Citation for published version:

Jones, S & O'Neill, E 2010, Feasibility of structural network clustering for group-based privacy control in social networks. in *SOUPS '10 Proceedings of the Sixth Symposium on Usable Privacy and Security.*, Article 9, ACM International Conference Proceeding Series, Association for Computing Machinery, New York, 6th Symposium on Usable Privacy and Security, SOUPS 2010, July 14, 2010 - July 16, 2010, Redmond, WA, USA United States, 1/01/10. <https://doi.org/10.1145/1837110.1837122>

DOI:

[10.1145/1837110.1837122](https://doi.org/10.1145/1837110.1837122)

Publication date:

2010

Document Version

Publisher's PDF, also known as Version of record

[Link to publication](#)

© ACM, 2010. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in Jones, S & O'Neill, E 2010, 'Feasibility of structural network clustering for group-based privacy control in social networks'. in *SOUPS '10 Proceedings of the Sixth Symposium on Usable Privacy and Security.* ACM International Conference Proceeding Series, Association for Computing Machinery (ACM), New York, pp. Article9, 6th Symposium on Usable Privacy and Security, SOUPS 2010, July 14, 2010 - July 16, 2010, Redmond, WA, United States, 1 January
<http://doi.acm.org/10.1145/1837110.1837122>

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Feasibility of Structural Network Clustering for Group-Based Privacy Control in Social Networks

Simon Jones, Eamonn O'Neill

Department of Computer Science,

University of Bath,

Bath,

BA2 7AY, UK

s.jones2@bath.ac.uk, eamonn@cs.bath.ac.uk

ABSTRACT

Users of social networking sites often want to manage the sharing of information and content with different groups of people based on their differing relationships. However, grouping contacts places a significant configuration burden on the user. Automated approaches to grouping may have the potential to reduce this burden, however, their use remains largely untested. We investigate people's rationales when grouping their contacts for the purpose of controlling their privacy, finding six criteria that they commonly considered. We assess an automated approach to grouping, based on a network clustering algorithm, whose performance may be analogous to the human's use of some of these criteria. We find that the similarity between the groups created by people and those created by the algorithm is correlated with the modularity of their network. We also demonstrate that the particular clustering algorithm, SCAN, which detects hubs and outliers within a network can be beneficial for identifying contacts who are hard to group or for whom privacy preferences are inconsistent with the rest of their group.

Categories and Subject Descriptors

H5.3. Group and Organization Interfaces; Asynchronous interaction; Web-based interaction.

General Terms

Algorithms, Design, Human Factors.

Keywords

Social networks, social media, privacy, content sharing, group-based access control, network structure, tie strength, automation.

1. INTRODUCTION

Social network sites (SNS) allow their users to disseminate information and digital content across social networks. However, these services typically treat all of a user's contacts equally, enforcing a monolithic notion of 'friendship' and ignoring the complexity and diversity of real world relations and the different roles they play. The varying nature of these relationships has implications for privacy in particular. Users may feel that they are compromising their privacy if they do not have the ability to allow

different people varying degrees of access to their information.

Specifying unique privacy settings for each individual within a user's social network provides the greatest level of control, however, it is common for a user to have hundreds of contacts. This makes it infeasible to accurately configure individual privacy settings every time one chooses to share a piece of content. Some social media sharing services address the need for granular privacy controls by allowing users to compartmentalize their social network, grouping contacts with whom they might share information similarly. The goal is to reduce the overhead of privacy management by reducing the granularity of disclosure decisions from individuals to groups. Managing groups may itself present a configuration burden, although Lederer et al. [16] suggest that good design practices can lessen this burden.

The contribution of this paper is threefold: it (1) identifies what information humans use when they conceptually group 'friends' for privacy purposes; (2) tests an existing network analysis algorithm to see how well it can accomplish this grouping and lessen the burden on the user; and (3) examines whether automated grouping is a viable approach to privacy management, given that personal privacy policies may vary depending on numerous factors such as the current context, need and activity. These contributions are timely as many users of social network sites are obliged to deal with complex privacy management decisions resulting from the integration of multiple social groups inherent in the expansion of such sites.

Contribution (1) is described in Section 3, which reports our collection and analysis of a dataset of 15 egocentric Facebook networks (i.e. networks centered on a single user), with a combined total of 3000 contacts and over 15,000 links. We used a card-sorting and interview study to identify the factors that people considered when creating groups with which to classify their contacts for the purpose of controlling privacy. Contribution (2) is described in Section 4, in which we investigate whether an approach to automating the creation of such groups using network clustering can take account of factors commonly considered by humans when grouping, and examine how closely algorithmic and human groupings match. Contribution (3) is described in Section 5, in which we report a questionnaire study of our participants' willingness to share items of their personal content with individuals within their network, and examine how well their preferences correlate with the groups created in Section 4.

In Section 6 we discuss our findings and make recommendations for how we might design systems that assist users in grouping contacts for content sharing and that support privacy-sensitive

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.

Symposium On Usable Privacy and Security (SOUPS) 2010, July 14-16, 2010, Redmond, WA, USA.

disclosures in social networks where multiple social groups co-exist.

2. RELATED WORK

The networks of ‘Friends’ which users maintain on SNSs often consist of contacts associated with distinct facets of their lives. Sites such as Facebook initially tailored content and user experience to a certain demographic, but now generate broader appeal and support users in maintaining diverse social relationships that span different aspects of their lives [6]. Users are expected to manage multiple social group memberships through a single system, which makes controlling privacy and online identity an interesting new challenge. Presenting information uniformly across connections from these various groups can be problematic, particularly for users who are privacy conscious. Goffman [9] observed that people attempt to maintain a great deal of control over their personas and minimize the appearance of characteristics that are contrary to an idealized version of themselves. Hewitt and Forte [11] suggest that SNSs, in which distinct social groups are co-present, challenge our ability to exercise such control and that a perceived loss of such control is part of many users’ experience on SNSs.

Skeels and Grudin [19] define *group co-presence* as “a situation in which many groups important to an individual are simultaneously present in one context and their presence is salient for the individual”. They found that as people tried to manage the co-presence of multiple groups within their network, they adjusted their posts to make them suitable for a broader audience but were often unhappy with this as a solution. Lampinen [14] asserts that limiting personal disclosures according to the “lowest common denominator” or resorting to other defensive strategies of social identity and privacy protection is a problem that merits attention from SNS researchers and designers. His study revealed that even the most carefree SNS users had attempted to manage group co-presence even when it was not explicitly supported by the system. This required dividing the platform into separate spaces, performing self-censorship and using suitable channels of communication, establishing and continually managing group identities to facilitate more contextual sharing, or relying on the goodwill and discretion of other users.

Several studies have demonstrated users’ desire to create groups of contacts that act as a mechanism for multi-level access control when sharing content, reducing the burden of employing such strategies [6][19]. boyd [2] found that in addition to SNS users being concerned about whether a particular piece of information or content would be suitable for a particular audience, they are also concerned with contacts from distinct contexts or groups being able to reach out and interact with one another. Group-based access control could offer a solution that prevents the inadvertent flow of information between groups within a user’s network.

Lederer et al. [15] found that people base decisions about sharing information more on the identity of the recipient than on the situation within which the information was sought. Similarly, Davis et al. [5] established that people decide with whom to share information based on the type of relationship (e.g. spouse, friend, peer etc). Jones et al. [13] and Olson et al. [18] showed that people want to specify groups and basic categories centered on these relationships, for which they can specify an appropriate privacy setting. However, it has also been noted [16] that managing groups can be a significant burden for the user, particularly as the number of contacts and relationship types expands with the growth and popularity of the service. In

addition, Ackerman & Mainwaring [1] stress that, although valued, privacy is not the users’ primary task and that making it an explicit task for the user can be problematic. Designing systems that reduce privacy violations without significant configuration effort from the user is therefore an important objective.

