Providing Awareness, Understanding and Control of Personalized Stream Filtering in a P2P Social Network

Sayooran Nagulendra and Julita Vassileva

MADMUC Lab
University of Saskatchewan
Saskatoon, Canada
{sayooran.nagulendra,julita.vassileva}@usask.ca

Abstract. In Online Social Networks (OSNs) users are often overwhelmed with the huge amount of social data, most of which are irrelevant to their interest. Filtering of the social data stream is the way to deal with this problem, and it has already been applied by OSNs, such as Facebook. Unfortunately, personalized filtering leads to "the filter bubble" problem where the user is trapped inside a world within the limited boundaries of her interests and cannot be exposed to any surprising, desirable information. Moreover, these OSNs are black boxes, providing no transparency of how the filtering mechanism decides what is to be shown in the social data stream. As a result, the user trust in the system can decline. This paper proposes an interactive method to visualize the personalized stream filtering in OSNs. The proposed visualization helps to create awareness, understanding, and control of personalized stream filtering to alleviate "the filter bubble" problem and increase the users' trust in the system. The visualization is implemented in MADMICA – a privacy aware decentralized OSN, based on the Friendica P2P protocol. We present the results of a small-scale study to evaluate the user experience with the proposed visualization in MADMICA.

Keywords: Online communities, Social networks, Social visualization.

1 Introduction

Today, with the enormous growth of Online Social Networks (OSNs) such as Facebook and Google+, millions of users are sharing social updates with friends and followers creating a "fire hose" of data in real-time. The updates vary from personal news (such as what's on their mind, what they are doing, what they are thinking of) to global news (such as news about politics, science, sports, technologies, etc.). If we consider the social data stream of a single user from her friends, only a fraction of it is relevant and interesting and the rest of the stream results in social data overload to the user. Personalized stream filtering mechanisms aim at solving these challenges of social data overload by presenting the user with the most relevant content. Social media sites such as Facebook, Digg and YouTube have already implemented personalized stream filtering which presents the most relevant content to users while reducing

P. Antunes et al. (Eds.): CRIWG 2013, LNCS 8224, pp. 61–76, 2013.

the social data overload. However, these systems are black boxes and provide no transparency or explanation, so users do not have any idea about what social updates that are hidden in the social data stream by the system and why they are hidden. As a result the user trust in the system can decline. Moreover, while attempting to personalize the stream with relevant content, in a long run the user can be trapped inside a world within the limited boundaries of her interests. This is called "the filter bubble" problem.

There are three key research questions that we are interested in answering by this research.

1. Is Visualization of the Filter Bubble an Effective Technique to Create Awareness, Understanding and Control of Personalized Stream Filtering?

The main purpose of personalized stream filtering is to reduce the social data overload by presenting only the relevant content. But showing what is hidden and filtered away in the stream can increase the social data overload problem. Therefore the main challenge is to find an effective visualization technique that can be seamlessly integrated into the activity stream without contributing additionally to the social data overload. What is the right amount of detail to expose in the hidden filtered social data and its explanation? How do we organize these hidden filtered social data? What type of visualization is effective to display the hidden social data stream? These issues can be explored through theoretical design and experiments with users.

2. Can a Visualization of Personalized Stream Filtering Increase the User's Trust in the Personalized Stream Filtering?

There is the possibility that some of the hidden filtered social data are being wrongly classified as undesirable. We believe that showing hidden filtered social data will provide transparency of the personalized stream filtering to the user and explaining them will build the users' confidence and will increase the user acceptance of the system.

3. Can a Visualization of Personalized Stream Filtering Alleviate "the Filter Bubble" Problem?

As the activity stream is personalized according to the user's interests, the user will ultimately only see activities related to her interest and will have no opportunity of discovering new interests. This will lead to "the filter bubble" problem where the user is trapped in a world filled with only items matching her interests. By exposing (some of the) hidden filtered social data, the user will become aware of the model that the system has of her, and may consciously decide to explore items from other areas by changing interactively her model and it will open the avenue for discovering new interests.

This paper proposes an interactive method to visualize the personalized stream filtering in Online Social Networks to create awareness, understanding, and control of personalized stream filtering to alleviate "the filter bubble" problem and increase the users' trust in the system.

2 Related Work

The social data overload problem is commonly solved by filtering out the irrelevant data. The problem of filtering out irrelevant data and providing personalized recommendation of data are addressed by Recommender Systems (RSs). RSs adapt to the needs of an individual user and provide personalized suggestions of most relevant information [12]. The personalized suggestions help users to make decisions on various types of items such as what book to read, what movie to watch and so on. Tandukar & Vassileva [15] developed an interest based filtering model which recommended relevant social data in the activity stream while filtering out the irrelevant social data to reduce the social data overload problem in a P2P social network. RSs provide recommendations using specific techniques based on background data, input data and algorithm. Collaborative filtering and content-based filtering are the main techniques used in most RSs [1]. Content-based filtering generates recommendations of new information using the history of information and the ratings previously given by that user. In collaborative filtering, recommendations are generated using only information about rating profiles for different users. Peer users with a similar rating history as the current user are identified and used for recommending new information.

