Accepted by IEEE Pervasive Computing, 2014

Toward a Group-Aware Smartphone Sensing System

Bin Guo[1], Zhiwen Yu[1], Daqing Zhang[2], Huilei He[1], Jilei Tian[3], Xingshe Zhou[1]

[1] School of Computer Science, Northwestern Polytechnical University, P. R. China
[2] Institut TELECOM SudParis, France
[3] Nokia

{guob, zhiwenyu, zhouxs}@nwpu.edu.cn, daqing.zhang@it-sudparis.eu, jilei.tian@nokia.com

Abstract

Nowadays, social activities are becoming more and more popular and important to human life. As the number of contacts increases, the implicit social graph becomes increasingly complex, leading to a high cost on social activity organization and activity group formation. This paper presents GroupMe, a group-aware smartphone sensing system that supports group management and activity organization in real life applications. We first present a systematic methodology that can steer the development of mobile group awareness applications. A multi-granular group model is then proposed. Based on the methodology and the model, we present our approaches that support a variety of user needs on group formation and management, including closed group suggestion, open/opportunistic grouping, and new group member suggestion. Experimental results verify the effectiveness of the proposed approaches.

Keywords: smartphone sensing, social graph, mobile group awareness, social activity organization

1. Introduction

A society acquires many of its characteristics from interpersonal and group interactions. In modern societies, people participate in various daily social activities. Depending on the distinct nature of an activity, different crowds of people are involved. We define people participating in a social activity as a *group*. For example, groups at universities can be referred to as project teams, dining partners, and co-players, among others. The reason for the formation of distinct groups for different activities is that human tend to be with a similar group of people to participate in certain activities. The interaction in groups greatly affects our work or living performance; thus, developing tools that can support group-related activities is important.

Technology already plays an important aspect in supporting group activities and interaction. More than two decades ago, the groupware concept has been proposed and applied in workplaces [1], which mainly investigates how groups work and seeks to discover how technology could enhance collaboration over existing groups. It has, nevertheless, paid little attention to how groups are formed. Several studies have recently been devoted to group management [2, 3] in online communities, such as emails, Facebook. Compared to online interactions, real-world group activities are more difficult to be captured. For instance, there basically lacks a preexisting infrastructure for physical activity logging and mining. Furthermore, activity organization in the real world is often influenced by various contexts like location and nearby people, which should be jointly considered in group service design. Rapid development in sensor-equipped smartphones has brought unprecedented opportunities for pervasive social sensing [4]. Thus, in the present study, we intend to explore the unique capabilities of smartphones to achieve real-world group awareness. We use the scenario to illustrate our motivation below.

*A university campus is a typical socially active environment. Harry, a second-year graduate student, initiates or participates in various social activities daily. In many cases, the formation of activity groups is roughly fixed or **closed**. For instance, Harry often takes his lunch with a subset of his contacts (with B, C, D on one day and B, C on another day). He creates groups by activity types to facilitate activity organization, but this approach is tedious and time-consuming. Thus, Harry realizes that he wants a tool that can automatically identify and manage his groups. However, merely having groups remains insufficient, and choosing from a large number of raw groups becomes burdensome (e.g., B,C,D and B,C) when initiating an activity. It is beneficial if highly relevant raw groups can be merged into logical units (e.g., combining the two raw groups into a logical one). In addition, Harry wants a tool that suggests highly relevant groups when an activity is initiated according to his contexts.*

*Foraging new connections (or **open** groups) is a social nature of human being. Harry at times plans to extend an activity group by inviting new members, and wants the tool to suggest like-minded, but non-linked people when organizing an activity. Compared with routing-like activities, groups for some activities are formed in a highly dynamic or opportunistic manner. For instance, Harry may choose to form opportunistic groups to achieve short-term, planned goals, e.g., grouping a number of skiing lovers from the campus to visit a ski resort during the weekend.*

Several challenges arise as a system is developed to support the above requirements. The key challenge is *how to develop a fundamental methodology that supports context-aware computing regarding to real-world group activities*. From the scenario we find a variety of demands for group management, such as closed/open group formation, group suggestion, and so on. It is beneficial to study a generic methodology that can steer the development of varied group-aware applications. The next challenge is *how to accurately model and extract human groups*. People usually participate in several groups, which are often overlapping and nested. Extraction of logical groups by merging large number of raw groups also needs to be explored. The third challenge is *how to meet diverse user needs on group formation*. For closed groups, the issue is how to recommend highly relevant groups to the user when forming an activity. For open groups, the issue refers to new member suggestion and opportunistic group formation.