Researchers have suggested using automated algorithms that use information such as network measures or tie strength to automatically determine distinct groups within a social network, however we have not found any work which builds and tests such an algorithm in order to analyze how well it performs. Gilbert and Karahalios [8] suggest that privacy controls based on tie strength may help to segment a user’s social network into meaningful groups. For example, a system could decide which contacts fall into trusted and untrusted categories and restrict content accordingly. In order to distinguish between strong and weak ties, Gilbert and Karahalios examined activity networks in Facebook, i.e. interactions between members of a social network rather than merely ‘Friend’ connections. They showed that tie strength may be modeled with high accuracy based on these interactions.

Other work has proposed that network measures might reasonably detect distinct groups within a social network. There has been considerable research into partitioning networks into such clusters by algorithmic analysis of the network structure [e.g. 4, 21]. Most of these algorithms cluster vertices within the social network such that there is a dense set of many ties within each cluster and few ties between clusters. A network with this property is said to be highly modular. Widely used modularity-based clustering algorithms, such as the CNM algorithm [4], cluster vertices in the network such that modularity is maximized.

In a study by McCarty [17], network clustering generated clusters that were subsequently verified as meaningful by their respective network owners. McCarty suggests that the reason for participants being able to recognize clusters as meaningful groups is that “frequently members of one group do not know members of another. For example, it is not uncommon for there to be no network ties between family and co-workers”; i.e. there is a dense set of ties within a group and few ties between groups. We speculate that such clustering might reasonably predict groups that users wish to create for controlling their privacy within a social network, separating groups representing distinct contexts and relationship types from one another.

Although many contextual factors can affect privacy decisions, the findings from Lederer [15] suggest that the primary index for such decisions should be the identity of the recipient. The precise context surrounding the disclosure is secondary to this and has less influence on the overall decision. Standard network clustering algorithms are oblivious to such context and operate only on the structure of the network, i.e. the individuals in the network and their links. Thus, algorithmic network clustering may provide a suitable starting point or set of defaults to reflect the primary index in grouping potential recipients of shared information, provided that allowance is made for adjustment of these defaults as required.

Although we share with some related work a goal of developing approaches for the automated detection of co-present groups within social networks, we begin by trying to understand how people manually group contacts for the purpose of controlling privacy. This understanding can then be used to inform an automated approach to assisting such grouping.

3. EXAMINING GROUP-BASED PRIVACY

The first part of this study had three goals: to understand how group-based privacy controls in a widely used social networking service are currently being used; to understand the factors that influence how people group their contacts; and to build a dataset of privacy-based groups, created by participants, against which we could compare the output of automated approaches to grouping.

Participants were recruited through university mailing lists within several departments, and messages advertising the study on a Facebook page used to recruit participants for academic studies. Participants were offered entry into a prize draw for a £25 book voucher as an incentive.

We had 15 participants in total, all of which were Facebook users, 8 male and 7 female, mean age 27 with a range from 19 to 43 years. 6 of the participants were university students and the remaining 9 were in full-time employment. All participants stated that they logged into Facebook at least once a day and had used it to share personal content and information. Our participants had a mean of 200 Facebook ‘Friends’ ($SD=61.4$), with a range from 99 to 312. Our study involved data collection using a custom Facebook application, a card sorting exercise, an interview and an online questionnaire. Our application required no more access to personal information than any other Facebook application and complied with the regulations set by Facebook. The following sections describe each stage of the data collection in more detail.

3.1 Facebook Data Collection

Each participant began by using our application to retrieve the set of all her Friend connections. These undirected links (Figure 1a) grant users access to each other’s profile and shared content. Restrictions can be applied to these connections by placing a Friend into a particular group (referred to as ‘lists’ in Facebook) for which permissions have been configured. Facebook attempted to provide automatic grouping of Friends into lists using data from the optional “How do you know this person?” field and from information about which institutions users indicated they belong to (although this feature is now deprecated in the current version). We also retrieved the names of the lists that each user had within her account and the contacts contained within them in order to question participants about their usefulness and whether they had relied on the automatic grouping by Facebook or manually adjusted these groups themselves.

Figure 1a. Direct ‘Friend’ Ties 1b. Friend-to-Friend Ties

We also used our Facebook application to retrieve a list of all connections amongst a participant’s Friends (Figure 1b), effectively telling us who knows whom within the user’s immediate social network. We used these data to build a graph of the social network of the user’s Friends on which we could run a network-clustering algorithm in an attempt to identify distinct

groups, based solely on the structure of the network in terms of edges and vertices.

3.2 Card Sorting Task and Interviews

In order to investigate how users segment their social network into groups we carried out a manual grouping exercise. Using xSort¹, a card sorting software application, participants were asked to sort cards representing their Friends, as if they were *grouping them for controlling their privacy* on Facebook. They were encouraged to think about with whom they would share information similarly and create groups such that they could be used to grant or deny access to information or content within their Facebook account. Friends’ names were imported to the card sorting application from the data collected on Facebook. The card sort was ‘open’, allowing participants to create an unlimited number of groups, sort according to any criteria they deemed suitable and describe each group with an appropriate label.

Upon completion of the card sorting exercise, participants were interviewed, with each interview session lasting roughly 30 minutes. Approximately 15 minutes was allocated to discussing current practices when sharing content and their use of the group-based privacy controls provided in that version of Facebook. The other 15 minutes was allocated to discussing the rationale behind the groups that the participant had created in the card sorting exercise. The first part of the interview was semi-structured with participants answering questions such as: “Who do you share content with on Facebook?”; “How do you decide what to share and what not to share on Facebook?”; “Have you ever regretted sharing anything with anybody?”. Participants were prompted to discuss in depth any interesting points relating to controlling the disclosure of content and the distinction between groups and individuals in their network. In the second part of the interview the participants were instructed to talk about each of the groups that they had created in turn. They were asked to explain the label given to the group, the criteria that group members conform to in order to qualify as a member of the group and the rationale for creating the group. They were also asked whether they had difficulty grouping any particular contacts and to identify these contacts.

Data collected during the interviews were validated through member checking. Notes taken by the interviewer during the interview were summarized and shared with the participant in order for them to affirm that their explanations were completely and accurately reflected. Interviews were recorded, transcribed and coded using a grounded-theory analysis where emergent themes were iteratively refined. An emphasis was placed on what emerged from the data, rather than using a hypothesis-driven approach for analysis.

3.3 Interview and Card Sort Results

3.3.1 Current Use of Group-Based Privacy

The following section describes the findings from the first part of our interview in which we questioned users about their use of group-based privacy controls for sharing content on Facebook.

¹ www.xsortapp.com

Table 1. Card Sorting Results

PARTICIPANT	NETWORK SIZE	NO. OF GROUPS	MEAN GROUP SIZE	STDEV GROUP SIZE	EXAMPLE GROUP LABELS	
					LARGEST GROUPS	SMALLEST GROUPS
1	286	20	14.3	25.2	Old School Mates, Work Mates, Loughborough Mates	Best Mate, Hairdresser, Ex-Boyfriend
2	205	11	18.5	18.5	Old Uni/School Friends, Don't Know/Don't Care, Lost Contact	Family, Know Me Best, Regular Contact
3	126	5	25.2	8.1	Friends, Not Close	Best Friends, Don't Like, Barely Know
4	202	15	13.5	10.5	Ski Season, School Mates, Uni Acquaintances	Work, Random, EngD
5	240	16	15.0	19.3	University PhD, School Friends, University	Work Colleagues, San Diego, Met at Gig
6	147	11	13.4	15.1	Secondary School, Loughborough, Home Friends	Girlfriend and Family, Closest Friends, Primary School
7	175	9	13.4	8.4	Department Friends, High School, Friends of Friends	Extended Family, Distant Family, Work Colleagues
8	200	17	14.2	15.1	Friends - Abroad, Friends - My City, Undergraduate Friends	Family, Bahrain, Who Are They?
9	312	20	16.3	21.8	Waterpolo, Department, School	Family, WH Smiths, Bristol Waterpolo
10	118	4	29.5	19.2	University Friends, Non-University Friends	Department, HCI Group
11	269	14	19.2	22.1	Thai Society, Bath, CS Bath	Family, The Den, Travel Buddies
12	206	16	12.8	13.0	Bath, School, Glasgow Friends	Unilever Colleagues, American Friends, Amsterdam Friends
13	99	10	9.8	5.2	UKC, Greek Dancing, Friends	School, Canterbury College, Nottingham
14	224	16	14.1	18.2	Old School Friends, Friends of Friends, Music Society	Colleagues, Supervisor, Close Friends
15	195	9	21.7	15.0	Work Friends, Old Friends, Best Friends	Ex's, Friends of Friends, Boyfriend
MEAN	200	12.9	16.7	17.1		

The majority of our participants, despite being seasoned Facebook users, were unaware that the group-based privacy control features even existed. Of those who were aware, only two had taken the time to manually organize their contacts into groups.

