

Talking in Circles: Selective Sharing in Google+

Sanjay Kairam¹, Michael J. Brzozowski², David Huffaker², Ed H. Chi²

¹Computer Science Department
Stanford University
sanjay.kairam@gmail.com

²Google Inc.
Mountain View, CA
{zozo, huffaker, edchi}@google.com

ABSTRACT

Online social networks have become indispensable tools for information sharing, but existing ‘all-or-nothing’ models for sharing have made it difficult for users to target information to specific parts of their networks. In this paper, we study Google+, which enables users to selectively share content with specific ‘Circles’ of people. Through a combination of log analysis with surveys and interviews, we investigate how active users organize and select audiences for shared content. We find that these users frequently engaged in selective sharing, creating circles to manage content across particular life facets, ties of varying strength, and interest-based groups. Motivations to share spanned personal and informational reasons, and users frequently weighed ‘limiting’ factors (e.g. privacy, relevance, and social norms) against the desire to reach a large audience. Our work identifies implications for the design of selective sharing mechanisms in social networks.

Author Keywords

Social media; group-based access controls; social networks

ACM Classification Keywords

H.5.3 Group and Organization Interfaces

INTRODUCTION

Social media, from blogs to forums to online social networks (OSNs), have become a major venue for social interaction and communication. In the United States, over 65% of adult Internet users are currently active on social networking tools [17]. This phenomenon is global; at the start of 2012, Facebook, the largest OSN, claimed over 800 million active users, 75% of whom were outside the US [7].

A key affordance of these tools is interaction with audiences to which users may not have had contact offline; often these interactions accompany disclosures of personal information. Research on social media over the past decade has shown that users are revealing a startling amount of personal information through blogs [14, 25], profiles [10], and status updates [19]. In situations where readership may not be clearly delineated, users often underestimate the size of their potential audience [1], in many cases leading to tension [8, 23], regret [26], or tangible negative consequences in one’s personal or professional life [25].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI’12, May 5–10, 2012, Austin, Texas, USA.

Copyright 2012 ACM 978-1-4503-1015-4/12/05...\$10.00.

If users are aware of the consequences of ‘over-sharing’, then why continue to share so much? This question constitutes what Barnes has called the ‘privacy paradox’ [4]. In this paper, we study individuals who are actively sharing while successfully handling privacy concerns. We observe their behavior to identify strategies for improving sharing precision through *selective sharing*.

We base our analyses on usage data from the field trial of Google+, a social networking service introduced in 2011. Through the incorporation of ‘Circles’, this service is designed around selective sharing and consumption of information. In the following section, we provide some brief theoretical background on needs for selective sharing and prior research on challenges for users sharing using existing social technologies. We then provide a brief introduction to how these needs are addressed in Google+.

BACKGROUND

In applying Altman’s theory of privacy boundary regulation [3] to the online world, Palen and Dourish argue that the disclosure of information is a necessary consequence of participating in a social world [22]. In the offline world, some information (such as presence) may be freely disclosed to the general public, while other information (such as personal secrets) may only be revealed to a privileged few. Using Goffman’s dramaturgical metaphor for interpersonal communication [9], social interaction provides a stage and an audience for whom we tailor self-representations, disclosing what we see fit.

Prior work has revealed the existence of ‘facets’ [8] or ‘modes’ [20], representing the various audiences for whom we must fluidly adapt our performances over the course of our day. Generally, in offline settings, information disclosed is ephemeral and the audience evident, enabling a tight loop between what we share and with whom we share it. As Palen and Dourish identify, moving interactions online means that content shared may persist beyond the scope of a conversation and reach unintended audiences.

Many existing tools for online communication are designed to enable ‘public’ sharing to wide groups. Rather than focusing on the natural ‘facets’ of a user’s life, many social networks, for instance, grant default access to content to a user’s ‘followers’ or ‘friends’, restricting those who fall outside this boundary. While this model may be simple for users to understand, it ignores the ‘natural’ groupings which may exist within these networks [16], making selective control over content difficult and prone to error [24].

Difficulty in disclosing information selectively to various life facets can lead to what Marwick and boyd have called ‘context collapse’ [18]. Studies of OSN use in the workplace reveal that this leads to very real tensions for users of such systems [2, 5, 23]. As a result, users may engage in drastic behaviors to avoid future privacy violations, such as self-censoring or ceasing sharing altogether [24, 26].

Tools such as email fall on the other side of this spectrum, easily enabling ‘targeted’ sharing of specific items to individuals. By essentially requiring access controls to be set for each item shared, such systems provide a high degree of control over access to content. While prior work has shown that particular sub-groups of email contacts may be common across many users [20], the creation and maintenance of persistent groups for access in these systems can often be cumbersome or difficult.

Introducing Google+

Google+, a social networking service launched into field trial by Google in 2011, was designed to address some of these problems. To aid users in selectively disclosing information to common sub-groups of their network, Google+ introduced ‘Circles’, an intuitive mechanism for organizing contacts. Similar to ‘lists’ or ‘groups’ available in networks such as Facebook or Twitter, Circles are user-created groupings of contacts which may be overlapping or hierarchical, allowing users a great deal of flexibility in organizing their networks. Each time a user shares a piece of content, he or she makes a contextual decision about the audience with whom to share it; content can be shared publicly or shared selectively to one or more circles. Users can also share to ‘friends of friends’ by selecting an ‘Extended Circles’ option, though this feature is not analyzed in detail in the present work.