To address these challenges, we propose a group-aware smartphone sensing system called GroupMe. The system exploits smartphones to capture human interaction, assist group management, and support activity organization in real-world scenarios. Specifically, our contributions include the following:

- *A systematic methodology for mobile group awareness*. It consists of a four-phase lifecycle and a context ontology for mobile groups. To the best of our knowledge, this is the first work that investigates the generic methodology for group awareness.
- *Multi-granular group modeling*. We propose a social graph-based model that can characterize the network of social activity participation at multiple granularities.
- *Varied real-world group formation*. Algorithms are proposed for a combination of closed and open group formation. Closed groups can be suggested by considering the surrounding contexts and the adhesion of each group to a user. Both centralized and opportunistic, delay-tolerant methods are explored for open group formation.

2. Related Work

Effective grouping of human contacts is crucial for interpersonal communication. ContactMap provides an editable visualization of personal contacts, spatially organized and colored by group membership [5]. Researchers from Google have proposed a friend-suggestion algorithm that can generate a recipient group when composing e-mails, given a small seed set of contacts [3]. MacLean et al. developed a social group browser called SocialFlow [2], which can show social groups automatically mined from e-mail data. These systems can extract social groups from online interactions, but fail to address group management and activity organization in the real-world.

This paper discusses two approaches to open-group formation: *suggesting new members to existing groups* and *building opportunistic groups*. Existing studies on the former mainly concern recommending friends to individuals based on user similarity (e.g., common profile). In our study, user similarity is measured by the interaction history of these individuals with varied real-world activities. We further propose an improved similarity metric to address the data sparsity problem. Several studies have focused on opportunistic-group formation. For example, Flocks [6] is a system that supports dynamic group creation on the basis of user profiles and physical proximity. MobilisGroups [7] is a location-based group creation service that allows users to initiate social events on the map and recruit nearby participants. Despite the facility of group formation in real-world settings provided by these systems, these systems mainly aim to group people already located nearby and not to recruit like-minded contacts who are not yet gathered but should be. We propose a delay-tolerant approach to organizing planned activities, which varies from these instant grouping methods.

3. A Systematic Methodology for Mobile Group Awareness

Context-aware computing has developed as an important research branch of pervasive computing since late 1990s. Its objective is to make the pervasive systems and services more intelligent by considering the relevant context which is not taken into account in the system design before. Afterwards, with the development of social computing, the social aspect of context caught the attention of the research community [8]. However, there still lacks a generic methodology that drives the development of context-aware group systems.

The building blocks of context-aware systems [9, 10], such as context ontology, context query/reasoning, and context filtering, have been developed. But there are new features that should be further studied for group-specific context awareness. First, a group activity, by its nature, has a running lifecycle, which should be explored to logically link different group-aware applications. Second, group activities introduce new types of contexts that function at different phases of the lifecycle.

