintaining presence within the group. In the case of online communities, this means more contributions and member interaction.

Social presence can be established in many ways. As Short et al. [12] stated in their seminal work on the subject, one of the primary means of increasing social presence is by increasing *salience*, the conspicuous prominence of the communication partners in the awareness of any given community member.

Walther and Parks [15] found that adding timestamps to asynchronous computer mediated communication was a sufficiently significant social cue to affect judgments of affection by participants, yielding significantly more enjoyment of the medium since participants perceived the connection as being more "real". Explained in social presence terms, adding a dimension of time to the presence of community members and their contribution provided a sufficiently significant increase in the salience of these members, which was also sufficient for increasing the perception of presence. Thus, we know that the added dimension of time can increase social presence, in turn raising contributions.

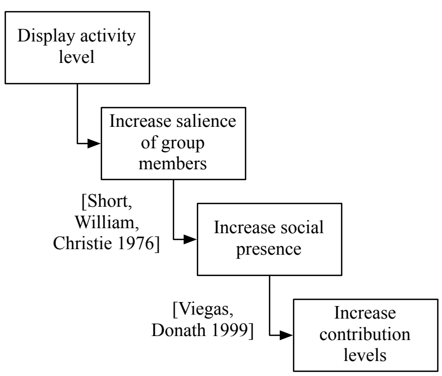

Figure 1. Causal model for increasing contributions to an online community.

Putting these theories together, we find that in order to increase the number of contributions made to a community, we should seek to increase a user's *motivation* to make contributions to their community. One way to increase a user's motivation to make contributions, is to increase the salience of other community members, making the communication feel more like face-to-face communication. This is done by making communications more synchronous, which has the effect of increasing the social presence of the members of the user's community. Increasing social presence has the direct benefit of increasing the number of contributions made by the user. Thus, in order to bootstrap this process, we introduce the concept of an *activity level display*, which is a display of the activity level of a user's online community. This display is designed to make communication between members more synchronous, thus increasing the salience of other community members, and increasing the amount of social presence of those members as felt by the user, and finally increasing the contributions made by the user. A diagram of this process is presented in Figure 1.

2.2 Making blogs more synchronous

All virtual communication mediums can be characterized by their level of synchronous communication. We define the notion of *synchronicity* for a communication medium as the medium's position in the spectrum depicted in Figure 2. This figure indicates the level of synchronicity for several types of communications mediums.

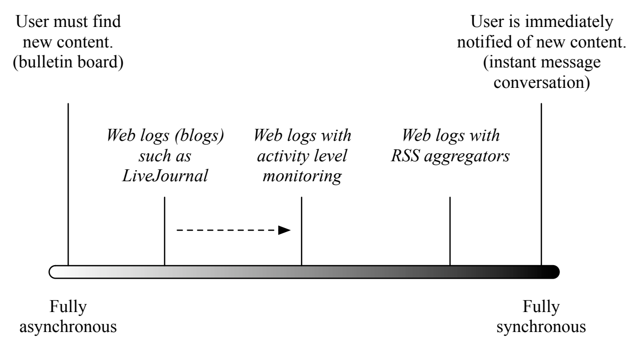

Figure 2. Synchronicity levels of different online communities.

As can be seen in Figure 2, on the one end of this spectrum lie bulletin board systems, such as message boards and newsgroups. These are highly asynchronous communication mediums as it is up to the reader to check them for new updates. On the other extreme lie nearly-synchronous mediums such as instant messaging, where parties communicate with each other in real time (e.g., Unix *talk*, where any typed character appears on the other party's screen) or close to real-time (e.g., AIM, where entire sentences are composed before they are sent).

Developers and virtual community designers add various features to manipulate the level of synchronicity of these mediums, and thus influence the salience and social presence of the communication. For example, some bulletin board systems automatically dispatch emails to the author of a certain post whenever somebody else replies to that post.