Both of these users had done so only when the instant messaging service integrated into Facebook began to utilize the groups. “Although I knew it would be useful to have people grouped for the purposes of controlling access to my content, I was only really motivated to do it when it became a part of the chat list too. It was messy having all of my contacts in one long list, so I grouped them then”. Thus, controlling privacy was not enough motivation to group contacts but when other features started to utilize the groups, e.g. organizing the chat list, sending messages to whole groups, filtering content received from groups, users were more motivated to create them.

The overwhelming consensus amongst our participants was that the task of organizing contacts into groups required too much time and effort to be worthwhile. The mean time to complete the card sort was 27 minutes. All participants stated that they would not spend that much time configuring the groups for their Facebook account, especially as they felt that they are impeded by the interface for doing so to a greater extent than with our card sorting application. One participant supported Lederer et al.'s [16] claim that good design practices, such as allowing contacts to be bootstrapped into groups as they are added, might reduce the

configuration burden: “I think if [a grouping mechanism] was there when you added contacts and it was easy to do then I'd use it, but I don't think I could be bothered to go through all my contacts now to group them.”

A previous version of Facebook attempted to automatically create groups and organize contacts into them, relying on information explicitly provided about the nature of the relationship or the organization to which a person belongs, resulting in groups such as “Family”, “Work”, “Grad School”, “High School” and “Elementary School”. We found that on average only 20% of Friends were placed into the groups by Facebook and that none of our participants had used these groups to restrict access to particular content. They typically regarded them as “too incomplete to be useful”. None of our participants had attempted to finish populating the groups automatically created by Facebook, other than a group called ‘Limited Profile’. This group has highly restrictive privacy settings by default and acts as an easy way for users to segregate contacts to whom they wish to give very limited access to their information.

3.3.2 Card Sort Results

Table 1 shows a summary of the data collected from the card sorting exercise. All participants completed the exercise, creating a mean of 12.9 groups (SD=4.93). The mean group size was 16.7 contacts (SD=5.26). The table also gives some examples of the

labels given to these groups by participants. Although these labels give an overview of the card sort, in many cases they do not reveal details about how or why the group was created. We were careful to avoid relying on our own interpretation of these labels when attempting to identify the criteria that participants had considered when forming groups. Participants gave a description of the considerations they had made for each group, which were recorded by the interviewer and subsequently validated by the participant to ensure accuracy.

The following section describes the commonly considered criteria, which were uncovered from the card sorting exercise and interview sessions. These were exposed by first open-coding the interview data and developing initial categories that were iteratively refined in order for core criteria to emerge.

3.4 Grouping Criteria Used in the Card Sort

Our analysis of the interviews with the participants about how they had performed the card sorting to group their contacts as they would for controlling privacy revealed the following 6 commonly considered criteria:

- Social Circles & Cliques
- Tie Strength
- Temporal Episodes
- Geographical Locations
- Functional Roles
- Organizational Boundaries

3.4.1 Social Circles & Cliques

Almost all (13/15) participants considered the cliques and social circles within their network when grouping contacts. Social circles and cliques refer to tightly knit groups, e.g. a group of friends that are highly connected to one another.

Large groups formed by the participants were often subdivided by them in order to separate distinct social circles: “Two of the groups that I’ve created contain my university friends, but I’ve divided the two separate friendship groups that I belong to. The personalities of people in each of the groups are quite different and I think I probably behave a little differently towards each group, even though I consider myself just as close to the people in both of them”. This participant sometimes shared content and information differently with each of the groups in order to comply with the “expectations” of each group. “With one group I tend to joke around more often, with the other group I think I’m a bit more sensible”. Users are not only aware of social circles and cliques within their wider social network but are influenced by them in their information and content sharing choices, with an inclination not to share information or content – or indeed behavior – that does not conform to the social norms or values of a particular social circle.

The tight-knit nature of social circles and the ease with which information could be disseminated amongst its members also contributed to participants’ reasons for grouping them together: “The likelihood is that if I reveal something to one person in this group then the other people would find out about it from them anyway. I would either restrict something from all of them or none of them”.

3.4.2 Tie Strength

Eleven participants considered factors relating to their relationship with a contact, such as the closeness, emotional intensity, level of

trust and frequency of communication, when assigning them to a group. These factors can be taken as indicators of the strength of an interpersonal tie [10].

It was common for participants to create groups of ‘Close friends’ or ‘Best friends’. These were either distinct groups or were sub-groups within a larger group. For example, participants either created a group based on strong ties, disregarding other factors, or they divided a group formed using other factors to produce two sub-groups, populated by strong ties and weaker ties respectively. One participant justified her reason for grouping strong ties thus: “Even though [Facebook] classes contacts as ‘Friends’, they’re not actually all my friends. Sometimes they’re just people that I know, perhaps not even very well at all. I don’t want them knowing everything about me. But it’s different for close friends. I’d be willing to share far more with them”.

It was also common for participants to create groups exclusively containing weak ties, such as ‘People I hardly know’, ‘Acquaintances’, ‘Friends of friends’ and ‘People that randomly added me’. Although the labels given to some of these groups did not always explicitly suggest tie strength, the reasoning behind the group’s creation was to group people with whom they had weak ties. For example, one participant had some contacts who had been “randomly added” to her set of Friends but who subsequently became strong ties. They were not grouped with the majority of “randomly added” contacts, because this latter group was only for weak ties.

One participant told us that she had based her groups on the different levels of trust that she placed in people and on her assessment of whether people would “judge” her based on the content she uploaded. “They might judge personal content in a negative way. It’s not considered very cool to have lots of photos with your Mum and Dad, but my close friends and family enjoy seeing those photos. I worry that people that I don’t know very well might see things like that and laugh at me”. Others were more concerned about other components of interpersonal tie strength such as the frequency or amount of contact. They felt that it would be important to share content differently with people with whom they did not interact a large amount or very often versus people with whom they interacted a lot and frequently.

Participants isolated individuals whom they described themselves as being exceptionally close to, or even exceptionally cautious of, to be able to specify unique privacy preferences for them. Typically these were very close friends with whom participants felt they could share everything, and untrusted or disliked individuals with whom they wished to share very little. Hence, although participants are often able to identify *groups* of contacts for privacy purposes, there is still a need to specify privacy settings for individuals. Interfaces that allow group-based privacy control should also allow for unique settings to be applied to exceptional individuals.