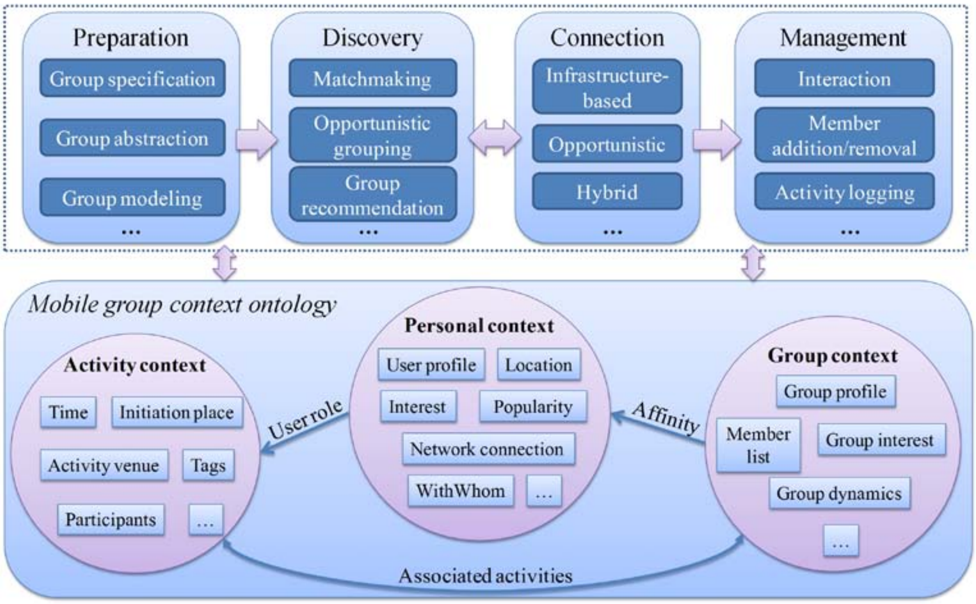

Fig. 1. Mobile group awareness: the lifecycle and supporting contexts.

(1) The lifecycle. In general, the lifecycle of a mobile group activity consists of four phases: preparation, discovery, connection, and management (as shown in Fig. 1). The key functionalities of each phase are described below:

- *Preparation*. Specification of group activities according to user needs; modeling of groups; extraction of logical groups from raw groups.
- *Discovery*. Discovery of people and devices for forming new groups; recommending closed groups for an activity; discovery of like-minded neighbors to form open, opportunistic groups.
- *Connection*. Connecting individuals to form groups, leveraging infrastructure-based, ad hoc, or hybrid networking techniques.
- *Management*. Adding individuals to groups or removing individuals from groups; managing interactions among group members; activity record logging and storage.

(2) The context ontology for mobile group awareness. In terms of the application domain, contextual information can be of different types [9]. In mobile group-aware systems, the context refers not only to the information about individual, but also about the social activity and related groups. Personal, group, and activity contexts are all indispensable to intelligent decision making in each phase of the mobile group lifecycle. For example, if one intends to form a group by discovering nearby users with certain attributes, personal context such as user location and interest would be essential; if one wants to suggest groups for organizing an activity, personal context such as WithWhom (companions) and group context such as group affinity to the initiator might be needed. We elaborate the three types of contextual information at the bottom of Fig. 1.

- ***Personal context***. It is defined as all the relevant information about an individual's situation, can be either static or dynamic. While dynamic personal contexts (e.g., location, WithWhom) change from time to time, static ones (e.g., profile, interest) change slowly with time.
- ***Group context***. It refers to the information that is related to a logical group, such as group profile (e.g., closed/open, cohesion), members, and group dynamics (e.g., the strength and evolution of interpersonal interaction in groups).
- ***Activity context***. Each real-world social activity includes an initiator, the initiation place (*I-Loc*) and time, the activity venue (*A-Place*), and a list of participants (*MemL*). Users may provide *tags* for an activity (e.g., dinner, meeting). We use the social activity logging (SAL) repository to keep activity relevant information. When an activity is initiated, the initiator sends an invitation message (e.g., SMS) to a group of contacts. The message is categorized into two types: SA_{in} and SA_{out}. All invitation messages are kept in the initiator's SA_{out} box, whereas received activity requests are kept in the SA_{in} box.

There are also associations among different type of contexts, such as user role (initiator or participant) in a group activity, the affinity (e.g., interaction frequency) of a group to a user, and the link between a group and its associated historical activities.

4. Group Preparation and Modeling

From this section we intend to make a showcase study to the methodology. Specifically, we will present the key techniques used to implement the scenario depicted earlier. A multi-granular group model in the preparation phase will first be presented.