In addition, specific individuals can be referenced through inclusion of a ‘+Mention’. Mentioning users individually in this way can be used to target content directly at those users. Use of a ‘+Mention’ when sharing to a wider audience can be a means of notifying the mentioned user and directing their attention to the post, similar to a ‘@mention’ on Twitter. In addition to providing a simple mechanism for selective sharing, Circles also support selective consumption. Users can filter their incoming stream in order to see content being shared only by members of a specific Circle. While sharing with specific groups is possible in other OSNs, we study Google+ here as it represents the first widely-deployed system with such a ‘selective sharing’ mechanism built in from the start.

OVERVIEW OF THE PAPER

In the following section, we examine large-scale usage data collected during the field trial of Google+. We first look at aggregate sharing patterns to understand the extent to which users are actively engaging in public sharing, targeted sharing to individuals, and selective sharing using Circles. We then look at names frequently given by users to custom Circles with the aim of investigating the various ways that

users separate their networks for the purposes of selective sharing and consumption.

In the subsequent section, we report findings from a qualitative study of users sharing comfortably on Google+. We analyze the behavior of these ‘expert’ users through data collected via surveys and interviews in order to understand how they use the product to successfully tailor self-presentations for various parts of their network through selective information sharing. We focus specifically on how they organize their contacts into Circles, what factors motivate them to share, and contextual considerations for choosing an audience for a particular piece of content.

QUANTITATIVE LOG ANALYSIS

In this section, we conduct analysis using usage logs from the field trial of Google+ to examine overall user sharing patterns and characterize how users organize their networks based on names assigned to Circles they have created.

High-Level Sharing Patterns

To understand the extent to which users were engaging in selective sharing with Circles, as opposed to public or targeted sharing, we first examined logs of per-user sharing choices aggregated over the week of July 20-26, 2011.

Data

User IDs were hashed and sampled randomly from the population of users who had shared at least once during the data collection period, providing a representative sample of 100K active users. To preserve user privacy, we did not conduct any analysis concerning specific Circles, but rather restricted our analysis to public (posts shared using the ‘Public’ setting), selective (posts share to one or more Circles), and targeted (posts shared using a ‘+Mention’) sharing. In order to compare our results to those found in the later qualitative study, we focused on users sharing via the desktop web application, resulting in a data set of 64,005 users who shared content during this week.

As shown below in Figure 1, the number of items shared weekly per user roughly follows a power-law distribution. Using Clauset et al.’s approach [6], we estimate the parameter α as 2.578, meaning that the distribution matches those of a variety of man-made and naturally occurring phenomena, giving us confidence that our data was representative of the wider population

Results

In Table 1 below, we report the extent to which users engaged in public sharing, selective sharing (Circles), and targeted sharing (+Mentions). In the first column, we see that about three quarters (74.8%) of users who share during the period engaged in selective sharing using Circles at least once. Portions of users engaging in public (33.9%) and targeted (10.3%) sharing were smaller, but still significant, indicating that at least some users were employing multiple sharing options. The second column displays percentages of all items shared using these options; here, we see that 67.6% of all items were shared using Circles and 33.8% of

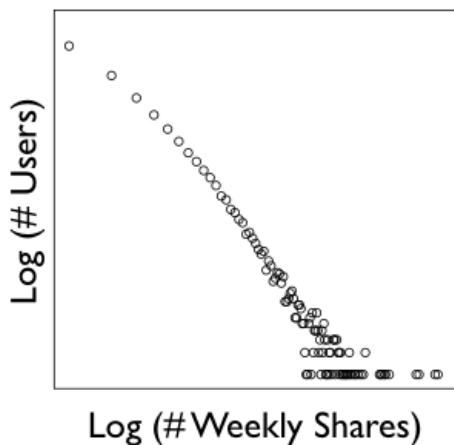


Figure 1. Log-log frequency plot of # weekly shares per user.
 Note. Axes obscured in order to not reveal total usage.

items were shared publicly (these percentages add up to more than 100% since some users combined sharing options for posts). We note that sharing a post both publicly and to Circles functioned the same as sharing the post publicly alone; this pattern in the data signaled that a small portion of users may have been mistaken about the function of Circles and were expressing additional needs for selective sharing aside from access control.

Patterns in Custom Circle Names

Users joining Google+ are initially presented with a set of four default Circles: *Family*, *Friends*, *Acquaintances*, and *Following* (intended primarily for consumption). Users are then able to rename these Circles or create additional Circles to serve their needs. In order to identify factors governing how users organized their networks, we examined the names frequently used for custom Circles. Our intuition here was that frequently occurring patterns in these names might reveal widespread needs that users had for selective sharing and consumption of information.

Data

Our data was drawn from a snapshot taken on Sept. 6, 2011 of the names of all non-empty, user-created Circles (thus excluding all Circles with names matching the defaults). We focused our analysis on the top 1,000 most common names in order to identify common needs via frequently occurring patterns; this method also avoided the possible privacy risk of revealing personal or idiosyncratic names. Capitalization and punctuation differences were ignored when counting names, but no stemming was performed. The frequency of occurrence of Circle names very closely approximated a power-law distribution with $\alpha = 1.99$ – similar to observed distributions for English word

Sharing	% Users	% Items
Public	33.9%	33.8%
Selective	74.8%	67.6%
Targeted	10.3%	5.8%

Table 1. Frequency of types of sharing conducted over a 1-week period in July 2011. The first column reports percentages of users choosing each option at least once during the week; the second column reports percentages of total items shared by the same set of users.

frequencies, such as the occurrence of unique words in *Moby Dick* [6]. The names in this data set covered 34.5% of all non-empty Circles in users’ accounts at snapshot time.