Web logs (blogs), by the same token, have evolved in their use and feature set to accommodate different levels of synchronicity. In addition to sending notification emails for comments (as in LiveJournal [7]), the salience of the blogging medium has also been manipulated by features such as *Trackback* [13] and *Instant-Gratification* [4]. Software tools known as *RSS aggregators* are also used to notify users of activity in blogs. From talking with several members of the web log community, we found that many bloggers choose not to use any of these features because of their distracting nature, which counteracts the benefits of added synchronicity. Our proposed manipulation provides a slightly lesser degree of synchronicity, but does not incur the cost of distraction.

2.3 Blog community testbed

LiveJournal [7] is an online community made up of thousands of bloggers. In fact, it can be though of as a collection of communities, since each user has their own virtual community which consists of themselves and other bloggers whom they regard as friends. In addition, there are shared community blogs centered around specific subjects. These features make LiveJournal an ideal testbed for experimentation. It provides the public with a mechanism for remotely collecting blog usage data, and the site's user base is typically active and enthusiastic.

3 Ambient displays

We designed our activity level display as an *ambient* display for two main reasons. Firstly, we only have a single piece of information to present to the user: the activity level. Secondly, we wanted a display which would be minimally distracting and easy to read at a quick glance. Ambient displays satisfied both of these requirements.

3.1 Background

An ambient display is one that delivers non-critical information to a user in a passive manner. Users perceive this information in their periphery, gaining an overview of the information without requiring the shifting of focus from their other activity. They may gain additional insight by explicitly focusing on the display and possibly interacting with it. It has been shown [10] that designers can successfully communicate additional information to users through ambient displays. In today's world, where the amount of available information is increasing but the human ability to focus remains constant, ambient displays provide an opportunity to communicate information to users in a way that doesn't inundate them. Some argue that because ambient displays are located in the 'everyday spaces' of the user's world, aesthetic quality should play a particularly large role in their design [8].

By far, the most prevalent use of ambient displays has been the abstraction of real-time information and the visualization of this data to users. Recently, such displays were also used to display trends of this data over time rather than just its instantaneous value [1].

Because relatively few ambient displays have been designed and implemented, there is still a great deal of experimentation and an abundance of new ideas developing in the field. The *Microsoft Sideshow* [2] team has come up with a list of six design principles for designing their ambient disp.la. The team postulated that a Sideshow must: *(i)* always be present, *(ii)* minimize visual motion in order to minimize distraction, *(iii)* provide information that is highly relevant to the user, *(iv)* be extensible so others can leverage the ambient platform, *(v)* support quick drilldown and escapes and *(vi)* be scalable.

In order to determine whether the modality of the display has any effect on contributions, we decided to design both physical and computer-based activity level displays. In designing these displays, we followed the Sideshow rules as guidelines when applicable. Namely, our displays are always present in the user's environment, only a single color is used to display information (thereby minimizing distraction from complexity), and the information provided is of interest to the user. Since we have two types of displays, we make the following distinction: our *ambient display* is a physical device in the user's environment, and our *peripheral display* is a display on the user's computer screen, situated in the periphery of their vision. Both displays display exactly the same information. The next section discusses the physical display we used in our experiment.

3.2 Orb device

The Ambient Orb device [1] is a frosted glass orb which glows a particular color to indicate a piece of information. It operates on AC power, and communicates through a national wireless network[1] in order to receive updates as to what color should be displayed. It can display any color in the spectrum, with a specific gradient defined for each application. The orb device is designed to sit in a visible location, and receives updates once every 10 minutes from the wireless network. A picture of the orb device is given in Figure 3.

3.3 Ambient Display of Community Activity

Our activity level displays show a measure of the number of blog entries and comments on entries in a user's com-

[1]This network is called the Ambient Information Network, and the orb's manufacturer claims to have 90% coverage in the United States.

Figure 3. The Ambient Orb Device.

munity (i.e., a user's journal and the journals of his or her friends). In order to present this information to the user, we map activity levels to colors and intensities, where brighter colors denote higher levels of activity. To handle differing sizes of personal communities, the level of activity displayed is actually a ratio of the activity level for the current time period, to the average activity level over all time periods. Time is discretized into fixed sized chunks (buckets) in order to sample activity over meaningful periods of time.