4.1. Social Activity Logging and Group Modeling

We use the *social graph* to characterize the structure of the network of social activity participation. Edges are formed by sending or receiving activity requests. We employ the egocentric network method used in [3], in which a message sent by a user to a group of contacts is regarded as one that forms a single edge (a *hyperedge*). The edge is directed, represented as *in* and *out* edges (corresponding to SA_{in} and SA_{out} in SAL). Each hyperedge is referred to as an *explicit/raw group*. Figure 2 gives an example of *A*'s social graph, where three raw groups are involved (e.g., *G1* to *G3*). As presented earlier, the social graph of a person often consists of a set of overlapping (e.g., groups *G1* and *G3*) and nested groups (e.g., groups *G1* and *G2*).

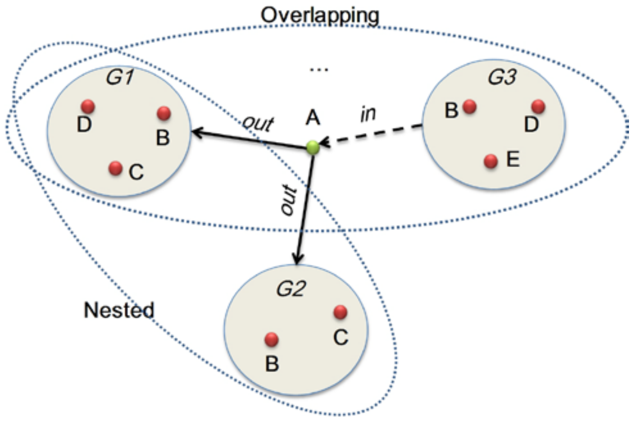

Fig. 2. An example of social graph.

4.2. Group Abstraction

People can participate in various social activities, and different social activities usually link different group instances. This relation results in a large number of groups in the initial social graph. We introduce the *group abstraction* process to eliminate minor subsets of groups by merging highly nested or overlapping groups into *implicit/logical groups*. We refer to the merging of nested groups as *group subsumption* and the merging of overlapping groups as *group integration*.

Group Subsumption. Given two nested groups, *G1* and *G2* ($G2 \subset G1$). The two groups can be subsumed if they are highly nested. We refer to MacLean *et al.*'s *information leak* metric for group nesting evaluation [2]. The value of information leak is determined by two factors: *similarity of the two groups* and *the ratio of the number of social activities held by each group*. We thus define a new parameter *subrate* to evaluate whether two groups can be subsumed, expressed in Eq. (1).

$$subtrate(G1, G2) = \frac{|G1| - |G2|}{|G1|} \times \frac{num(G2)}{num(G1)}, \quad when\ G2 \subset G1 \tag{1}$$

where $|Gi|$ refers to the number of members of group *Gi*, and $\frac{|G2| - |G1|}{|G2|}$ characterizes the similarity of two groups; *num(Gi)* refers to the number of social activities held by *Gi*. If the *subtrate* value is below a predefined threshold, the two groups can be subsumed.

Group Integration. The two overlapping groups can be integrated if they are very similar. We use the *Jaccard* metric to measure it (see Eq. 2). The two groups can be integrated if their similarity exceeds a threshold.

$$intrate(G1, G2) = \frac{|G1 \cap G2|}{|G1 \cup G2|}, \quad when\ overlap(G1, G2) \tag{2}$$

Group abstraction results in a set of logical groups according to user–group interaction history, which facilitates the management of groups.

5. Mobile Group Discovery and Management

People have varied needs on group discovery and management in real-world settings. Based on the methodology and the group model, this section presents our efforts to support closed and open group formation, corresponding to the distinct requirements defined in the scenario.

5.1. Mobile Group Recommendation

This application can recommend highly relevant groups to users, with two major factors considered: *context of the user* and *affinity between the user and the user groups*.

(1) Context-Aware Group Filtering. The various contexts that are identified when users initialize activities are used to filter irrelevant logical groups.

- *Time*: we divide the initiation time into four logical periods such as *morning* (6:00 to 11:00) and *noon* (11:00 to 13:00).
- *Location*: the place where the user initiates an activity. It can be obtained by *in-phone* GPS positioning or Wi-Fi indoor positioning techniques.
- *WithWhom*: nearby friends who are often co-initiators or members of an activity. We use *WithWhom (i)* to indicate that a number of *i* contacts are together with the initiator. This context can be obtained using the Bluetooth ID of user mobile phones.