Many researchers have worked on developing new RSs and improving the accuracy of their filtering algorithms. However the ultimate measure of success in this area is the user acceptance and trust of the recommendations and with respect to this measure there is still a lot of work that needs to be done [6]. The standard performance measures for RS are good when it comes to testing the recovery of missing data by RSs. But they cannot provide a valid method to test whether recommended data are valuable and previously unknown to the user. Providing a better user experience with RSs can increase user acceptance of recommendations. So user experience is becoming one of the most important current areas of research in RSs. The RSs must adapt and understand the needs of the users at different stages and provide not only valuable recommendations to the users, but also, as proposed by Chen & Pu [10] explanation interfaces which turn to be very effective in building the users' trust in the RSs. Previous research shows that explaining recommendations can increase the transparency of RSs and the users' trust in RSs [4, 18].

Explaining the rationale behind the recommendation is an important aspect of recommender systems. Explanations provide users with a mechanism for handling errors that might come with a recommendation. When we consider how we accept the recommendations provided by other humans, we recognize that other humans are imperfect recommenders. In case of the recommendations suggested by a friend, we might consider the quality of previous recommendations by the friend or we may compare that friend's interests with our interests in the domain. However, if there is any doubt, justification of the recommendation is needed and we let the friend explain it. Then we can analyze the explanation and decide whether to accept the recommendation or not [13].

Tintarev and Masthoff [17] describe three motivations for explanations in recommender systems: (1) transparency, which exposes the underlying logic of forming the recommendation so that the user can trust the system; (2) trust, which enables the user to consider the recommendation regardless of its accuracy level, and (3) scrutability, which enables the user to provide feedback on the recommendation to the system, so that the system can improve the future recommendations. Previous work on expert systems and automated collaborative filtering systems has shown that explanations can provide considerable benefit [13]. Work related to explanations can also be found in many other domains such as psychology, philosophy and cognitive science. Incorporating an explanation feature in recommender systems provides several benefits to users. It removes the black box from around the recommender system, and provides transparency. Herlocker et al. [3], mention some benefits provided by explaining recommendations such as: justification, user involvement, education and acceptance. Johnson & Johnson [4] have done research on explanations in human-computer interfaces.

The way recommendations are presented is critical for the user acceptance of recommender systems. Visualization techniques can be deployed to provide an intuitive "at a glance" explanation for recommendations and can also motivate the user to accept the recommendation. Presenting the recommendations in a ranked list according to their recommendation score is the most simple and commonly used visualization technique. Webster & Vassileva [19] proposed an interactive visualization of a collaborative filtering approach in RSs that allows the user viewer to see the other users in her "neighborhood", who are similar to her, and also to change manually to degree of influence that any of the other users can have on the recommendations of the viewer.

As a result of personalized filtering, the user can be trapped inside "the filter bubble" - a term introduced by Eli Pariser [9] to denote a limited scope of information defined by the user's interests and isolated from anything that doesn't belong to this scope. Resnick et al. [11] discuss the dangers of isolating users in filter bubbles and outline some strategies for promoting diverse exposure.

As discussed above, some approaches for increasing the transparency and the users' trust in RSs involve explanations or making the mechanism of recommendations visible to the user. Yet there haven't been approaches to visualize or explain the filter bubble problem. We propose an interactive visualization that presents a metaphoric view of the recommended and the hidden filtered social data in the personalized stream filtering in OSNs. The purpose of the approach is to alleviate the filter bubble problem and increase the users' trust in the filtered stream. The next sections present the design of the visualization and the results of a small scale user study with exploratory purpose.

3 Proposed Visualization

To achieve the goal of creating awareness, understanding, and control of personalized stream filtering in an OSN to alleviate the filter bubble problem and increase the users' trust in the system, we propose a visualization that metaphorically explains the filtering mechanism and provides means of control over certain parameters of the filtering for the users.

3.1 Visualization Design

The visualization is based on a bubble metaphor to make the effect of the personalized stream filtering in OSNs more understandable for the users (see Fig. 1). It divides the space of the screen in two parts - outside and inside the bubble. The items that are inside the bubble are visible for the user, those outside the bubble are those that have been filtered away and are invisible in the stream (but they are shown in the visualization). The visualization provides two alternative points of view: one focusing on the user's friends and one focusing on the categories of the social data originating from them in the OSN.