Coding

In coding circle names, we adopted the following conservative approach, drawing on Hsieh’s definition of ‘conventional content analysis’ [13]. We first looked at patterns in the 50 most frequently occurring Circle names to develop four labels. One author then coded the entire list (the top 1,000) using these labels, coding all uncategorized Circles as ‘other’. This method ensured categorization of the most common names, and thus wider coverage overall.

Non-English names were translated using Google Translate; ambiguous translations were cross-referenced against occurrences on pages found through Google search. While this approach may seem overly reductive, the brevity of Circle names made inferring user intent difficult in many cases. As Circle names are created for private, rather than public use, names often contained abbreviations, idiolectic slang, or terms with multiple or ambiguous meanings. A second author independently categorized the first 100 names to assess inter-rater reliability; Cohen’s $\kappa = 0.91$ indicated high agreement on label assignments.

Labels developed from the data fell into two higher categories: *life facets* and *tie strength*. The first category, *life facets*, contained two labels, ‘work’ and ‘school’. Circle names labeled ‘work’ contained synonyms or translations for words such as *work* or *professional* or denoted business relationships (e.g. *employees*, *customers*). Names labeled as ‘school’ contained synonyms or translations for words such as *school* or *college* or academic relationships (e.g. *professor*, *students*). Specific company and school names were excluded, as it was impossible to resolve ambiguity over users’ relationships with those names (e.g. Users may create a *Google* Circle for sharing content about Google products rather than for sharing with Google employees).

Rank	Work	School	Strong	Weak	Other
10	74.1%	13.8%	7.2%	0.0%	4.9%
100	47.5%	16.0%	9.0%	3.5%	24.0%
1000	31.4%	12.6%	8.8%	4.0%	43.1%
All	10.8%	4.3%	3.0%	1.4%	-----

Table 2. Categories for the top 10, 100, and 1000 most common names for user-defined Circles. Percentages reported in bottom row represent strict lower bounds on category representation across the entire data set.

The second category, *tie strength*, included ‘strong ties’ and ‘weak ties’. Circles were deemed ‘strong ties’ if their names applied modifiers such as *close*, *best*, or *inner* to common relations such as *friends* or *family*. We also included specific relationships such as *girlfriend* or *husband* and common slang for close ties (e.g. *BFFs*). Circles labeled as ‘weak ties’ included words such as *other*, *random*, or *extended* or words otherwise implying limited knowledge about Circle members (e.g. *unknown*, *don’t know*, or *WTF*).

Results

In Table 2, we report results from labeling the top 10, 100, and 1000 most commonly used Circle names. Each row presents a different view of the representation of the various labels in the data. The first row indicates that the ‘work’ label covers 74.1% of all Circles with names matching the top 10 most frequently occurring user-defined names. We note here that ‘work’, ‘school’, and ‘strong ties’ labels cover almost all of the Circles in this group (the 10th most common name, categorized as Other, was the empty string).

Expanding our view to the top 100 names, we see that Circles pertaining to ‘school’ and tie strength occupy a greater portion of this data set, while an increasing portion of the names become difficult to categorize. Given our knowledge that a large portion (34.5%) of all custom circles are covered by the top 1000 most commonly used names, we can estimate that Circles pertaining to “work” and “school” represent at least 15.1% of all Circles, indicating that these life facets are commonly important dimensions for selective sharing and consumption of content. Circles pertaining to tie strength occupy at least 4.4% of all custom Circles, an amount which is smaller but still substantial.

The prevalence of Circles focused on professional life is consistent with prior hypotheses about specific sharing needs relating to workplace OSN use [2, 5, 23]. To examine if this pattern was consistent across cultures, we looked at the most common Circle names in 10 languages popular amongst users of Google+, finding that the most common name in each translated into either ‘work’ or ‘colleague’.

Though the coding approach used was fairly conservative, it provides the advantage of being rather unambiguous while covering a large percentage of names across the data set. Still, we also discovered tremendous diversity in the

‘other’ common Circle names which invite further study. Observing the top 10 most common Circle names categorized as ‘Other’ (*church*, *music*, *google*, *tech*, *twitter*, *photographers*, *celebrities*, *news*, and *relatives*) hints that Circles based on topical interest may be common as well.

Discussion

Based on log analysis of sharing patterns and contact organization, we make two observations about needs for selective information disclosure and consumption. First, we find that users are actively engaging in a mix of public, targeted, and selective sharing; at least a small portion of users are utilizing different options for various pieces of content. In addition, we find that frequently occurring names for user-defined Circles reflect a need to manage content according to *life facets* (specifically professional life), *tie strength*, and possibly, *topical interest groups*.

Our analysis focused on data which both ensured broad coverage of Google+ usage and avoided product confidentiality or privacy concerns. While additional statistics about Circles and sharing might augment these findings in the future, we believe that the patterns discussed still provide ample insight into selective-sharing practices. Keeping these high-level findings in mind, we move on to data collected through qualitative study of active Google+ users, through which we aim to gain a more detailed understanding of the contexts in which these patterns occur.

QUALITATIVE USER STUDY

In this section, we describe the results of a survey and interview study conducted with active Google+ users.

Method

As our aim was to understand why users were actively sharing with different groups, we chose to focus our analysis on active users who had familiarity with the product and a history of shared content. Participants in all stages of this study were recruited from a pool of approximately 2500 users compiled via snowball sampling through a survey posted on Google+ which collected demographic and general Google+ usage information. A subset of these users were chosen to receive our follow-up survey designed to probe user motivations for creating Circles and sharing content.