(2) Group Affinity Ranking. It is used to calculate the tie strength between a user and the logical groups. In addition to interaction frequency, two other factors are considered:

- *Recency*. Human relationship is evolvable and dynamic over time.
- *User role.* The social activities in which the user is the initiator are considered more important than those in which the user is merely a participant.

We define the affinity rank between user *Ui* and the logical group *Gj* as *affrank(Ui, Gj)*, which can be calculated by Eq. (3). Given *Ui*, the implicit group with the highest rank is finally recommended.

$$affRank(Ui,Gj) = \omega_{out} \sum_{Ai \in SA_{out} \wedge G(Ai)=Gj} \left(\frac{1}{2}\right)^{d_{now}-d(Ai)} + \omega_{in} \sum_{Ai \in SA_{in} \wedge G(Ai)=Gj} \left(\frac{1}{2}\right)^{d_{now}-d(Ai)} \quad (3)$$

where ω_{out} and ω_{in} represent the weights of the user roles in social activities, with the former being larger to represent the importance of initiator roles. We empirically use 1.5 and 1.0 in the current implementation. d_{now} and *d(Ai)* refer to the current date and the initiation date of activity *Ai*, respectively.

5.2. New Group Member Suggestion

To address user needs on new group member suggestion, an approach that combines community detection and user similarity measurement is proposed. With community detection, a set of communities can be detected over the activity participation network. We use this result to distinguish the friends (in the same community) and the non-friends of a user. The popular *GN* algorithm is used in this approach [11].

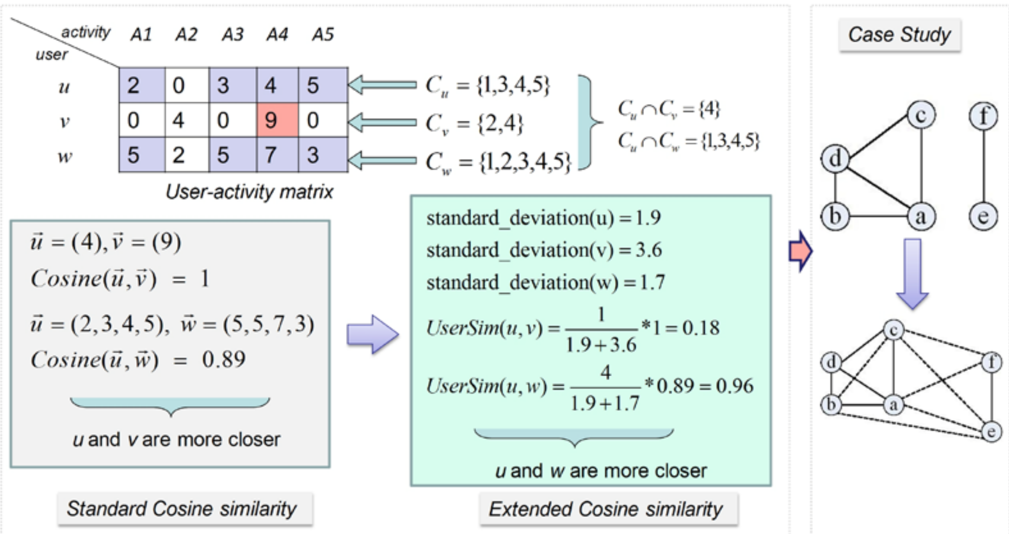

Fig. 3. User similarity measurement and new member suggestion.

To identify potential group members from the candidates in the non-friend user set, we need to measure the similarity among users. An extended *Cosine*-similarity method is proposed, which can calculate user similarity based on their activity participation history. We build the user-activity matrix according to the degree of user participation in each activity (see Fig. 3). The standard *Cosine* similarity between users can be derived from the matrix. However, one issue that needs to be addressed in similarity measurement is data sparsity, which may lead to poor prediction quality. Herlocker *et al.* [12] have suggested the addition of correlation significance weighting factors to devalue similarity weights according to the number of co-rated items. We use this method and extend the *Cosine* metric by introducing two factors:

- *Activity zone*. It denotes the number of activity types in which a user has participated. We represent the activity zone of user *u* as C_u and assume that if the common activity zone is larger for two users, the similarity between them is higher.
- *Behavior patterns*. The standard deviation (*SD*) is also considered to reflect user behavior patterns. Two users with a small variance are likely to appear more similar because they attend activities at regular patterns.

