Poster Abstract: Profile-Cast: Behavior-Aware Mobile Networking

Wei-jen Hsu[a] *wjhsu@ufl.edu*
Debojyoti Dutta[b] *dedutta@cisco.com*
Ahmed Helmy[a] *helmy@ufl.edu*
[a]Department of Computer and Information Science and Engineering, University of Florida, Gainesville, FL
[b]Cisco Systems, Inc., San Jose, CA*

In this paper we propose a new communication paradigm called profile-casting where destination nodes are identified by their (inferred) properties, not by node IDs. As an example case study, we use the mobility profile to show a design of profile-casting protocol based on the similarity in user mobility processes. The similarity-based protocol has good potential in both overhead reduction and delivery ratio improvement compared with simple message delivery protocols, such as flooding or random transmission. It is also flexible to operate at various points of the overhead-performance tradeoff.

I. Introduction

Recent years have witnessed significant growth in the adoption of portable wireless communication and computing devices (e.g., laptops, PDAs, smart phones) and large-scale deployments of wireless networks (e.g., cellular, WLANs). We envision that future usage of mobile devices and services will be highly personalized. Users will incorporate these new technologies into their daily lives, and the way they use new devices and services will reflect their personality and lifestyle. This opportunity opens up the door for novel paradigms such as *behavior-aware* protocols and services. Such services look into user behavior and leverage the underlying patterns in user activity to adapt their operations and have the potential to work more efficiently and suit the real needs of the users. One classical example is the data mining efforts from various online stores (e.g., Amazon.com) that provide personalized shopping offers based on browsing history. However, little attention has been directed towards leveraging behavioral patterns for services or protocol design in the mobile computing paradigm.

In this paper we focus on a new class of service named *profile-casting*. In this service, instead of targeting at a particular end-point or host, the message is to be delivered to *all* hosts with a certain property (i.e., those who match with the specified *profile* are intended receivers). Potential applications of such a service are notification or advertisement for a scoped group within the general population, or a matching service trying to find people with certain characteristics or interests. Note that the notion of *profiles* refers to the implicit, intrinsic properties to be discovered from the behavioral patterns of users. This distinguishes *profile-casting* from traditional multi-casting where users join multicast groups explicitly with their network identities.

We take user *mobility profile* as a case study to demonstrate the usefulness of the *profile-casting* paradigm. Mobility is a key aspect of wireless networks, and the long-run trends in mobility also reveal the social contexts, which can then be used to categorize users [1]. We show that by incorporating mobility profiles of users, our *similarity-based* forwarding protocol limits the scope of message delivery in delay-tolerant networks (DTNs) [2] to a specific behavioral group. Thus we avoid the high overhead of the epidemic routing [4] (less than half of the overhead with a little reduction in performance) and outperform random-walk based protocols in terms of delivery ratio (30% better under similar overhead).

II. Preliminaries

Delay tolerant networks (DTNs) [2] are networks characterized by sparse, time-varying connectivity, in which end-to-end spatial paths from source to destination nodes are often not available. Messages are stored in intermediate nodes and moved across the network with nodal mobility. One particular important decision to make for nodes in DTN is whether to forward a packet to other nodes they encounter (i.e., move into the radio range). Such decisions have implications on many aspects of how efficiently the routing strategies work, such as delay, overhead, and message delivery rate. In this work we utilize similarity of user mobility profiles to guide such decisions.

*The work was done when Debojyoti Dutta was at Univ. of Southern California.

III. A Mobility PROFILE-CASTing Protocol

In the paper we consider the scenario where the message sender is interested in forwarding messages to users with a similar mobility profile to the sender itself. Such a need arises, for example, when a student loses a wallet and wishes to send an announcement to other fellow students who visit similar places often as he does to look for it. Or, it can be used for location specific announcements such as power shutdown in parts of campus directed only to patrons of the specific area. Note that this application is different from *geo-casting*, which targets at the nodes *currently* within a geographical region as the receivers. Our target receivers are nodes with a certain mobility profile, *regardless* of their actual locations at the time the message is sent. To enable such *profile-based casting* services, it is important to have a descriptive representation for user mobility profiles and a measure of the similarity between users, to guide the message forwarding decisions.