A sub-group of users who completed the follow-up survey were asked to participate in semi-structured interviews aimed at understanding how users group their contacts into Circles. These interviews also included a critical incident study of posts that the user had recently shared, aimed at providing greater insight into decisions about whether to share content and about the audience with which to share it.

Survey Study

Participants. From the initial pool described above, 300 users were invited to participate in our follow-up survey via the following criteria: each had at least 4 Circles containing a total of at least 30 members, had shared at least 5 items, and had been visiting Google+ more than once a week.

168 (56.3%) users participated in our follow-up survey on Circle management and sharing preferences. In a few cases, follow-up responses could not be connected to initial survey responses; we thus report demographic information for the 162 surveys for which this was possible. Respondents were primarily (74.2%) male and had a wide range of ages ($\mu = 35.0$, $\sigma = 8.3$). Almost all participants (98.8%) used other social media tools at least weekly, with YouTube (89.5%), Facebook (82.7%), Twitter (72.2%), and LinkedIn (45.1%) being most popular. While occupations were skewed towards technology-focused positions, we observed a wide range of professions, including several lawyers, IT consultants, artists, analysts, and software developers.

Measures. Beyond the demographic and usage questions collected in the initial survey, this survey probed further about the frequency of use of various sharing settings and shared content of various types. A portion of the study was dedicated to a critical incident report about the item most recently shared in this stream; participants uncomfortable with discussing their most recent post had the option to discuss another item, with 14 of our participants doing so.

Interview Study

Participants. From the set of users who completed the follow-up survey, 12 were chosen to participate in a 1-hour remote interview about their Google+ usage. Interview participants were selected such that they lived in the United States in order to avoid language fluency issues and balanced by gender, but otherwise matched the demographics of the overall population initially surveyed.

Measures. Interviewees first provided a walk-through of how they use the product (10-15 minutes), and then answered questions about their management of contacts, including in-depth questions about their most heavily used Circles (10-15 minutes). The interview also included a critical incident portion similar to that administered in the surveys, in which participants discussed 3 or 4 posts recently shared (15-20 minutes). The remainder of the interview was used to collect general product feedback. Additional details are given below about questions central to the study and the methods used to analyze responses.

Sharing and Audience Choices

In characterizing reasons why users shared content and factors considered when choosing an audience with which to share, we utilized a grounded approach to code and categorize responses to the following two questions about the users' recently shared posts:

- Please describe this post and why you shared it.
- Please describe the audience that you chose and why you chose to share this content with them.

In coding responses to these two questions, we adopted an approach similar to that used in Naaman et al. [19]. For each question, two authors worked together to develop a coding scheme using 50 responses sampled from the data; each author then independently coded half of the remaining responses in order to identify additional labels which might apply. Participants often gave rich responses touching on multiple reasons; thus, labels were not mutually exclusive.

After agreeing on a complete set of labels, each author then went through and re-coded the full set of responses. Finally, these two authors discussed each coding discrepancy in order to reconcile it. These labels were then used to generate the higher-level categories discussed below. Because complete agreement was reached between coders, as in Naaman et al. [19], we do not report inter-rater reliability measures; however, secondary coding by the other two authors on 10% of the data revealed at least moderate agreement for each of the high-level categories.

It should be noted that the goal in each of these cases was to identify patterns in the self-reported 'reasons' why users made particular decisions. This goal is distinct from and serves to extend prior work in describing message content in similar types of social awareness streams [12, 15, 19].

Management of Contacts

During the interviews, users were asked to identify the Circles with which they 'interacted' most frequently, where the nature of interaction (e.g. sharing, consumption, conversation) was left open-ended. For each Circle identified, participants responded to two sets of statements describing the Circle, using one of two 5-point Likert scales. The first scale ranged in level of agreement from *Strongly Disagree* (1) to *Strongly Agree* (5), while the second scale ranged over the proportion of Circle members to which the statement applied, from *None* (1) to *All* (5).

Questions were informed by prior study of Facebook friend grouping [16], and were designed to test the extent to which Circles represented distinct social cliques, collections of weak or strong ties, common organizations, or connections to a common geographical location or episode in the user's life. In addition, in light of findings from our analysis of Circle names, we also included questions to probe whether Circles connected users on the basis of topical interest.

Results

We start this sub-section by summarizing general patterns reported by survey respondents about their use of Google+, focusing specifically on types of content typically shared and the types of sharing (public, selective, and targeted) usually chosen. Our intent here is not to characterize the general population, but to understand the behaviors and attitudes which aid early-adopting users in using the product successfully for sharing and filtering content.

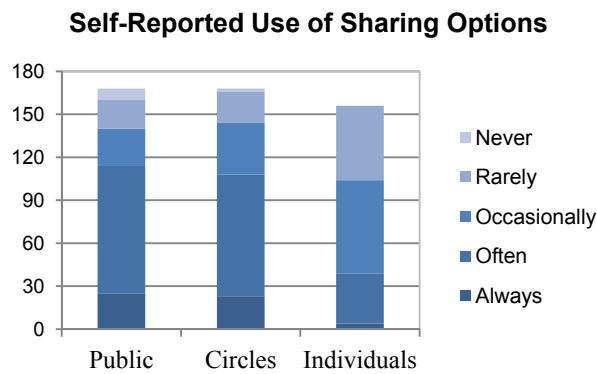


Figure 2. Frequency of use of various sharing options based on participant self-report.

The remainder of this sub-section is organized according to our three primary areas of interest: what reasons spur users to share content, what factors they consider when choosing with whom to share, and how they manage their contacts generally for the purposes of selective sharing and filtering.