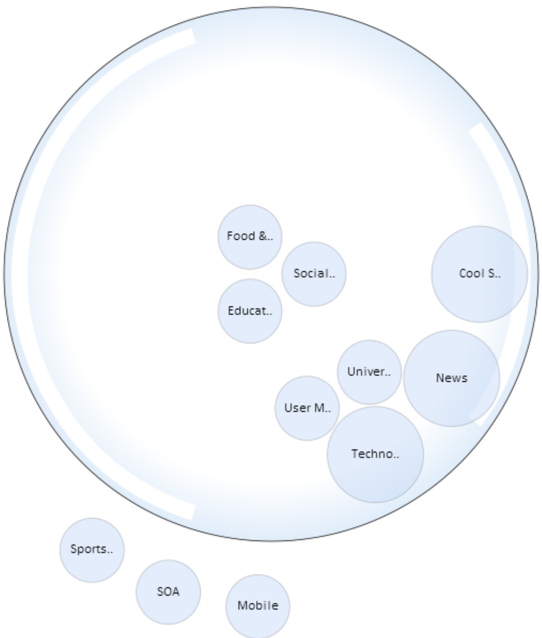

Fig. 1. Filter bubble visualization - category view

Fig. 2. Filter bubble visualization - friends view

Fig. 1 shows the "Category view", where in the bubble are shown the categories of social data that are not filtered away and can be seen in the user's stream (e.g. "Food & Health", "Education", "Cool Stuff", "News"). Outside the bubble are shown the categories of social data that are currently being filtered away from the user stream ("SOA", "Sports" and "Mobile"). The bubble shape design was chosen not only because it fits well with "the filter bubble" metaphor, but because it is scalable to accommodate more circles inside. Making the users aware of the different categories in which status updates are classified provides some transparency of the mechanism, which otherwise users won't be aware of. In essence, the category view summarizes what categories of social data the user is interested in and what categories of social data she tends to ignore in her stream. In addition, the abstract category view scales better than showing the specific updates and does not lead to an overcrowded view and cognitive overload. Upon clicking on a circle representing a given category, a small pop-up window shows the list of social updates from the stream that belongs to the category. In this way, for example, by clicking on the “Mobile” circle shown in Fig. 1, the user can see all the status updates from her OSN stream related to the “Mobile” category, that have been hidden from her. Thus we follow Shneiderman’s [14] visualization design principle “overview first, details on demand”.

The second view, called “friends view” (see Fig. 2), shows in a similar way the bubble, but instead of circles representing categories of social data, the circles represent the user’s friends who have posted the social data. If a friend’s circle is inside the bubble, then the social data from that friend are visible in the user’s stream, whereas if the friend circle is outside the bubble, the social data from that friend are hidden and not displayed in the stream. Since the filtering mechanism differentiates the filtered data both based on who the data comes from and the category of the data, the friends view displays the relationship that the user has with each of her friends with respect to a given category. Thus the user has first to select a specific category from a drop-down menu on the top of the screen (see Fig. 3), and then sees which of her friends are inside the bubble for this category, i.e. who she is connected with respect to the chosen category. These are the people whose social updates in the selected category the user is seeing in her stream; the updates in this category of the other friends who are outside the bubble are being filtered away from the stream. In order to provide a better understanding of what is happening in the personalized stream, the location of category/friend circle (inside or outside the bubble) represents the visibility of social data in your stream. For both views, the size of the category/friend circle denotes the number of social updates in a certain category or by certain friends, and it helps to understand the relative proportion of social updates that are visible versus those that are hidden as well as who is posting more social data and who is posting less.

Another feature of the visualization design focuses on giving users control over the stream filtering process. This is done by allowing users to drag and drop the category/friend circles inside and outside the filter bubble. Depending on which current view is selected (category view or friend view), dragging a category or friend circle inside the filter bubble enables the users to see updates from a category which appears interesting, but so far has been filtered away, or to strengthen the relationship with a friend whose social data have been not visible in the stream due to the personalized filtering process. In reverse, when users drag a category/friend circle outside the bubble, the social data belonging to that category or from that friend will not appear in the stream anymore. This helps the users to get rid of uninteresting social data and also to avoid spammers who flood the stream with uninteresting and unwanted social data.