3.4.3 Temporal Episodes

Six participants said that they had created groups that represented significant parts of their lives. One participant referred to these as “episodes”, each episode representing a certain period of time. Into each of these groups they placed contacts that they associated with that episode. Time scales ranged from hours (e.g. the time spent at a certain event) to months or years (e.g. the groups ‘Summer 2009’ and ‘My Childhood’). Temporal Episodes were often closely linked and sometimes conflated with Geographical Locations. For example, when explaining the rationale behind

Table 2. Card Sorting Criteria and Strategies

		PARTICIPANT															TOTAL
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
CRITERIA	SOCIAL CIRCLES & CLIQUES	✓	✓		✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓	13
	TIE STRENGTH	✓	✓	✓	✓	✓	✓	✓	✓					✓	✓	✓	11
	GEOGRAPHICAL LOCATIONS	✓			✓	✓	✓		✓		✓	✓	✓	✓		✓	10
	ORGANIZATIONAL BOUNDARIES				✓		✓	✓	✓	✓	✓	✓		✓	✓		9
	TEMPORAL EPISODES		✓		✓		✓	✓					✓			✓	6
	FUNCTIONAL ROLES	✓	✓		✓					✓		✓				✓	6

their groups, participants would interchangeably talk about periods of time by referring to particular locations where that time was spent and vice versa.

3.4.4 Geographical Locations

Ten participants reported compartmentalizing their contacts into groups associated with particular locations, for example where they first met or where they spent a significant amount of time throughout their relationship. Locations varied in scale from particular venues to entire countries. For example, participants created groups such as ‘Loughborough’, for anybody they had met in that town. ‘People I met on holiday in America’ was used to group people met in a certain country during a particular period of a few weeks, and ‘Met at a gig’ was used to group people met at a particular venue in the space of only a few hours.

When asked why temporal episodes and geographical locations were important to consider for controlling privacy, users expressed concern about their momentary actions having unforeseen consequences if viewed in different contexts. One participant told us, “There’s a time and a place for certain behavior. If I share content such as photos, I think to myself, would I have been comfortable acting like this in front of people if they were actually there at the time?” By compartmentalizing contacts associated with different times and places our participants seemed to be attempting to mitigate potentially negative effects of future public access and the loss of physical boundaries when migrating information about their lives to the digital domain.

3.4.5 Functional Roles

Although Facebook is primarily a *social* networking service, participants sometimes used it to foster non-social connections. Some contacts were added for professional networking and others were added for functional reasons, providing a particular use or service to the user. For example, one participant had added as a contact somebody whom she had encountered through a classified advertisement, which she had placed on Facebook to sell an item, merely as a way of bookmarking him for when she wanted to communicate about the transaction. Participants frequently said that they had asked themselves *why* they had added these ‘functional’ contacts to Facebook when deciding which group to place them in. They expressed concern about making personal information, particularly regarding their social activities, available to such contacts, however, they had not taken any measures to address such concerns.

3.4.6 Organizational Boundaries

Some participants grouped contacts based on the institutions or organizations that they belonged to, in order to be able to separate the professional and social aspects of their lives. They created groups representing particular companies, departments and different roles within their workplace.

One participant told us: “I do sometimes worry what these people might think when they see my profile. Not everything on there portrays me in a professional light, which is how I try to represent myself when I’m around them in person.” He added, “I keep meaning to alter my privacy preferences for these people, but I never get around to it”.

3.4.7 Combinations of Factors

When considering which group to place a particular contact in, the participants asked themselves questions relating to the 6 criteria, such as: which friendship group does this person belong to, how strong is our relationship, when did we meet, where did we meet, why did I add her as a contact, does he belong to a particular organization or institution. Participants used various combinations of these criteria, ranging from consideration of only a single criterion to consideration of all 6. In explaining their rationale for grouping contacts, most participants stated that they had used the criteria that felt most natural to them. Some stated that it was easier to group participants in a certain way but that they were aware of other approaches to grouping. For some participants it was easier to think about the strength of interpersonal ties and to separate trusted and untrusted contacts. For others it was easier to consider only organizational boundaries or functional roles of contacts because they were more ‘tangible’ concepts.

Table 2 shows the factors that were considered by each participant when they were organizing their contacts into groups for the card sort. Only participants 8 and 13 considered precisely the same combination of factors. Participants most commonly considered the formation of social circles and cliques within their network. The next most commonly considered factor was the strength of interpersonal tie between the participant and the contact.

4. TESTING AUTOMATED GROUPING

Having exposed a number of factors that people consider when creating groups for controlling access to their personal content and information, we investigated automating this grouping using a network analysis algorithm. We were interested in exploring how closely we could match humans’ grouping of their social networks

with automated grouping, with the ultimate aim of assisting users in grouping contacts, thereby reducing the configuration burden of managing privacy and content sharing.

Our card-sorting and interview findings suggest that a suitable automated approach should group contacts using criteria analogous to as many of the 6 identified factors as possible. An additional requirement for an automated approach to grouping is that minimal explicit input should be required from the user beyond their normal interaction with the system [1]. For example a user should not have to provide additional information about each of their contacts in order to assist the automation.

We believe that network clustering is a candidate for automated grouping for a number of reasons. First, network clustering relies only on the network graph of connections between individuals; no additional information is required. Forming these connections is a fundamental part of the user experience on a social network such as Facebook. Hogan [12] suggests that the norms of ‘friending’ on Facebook give rise to a more coherent depiction of personal ties in real-world networks than was previously possible through other means.

Secondly, performing network clustering on an egocentric network has the potential to recognize unique features of that network which may not be present in other users’ networks. Clusters/groups do not have to follow general rules (e.g. everyone must have a group of “Friends” and a group of “Work Colleagues”, a certain number of groups, or groups of a particular size). Instead, the groups created are exclusive to the individual user.

Finally, network clustering appears to capture groups that are formed analogously to several of the 6 criteria that we have identified. Social circles and cliques can readily be derived by algorithmic analysis of the network structure [3, 13]. The persistence of ties within the network also allows users’ network graphs to evolve over time, reflecting a personal history of relationships. Clusters of ties are likely to form as people transition between temporal episodes, distinct geographical locations and organizational boundaries and build new sets of relationships. Such clusters will also be identifiable algorithmically in the network structure.

4.1 The SCAN Clustering Algorithm

We used the SCAN algorithm [21] to cluster vertices (i.e. people) within each participant’s egocentric social network (see Fig 1b) into groups. This algorithm also detects and isolates two kinds of vertex that play distinctive roles: ‘hubs’ that bridge clusters and ‘outliers’ that are marginally connected to a cluster. We observed in the card sorting exercise that participants struggled to place some contacts into groups because they were either weakly associated with a group (outliers) or strongly associated with multiple groups (hubs).

Hubs are likely to represent people who belong to multiple social circles within the network. Consider the example of a work colleague who is also a friend of the family. Modularity based algorithms such as CNM [4], which consider only direct connections between contacts, are likely to place this person into *either* a cluster of work colleagues *or* a cluster of family members, depending on which cluster she has most connections to. However, from a human perspective a work colleague who is also a friend of the family may be viewed differently from other work colleagues who are not. She may be treated differently in terms of

the access she is granted to personal information and content. For example, being a family friend could justify giving this person access to a certain piece of content. Alternatively, she might be placed in groups with more restricted access to some personal content in order to prevent the dissemination of information from one cluster to another through her bridging effect.

The SCAN algorithm uses the neighborhood of each vertex as an additional clustering criterion. Vertices are grouped based on how they share neighbors. Thus, hubs and outliers are detected because they differ in terms of structure and connectivity from the vertices around them. Figure 2 shows an example of one participant’s egocentric network clustered using the SCAN algorithm.

Figure 2. SCAN clustering output for a participant’s egocentric network showing clusters, hubs and outliers.