4 Experimental Design

This section is written as if we had already performed the experiment. Some numbers which we do not have yet, such as the number of drop-outs, have been replaced with more general variables. These will obviously be filled in as we learn of them, but we wanted to give a general idea of the some of the statistics that would actually go in a conference paper.

In order to test our hypotheses, we performed the following experiment. Our primary independent variable is the display type that a participant is given, and it has three levels: *no display*, a *peripheral display* on the participant's desktop computer, or an *ambient display* which is the orb device placed on the participant's physical desktop. The orb device is discussed in Section 3.2.

Our experiment uses a within-subjects design, where participants are rotated through periods where they use each type of display. We correct for ordering effects by using two counterbalanced groups, with respect to the first display received by the participant.

Our primary dependent measure is *contribution*, which is operationalized as the number of posts and comments made by a participant. Activity levels are operationalized as the ratio of the number of posts and comments made by all users in a participant's community in the current time period[2], to the average number of posts and comments made in all time periods. Posts and comments made by a participant do not factor into the activity level calculation. Any other interactions with the LiveJournal web site, such as reading other friends' posts, or interactions which take place outside of the LiveJournal are not counted as contributions because they do not generate new content which others can consume (i.e., the community derives little or no value from a user who simply reads messages).

[2]For this experiment, we used time periods of 30 minutes.

4.1 Participants

Participants in our study consisted of [N students from Carnegie Mellon University] (X% male), and were all active LiveJournal users. [N participants dropped out of our study for X, Y and Z reasons.]. Potential participants were disqualified from participating in our study if they indicated that they used third party software to interface with *LiveJournal*. Examples of such software include *RSS aggregators* or other programs which monitor a participant's LiveJournal activity and inform the participant when new content has been posted. Such software directly interferes with our manipulation, as it provides the same (if not more) activity level information to the participant.

4.2 [Proposed] Method

Participants were first asked to sign informed consent forms, as our experiment requires monitoring a participant's behavior through the use of a software program. They were also provided with the software itself, and received help in installing and operating it, when necessary.

Each participant was monitored over five week-long periods, during which our experiment software (presented in Appendix A) recorded information on the activity in the participant's community. This information included the posts and comments made by the participant, as well comments and posts made by the participant on the participant's *friends list*. The information collected in the first period was used to determine base rates of activity and participation for that participant, so that the activity displayed to the participant was relative to the usual activity in the their community. This allows us to conduct the experiment with a heterogenous set of communities, and still measure relative changes in activity.

In each of the five weeks of the experiment, participants received one of the two activity display types (peripheral or ambient) or no activity display at all. These are presented in Table 1. Participants also filled out online questionnaries before the start of each phase of the experiment, designed to measure their satisfaction levels of using the current display, the degree of connectedness they felt with their community, and the amount of motivation they felt for contributing to their online community. A pre-experiment survey also measured several demographic variables such as age, gender, and blogging experience, and a post-experiment survey collected each participant's final thoughts on each of the display types. (*These surveys have not yet been written. They will be before we run the experiment.*)

Phase	Week	Group #1	Group #2
Pre-manipulation	#1	None	None
Manipulation	#2	Peripheral	Ambient
	#3	None	None
	#4	Ambient	Peripheral
Post-manipulation	#5	None	None

Table 1. Display types in experiment phases.

As shown in Table 1, participants were randomly assigned to one of two groups. No activity display was provided in the first week, although the monitoring software was used for the baseline activity calibration. In the first of the three manipulation weeks, participants in the first group received the peripheral desktop display, while those in the second group received the ambient orb display. The next week was a transition period designed to return contribution levels to their base rates, and thus no activity information is displayed; this also allows us to measure post-display effects. The fourth week provides participants with the second display option, followed by the last week where where activity levels are expected to return to normal as no activity display is provided.

4.3 Hypotheses

As part of our experiment, we test the following hypotheses:

H1: Presenting an activity level increases contributions.