3.2 Implementation

MADMICA [8] is an implementation of a privacy-aware decentralized (peer-to-peer) OSN using the Friendica open source framework [7]. MADMICA implements an approach to filtering social data, according to a model of the strength of the user's interests in different semantic categories overlaid over a model of their social relationships, which was originally developed and evaluated in a simulation [15]. The intuition behind the filtering approach is that two people can be friends, but not share the same level of interest in different topics or categories and not trust each other's judgment with regard to these categories. In essence, the filtering approach is based on a model of the user's interest in a finite set of categories of social data that is overlaid with a model of the strength of user interpersonal relationships (over each category). It consists of a matrix of relationship strengths (values between 0 and 1) between the user and each of her friends in different areas of interest. The model is updated based on implicit and explicit feedback from the user, based on the user actions over the social data (e.g. rating, commenting, forwarding or ignoring). The filtering of social data depends on the value of the strength of the relationship between the two users. The current relationship strength between a user and her friend in a given category is compared to a certain threshold value (currently a constant for all users in the OSN, but this could be personalized in the future) by the filtering algorithm to decide whether a new social update from this friend in the given category should be shown in the user's stream, or hidden. More details about the filtering approach can be found in [16].

MADMICA (http://madmica.usask.ca) is built with PHP, jQuery and MySQL technologies. The technology used to implement the visualization is HTML 5 with jQuery. The code can be run by any device on a browser without any plugin and can be adjusted to fit any size screen in a graphically pleasing manner [5]. The visualization is implemented in MADMICA as a plugin. This ensures that the modularized plugin architecture of MADMICA is preserved. So the user of each MADMICA node has the ability to turn off the plugin so the visualization. Users are notified in a side menu next to their stream with a message "Do you know this? N posts from your friends are hidden in your news feed based on your interest. Please, click on the bubble below to see them!". This creates awareness to the users that filtering is happening in the stream and some social data are not shown in the stream. When users click on the small bubble icon, the visualization plugin is loaded. When loading the visualization, all shapes are generated on the HTML5 canvas using KineticJS framework according to the data retrieved from the database. The visualization view is updated instantaneously and it always shows the category/friend circles according to the newest value from the user's relationship model. The default view is category view. Stored procedures have been used in MySQL to speed up the loading of visualization with necessary data.

The visualization can be viewed based on three different filters: bubble view, friends/category, and time period (see Fig. 3). This provides flexibility for the users to choose the desired view, and a time period of interest, since their interests in different categories and their relationships with friends are dynamic. The Bubble view filter consists of a dropdown menu that allows the user to select one of two views: category view and friend view. When the "category view" option is selected, a dropdown list is loaded in the Friends filter containing all the user's friends, so that she can individually select a friend and view all the semantic categories of social data that the user shares with this friend (i.e. shared interest) inside the bubble and those categories with respect to which the user and this friend do not share interest (outside the bubble).

In this view (selected category and selected friend), the circles representing the categories will appear positioned either inside or outside the bubble, based on the relationship strength value with that friend on that category.

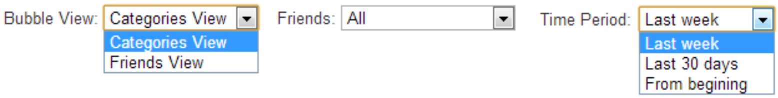

Fig. 3. Filters on the visualization view

To handle the problem of category name not fitting into the small circles, we display only a fraction of text inside the circle and the full text is shown as a tooltip when the mouse pointer is hovered over a category circle. To see the hidden social data for a particular category, the user has to click on the category circles which are outside the filter bubble and a pop-up window is loaded with the individual links to the social data (see Fig. 4). By clicking each link on that menu, the user can see the individual social data items which have been hidden in the stream by the filtering mechanism.

Fig. 4. Screenshot of hidden posts pop-up window

The default option in the Friends filter is “All” which shows just all categories represented inside or outside the bubble depending on the average relationship strength in these categories across all the friends of the user. On the other hand, when the “Friends view” option is selected in the Bubble view filter, the second filter changes to “Category”, allowing the user to pick a particular category of interest (the default is again “All”). The circles in this case represent the user’s friends and contain the friend’s avatar or photo, and the name of the friend appears as mouse is hovered over it (see Fig. 2). In this case, the second filter shows a dropdown menu showing all the semantic categories available in MADMICA. The user can select a particular category and see which of her friends (from those who have posted social data in the selected category) are inside or outside her bubble. Both views can be generated based on a time period filter. This filter comprises several options: “from beginning”, “last 30 days” and “last week” as the dropdown list labels. By default “last week” is selected when the visualization is loaded which shows the categories/friends circles in/by which social data were generated during the last week.

To add control of the filtering, we have added the drag and drop feature so that users can drag a category/friend circle inside and outside the filter bubble to show or hide data from this category/friend. When dragging a category/friend circle inside the filter bubble, AJAX (Asynchronous JavaScript and XML) requests are generated from the visualization and the corresponding model values for the interest based relationships are updated in the database. Similarly, when dragging a category/friend circle outside the filter bubble, another set of AJAX requests are generated to save the data. To let the users know about the results of the drag and drop action, a message is displayed to the user informing about whether the social data will be made visible or hidden based on the users' action.