4.2 Comparison of Algorithmic and Human Grouping

We compared the output of the SCAN clustering algorithm with the output of the card-sorting exercise in which our participants grouped their own social network. Our implementation of the SCAN algorithm performed clustering with ϵ values from 0.1 to 0.9 (the threshold for determining whether vertices share enough neighbors to be clustered together) and selected the clustering output with the highest modularity value. In order to see how closely the output of the SCAN algorithm matched the groups produced by each participant we used a formula designed to measure the similarity between two clusterings of the same dataset [20]. Let $C = \{C_1, C_2, \dots, C_m\}$ be the set of clusters produced by the clustering algorithm and $G = \{G_1, G_2, \dots, G_n\}$ be the set of groups produced by the participant in the card sorting exercise. The similarity S_{ij} of a cluster C_i and group G_j is p , the number of contacts common to both groups ($C_i \cap G_j$), divided by q , the total number of unique contacts in both groups ($C_i \cup G_j$). The overall similarity of a clustering C and grouping G uses all pair-wise comparisons of clusters and groups and is defined as:

$$\text{Sim}(C, G) = \sum_{i \leq m, j \leq n} S_{ij} / \max(m, n)$$

Table 3 shows the similarity of groups produced by our participants with algorithmically produced clusters. A value of 1 represents an identical output and 0 represents an output in which none of the contacts were grouped similarly. The modularity of the network clustering using SCAN is also shown.

On average, the SCAN algorithm produced clustering outputs that were 44.8% similar to the groups created by participants using the card sorting, with a range from 18.1 to 79.5%. Although this value is by no means high enough for the clustering method to be

used to automate groups within a social networking service, it corroborates that the presence of social circles and cliques within the network structure is a significant factor when manually organizing contacts into groups. Using only data relating to the structure of the network approximately half of a user's contacts can be placed into what the user considers suitable groups for controlling privacy. However, this may not reduce the burden on users in grouping their contacts, as they would incur the extra cost of having to identify and adjust contacts that are incorrectly grouped.

Table 3. Clustering Modularity and Similarity Values

PARTICIPANT	SCAN CLUSTERING MODULARITY	SIMILARITY BETWEEN HUMAN & ALGORITHMIC CLUSTERING	SIMILARITY (AFTER MERGING STRONG & WEAK TIE CARD GROUPS)
1	.4032	.4100	.7454
2	.7102	.5222	.6788
3	.1625	.2849	-
4	.5783	.4111	.6166
5	.6085	.4812	.5922
6	.3679	.4416	.6071
7	.6877	.4374	.5383
8	.1810	.3085	.6282
9	.4562	.5804	-
10	.0220	.1919	-
11	.4842	.4411	-
12	.6861	.5777	-
13	.7952	.8004	.8892
14	.3417	.3979	.6719
15	.3730	.4384	.7306
MEAN	0.4572	0.4483	0.6698

4.2.1 Tie Strength Divisions

Within our data we observed that single clusters produced by the algorithm were often similar to the combination of two groups produced by the human. For example, Cluster 14 of Participant 1's network was 45% similar to the group that he had created and called 'DeMontfort Friends', and 37% similar to his group called 'Close DeMontfort Friends'. In order to produce an algorithmic clustering output that was more similar to the human output, the algorithm would have needed to divide Cluster 14 into two corresponding clusters, however, the names of each group and the participant's expressed rationale for creating them suggest that knowing how to perform this division relies not on analysis of the network structure but on consideration of tie strength within a cluster. The only distinction between members of the two groups is their 'closeness' to the participant.

Our card sorting and interview data allowed us to identify where users had overtly created a single group and then divided it with a consideration for weak and strong ties. In cases for which we had documented evidence of this behavior (10 out of 15 participants) we were able to reconstruct the original, undivided group and re-calculate the similarity measures to show the effect of tie strength. The final column in Table 3 shows the re-calculated similarity values where applicable. The mean average similarity for these participants was 0.6698, suggesting that distinguishing between weak and strong ties in the automation process could improve the similarity between its output and the groups created by

participants. A paired t-test revealed that this improvement in similarity was statistically significant. ($t = 6.8917$, 9 d.f., $p < 0.01$).

Tie strength data for every connection between a user and her contacts would be required in order to augment the network clustering algorithm to split groups according to tie strength. The work of Gilbert et al. [8] suggests that such a process could use an analysis of the interactions between members of a social network to derive the strength of ties, however we were unable to capture sufficient data about these interactions to replicate such a model within our study.

4.2.2 Analysis of SCAN Algorithm Output

We found that the detection of 'hub' and 'outlier' vertices provided some additional value. When interviewing participants about how they had carried out the card sorting exercise, we asked them to identify any particular contacts that they had difficulty placing into a group. We investigated how the SCAN algorithm had dealt with these contacts, 16 in total from all participants. Nine of the 16 were identified as hubs and a further 3 were identified as outliers. Thus, hubs and outliers identified by SCAN represented 75% of the problematic contacts. In an automated grouping service the algorithm could, for example, flag them and prompt the user to deal with them manually.

Within many of the participants' social networks there were vertices that were disconnected from any of the participant's other contacts. We found that for 7 of the participants the set of isolated vertices closely matched (>75%) a group of contacts that they had created during the card sorting exercise. These were typically the groups that participants had created for weak ties, with labels such as 'People I barely know', 'Who are these people?' and 'People I met online'. If a person has no mutual friends with one of her contacts, i.e. that contact is a disconnected vertex in the network, that contact may not be well known to the user. The similarity of the disconnected vertices detected by the clustering algorithm with weak-tie groups created by the participants suggests that people may often group such contacts together.

As well as being unable to divide single clusters based on a division in tie strength, we also noticed that in some cases the clustering algorithm produced clusters which, had they been merged, would have been more similar to participant created groups. Contacts that a participant perceived as a single group, for example "My Girlfriend's Friends", often consisted of multiple distinct cliques. The distinction between cliques may be more apparent to the members of the group, however it is unimportant to the participant for the purpose of controlling his privacy. Merging clusters would produce a single cluster that is more similar to the group created by the participant, however, any algorithm that works solely on network structure cannot know which groups to merge. A guess could be made based on the connectedness of clusters but it would be prudent to allow user intervention before performing such merging.

This feature was observable in the comparison of human and algorithmic outputs for many participants' networks, particularly those with a low modularity value. By definition, low modularity tells us that there are clusters which are distinct but which still have many edges between them. An analysis using Pearson's correlation coefficient indicated a statistically significant linear relationship between network modularity (as measured by the SCAN algorithm) and the similarity between human and algorithmic clustering, $r(15) = 0.837$, $p < 0.01$. That is, the more highly modularized a network, the more the algorithm's clusters

Figure 3 a. Modularity vs Similarity graph; b. Comparing number of groups produced by algorithm and card sort.

are similar to the groups produced by the human (see Fig 3a). A possible explanation is that people who have clearly distinct social circles and cliques in their social network structure are more inclined to use this as a criterion for grouping their contacts. Networks with less distinct social circles and cliques produce algorithmic clusters that are dissimilar to those produced by humans. In these circumstances, participants may rely on other criteria for grouping, e.g. tie strength. The networks of participants 3 and 10 have the two lowest modularity values and these were the only participants not to consider cliques and social circles when manually organizing their contacts into groups (see Table 2).

We found no significant effect of network size (number of contacts) on either the modularity or the similarity of the algorithmic and human clustering. This suggests that the approach of identifying distinct groups using network clustering will scale from small to large networks. There may be an upper or lower threshold at which the method becomes less effective but we did not find this within our range of 99 to 312 contacts.

An analysis using Pearson's correlation coefficient indicated a statistically significant linear relationship between network size and the number of groups created by participants, $r(15) = 0.8308$, $p < 0.01$. That is, participants with more Friends created more groups within which to place them. Similarly, Pearson's correlation coefficient indicated a statistically significant linear relationship between network size and the number of clusters found by the SCAN algorithm, $r(15) = 0.6542$, $p < 0.01$. This suggests that as people's networks increase in size they do not merely increase the size of the social groups that already exist but also add new social groups to their network.