Our main hypothesis is that participation in an online community will be increased when a participant is presented with activity level information. This hypothesis follows directly from the causal model presented in Figure 1. It has been shown that increases in the salience of the other members in a participant's community directly translate into increases in social presence [12], and increases in social presence translate to increases in contribution [14].

H2: The manipulation effect is stronger for an ambient display than for a peripheral display.

The ambient orb device with which our participants are provided is likely to be more prominent in their environment than the peripheral desktop display. The desktop display can be seen only when the participant actually looks at the monitor and focuses on the display (which is located in the system tray). It is not seen when hidden by other windows, playing full-screen games, when the screen saver is running, and in many other cases. The orb, on the other hand, is constantly on. Whether the participant is watching the screen or not, he or she is likely to notice changes in the orb's intensity in the periphery of his or her vision. Thus, even when the participant has gone to bed (assuming the orb is kept in the same room), the participant will be aware of the activity in his or her community.

H3a: Our displays provide enough information to stimulate a coordination effect.

This assumes that display updates are highly detectable, because the update time period is enough time for activity to significantly change. Since we do not yet have the orb devices, we are unsure of how detectable updates to the orb's color will be.

A *coordination effect* occurs when the difference between the time a participant comments on a post and the time when the post was originally made is significantly lower in the activity display conditions than in the no-display condition. We expect that a feedback loop will be created in the presence of an activity level display. Our display updates are highly detectable to participants, such that when the display is updated after the update period, a noticeable color shift takes place. This detectable shift in color would then prompt the user to check the latest activity in their community, which would potentially lead to the creation of a new contribution. After the new content is created, other members of the community have one update period in which to post replies to the participant's contribution, and have the participant see the activity in the next time period. However, generation of a large number of replies during this period[3] will increase the amount of activity seen by the participant in the next time period, again prompting the participant to check the new activity, and thus completing the feedback cycle. Thus, a coordination effect is present in this cycle, since the participant is given enough information from the activity level updates to know when new content is available, giving the participant more opportunities to make contributions. In order to negate the confounding effects of coordination on the actual mechanism behind increases in contribution (motivation to contribute), we will analyze contributions in terms of the number of contributions per unit of community activity.

[3] Perhaps because the participant goes through several synchronous rounds of reading and contributing.

H3b: Our displays do not provide enough information to stimulate a coordination effect.

This assumes that display updates are not highly visible, because the update time period is not enough time for activity to significantly change. This hypothesis is mutually exclusive to H3a, and really depends on our evaluation of the orb, and what activity level changes really look like.

A *coordination effect* occurs when the difference between the time a participant comments on a post and the time when the post was originally made is significantly lower in one condition than another. For such an effect to occur, we would see that the difference in timestamps between posts and comments was significantly different in a display condition than the no-display condition. However, we do not believe such an effect will occur because the activity level is updated only once every 30 minutes[4], and we have noticed that changes in the display are gradual enough that participants will be unable to tell exactly when new content arrives. Rather, they simply have information as to the relative amount of activity which has occurred in the past 10 minutes.

H4: Participants will feel motivated to make contributions because of feelings of connectedness to their communities.

Most participants will have increased contribution rates, and contribution increases that are not the result of a coordination effect will be the result of motivations to contribute to their LiveJournal community. We expect to see participants report these motivations, as well as greater feelings of connectedness to his or her LiveJournal community.

H5: Participants with larger communities will hold stronger feelings of connectedness to their communities after being presented with an activity level display.

Participants with larger communities will experience more variance in their display because of the larger volume of activity in their communities. This increased variance will make changes in the activity level more noticeable in larger communities than smaller communities (where activity is more infrequent). In general, increased awareness of these changes should cause stronger feelings of connectedness to the community, because of the increased salience of the community's interactions. Thus, we expect that participants with larger communities will see more noticeable activity level updates, and hence will have stronger feelings of connectedness to their community.

[4]The orb device is on a fixed 10 minute update interval. Thus, every two out of three updates will not change the orb's color.

5 Results

Note: at the time of writing this paper, we have not yet performed the experiment and do not have the results. The following describes our expected results.