4 Evaluation

A qualitative study was carried out to evaluate the usability and user acceptance of the visualization and whether it achieves its goals of providing awareness, control and trust in the filtering mechanism in MADMICA. The subjects were 11 graduate students from our research lab who used the MADMICA system instead of Facebook to share interesting and research relevant links over a period of three weeks in March 2013. All participants were international graduate students (six female and five male) from various parts of the world (the Middle East, Asia, and Africa), with computer science background and all were very experienced users of social networks (Facebook).

4.1 Hypotheses

The goal of this small-scale user study was to find out if the visualization is usable, if it creates awareness and understanding of the personalized stream filtering mechanism and ability to control it to alleviate the filter bubble problem and if as a result it helps to increase the users' trust in the filtering. So the evaluation aims at testing the following hypotheses.

1. The visualization creates awareness, understanding and sense of control of the personalized stream filtering mechanism to alleviate the filter bubble problem.
2. The visualization increases the user's trust in the personalized stream filtering.
3. The visualization of filter bubble increases the users' satisfaction with the system.

4.2 Experimental Setup

Due to the small number of users and the fact that the users were lab students and knew each other well, privacy wasn't an issue, so for efficiency sake, we hosted only one peer node to support all the participants. Each participant was asked to register and create a profile on MADMICA. Then the participants added each other as friends and started sharing anything they found interesting with their colleagues over the course of 3 weeks. We chose 11 semantic categories to classify the social data (the classification into one of the categories had to be done manually by the user when sharing something new with their friends), but allowed users to create their own categories (subject to approval by an administrator). The categories were chosen based on the main research areas in our lab, such as, education & mentoring, user modeling, mobile technologies, social computing, SOA, and common interest areas, such as food & health, news, sports & games, technology, university news and cool stuff.

To keep the participants engaged and motivated to be active in the network throughout the study period, we provided monetary rewards for participation in the study. Also in the second week of the study, to boost user activity with respect to the visualization, a notification was posted on the main page of MADMICA, to remind users to check the visualization of the hidden and visible social updates.

At the end of the study, the participants were asked to answer a questionnaire. As this was a qualitative study, the questionnaire had mostly open ended questions enabling participants to provide free feedback and describe their own ideas or suggestions. Responses for the few closed questions in the questionnaire were given on a 10-point Likert scale. Both types of questions focused on finding out about the user experience related to the proposed visualization and about the usability of the visualization. All participants completed the final questionnaire.

In addition to the questionnaires results, the usage of visualization of filter bubble was tracked by the system in order to collect data about users' actions on the bubble such as viewing the filter bubble visualization, dragging category/friend circle inside the filter bubble and dragging category/friend circle outside the filter bubble.

4.3 Results

Based on the tracked data, the number of users who performed actions on the visualization, such as clicking on the bubble, dragging a category/friend circle inside, and dragging a category/friend circle outside was plotted for each day throughout the experiment (see Fig. 5). In the first week of the experiment, 19 click actions, 4 drag out actions and 12 drag in actions have been recorded. During the second week when the popup was introduced, the number of click actions has dramatically increased to 28 and while the number of drag outs remained unchanged, the number of drag in actions has doubled as the previous value. In the last week of the experiment, although there is a small decrease in the number of all actions compared to the previous week (21 click actions, 2 drag out actions and 19 drag in actions), the number of actions has increased comparing to the first week.

The questionnaire included a number of closed questions that we asked to get some quantitative data on important aspects of the visualization. A subset of those closed questions focused on evaluating user experience with filter bubble visualization. The results are summarized in Table 1. On average, most of the participants answered above 5 on the scale of 1 (very low) to 10 (very high).

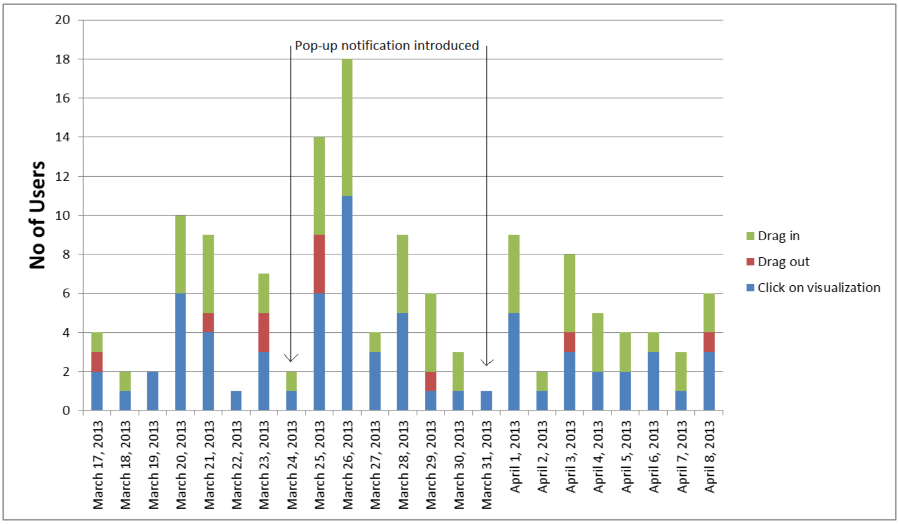

Fig. 5. Number of users performing three types of actions in the visualization

Table 1. Results of closed questions related to user experience of filter bubble visualization (#participants/percentage of all participants)