General Patterns of Use

Survey participants were asked to report on how frequently they shared various types of content. Sharing of URLs and photos was common: 134 of our 168 respondents (79.8%) indicated that they shared URLs ‘often’ or ‘always, and 138 (82.1%) reported sharing photos at least ‘occasionally’. Sharing of videos and locations was less prevalent, with 76 respondents (45.2%) sharing videos and 114 (67.9%) sharing locations ‘rarely’ or ‘never’.

Participants also reported on how often they utilized various sharing options. Figure 2 summarizes participants’ self-reports about how often they engaged in public, selective (Circles) and targeted (individuals) sharing. 144 of our respondents (85.7%) reported sharing to Public at least occasionally, hinting at the possibility that these active users were more comfortable sharing publicly than the general population (see Table 1). Patterns for public and selective sharing were generally similar, but participants reported sharing to individuals less frequently overall.

Sharing: Why do users share content?

Of the 168 users who participated in the survey, 166 submitted responses which could be coded to the query ‘Please describe this post and why you shared it.’ Results are organized below according to categories developed based on responses to the survey. As mentioned previously, these categories were not designed to be mutually exclusive. Quotes from surveys and interviews provide greater detail for some observed phenomena; interview subjects are designated using a unique identifier [P*], while survey quotes are unmarked. Typographical errors in user quotes are corrected using {brackets} for clarity.

Inherent Value of Content. Based on survey responses, the reasons for sharing most frequently referenced by users (98/166, or 59.0%) referenced the inherent value of the content itself. Half of these cases (49) included a stated

assumption that others would value the content as well, indicating that these users are actively reflecting on the consumption needs of others. Content was most frequently deemed valuable because it was topically ‘interesting’ or generally ‘exciting’ or ‘cool’ in some way (48/166, or 28.9% of responses). One user described why he shared a post containing a book recommendation:

I shared this content because, as listed in the post, I have been reading this particular book and I love it! I am very much into the subject matter and wanted to share this information so that those that are following me would know that it was a good book (according to me) and to get into a discussion started about the subject matter.

Another common reason for deeming content as valuable was because it was ‘funny’, ‘amusing’, or ‘silly’ (24/166, or 14.5%). Content in this category included quotations, personal stories, and humorous pictures and videos. One user reported sharing a quote because it was “an amusing and accurate generalization of the activities that make up our lives”, while another who shared a humorous picture simply said “I re-shared it because I thought it was funny as hell!” [P12]. 16/166 users (9.6%) indicated in their responses that they shared content because it was ‘informative’ or ‘useful’. Half of these posts were about Google+ itself, perhaps an artifact of the novelty of the service, though prior work [12] has shown that this type of meta-commentary is evident in other social systems as well.

Finally, a few respondents reported sharing content because it was positive, though not explicitly useful or funny (8/166, or 4.8%). These responses described content that was ‘heartwarming’ or ‘enjoyable’. One user, who had shared a YouTube video of a man singing old TV theme songs indicated that it was “very nostalgic” and thought that “other people...might enjoy it in the same way.” An interview participant who shared a story about an elderly knitter had this to say: “It just struck me as such a sweet story...it wasn’t just something that would appeal to knitters, it would appeal to people in general” [P9]. Motivation to share in these cases generally seemed connected to a desire to elicit positive emotion in others.

Sharing about Self. The other theme which emerged from users’ responses was a desire to share about oneself. Of the 43 (25.9%) responses in this category, only 4 explicitly or implicitly mentioned considerations about whether others would find the content useful, suggesting that the underlying motivations in these cases differed fundamentally from those in the prior category. Within this group, the most frequently referenced reason for sharing was a desire to share one’s personal experiences with others (19/166, or 11.4%). These ranged widely from sharing brief updates (“just updating my friends on what I was up to”) to long posts with strong emotional content. One user, who had posted about a long struggle with weight management, revealed a desire to share her story with her audience:

I am feeling very frustrated about my (lack of) weight loss, and I wanted to explain a bit of my background to my followers, and ask the collective to help me a bit when my self-discipline is lacking, {especially} about skipping Weight watchers meetings.

A number of survey respondents (16/166, or 9.6%) indicated a desire to share personal opinions or commentary; in the majority (10) of these cases, users were commenting on technology products, which again may have been indicative of the early-adopting population studied.

Discourse. An additional theme bridging the two primary categories listed above was a desire to start or participate in a discussion (28/166, or 16.9%). This frequently focused on soliciting help or information from others. These ranged from immediate information needs (“I need to be helped out by getting an information about printing several documents at once.”) to general topical inquiries (“I have been checking to see when the app Instagram would be coming out for Android and could not find a definitive answer.”).

In other cases (13/166, or 7.8%), participants were motivated to start or participate in a discussion on a topic of interest. In many of these cases, posts focused on general interest topics (e.g. news, current events). One participant living in Virginia, who was interviewed the day after the 2011 Virginia Earthquake, described how she shared in order to solicit additional information about the event:

When we had the earthquake...I geo-tagged it hoping that people around me could see it as well and maybe...contribute to it as far as whether they felt it or not. [P5]

In other cases, discussions were specifically relevant to a sub-group of users. One respondent gave the following context for a post about a knitting project:

There is a knit-along going on specifically for g+ so I was sharing my project status with the other knitters. Many of the other knitters are sharing project photo, so I wanted to add mine too.