After performing numerous statistical tests on our data, some of which will go in this section, we present justifications for our hypotheses. These are not designed to be exhaustive, as there are many possible outcomes. Rather, they attempt to describe one particular outcome for each case of validation/non-validation.

If there is evidence supporting H1: We found that participation levels were increased when a participant had information regarding the activity of their LiveJournal community. Further, participants reported that they felt closer to their community of friends, and that they were more motivated and compelled to contribute more often when they had the activity level information than when they did not. This validates the causal model presented in Figure 1, that activity levels are increased because of an increase in social presence (and not another mechanism such as a coordination effect).

If there is no evidence supporting H1: We did not find that participation levels were increased when a participant had information regarding the activity of their LiveJournal community. However, participants did report that they felt closer to their community of friends. From this we conclude that the displays were unable to convey enough information to coordinate participants with specific occurrences of activity, and thus a different type of display which could provide a more fine-grained display of activity would be more appropriate for increasing contribution levels.

If there is evidence supporting H2: Participation increases were greater for the ambient display than for the peripheral display. Participants reported that they paid more attention to the ambient display because it was visible at all times, whereas the peripheral display was not always visible to them.

If there is no evidence supporting H2: Participation levels were not significantly different between the ambient display and the peripheral display. This could either be because participants only monitor both types of displays when using the computer (especially in the case where the orb device is placed near the computer), or because there simply was no benefit in moving the display from the computer desktop into the physical world.

If there is evidence supporting H3a: Participants experienced a coordination effect when using the activity level displays. The average delay between a post made by a participant's friend, and a comment made by the participant was $[X]$ under the display conditions, and $[Y]$ under the no-display conditions $[(p < Z)]$. *Of course, we will have fancier statistical testing here, but this is the basic idea.*

If there is evidence supporting H3b: Participants did not experience a coordination effect when using the activity level displays. This is due to the activity level displays not providing a sufficient amount of information as to the amount of new content available during display updates. Participants reported that they were unable to determine exactly when the display was updated under periods of constant activity, since the color change was not radical enough, and thus they were unsure of whether or not any new content had actually arrived.

If there is evidence supporting H4: Participants reported that they felt more connected to their community, and felt more motivated to make contributions, under the display conditions.

If there is evidence supporting H5: Participants with larger communities experienced greater feelings of connectedness with their community of friends than participants with smaller communities. Participants with larger communities reported that they noticed their display changing color much more often than participants with smaller communities. The exact reasons for why this occurs will be made clearer by the results of our post-manipulation survey.

6 Conclusions

Our conclusions and future research directions depend on the results of the experiment.

7 Acknowledgements

We would like to thank Dr. Bob Kraut and Dr. Paul Resnick for their valuable feedback, and David Rose of Ambient Devices, Inc. for providing the orb devices.

A Software Architecture

In this appendix we present an overview of the architecture and the algorithms used by our system.

A.1 Architecture

Our system follows a standard client/server architecture, with our client software running on the participant's machine, and the experiment server running on one of our systems. The client has two roles. One is to provide an ambient or peripheral display of the activity level in the community. Another is to collect information about the activity in the community from the community server, process it, and send it to the experiment server. The experiment server maintains a repository of information over the experiment period, including all the activities and contributions that took place in the community. We collect all of this information so it will be traceable and we can analyze it if the need arises. The server also calculates a value for the community activity level, which the client requests and uses to update its display. This architecture is presented in Figure 4.