5. INDIVIDUAL VS. GROUP PRIVACY

Our final phase of data collection consisted of a questionnaire integrated into a Facebook application that was presented to our participants approximately 3 weeks after the interview and card sorting session. Participants were asked to select a single item of information or content from any part of their Facebook profile which they either a) had shared but not with all of their contacts, b) had shared with all of their contacts but had actually wanted to share with only a subset of their contacts, or c) had not shared on Facebook because they did not want certain contacts to have access to it.

Participants were then presented with a list of 100 of their Facebook contacts in a randomized order (99 in the case of the participant with fewer than 100 contacts). These contacts were a

stratified sample from the groups created in the card sorting exercise, although this was not explained to them. They were asked to use a sliding scale to indicate their level of 'willingness-to-share' the particular item of content that they had selected with each of these contacts. The scale ranged from 'Not at all willing' to 'Very willing' and was initially set at a central, neutral position. Two of the participants were unavailable to complete the questionnaire. The remaining 13 provided complete responses for all 1299 of the sampled contacts from their respective networks.

Thus, we could analyse the groups created by the humans and by the SCAN algorithm for uniformity in the willingness-to-share values of their members for content taken from a range of contextualized information sharing situations. By contextualized, we mean that they are not hypothetical or generalised but specific instances with known context.

Since participants had been asked to group their contacts before focusing their attention on a specific item of content, these groups were intended for generic information sharing management; participants created groups they considered would be useful in a range of sharing situations. Such groups could be inappropriate for managing the sharing of particular content with a given contact given that privacy related decisions may vary depending on contextual factors specific to an individual. We cannot test the efficacy of groups against all conceivable content sharing contexts. However, if willingness-to-share values are uniform within groups for a range of content sharing instances that participants have identified as being privacy sensitive, we can make a case for grouping as a viable approach to reducing the burden of privacy management.

We also examined the properties of contacts whose willingness-to-share values differed markedly from the rest of their group. If willingness-to-share values are not uniform within groups, we can investigate whether the contacts whose values are outliers within the group exhibit any particular structural properties within the network that might help to identify them. They could then be excluded from the group for particular information sharing purposes.

5.1 Selected Content Items

Our participants selected items from a wide range of content and information on which to base their responses. These were: home addresses, contact details such as phone numbers or email addresses, specific status updates relating to their work life, organizational details about private events, messages containing secrets or sensitive information and particular photographs taken on holiday or while out partying. Participants were asked to pick specific instances of content or information rather than a particular type of content, e.g. status updates or photographs, in order to avoid generalization or responses to hypothetical situations.

One participant was unable to find an example of content that he had been unwilling to share uniformly with all of his contacts and therefore did not complete the questionnaire in full. Two completed the questionnaire but gave entirely uniform willingness-to-share values for all contacts in the sample. When we questioned these participants about their responses, they told us that they were willing to share the selected items with their entire network except for one particular individual. In both cases this individual had not been included in our sample population.

5.2 Individual vs. Group Privacy Results

The bar chart shown in Figure 4 exemplifies the data collected. A single chart corresponds to the data from a single participant sharing a particular item with their contacts. Each bar in the chart represents the willingness-to-share value assigned to a particular contact. Contacts are ordered according to the group to which they were assigned. Each shaded region delimits a group. The y-axis ranges from -50, representing 'Not at all willing', to +50, representing 'Very willing'.

Figure 4. Example willingness-to-share data for a single participant sharing a single item of information/content

We used these data to help answer the following questions: are the groups that participants have created accurate and useful for sharing content; are the groups created by the algorithm more or less useful than those created by the participant; to what extent does information leakage or undesirable sharing occur; and can the structural properties of the network provide any insight into when and why this happens?

We measured the extent to which undesirable sharing outcomes occurred when privacy preferences were specified at the group level. We define an undesirable outcome as the user either unwillingly sharing an item with a contact or unwillingly denying a contact access to an item. A group can be used either to grant or deny access to the item for all the group members. An undesirable outcome can therefore occur when an individual contact's willingness-to-share value conflicts with the decision to grant or deny access to their group.

Since willingness-to-share values were denoted on a sliding scale, we began by creating a bipolar classification of values as either 'Yes' or 'No' decisions to grant access to the item to that individual contact. This approach oversimplifies privacy related decisions to some extent, and assumes that there is no margin of neutrality within which users are unconcerned about the effects of sharing their content. However, it serves to indicate the extent to which conflict occurs within groups, if we assume that sharing when only slightly unwilling to share and denying access when slightly willing to share are undesirable outcomes.

If users of group-based access controls grant or deny access to a group within which there is a conflict, they may not necessarily choose the option that minimizes the number of undesirable outcomes. For this reason, we consider best and worst case scenarios. We define best case as the lowest number of contacts which contravene the overall group decision and worst case as the highest number of contacts which contravene the overall group decision. For example, consider a group of 5 contacts within which you *would* share content with 4 and *not* with the remaining

1. The worst case scenario is that you choose not to share with the group, denying access to 4 contacts that you wish to share with. The best case scenario is that you grant access to the group but therefore grant access to 1 contact you did not wish to share with. By simplifying our model of privacy-related decisions, we assume that inadvertently granting and denying access are both equally undesirable. In reality, of course, the consequences of unwillingly granting access to one individual may outweigh the benefits of sharing with the others. However, despite its limitations, this model enables us to approximate the efficacy of group-based privacy control.

The following sections report the extent to which information leakage and over-restriction occurred within the groups. This analysis was performed for groups created both by participants and by the algorithm. We were interested in determining whether groups created by participants explicitly, but generically for privacy control are sufficiently useful and whether it is worthwhile striving to automate their creation. We also wish to test how useful the algorithm is in its current form, which considers only structural network clusters.

5.2.1 Best and Worst Cases

For the groups created by humans, the average *worst case* scenario across all participants is that 22.2% of contacts are either undesirably given or denied access when permissions are set at group level ($M=0.222$, $SD=0.036$). In other words, 77.8% of contacts are correctly granted or denied access.

The average *best case* scenario across all participants is that 9.2% of contacts are either undesirably granted or denied access ($M=0.092$, $SD=0.025$), meaning that 90.8% of individual contacts are correctly granted or denied access when permissions are set at group level.

For groups determined by the algorithm rather than by humans, the average worst case scenario across all participants is that 61.7% of contacts are either undesirably granted or denied access ($M=0.617$, $SD=0.094$), meaning that 38.3% of individual contacts are correctly granted or denied access when permissions are set at group level.

The average best case scenario across all participants is that 23.9% of contacts are either undesirably granted or denied access ($M=0.239$, $SD=0.081$), meaning that 76.1% of individual contacts are correctly granted or denied access when permissions are set at group level.

5.2.2 Improving the Model of Privacy Decisions

The values in Section 5.2.1 represent a baseline measure of how well the groups perform as mediators of privacy preferences. Modeling indifference/neutrality to sharing particular content is only likely to improve this performance. For illustrative purposes we created a margin of 'Neutral' values (-16.6 to +16.6 i.e. the central third of the slider scale) such that low positive and negative values were classified as 'Neutral'. By restricting 'Yes' and 'No' decisions to values that are more polarized to either end of the willingness-to-share scale, we restrict undesirable outcomes to those that are likely to have higher cost, e.g. a conflict in which a participant is 'Very willing' to share with one contact and 'Very unwilling' to share with another contact in the same group.

In reality the boundaries between neutrality and choosing to share or restrict access to content are far more fuzzy and subject to individual variability between participants as well as the specifics of the content and its surrounding context.

When neutrality is considered, for the human-created groups the average worst case scenario is that 5.7% of contacts are undesirably granted or denied access ($M=0.057$, $SD=0.019$), meaning that 94.3% of individual contacts are correctly granted or denied access when permissions are set at group level. The average best case scenario is that 1.9% of contacts are undesirably granted or denied access ($M=0.019$, $SD=0.012$), meaning that 98.1% of individual contacts are correctly granted or denied access when permissions are set at group level.