We further define the extended similarity between *u* and *v* as *UserSim(u,v)*, expressed in Eq. (4). Users from the candidate set who exhibit high similarity (the value is above a threshold called *SimThres*) to the user will finally be recommended. Figure 3 shows that distinct results when applying the two different metrics.

$$UserSim(u,v) = \frac{|C_u \cap C_v|}{SD(\vec{u}) + SD(\vec{v})} * Cosine(\vec{u},\vec{v}) \tag{4}$$

5.3. Mobile Opportunistic Grouping

In addition to the centralized methods mentioned above, the GroupMe system also supports the formation of opportunistic groups using opportunistic IoT techniques [13]. Opportunistic IoT addresses information dissemination and sharing within and among opportunistic communities that are formed based on the movement and opportunistic contact nature of humans.

For example, if a user called Harry plans to organize a skiing activity on a weekend, he can post a grouping request to GroupMe to recruit potential nearby participants (i.e., basketball fans). As Harry's activity continues, his request is shared by people in the vicinity. Match-making (each user keeps a list of interests) is then executed and matched users are added to the list of group members. Given roughly fixed moving range and mobility pattern of Harry, to increase the number of grouping request receivers and hasten request dissemination, GroupMe employs other mobile nodes as "group brokers" to help store and forward Harry's request. Broker election is also conducted in opportunistic communities based on user popularity, which measures the number of people the broker can possibly meet within a given period. The "broker-switch" action will be performed once there is a more effective broker. Dissemination ends when *i*) the required number of participants is obtained, or *ii*) the pre-specified request dissemination time has expired. For instance, Harry hopes that the group can be created within three days. We define it as the *group formation expiry time*.

The crucial issue for opportunistic grouping is to design an appropriate broker-selection protocol to facilitate data dissemination. Two social metrics are used to measure the usefulness of candidate brokers: *popularity* and *effectiveness*. As a basic metric, *user popularity* chooses a new broker simply based on the predicted number of contacts the user may encounter within a given period. As an improved metric, *user effectiveness* additionally uses the contextual information obtained during the group creation process. Specifically, each user maintains a list of users that she is likely to meet, and the current broker maintains a list of already encountered users (i.e., the context). We then calculate the difference-set (*DS*) of the two user lists, the size of which is used to measure the effectiveness of a user in the broker election process. If the *DS* of a new encounter is higher than the current broker, *broker switch* will happen.

6. System Evaluation

As an intelligent system based on social interaction history mining, data collection becomes the basis for system performance evaluation. In the current stage, a combination of two methods can be applied for data collection: *smartphone logging* and *online blogging*. The former logs user activities occurring at any time and place by using a customized mobile application. The latter requests users to manually record their daily social activities on an online blogging Web page. Forty more students from the department of computer science were recruited to contribute data which were collected from March 2012 to July 2012.

We collected about 1,000 data records, including 22 types of predefined activities. The activities are broadly divided into three types: *working* (e.g., meetings), *relaxation* (e.g., parties, shopping) and *sporting*. The two most popular initiation places are the *lab* and the *student dormitory*. Experiments based on the dataset are presented below.

(1) Mobile Group Recommendation. To validate the effectiveness of the group recommendation algorithm, we employ two generic criteria: *Precision and Recall*. In the experiments, we chose 500 SAL records as the training set, and 100 as the test set. The *MemList* in the test records was regarded as the ground truth.