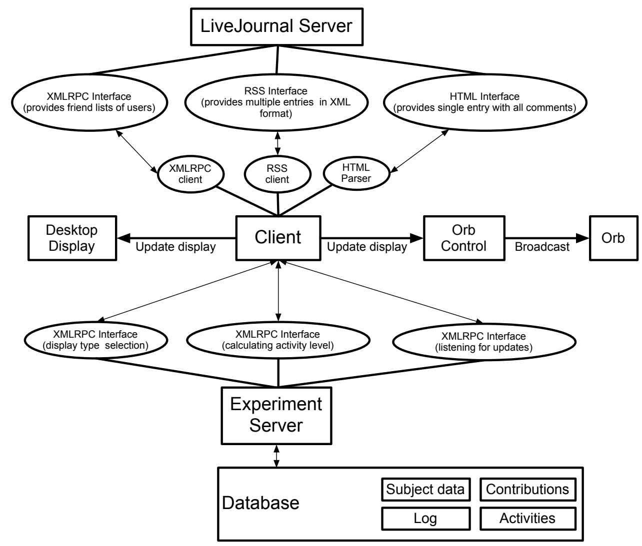

Figure 4. Structure of our system.

One may wonder about the fact that the client is responsible not only for displaying the activity level, but also for collecting activity information from the community and passing it on to the experiment server. After all, if the client eventually receives the activity level information from the server, why would it be responsible for collecting it in the first place?

There are two reasons for this. First, some features of the community server, such as retrieving the list of friends of a LiveJournal user, require the user of a password. By having the client software collect the information, the user never has to supply us with their password and we do not have to store it on the server. It is only stored in the non-persistent memory of the client software. Also, because of the current limitations of the LiveJournal software, we have to open a lot of HTTP connections to the community, and also parse a lot of HTML files looking for comments. By having each client carry out these time consuming operations for itself and just reporting the results asynchronously to the experiment server, we are able to accommodate multiple participants at the same time despite the limited power of the server.

Nevertheless, before conducting the experiment we will try an alternate architecture where the server carries out all the activity. This will reduce the dependency on the client software being active during the entire five weeks of the experiment, since there is currently a risk that the participants will not keep it running during the weeks when no activity display is provided, or when the orb is used.

A.2 Operation

Let us explore the operation of the client, which is presented in the sequence diagram in Figure 5.

When the client program is activated for the first time, the participant is presented with a login screen where she provides her subject ID, as well as her LiveJournal user name and password. The subject ID is used to contact the experiment server to request the display instructions for that day[5]. If the display instructions for that week call for the use of a desktop display, an appropriate widget is instantiated.

Routinely (as specified by a customizable interval) the client polls the experiment server for the current activity level. This information is used to update the display widget, or to broadcast instructions to the peripheral display device (the orb).

As mentioned, the client is also responsible for collection activity information from LiveJournal and updating the experiment server. Because at the time of writing the software LiveJournal did not provide APIs for easily carrying out this task, the process is elaborate and is thus described.

The client starts by using LiveJournal's XML-RPC protocol to request the updated list of friends of the user. This list is returned and processed. The client proceeds to examine the journals of each of these friends for activity, as follows.

LiveJournal provides a collection of a user's recent[6] entries in a convenient XML format based on the RSS protocol [11]; this is intended for use by news aggregators, and was a convenient way to get a series of recent posts. The client receives standard RSS entries for each journal entry, and processes each of the new entries. By processing we mean reporting the new activity to the experiment server for tracking, as well as for searching for new comments on this entry.

At the time of writing our software, LiveJournal does not provide convenient programmatic access to the comments on journal entries. We need these comments for tracking contributions and activity, and therefore have to infer them from the entry's HTML page. LiveJournal provides access to each journal entry along with its comments using a unique URL, which is available as part of the RSS feed we got earlier. For each journal entry, the client requests that page from the server, parses the HTML and looks for all the comments on that page, collecting information on who posted them and when. Each comment is reported to the experiment server as a contribution (if it is a comment the participant has made on another journal) or an activity (if someone else made the comment). The experiment server receives the reports and records the new ones in the database.

After the client finished checking for updates and comments on the journals of each of the participant's friends, it does the same for the participant's own journal. Each new post is reported to the experiment server as a contribution, and comments made by others are reported as activities.

A.3 Stored information

Our server database records all blog posts made by each subject and their LiveJournal friends, as well as all comments on those posts, for the 5 week period of the experiment. Each item includes the LiveJournal ID of the submitter, the post URL, and the time of submission. Lastly, data is kept for each unique participant; their LiveJournal username and password are not kept on the server, but their subject ID is, along with their schedule of display types.