Evangelism. Another less-common, but interesting, category of reasons for sharing was a desire to raise awareness on behalf of another person or cause. The majority of the (11/166, or 6.6%) survey responses which received this label included posts intended to spread information about specific events or work being conducted by others so that recipients could appreciate, attend, or contribute. Others sought to spread awareness of general topics, such as this photographer who shared a post about undue harassment that photographers were receiving from police officers when photographing government buildings:

My hope was to raise awareness of the issue, and at some point, I invite people to duplicate the experience in their town. A little bit of activism attached to this one. [P12]

Sharing Options for Recent Posts

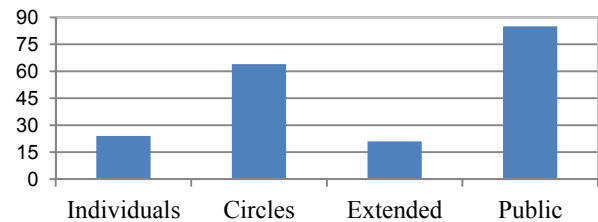


Figure 3. Sharing options used for the ‘most recent’ post in users’ streams. As participants had the option to select an alternate post (14 users indicated doing so), reported results may be biased towards publicly shared content.

Selecting an Audience: How do users choose?

Of these respondents, 165 also reported the sharing options they chose and provided a response which could be coded to the query ‘Please describe the audience you chose and why you chose to share this content with them.’ Sharing options chosen for these posts are displayed in Figure 3 above. Here, we see that the majority of posts discussed in this study were shared publicly (85/165, or 51.5%). This may be due in part to the fact that the population studied was fairly comfortable with public sharing (see Figure 2), and also to the fact that private posts may have been more likely to contain sensitive information that users would not have chosen to discuss as part of the study. In addition, a number of users chose multiple options when sharing, including several who shared both to Public and Circles.

Again, results are organized below according to the categories developed based on survey responses; quotes from interviews are designated using the same unique participant identifiers as in the prior section.

Privacy. Concerns about privacy were common for users considering whether to limit the audience for their posts (36/165, or 21.8%). Most users (24) with responses in this category focused on risks to their own privacy when choosing an audience. These users tended to subscribe to one of two overall ‘policies’ regarding shared content.

Followers of the first policy, which we call *public-by-default*, tended to share content, or at least most types of content, widely unless they had a reason not to do so (13/165, or 7.9%). Most explanations included general statements like “I prefer to share as much publicly as I can” or contextual statements such as “There was nothing specific or private about it so I shared it publicly.” Followers of the second policy, which we call *private-by-default*, tended to have overarching reasons for placing restrictions on shared content (5/165 = 3.0%). Often these restrictions arose from concerns about specific individuals or groups having access to content shared. One user explained her decisions to limit sharing to Circles, saying:

I almost always limit my posts for privacy reasons. There are people from my past, including my ex-husband, who don’t need to know what I’m up to.

Another user demonstrated the use of different policies depending on time of day, restricting posts specifically during work hours due to concerns over social media use:

I have people from my work who are {present} in 'Extended Circles', who I don't want to see my posts in 'work hours'. So {usually} I prefer 'Extended' but just in non-working time."

Fewer users (6/165 = 3.6%) expressed concerns based on the possibility of privacy risks to others. A specific source of concern here was the sharing of photos, illustrated by one interview participant who appeared to be very comfortable sharing personal information, but carefully considered options when sharing photos of her daughter:

I'm not overly concerned about pictures of my kids being on the Internet, because I have an active online life...but on a post like this where I'm just sharing pictures of my daughter, it's probably one that I would not share publicly. [P7]

Relevance. Users also frequently considered whether content shared would be relevant to particular parts of their audience (38/165, or 23.0%). Often, these users targeted content to those who might be interested, in some case utilizing Circles created for this purpose. One interview participant described how she frequently used a topical Circle to selective share nutrition and health content:

The Circle I use the third-most is a subject-based Circle, and it's about paleo-primal health and exercise stuff and the science behind it. I shared this with that Circle and a couple of individuals who aren't in that Circle but are interested in this stuff. [P3]

In other cases, users chose an audience by excluding those who might not find the shared content relevant, referencing a desire to avoid 'spamming' or 'polluting' others' streams. One user who shared a piece of technology news described only sharing it to people from his school:

The subject of the post is related to our studies. I didn't share it with anyone else because I didn't want to {pollute} their stream with things that probably won't {interest} them.

A particular sub-group of these users (9) mentioned limiting their audience because the content required shared context to understand. These posts ranged from inside jokes ("I limited it just to her because it was {the} kind of inside joke that only she would understand") to shared projects ("Since the message was related to our tour, {I} had shared it individually") to family matters ("This was a family {event} so I kept it to my friends and family").

Social Norms. A smaller number of participants referenced concerns about whether content might be appropriate for a particular, or unknown, audience (13/165, or 7.9%). One user who shared a humorous animal picture explained that

"since it wasn't too controversial in any way, there was no need not to share it publicly." Another user discussing a potentially risqué joke shared, hinting at a general policy for assessing whether content was safe to share publicly, explaining "It's the kind of joke I wouldn't be ashamed to show my mother so I can share it with everybody..."