Question	**Results**									
Aesthetically pleasing	Very Low									Very High
	1	2	3	4	5	6	7	8	9	10
						1/9	3/27	3/27	3/27	1/9
Friend View	Unhelpful									Helpful
					2/18	1/9	3/27	3/27	1/9	1/9
Category View	Unhelpful									Helpful
					1/9		3/27	3/27	2/18	2/18
Awareness about hidden posts	Inadequate									Adequate
					2/18	1/9		4/36	3/27	1/9
Arrangement of information on screen	Illogical									Logical
		1/9			1/9		4/36	2/18	2/18	1/9
Manipulation of interest/friend circles(dragging in and out)	Difficult									Easy
				1/9	1/9	2/18	2/18	1/9	2/18	2/18
Finding interest not inside your filter bubble			1/9		1/9	1/9	2/18	2/18	3/27	1/9
Discovering new interests					2/18		2/18	2/18	5/45	

Table 1. *(continued)*

Discovering the interests of friends					1/9	1/9	3/27	2/18	4/36	
Discovering the areas your friends are most interested			1/9		2/18		1/9	3/27	4/36	

A set of close-ended questions with Likert scale (1-10) shown in Table 2 were asked to evaluate the users' trust in the system. The results are summarized in Table 3.

Table 2. Closed questions for trust in the system with Likert scale

#	**Question**
Q1	Trust in the System before using the filter bubble:
Q2	Trust in the System after using the filter bubble:
Q3	Trust in the System after seeing the hidden posts:
Q4	Level of transparency in filtering provided by the system:

Table 3. Results of closed questions for trust in the system (percentage of participants who chose on a 10-level Likert-scale)

#	Very Low									Very High
	1	2	3	4	5	6	7	8	9	10
Q1				18	36	9	18	9	9	
Q2					9	18	9	36	27	
Q3				9	18			18	36	18
Q4					9		18	45	18	9

The questionnaire contained a set of questions aimed at evaluating the users' awareness and understanding of the personalized stream filtering mechanism and the filter bubble visualization. Ten (91%) participants reported that they used the filter bubble visualization and one participant reported that s/he didn't use it. To the open-ended question "What do you think the visualization represents?", nine out of the ten participants who used the filter bubble visualization (90%) responded that they thought it represented their interest categories of social data that were displayed in their stream. Some excerpts from the answers follow: "*Shows my interests to different categories (category view) or to posts of friends (friend view)*", "*It represents my* *interest and posts I will receive", "It represents my interest category and that of others that is filtered from me"* and "*It reflects the interest a person showed in certain category of posts*". One participant mentioned specifically about the position of friend/category circles "*inside the bubble is the categories of the news I like while the hidden news belong to the categories outside the bubble, if friend view is selected, the same as category but for friends"* .

For the question "What do you think about the category view in the visualization?", three participants (27.27%) commented on what they understood about the category view: "*Category wise news/posts*" and "*I think category view is useful to visualize my choice of posts and help me to somewhat sort the posts I want to have a look on my wall.*" The remaining eight participants (72.73%) commented positively on the aesthetic aspect of the category view (e.g. *"nice, compact visualization", "good, and easy to use"*).

For the question about what participants thought about the friends view, three participants (27.27%) reported that they didn't use the friends view. Two participants (18.18%) said that it's an unnecessary view and they interpreted it wrongly. Three (27.27%) reported that it was useful to avoid friends' social data in which they were not interested. Three participants (27.27%) said that it was a good and useful visualization. To a control question asking them to indicate a preference to one or the other view, all of the participants replied that they preferred the category view over the friends view. Five participants (45.45%) were happy with the current views and didn't suggest any other useful views. The remaining six participants (54.55%) suggested several other useful views, such as *"a mixture of both", "more subcategories! But I wonder about the tradeoff with the simplicity", "time view! Popular view!", "By Date and week, and popular post -by like and comments"*, and so on.