[5] The server is polled on regular intervals to check if the display instructions changed.

[6] The number of recent entries is indeterminate by RSS, but typically includes all articles from the past week.

References

[1] Ambient devices homepage. `http://www.ambientdevices.com`.

[2] J. J. Cadiz, G. Venolia, G. Jancke, and A. Gupta. Designing and deploying an information awareness interface. In *Proceedings of the 2002 ACM conference on Computer supported cooperative work*, pages 314–323. ACM Press, 2002.

[3] Y. Chen, S. Kiesler, J. Konstan, R. Kraut, P. Resnick, J. Reidl, and L. Terveen. Designing on-line communities to enhance participation – bridging theory and practice, 2003. A research proposal to the National Science Foundation, Carnegie Mellon University, Pittsburgh, PA.

[4] D. Grigsby. Instant gratification. `http://www.unpossible.com/blogtraffic.html`.

[5] R. Grinter and M. Eldridge. Wan2tlk?: everyday text messaging. In *Proceedings of the conference on Human factors in computing systems*, pages 441–448. ACM Press, 2003.

[6] R. E. Grinter and L. Palen. Instant messaging in teen life. In *Proceedings of the 2002 ACM conference on Computer supported cooperative work*, pages 21–30. ACM Press, 2002.

[7] Livejournal homepage. `http://www.livejournal.org`.

[8] J. Mankoff, A. K. Dey, G. Hsieh, J. Kientz, S. Lederer, and M. Ames. Heuristic evaluation of ambient displays. In *Proceedings of the conference on Human factors in computing systems*, pages 169–176. ACM Press, 2003.

[9] L. Markus. Towards a critical mass theory of interactive media: Universal access, interdependence, and diffusion. *Communications Research*, 14(2):491–511, 1987.

[10] W. Prinzmetal and A. Wilson. The effect of attention on phenomenal length. *Perception*, pages 193–205, 1997.

[11] Rss 2.0 specification. http://blogs.law.harvard.edu/tech/rss.

[12] Short, William, and Christie. *The social psychology of telecommunications*. Wiley, 1976.

[13] Trackback - a beginner's guide. http://www.movabletype.org/trackback/beginners/.

[14] F. B. Viegas and J. S. Donath. Chat circles. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 9–16. ACM Press, 1999.

[15] J. B. Walther and M. R. Parks. Cues filtered out, cues filtered in: Computer-mediated communication and relationships. In I. M. L. Knapp and J. A. Daly, editors, *Handbook of interpersonal communication*, pages 529–563. Sage, Thousand Oaks, CA, 3rd edition, 2002.