Table 1. Performance evaluation of the group recommendation approach

Context Groups	Precision	Recall
Baseline (Top-5)	22.9%	39.6%
Baseline (Top-10)	14.6%	44.8%
No Group Abstraction + Time +I-Loc	30.2%	35.1%
Time + I-Loc	58.2%	74.6%
Time + I-Loc + WithWhom (1)	68.1%	94.7%
Time + I-Loc + WithWhom (2)	81.1%	98.7%

Many contexts derived by mobile sensing are used to filter irrelevant groups. To evaluate the effects of different contexts, we have chosen four different groups of contexts, with *Time* and *I-Loc* as the basic group as well as *WithWhom(1),* and *WithWhom(2)* as additional elements in the other context groups. We have also introduced a baseline method, which recommends the top-*k* contacts (calculated based on co-activity-participation frequency) to the initiator in activity organization. Experimental results in Table 1 indicate that our method performs better than the baseline method. A bigger *k* (e.g., 10) in the baseline method can increase the system recall but decrease its precision. It proves the social phenomenon that human tend to be with a similar group of people to participate in certain types of activities, while not with the best-connected people for any types of activities. Another finding is that the *WithWhom* context performs more efficiently than the other two contexts, which means that identification of co-initiators of an activity results in a better group suggestion performance. Group abstraction is another contribution to group management, which can eliminate noisy groups and merge relevant raw groups with logical units. The experimental results suggest that group abstraction can greatly enhance the system performance (with the precision increases from 30.2% to 58.2%).

(2) New Group Member Recommendation. To verify the usefulness of the member recommendation approach, we analyze the evolution of activity participation-based social communities. We first detected the communities in April and provided new group member suggestion based on them. The result of community detection in July was used as the ground truth to validate the effectiveness of recommendations. We chose six first-year graduate students (labeled *a* to *f*) for a case study (as shown in Fig. 3), where two communities were formed in April: (*a*, *b*, *c*, *d*) and (*e*, *f*). The suggestions for *a* was *e*, and that for *f* was *a*, *c* in April; connections between them were observed in July.

We also conducted a large-scale study to measure the effectiveness of the extended cosine similarity metric. Douban[1] is an activity-based social network that can guide users to attend offline activities through online advertising. Users can give a 'want to attend' tag to the interested activities that they plan to attend, by which we can obtain a user-activity matrix to each user. There are ten more predefined activity types, such as exhibition, music, sporting, gathering, and so on. A dataset of 15,050 users and 45,561 activities was crawled during Jan. to Apr., 2013. We extracted 10,000 social links among the users from the dataset, which were used as the ground truth to measure the accuracy of activity-based link prediction (i.e., new link recommendation). The results under three different *SimThres* (40%, 60%, and 80%) are shown in the left of Fig. 4. We can find that the *extended* metric performs much better on link prediction than the *standard* metric under different settings.

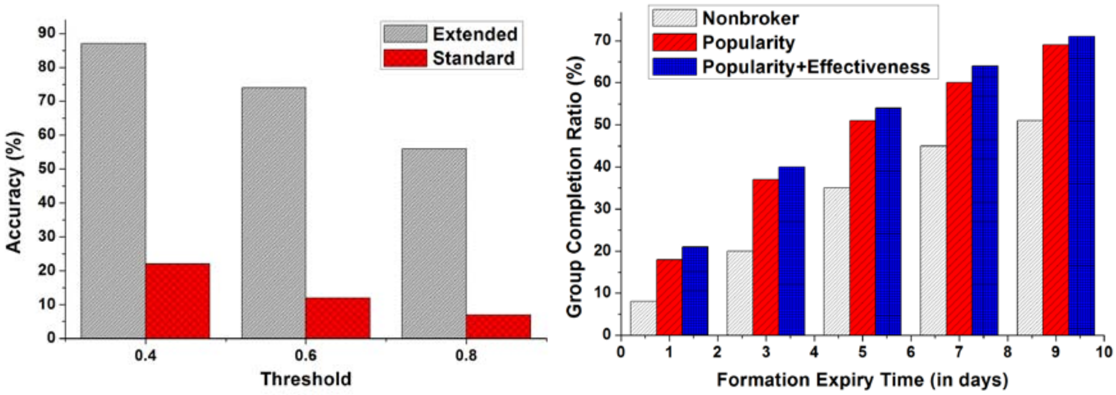

Fig. 4. Experiments results for member suggestion (left) and opportunistic grouping (right).