Distribution. In addition to considerations about why to limit one's audience, a large number of respondents specifically referenced reasons for maximizing the potential audience for content (71/165, or 43.0%). In the majority of these cases (49/165, or 29.7%), users were motivated to share content widely in order to make it more accessible to others who might want to consume it. One interview participant who had shared a photo while engaging in an outdoor activity did so publicly in the hopes of enabling serendipitous interactions with others:

I really like finding nice pictures in my stream...[I] put this out so more people can see it...Somebody could be on their phone nearby and say 'Hey, that's really cool!' I think it'd be nice if there were enough people doing the same thing that when I shared that picture, someone else on their phone nearby could be doing a similar thing and share their picture from across the lake. If I don't share pictures like that, then how's that experience going to happen otherwise? [P9]

In 14/166 (8.5%) cases, responses indicated a desire to publicize content in order to expand one's own audience, specifically with the aim of encouraging others to add the user to their Circles. Another reason for maximizing one's audience, related to the social search behavior discussed in the prior section, was to 'cast a wide net' for potential help or assistance (11/165, 6.7%). The user who posted about weight management explained "I decided to spread my net as wide as I could {without} taking it public, You {never} know when you will find 'I've been there' supporters," hinting at the balance that users must strike when weighing the benefits of various sharing options.

Circles: Organizing contacts for sharing and consumption

The majority of interview participants discussed three Circles; two discussed only two due to time constraints. This resulted in a data set of 34 Circles. Aided by a *k*-means clustering algorithm using data from participants' responses to the Likert-scale questions discussed earlier, we manually identified 4 major types of Circles discussed by users:

Inner Circle. The first cluster contained Circles which represented small, tightly-knit groups composed of members who were strongly tied to the user. These Circles had names such as *Family*, *Inner Circle*, *Tell-All*, etc. When discussing these Circles, participants tended to describe them as groups of users to whom they could say anything: "These people pretty much know the 'unfiltered' me" [P1]; "I don't have to be discreet in any way with them" [P2].

Structured Groups. Circles in the second cluster often represent structured groups. Members of Circles in this cluster often shared common interests, belonged to common organizations (either offline or online), and were likely to know one another. Examples included *Church/Ministry* (“Pretty much anyone who has anything to do with my church or Christian ministry” [P11]), *Google Big Wigs* (“Everyone that I’ve heard that’s on G+ that works for Google” [P12]), and *eGRC* (an industry-specific Circle: “Employees, clients, former clients...” [P3])

Interest Group. The third cluster contained Circles which grouped members with common interests but without other links such as a common location or organization. Examples of these included *Knitters* (“I know them from Ravelry...I know some of them by name, most of them by screen name” [P7]) and *Photogs* (“I pretty much throw any photographer in there” [P12]).

Catch-All. The fourth type of Circle appeared to collect users who were unrelated or loosely-related. The intent behind creating these Circles was clear from some of their names: *Possible Friends*, *Acquaintances*, *Everything Else*.

DISCUSSION

In this paper, we analyzed data from the Google+ field trial to examine how expert users utilize the system to engage in *selective sharing*. We believe that insights generated from this work will aid in understanding needs of users of OSNs with respect to the following:

Contact Management. Analysis of common names for user-created Circles reflects a need to manage content separately for different **life facets** (with a focus on ‘professional’ life) and according to **tie strength**. These findings were strengthened by further qualitative study with expert users; specifically, we found them managing contacts by tie strength, creating ‘Inner Circles’ for sharing personal information with strong ties and ‘Catch-All’ Circles for weak ties (with whom they might be more cautious). This finding suggests that metaphors from the sharing of physical space in interactions [11] may translate online. Log analysis and interview data suggested that **topical interest** may also be an important grouping factor, and future study may reveal the extent to which this is true.

Future systems could explicitly aid in Circle management by suggesting members and proactively help in creating Circles for managing different aspects of life. We note that the heuristics for grouping which emerged from the needs of active users differed in some ways from those identified in prior simulated studies of friend grouping [16], revealing that groupings may evolve as a result of longitudinal engagement with the system. Users in our study often referenced future plans to modify Circles based on sharing or consumption challenges experienced during use, suggesting that the complexity of these real-life boundaries might make them difficult to infer automatically.

Reasons for Sharing. Based on expert users’ responses, we have identified two primary ‘reasons’ for sharing. The first is the **inherent value of content** (e.g. it may be useful, funny, interesting, or otherwise pleasant to others). The other is a **desire to share about oneself** (e.g. expressing an inherent desire to relate one’s personal experiences and opinions to others). Less prevalent reasons included engaging in **discourse** and **evangelism** on the part of others. We observe a large overlap with categories (‘Informers’ vs. ‘Me-formers’) identified by Naaman et al. based on receiver characterization of message content on Twitter [19]. While it remains to be seen whether sender motivations match receiver perceptions, our findings coupled with those from prior work suggest that these reasons may generalize across public and semi-public OSNs.

System designers may wish to explicitly create feedback mechanisms to help users manage self-presentations. While companies (e.g. Klout¹) are pursuing explicit reputation scores, there exists room for innovation in self-analytics, such as visualizing for users types of content shared (e.g. proportion of messages about oneself) or responses and re-sharing of contributed information. These analytic capabilities might allow users to reflect on their sharing patterns and adjust them to suit their own needs, as well as those of their expected audiences.

Selective Sharing. We find that users engaging in selective sharing weigh four primary factors when choosing audiences for content. Three of these categories – **privacy**, **relevance**, and **social norms** – represent reasons why a user might choose to limit access, but we also found users expressing reasons for increasing content **distribution**.

These findings suggest revising our understanding of how users consider privacy generally in OSNs. Rather than framing decisions about privacy as ‘boundaries’ based on who is ‘allowed’ to view content, users may instead be balancing reasons to limit or to distribute content. Without proper controls for selective sharing, this can lead to previously under-explained behaviors such as over-sharing (e.g. ‘context collapse’ [18]) or fear of sharing (e.g. self-censoring [24, 26]). The design of proper selective sharing controls, however, allow users to balance sender and receiver needs and adapt these controls to different types of content. While similar controls may be available on other networks (e.g. ‘Lists’ on Facebook), making them central to the sharing model may be key to widespread user adoption.