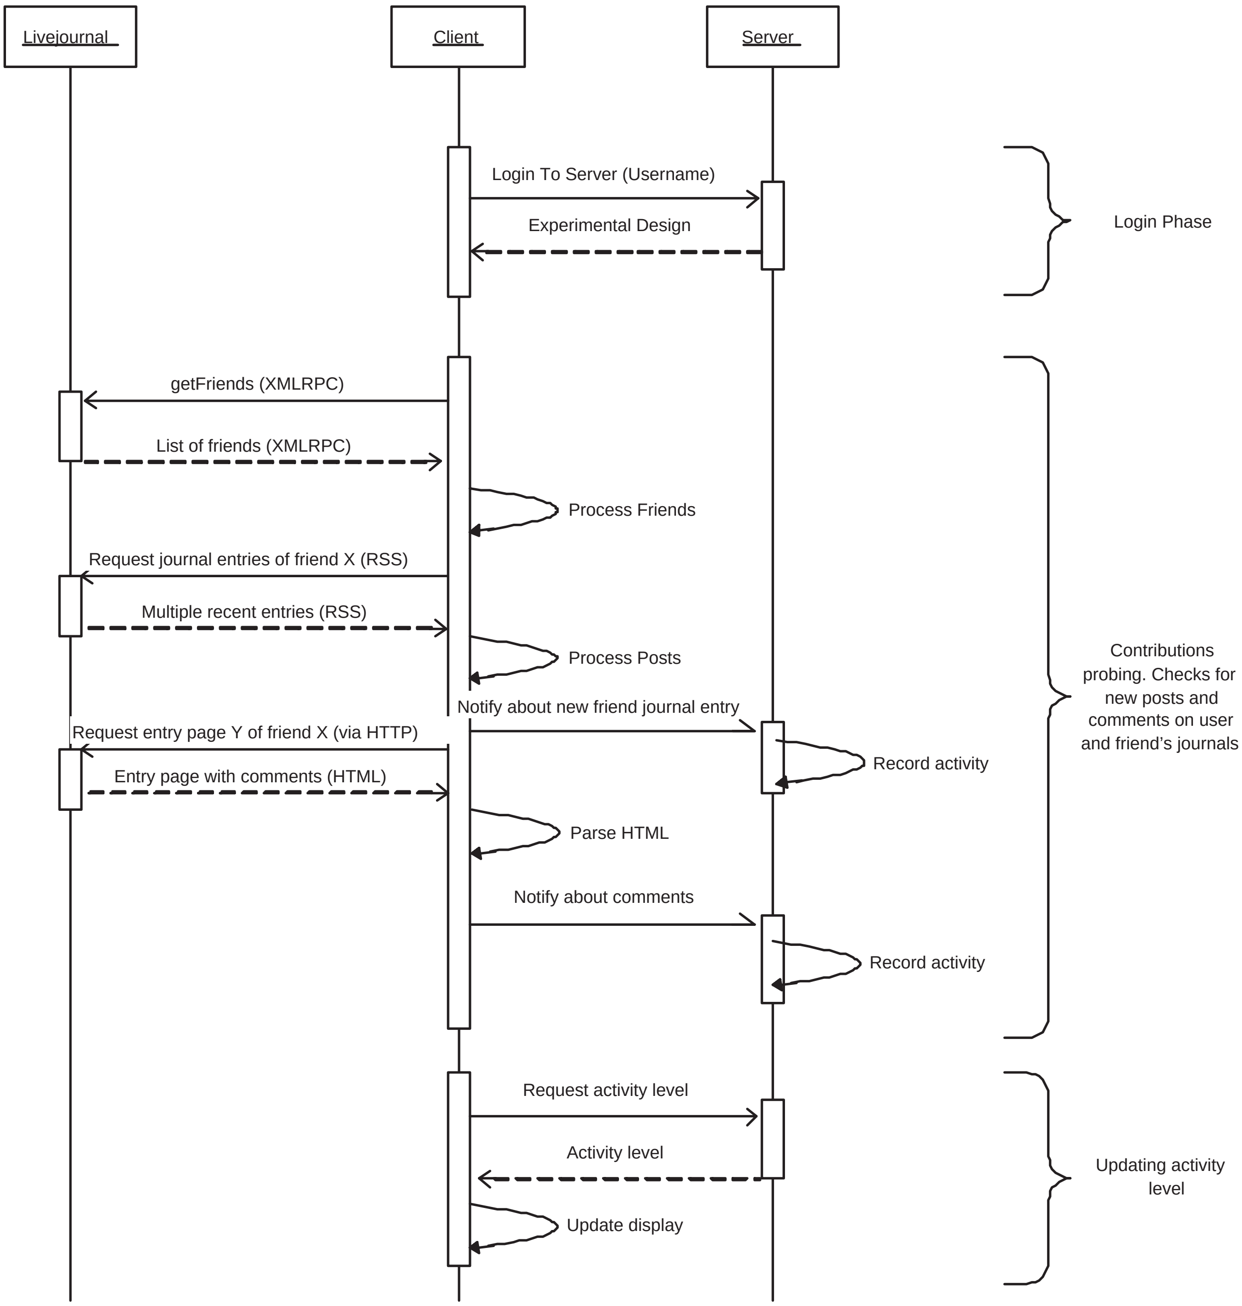

Figure 5. Sequence diagram for our system.