(3) Opportunistic Grouping. We used the MIT Reality Mining [8] dataset to evaluate the performance of opportunistic grouping. To make the dataset more manageable, we extracted twelve-week of collocation data. Specifically, the first eight weeks were used as the training dataset while the last four weeks were used as the testing dataset. In the experiment, we randomly generated 100 open group-formation tasks, where the task initiators and the start time of these tasks were selected randomly.

We tested opportunistic grouping under different group formation expiry time. Figure 4 (right) shows the experiment results of the three group-formation methods. We

1 http://www.douban.com, a service similar to MeetUp (http://www.meetup.com/)

measured them by calculating the group completion ratio (the average ratio of successfully completed tasks). The results indicate that better performance can be achieved when both social features are leveraged in the broker-selection approach.

7. Conclusion and Future Work

With the prevalence of sensor-enriched smartphones, GroupMe facilitates group awareness and activity organization in the real world. Historical group activity data sensed from the physical space can be used for group recommendation and open-group formation. Besides the functional purposes presented in this paper, we should further define the group-awareness objectives from the social perspective. For example, people often participate activities to strengthen the social ties particularly for those getting weak, while not only for the ones strong. It is thus needed to find new social metrics to measure and suggest human groups.

People apparently involve in both online and offline communities, and their behaviors within different spaces have been proven correlated [14]. Group awareness problems become more interesting and challenging when both online and offline activity data are considered. Questions on the online/offline convergence arise: whether online human groups often organize group activities offline, whether member suggestion for a real-world activity is possible by using online interaction data, and whether offline activities can benefit friend recommendation and information sharing in online social networks.

Acknowledgement

This work was partially supported by the National Basic Research Program of China (No.2012CB316400), the National Natural Science Foundation of China (No. 61332005, 61373119, 61222209), and the Program for New Century Excellent Talents in University (No. NCET-12-0466).

References

[1] A.C. Ellis et al., “Groupware: some issues and experiences,” *Communications of the ACM*, vol. 31, no. 1, 1991, pp. 39-58.

[2] D. MacLean et al., “Groups without tears: mining social topologies from email,” *Proc. of IUI’11*, 2011, pp. 83-92.

[3] M. Roth et al., “Suggesting friends using the implicit social graph,” *Proc. of the KDD’10*, 2010, pp. 233-242.

[4] N. J. Yuan et al., “We Know How You Live: Exploring the Spectrum of Urban Lifestyles,” *Proc. of ACM COSN’13, 2013*.

[5] S. Whittaker et al., “ContactMap: Organizing communication in a social desktop,” *ACM Transactions on Computer-Human Interaction*, vol. 11, no. 4, 2004, pp. 445-471.
[6] E.G., Boix et al., “Flocks: enabling dynamic group interactions in mobile social networking applications,” *Proc. of SAC’11*, 2011, pp. 425–432.
[7] R. Lubke et al., “Mobilisgroups: Location-based group formation in mobile social networks,” *Proc. of PerCom Workshops*, 2011, pp. 502–507.
[8] N. Eagle, et al., “Inferring friendship network structure by using mobile phone data,” *PNAS*, vol. 106, no. 36, 2009, pp. 15093-15094.
[9] G. Adomavicius, A. Tuzhilin, “Context-Aware Recommender Systems,” Recommender Systems Handbook, Springer, 2011, pp. 217-253.
[10] T. Gu, H.K. Pung, D. Zhang, “Toward an OSGi-based infrastructure for context-aware applications,” *IEEE Pervasive Computing*, vol. 3, no. 4, 2004, pp. 66-74.
[11] M. Girvan et al., “Community structure in social and biological networks,” *PNAS*, vol. 99, no. 12, 2002, pp. 7821-7826.
[12] M.R. McLaughlin et al., “A collaborative filtering algorithm and evaluation metric that accurately model the user experience,” *Proc. of SIGIR’04*, 2004, pp. 329-336.
[13] B. Guo et al., “Opportunistic IoT: Exploring the Harmonious Interaction between Human and the Internet of Things,” *Journal of Network and Computer Applications*, vol. 34, no. 6, 2013, pp. 1531-1539.
[14] X. Liu et al., “Event-based social networks: linking the online and offline social worlds,” *Proc. of KDD’12*, 2012, pp. 1032-1040.