For the algorithmically created groups, the average worst case scenario is that 42.4% of contacts are undesirably granted or denied access ($M=0.424$, $SD=0.145$), meaning that 57.6% of individual contacts are correctly granted or denied access when permissions are set at group level. The average best case scenario is that 13.2% of contacts are undesirably granted or denied access ($M=0.132$, $SD=0.048$), meaning that 86.8% of individual contacts are correctly granted or denied access when permissions are set at group level.

5.2.3 Network Properties vs. Group Uniformity

We analyzed the data to determine whether the number of undesirable outcomes was correlated with factors such as the size and modularity of the network and number of groups created. We also hypothesized that participants with coarse-grained groups would experience more information leakage and over-restriction, due to their not separating their contacts with sufficient granularity to manage their privacy settings effectively.

Tests of Pearson's correlation revealed no significant relationship between the number of undesirable outcomes and network size; $r(10)=0.0464$, $p>0.05$ n.s., number of groups; $r(10)=0.3042$, $p>0.05$ n.s. or modularity; $r(10)=0.3769$, $p>0.05$ n.s. There is no evidence to suggest that participants are better positioned to create more useful groups for controlling privacy if they have networks of a particular size or modularity, or if they create more or fewer groups. However, our hypothesis with respect to granularity holds. Participants creating groups with a higher group to contact ratio (i.e. finer granularity) had fewer undesirable outcomes when sharing content with groups, $r(10)=0.6355$, $p<0.05$.

5.2.4 Privacy Settings for Hubs and Outliers

We noted in Section 4.2.2 that hubs and outliers accounted for 75% of the contacts that participants struggled to assign to groups. We hypothesized that hubs and outliers might also give participants difficulty with content sharing decisions, either because of their ability to bridge multiple groups or their weak association with a particular group. By using the SCAN algorithm we were able to identify these structural hubs and outliers, allowing us to test for a correlation between contacts whose willingness-to-share values were numerically distant from the rest of their group and the structural network properties of these anomalous contacts. We identified all contacts for whom willingness-to-share values were more than 2 standard deviations from the group mean and noted whether they had been classified as group members, hubs or outliers by the SCAN algorithm.

An unpaired t-test showed a significantly larger percentage of outliers within the set of contacts whose willingness-to-share values are more than 2SD from the group mean ($M=0.337$, $SD=0.331$) than outliers within the entire network ($M=0.069$, $SD=0.059$), ($t = 2.3893$, 18 d.f., $p < 0.05$). That is, on average fewer than 7% of contacts throughout a network were structural outliers, but amongst contacts whose willingness-to-share values

were not uniform with the rest of their group, over a third were structural outliers.

Identifying outliers to the user could act as a useful "just-in-time" feature, helping to prevent undesirable outcomes when sharing information and content with a group of contacts. We found that removing algorithmically identified outliers from the human-created groups improved the accuracy of the groups, with between 84.4% (worst case) and 94.7% (best case) of contacts being correctly granted or denied access [cf. Section 5.2.1.]

An unpaired t-test comparing the percentage of hubs within the set of contacts whose willingness-to-share values were more than 2SD from the group mean ($M=0.044$, $SD=0.133$) and hubs in the entire network ($M=0.027$, $SD=0.039$) revealed no significant difference ($t=0.384$, 18 d.f., ns). The contrast between this finding and the corresponding significant finding for outliers may imply that our participants were less concerned about the effect of sharing with contacts who bridge multiple network clusters than sharing with contacts who are weakly associated with a cluster.

6. DISCUSSION

The results from our interviews suggested that participants were generally unwilling to organize their contacts into groups on Facebook because of the significant burden it placed on them. A few users were motivated to create groups only when they could also be used for a number of other purposes, such as organizing chat lists, sending group messages and filtering content received from groups. This finding suggests that working towards systems that reduce users' configuration burden by automating the formation of groups is potentially valuable.

Our participants were positive about the possibility of having an automated method for grouping contacts. While they were aware that it could be far less time consuming, they also noted that the groups would have to be both meaningful and complete in order to be useful. Facebook's previous attempt to automate the creation of groups of contacts did not adequately meet the needs of its users for two reasons: groups were often incomplete and the mechanism for grouping did not reflect the criteria that users commonly consider when grouping their contacts.

Through our card sorting exercise we uncovered six criteria that users considered when manually sorting contacts into groups for controlling privacy: social circles and cliques, tie strength, temporal episodes, geographical locations, functional roles, and organizational boundaries. These corroborate prior findings and validate prior contextual analyses of copresent privacy practices in social network sites. Skeels and Grudin [19] found that the multiplicity of groups on Facebook is both temporal, bringing together groups from different stages of the individual's life, and spatial, bringing together people who might live in different geographical locations. We corroborate these findings and show that users are also concerned about the coexistence of groups that vary in tie strength, span organizational boundaries and have different functional roles.

Our findings convince us of the pitfalls of striving for 100% automation, for several reasons. We found that different people considered different grouping criteria to varying degrees. Basing our grouping algorithm on analysis of egocentric networks allows for individualized outputs that incorporate unique features of a user's network. However, in order to better account for this individual variability users may need to specify which criteria the algorithm should account for.

While one might not expect a single, off-the-shelf algorithm to produce accurate groups for such a complex task, we have shown that the SCAN algorithm can go some way towards identifying groups that the user considers useful for controlling privacy. We do not wish to paint these results in too positive a light and we highlight the need for improvement by more accurately emulating the grouping process of humans. We have shown that a significant improvement over structural network clustering can be achieved by the additional consideration of tie strength. On average, for participants who created groups using criteria analogous to network clustering and then divided groups based on tie strength, 66.9% of contacts could be accurately grouped using automation. Some participants felt that they required particular groups to be divided into subgroups of strong and weak ties, whereas others did not. An algorithm with a model of tie strength might be able to automatically divide clusters into groups of strong and weak ties, however, it would not know which clusters the user would like to divide in this way. This illustrates the need for user intervention at certain points of the automation process to make decisions based on information unavailable to the machine and to ensure that control ultimately lies with the user.

We established that there is a significant linear relationship between the modularity of our participants' networks and the similarity of the groups created by humans and by the algorithm. Our results suggest that network modularity influences whether users consider the formation of network clusters when grouping contacts. The participants with the lowest network modularity did not consider the criterion that is most strongly associated with network clusters, social circles and cliques, and instead based their grouping entirely on other factors such as tie strength.

Research by both boyd [3] and Fisher [7] has demonstrated that people are very good at reading and interpreting their own social graph. Our findings suggest that users may also have a basic awareness of the modularity of their own social graph when considering how to group contacts. Participants with clear, distinct network clusters tend to group their contacts such that these clusters are reflected in their groups. Those who are unable to separate clusters easily instead resort to grouping using other criteria. Network clustering as an automated approach to group creation for privacy control is therefore most feasible for participants with high network modularity.

We also found that participants had difficulty grouping particular contacts during the card sorting exercise, either because they had a weak association with the created groups or they had strong associations with multiple groups. We found that the majority of problematic contacts were identified as hubs or outliers by the SCAN algorithm. This implies that our participants are somewhat aware of a contact's structural relationship with clusters in their network. When analyzing social network structure, we recommend using clustering algorithms such as SCAN that identify hubs and outliers within the network. While an automated approach to group creation would not be able to group these contacts reliably, it could at least recognize which contacts it should flag for the user to deal with, either when groups are initially created or when content is shared.

Our analysis of participants' willingness to share specific, privacy sensitive items of information/content with individuals from their network revealed that the detection of outlier vertices could also be advantageous. The concentration of structural outliers within the set of contacts whose willingness-to-share values were more than 2SDs from their group's mean was significantly higher than

their presence in the network as a whole. The same was not true for hubs, who had willingness-to-share values relatively uniform to the rest of their respective groups.