The design of selective sharing controls offers a great opportunity to help users manage their self-presentations to multiple audiences. One might desire to share humorous content with family or friends while promoting a more professional image for clients and teammates. In the context of prior work on profile-sharing in OSNs [10], we observe that users may also benefit from tailoring of profiles to

¹ <http://klout.com>

show different information to each Circle. While our study focused on identifying factors considered when choosing an audience, future work will prove how users balance these considerations to ultimately make these decisions.

Models combining life facets, tie strength, and topical interest will be key to refining audience management, through Circles in Google+ or similar feature in other OSNs. A combined model might use network structure to identify strong ties or post content to estimate topical interest. In the competitive attention economy of the social web, helping messages reach appropriate audiences is an important design challenge. In the short-term, at least, it would seem that this challenge may be best met through a combination of user control and system features.

CONCLUSION

In this paper, we study Google+ to provide a first empirical study of behavior in a network designed to facilitate selective sharing. As this study focused on active users of a relatively new product, an obvious limitation of this work is that it may not generalize to the wider population and users of other networks. However, much of our findings both match with and help to explain prior findings on sharing across a variety of OSNs. As social networks move from a 'boundary' model to a more nuanced selective-sharing model, our study here points to the complexity of user behavior in this brave new world. We hope the results will inform other designers of social sharing mechanisms and ultimately help users manage their information better.

ACKNOWLEDGMENTS

We would like to thank the entire Google+ team, particularly Makoto Uchida and Alex Fabrikant for their support and help with data collection.

REFERENCES

- Acquisti, A. & Gross, R. Imagined communities: Awareness, information sharing, and privacy on the Facebook. *Lect. Notes of Comp. Sci.* 4258: 36-58. 2006.
- Adams, P. "The Real-Life Social Network". <http://www.slideshare.net/padday/the-real-life-social-network-v2>. 2010.
- Altman, I. *The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding*. Brooks/Cole Pub. Co, Monterey, CA, USA. 1975.
- Barnes, S.B. A privacy paradox: Social networking in the United States. *First Monday*, 11 (9). 2006.
- Binder, J., Howes, A., & Sutcliffe, A. The problem of conflicting social spheres: effects of network structure on experienced tension in social network sites. *Proc. CHI '09*, 965-974. 2009.
- Clauset, A., Shalizi, C.R., & Newman, M.E.J. Power-law distributions in empirical data. *SIAM Review* 51, 661-703. 2009.
- Facebook, accessed 1/10/2012. <http://www.facebook.com/press/info.php?statistics>.
- Farnham, S.D. & Churchill, E.F. Faceted identity, faceted lives: Social and technical issues with being yourself online. *Proc. CSCW '11*, 359-368. 2011.
- Goffman, E. *The Presentation of Self in Everyday Life*. Garden City, New York: Doubleday. 1959.
- Gross, R. & Acquisti, A. Information revelation and privacy in online social networks (The Facebook case). *Proc. ACM WPES '05*, 71-80. 2005.
- Hall, E.T. *The Hidden Dimension*. New York: Doubleday. 1966.
- Honeycutt, C. & Herring, S.C. Beyond microblogging: conversation and collaboration via Twitter. *Proc. HICSS '09*. 2009.
- Hsieh, H.F. & Shannon, S.E. Three approaches to qualitative content analysis. *Qual. Health. Res.*, 15 (9), 1277-1288. 2005.
- Huffaker, D.A. & Calvert, S.L. Gender, identity, and language use in teenage blogs. *JCMC*, 10 (2). 2005.
- Java, A., Song, X., Finin, T., & Tseng, B. Why we Twitter: understanding microblogging usage and communities. *Proc. WebKDD/SNA-KDD '07*. 2007.
- Jones, S. & O'Neill, E. Feasibility of structural network clustering for group-based privacy control in social networks. *Proc. SOUPS '10*, 1-8. 2010.
- Madden, M. & Zickuhr, K. Pew Internet, September 2011. <http://www.pewinternet.org/Reports/2011/Social-Networking-Sites.aspx>.
- Marwick, A. & boyd, d. I Tweet honestly, I Tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society* 13, 96-113.
- Naaman, M. Boase, J., & Lai, C-H. Is it really about me? Message content in social awareness streams. *Proc. CSCW '10*, 189-192. 2010.
- Nardi, B., et al. Integrating communication and information through ContactMap. *CACM* 45 (4), 89-95.
- Ozenc, F.K. & Farnham, S.D. Life "modes" in social media. *Proc. CHI '11*, 561-570. 2011.
- Palen, L. & Dourish, P. Unpacking "privacy" for a networked world. *Proc. CHI '03*, 129-136. 2003.
- Skeels, M.M. & Grudin, J. When social networks cross boundaries: A case study of workplace use of Facebook and LinkedIn. *Proc. GROUP '09*, 95-104. 2009.
- Strater, K. & Lipfort, H. Strategies and struggles with privacy in an online social networking community. *BCS-HCI '08*, 111-119. 2008.
- Viégas, F.B. Bloggers' expectations of privacy and accountability: an initial survey. *JCMC*, 10 (3). 2005.
- Wang, Y, et al. "I regretted the minute I pressed share": A qualitative study of regrets on Facebook. *Proc. SOUPS '11*, 1-13. 2011.