The last few questions in the series of open ended questions aimed at evaluating the controls given to the user in filter bubble visualization: whether they were used (we could verify the answers as we had collected usage data, shown in Fig. 5), and whether they were considered useful and usable. The first question was about whether participants dragged the category/friend circles inside the bubble. Nine participants (82%) stated that they have dragged the category/friend circles from outside the filter bubble to inside the filter bubble. In a follow-up question, those who answered "yes" for dragging inside, were asked about the effect that they noticed after dragging a category/friend circle inside the filter bubble. Eight participants (88.89%) out of the nine participants said that there is an effect after dragging a category/friend inside the bubble. In particular, four participants out of those eight said that their interest areas expanded and more social data appeared in their stream. Only one participant out of those who tried dragging the circle inside said that there was no effect after the action. Similarly, a question was asked about dragging a category/friend circle outside the filter bubble. Four participants (36%) stated that they had tried dragging category/friend circle outside the filter bubble and noticed a change in their stream; particularly social data got filtered away. Other seven participants (63.64%) stated that they hadn't tried dragging a category/friend circle outside the filter bubble.

4.4 Discussion

The results show that the participants were aware of the filtering. The following results provide enough evidence to support the hypothesis 1: Most of the participants (80%) showed understanding about the representation of filter bubble visualization, knowing that the system is filtering their data stream (82%). The majority (73%) said that the visualization helped them to understand the filtering mechanism and more than 50% of the participants said that the visualization provided adequate awareness about the hidden social data. The participants' understanding of the graphical language of the visualization, i.e. the meaning of circle position and size, however was not uniformly good. The results show that 63.64% of the participants believed that there is a meaning in the position of the category/friend circle with respect to the filter bubble, so it is evident that the majority understood the general metaphor of the visualization. Even though eight participants (72.73%) responded that there is a meaning to the size of the circles, only two participants understood that the size denotes the volume of social data represented by the category or originated by the user represented by the circle. The remaining participants had various wrong interpretations of the size. So the design needs improvement with respect to using the size of the category/friend circles as part of the graphical language.

From the results of the open ended questions related to the category view and the friends view, we can see that the category view was more effective than the friends view in creating awareness and understanding of the personalized stream filtering and also the category view seems to be the most preferred view. So the Friends view needs to be improved, or removed. The results to the open ended questions that aimed evaluating the control given to the user to manipulate the visualization show that the participants felt they had control over their stream and the filtering mechanism. Thus we have sufficient qualitative evidence in support of hypothesis 1.

Hypothesis 1 can also be supported by the results of the user actions graph (see Fig. 5). The graph in Fig. 5 depicts the user actions performed on the filter bubble visualization over the time period of the experiment. The beginning of the graph period can be marked as the learning phase where users get familiar with the drag and drop of category/friend circles. Then there is a sudden spike in user actions in the second week when we introduced a popup window to notify the users that social data are filtered away from the stream and to introduce the visualization allowing them to gain control of the filtering. After one week, when the necessary awareness about the visualization has been created, the popup notification was turned off. Even after the notification was turned off, from the graph in Fig. 5, still we could see users checking the filter bubble visualization and dragging the circles in and out. This shows that the filter bubble visualization has been used to control of personalized filtering. Interestingly, most of the actions were "dragging in" categories or people, which means the participants counter-acted the filtering mechanism. There were a few "drag out" actions throughout the experiment and they were targeted at one particular participant, the most active one in the group, who was probably perceived as a spammer at a certain moments of high traffic by some of his/her friends.

The study results also provide evidence to support the hypothesis 2. Comparing the results of Q1 and Q2 provides more clear evidence to support the hypothesis 2 i.e. most of the participants (63%) rated below 6 (on a scale between 1-very low, 10- very high) for their trust in the system before using the filter bubble visualization. After seeing the filter bubble visualization, 72% of participants rated above 6 for their trust in the system. Moreover, 72% of participants rated high (above 7) for their trust in the system after seeing the hidden posts provided by the visualization and most of the participants (72%) rated the level of transparency as high (above 7).

The results shown in Table 1 provide answers to questions about the general user experience with the system. Following some user experience design guidelines [2], we consider user experience dependent on whether the artifact is aesthetically pleasing, logically composed and easy to use. They support hypothesis 3, because 90% of the participants found that the filter bubble visualization is aesthetically pleasing by rating it above 6; 90% found that category view was helpful, and 72% have found that the friend view was helpful. In addition, 72% of the participants found that the visualization provided adequate awareness about hidden social data, 81% of participants found that the information on the screen was logically arranged, 63% of participants said dragging the category/friend circles in and out of the filter bubble was easy, 72% said finding an interest which is not inside their filter bubble was easy, 81% said discovering new interests and discovering the interests of friends were also easy and 72% said that discovering in which areas their friends are most interested was also easy. So the results in Table 1 suggest that the user experience with the MADMICA was enhanced by the visualization. Moreover, the results showed that users were aware that they are able to find interests outside of their filter bubble and thus discover new interests that they didn't display otherwise in their behavior. This clearly shows that users became more aware of the filtering mechanism due to the visualization and are interested, able and willing to manipulate it to ensure that they will not be trapped inside a bubble world within the limited boundaries of their manifested interests.