We attempted to establish whether the relatively static approach of creating groups to which privacy settings could be applied is feasible for accurately controlling the disclosure of information, given the dynamic nature of privacy. Lederer [15] finds that context is a secondary consideration, after the identity of the recipient. This suggests that groups based on the identities of contacts (i.e. the social circle to which they belong, strength of tie to the user, geographical location, organization etc) may at least provide a suitable default group substructure, to which adjustments can be made in a range of contextualized content sharing situations. Our analysis revealed that on average groups from the card sorting exercise resulted in between 77.8 and 90.8% of all contacts being assigned correct privacy settings when the settings were assigned at group level. We also demonstrated that this is likely to increase if we are able to account for a margin of neutrality in which users are unconcerned about the effects of sharing, however, more work is needed to understand the variance attributable to privacy intentions. There is also a substantial improvement if outliers are removed from groups altogether. Groups formed by the algorithm did not perform as well, with between 33.8 and 76.1% of contacts receiving the correct privacy settings. Full tie strength information for each of the contacts could allow an augmented algorithm to produce groups more similar to the card sorted groups and therefore improve their accuracy.

Although the groups that participants created could be used to provide accurate privacy settings for a fairly high percentage of contacts in a range of cases, there is still the possibility that a single unwanted recipient of sensitive content may nullify the benefits of the group-based approach. Allowing the user to easily adjust these groups and verify that they accurately reflect their privacy intentions is a problem that requires more attention.

7. CONCLUSION

This work takes a step towards providing group-based privacy controls for social networks that take account of how users naturally organize groups and which reduce the configuration burden for the user. We have uncovered six criteria that people consider when organizing their social network contacts into groups for the purpose of controlling privacy. We have also demonstrated that automated approaches which account for these criteria, such as detecting cliques within a social network and separating weak and strong ties, have the potential to reduce the burden of organizing contacts for social network users. Although our algorithm does not produce entirely accurate groups, this work is a valuable contribution to the debate about an emergent privacy related challenge. In the course of implementing and testing the SCAN algorithm we found that the detection of outliers within the network has strong potential to offer a real advantage in identifying potentially problematic contacts when using group-based sharing in a social network. Designers who wish to improve burdensome privacy controls could consider an automated approach to grouping contacts with the criteria we have identified. In our ongoing work we aim to investigate how automated tools might consider all six of these criteria.

8. ACKNOWLEDGEMENTS

Eamonn O'Neill's research is supported by a Royal Society Industry Fellowship. We would also like to thank Prof. Xiaowei Xu for providing the SCAN algorithm source code and David Pollington, James Irwin and Tom Lovett at Vodafone Group R&D for their input.

9. REFERENCES

- [1] Ackerman, M. and Mainwaring, S. (2005). Privacy Issues in Human-Computer Interaction. In L. Cranor and S. Garfinkel (Eds.), *Security and Usability: Designing Secure Systems that People Can Use*, 381-400, Sebastopol, CA, O'Reilly, 2005.
- [2] boyd, d. m. & Ellison, N. B. (2008). Social network sites: Definition, history and scholarship. *Journal of Computer-Mediated Communication*, 13(1), 210-230.
- [3] boyd, d., & Heer, J. (2006). Profiles as conversation: Networked identity performance on friendster. In *Hawaii international conference on systems science* (Vol. (HICSS-39)). Kauai, HI: IEEE Computer Society.
- [4] Clauset, A., Newman, M., and Moore, C. (2004). Finding community structure in very large networks. *Phys Rev E* 2004, 70:066111.
- [5] Davis, M., Canny, J., Van House, N., Good, N., King, S., Nair, R., Burgener, C., Rinehart, B., Strickland, R., Campbell, G., Fisher, S., and Reid, N. (2005). MMM2: mobile media metadata for media sharing. In *Proceedings of the 13th Annual ACM international Conference on Multimedia* (Hilton, Singapore, November 06 - 11, 2005). MULTIMEDIA '05. ACM, New York, NY, 267-268.
- [6] DiMicco, J. M. and Millen, D. R. (2007). Identity management: multiple presentations of self in facebook. In *Proceedings of the 2007 international ACM Conference on Supporting Group Work* (Sanibel Island, Florida, USA, November 04 - 07, 2007). GROUP '07. ACM, New York, NY, 383-386.
- [7] Fisher, D. (2005). Using egocentric networks to understand communication. *IEEE Internet Computing*, 9(5), 20-28.
- [8] Gilbert, E. and Karahalios, K. (2009). Predicting tie strength with social media. In *CHI '09: Proceedings of the 27th international conference on Human factors in computing systems*, New York, NY, USA, pp. 211-220. ACM.
- [9] Goffman, E. (1959). *The Presentation of Self in Everyday Life*. New York: Doubleday.
- [10] Granovetter, M. S. (1973). The Strength of Weak Ties. *The American Journal of Sociology*, 78(6), 1360-1380.
- [11] Hewitt, A. and Forte, A. (2006). 'Crossing boundaries: Identity management and student/faculty relationships on the Facebook', *CSCW06*, November 4-8.
- [12] Hogan, B. (2008). A Comparison of On and Offline Networks through the Facebook API (December 18, 2008). Available at SSRN: <http://ssrn.com/abstract=1331029>
- [13] Jones, Q., Gandhi, S. A., Whittaker, S., Chivakula, K., and Terveen, L. (2004). Putting Systems into Place: A Qualitative Study of Design Requirements for Location-Aware Community Systems, In *Proceedings of Computer-Supported Cooperative Work (CSCW '04)*, Chicago, ACM Press, pg. 202-211, 2004.
- [14] Lampinen, A., Tamminen, S., and Oulasvirta, A. (2009). All My People Right Here, Right Now: management of group co-presence on a social networking site. In *Proceedings of the ACM 2009 international Conference on Supporting Group Work* (Sanibel Island, Florida, USA, May 10 - 13, 2009). GROUP '09. ACM, New York, NY, 281-290.
- [15] Lederer, S. Dey, A. K., & Mankoff, J. (2002). "A conceptual model and a metaphor of everyday privacy in ubiquitous computing," Intel Research Berkeley, Tech. Rep. IRB-TR-02-017, 2002.
- [16] Lederer, S., Hong, J., Dey, A., & Landay, J. (2004). 'Personal privacy through understanding and action: five pitfalls for designers', *Personal and Ubiquitous Computing*, vol. 8, no. 6, 440-454.
- [17] McCarty, C. (2002). Structure in personal networks. *Journal of Social Structure*, 3.
- [18] Olson, J., Grudin, J., & Horvitz, E. (2005). A study of preferences for sharing and privacy. In: *CHI '05 extended abstracts on human factors in computing systems*, Portland, OR, USA, pp 1985-1988.
- [19] Skeels, M. M. and Grudin, J. (2009). When social networks cross boundaries: a case study of workplace use of facebook and linkedin. In *Proceedings of the ACM 2009 international Conference on Supporting Group Work* (Sanibel Island, Florida, USA, May 10 - 13, 2009). GROUP '09. ACM, New York, NY, 95-104.
- [20] Torres, G., Basnet, R., Sung, A., Mulkamala, S., and Ribiero, B. (2008). A Similarity Measure for Clustering and its Applications. *Proceedings of World Academy of Science, Engineering and Technology*, Vol.31 (JULY 2008) ISSN 1307-6884, pp.490-496.
- [21] Xu, X., Yuruk, N., Feng, Z., and Schweiger, T.A. (2007). SCAN: a structural clustering algorithm for networks. In *Proceedings of the 13th ACM SIGKDD international Conference on Knowledge Discovery and Data Mining. KDD '07* ACM, New York, NY. pp. 824-833.