5 Conclusion and Future Works

The paper proposes an interactive method to visualize the content-based stream filtering in a P2P Social Network. The proposed visualization helps to create awareness, understanding, and control of personalized stream filtering mechanism to alleviate the filter bubble problem and increase the users' trust in the personalization of the system.

The results of the small scale study show that the filter bubble visualization makes the users aware of the filtering mechanism, engages them in actions to correct and change it, and as a result, increases the users' trust in the system. Future work directions include finding a solution to the limited number of categories, and conducting a large scale user study.

References

1. Burke, R.: Hybrid Web Recommender Systems. In: Brusilovsky, P., Kobsa, A., Nejdl, W. (eds.) The Adaptive Web. LNCS, vol. 4321, pp. 377–408. Springer, Heidelberg (2007)
2. Garrett, J.J.: The Elements of User Experience: User-Centered Design for the Web and Beyond. Pearson Education (2010)
3. Herlocker, J.L., Konstan, J.A., Riedl, J.: Explaining Collaborative Filtering Recommendations. In: Proc. ACM Conference on Computer Supported Cooperative Work, CSCW 2000, pp. 241–250. ACM Press, New York (2000)
4. Johnson, H., Johnson, P.: Explanation Facilities and Interactive Systems. In: Proc. 1st International Conference on Intelligent User Interfaces, IUI 1993, pp. 159–166. ACM (1993)
5. Kaela, S.: The Role of HTML5 and Flash in Web Design (2012)
6. Konstan, J.A., Riedl, J.: Recommender systems: from algorithms to user experience. User Modeling and User-Adapted Interaction 22(1-2), 101–123 (2012)
7. Macgirvin, M.: DFRN – the Distributed Friends & Relations Network, `https://macgirvin.com/spec/dfrn2.pdf` (accessed August 2, 2012)
8. Nagulendra, S., Vassileva, J.: Minimizing Social Data Overload through Interest-Based Stream Filtering in a P2P Social Network. To appear in Proc. of the IEEE International Conference on Social Computing, SocialCom 2013 (2013)
9. Pariser, E.: The Filter Bubble: What the Internet Is Hiding from You. Penguin Press HC (2011)
10. Pu, P., Chen, L.: Trust Building with Explanation Interfaces. In: Proc. 11th International Conference on Intelligent User Interfaces, IUI 2006, p. 93. ACM Press, New York (2006)
11. Resnick, P., Munson, S.A., Garrett, R.K., Stroud, N.J., Kriplean, T.: Bursting Your (Filter) Bubble: Strategies for Promoting Diverse Exposure. In: Proc. of the Conference on Computer Supported Cooperative Work, CSCW 2013 Companion Proc., pp. 95–100. ACM (2013)
12. Resnick, P., Varian, H.R.: Recommender systems. Communications of the ACM 40(3), 56–58 (1997)
13. Ricci, F., Rokach, L., Shapira, B.: Introduction to Recommender Systems Handbook. In: Ricci, F., Rokach, L., Shapira, B., Kantor, P.B. (eds.) Recommender Systems Handbook, pp. 1–35. Springer, US (2011)
14. Shneiderman, B.: The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. In: Proc. 1996 IEEE Symposium on Visual Languages, pp. 336–343. IEEE Comput. Soc. Press (1996)
15. Tandukar, U., Vassileva, J.: Selective Propagation of Social Data in Decentralized Online Social Network. In: Ardissono, L., Kuflik, T. (eds.) UMAP 2011 Workshops. LNCS, vol. 7138, pp. 213–224. Springer, Heidelberg (2012)
16. Tandukar, U., Vassileva, J.: Ensuring Relevant and Serendipitous Information Flow in Decentralized Online Social Network. In: Ramsay, A., Agre, G. (eds.) AIMSA 2012. LNCS (LNAI), vol. 7557, pp. 79–88. Springer, Heidelberg (2012)
17. Tintarev, N., Masthoff, J.: Effective Explanations of Recommendations: User-Centered Design. In: Proc. ACM Conference on Recommender Systems, RecSys 2007, pp. 153–156. ACM (2007)
18. Wang, Y., Zhang, J., Vassileva, J.: Towards Effective Recommendation of Social Data across Social Networking Sites. In: Dicheva, D., Dochev, D. (eds.) AIMSA 2010. LNCS, vol. 6304, pp. 61–70. Springer, Heidelberg (2010)
19. Webster, A., Vassileva, J.: The KeepUP Recommender System. In: Proc. of the 2007 ACM Conference on Recommender Systems, RecSys 2007, pp. 173–176. ACM, Minneapolis (2007)