In our previous work [1], we analyze two large scale WLAN user association traces[5, 6] and represent user mobility in the form of *daily association vectors*. This representation captures the relative importance of locations on campus to each user (i.e., the *preference* in the user mobility process). Utilizing Singular Value Decomposition (SVD), we extract summarized mobility features from each user and propose a similarity metric to classify the whole user population into distinct behavioral groups. These groups correspond to users with distinct mobility profiles. In this work, we take these behavioral groups as the targets for mobility profile-based casting.

The details of our *similarity-based* profile-casting protocol is as follows: Users in the network are not aware of the centralized decision of user grouping based on the similarity in their mobility. Instead, they collect their own mobility history locally independent of one another, and when two users meet with each other, they exchange the summarized mobility profiles (i.e., the descriptive vectors with their relative importance (weights)) of their previous mobility history and decide whether they are similar on the spot[1]. If the similarity index is larger than a threshold, they exchange the message. The philosophy behind the protocol is, if each node delivers the message only to others with high similarity in the mobility profile, the propagation of the message copies will be scoped within a group of similar users.

[1]The similarity index is calculated as the weighted sum of inner products of the mobility profile vectors[1].

$$Sim(U,V) = \sum_{i=1}^{rank(U)} \sum_{j=1}^{rank(V)} w_{u_i} w_{v_j} |u_i \cdot v_j|, \qquad (1)$$

where u_i and v_j are mobility profile vectors from user U and V, and w_{u_i} and w_{v_j} are their corresponding weights.

IV. Evaluation and Comparison

We compare the performances of the following schemes with the *similarity-based* protocol : (1) *Flooding*: The nodes in the network are all oblivious to user mobility profiles and blindly send out copies of the message to nodes who have not received it yet. This scheme is also known as the *epidemic routing* [4]. (2) *Centralized*: In this ideal scenario, all nodes acquire the centralized knowledge of the behavioral group membership, and only propagate the message to others if they are in the same group. The message will never propagate to an unintended receiver. (3)*Random-transmission (RTx)*: The current message holder sends the message to another node randomly with probability p when they encounter, and never transmits again (i.e., only the node who last received the message will transmit in the future). Loops are avoided by not sending to the nodes who have seen the same message before. This process continues until a pre-set hop limit is reached.

We utilize the USC trace [5] to study the message transmission schemes discussed above *empirically*. We use the trace for user mobility and assume that two nodes are able to communicate when they are associated with the same access point. Note that the WLAN infrastructure is merely used to collect user location information, and the messages can be transferred only between the users without using the infrastructure, as in [3]. We split the WLAN trace into two halves. The first half of the trace is used to determine the grouping of users based on their mobility and we identify 200 groups of users with distinct mobility characteristics from the total of 5000 users. Then we evaluate the performances of the fore-mentioned protocols using the second half of the same trace. For each group with more than 5 members, we randomly pick 20% of the members as the source nodes sending out a *one-shot message* to all other members in the same group.

The performance metrics used are as follows: (1) *Delivery ratio*: The number of nodes receiving the message over the number of intended receivers (i.e., nodes belonging to the same group). (2) *Delay*: The average time taken to deliver the messages to recipient nodes. (3) *Overhead*: The total number of transmissions involved in the process of message delivery.

We choose *flooding (i.e., epidemic routing)* as the baseline for our evaluation and show the relative performance of the other protocols relative to that of the epidemic routing in Fig. 1. In the graph we see that *flooding* has the lowest delay and the highest delivery ratio as it utilizes all the available encounters to propagate the message. However, it also incurs significant overhead. The average delay, which is the lowest possible under the given encounter patterns, is in the order of days (3.56 days in this particular case). Profile-casting based on the *centralized* user grouping information, the ideal scenario, shows a great promise of behavior-aware protocols, as it significantly reduces the overhead while maintains almost perfect delivery ratio, with a little extra delay. However, it is not realistic to assume such centralized knowledge.

For the *similarity-based* protocol, its aggressiveness can be adjusted with the forwarding threshold of the similarity index. Experiment results show a significant reduction of the overhead (only 2.5% of *flooding*) at the cost of the delivery ratio if we set a high threshold such as 0.7 (i.e., sending almost exclusively within the same group). Setting a low threshold (e.g., 0.5) leads to a better delivery ratio (92% of *flooding*) but still cuts the overhead to 45% of *flooding*. For the *RTx* protocol, although the overhead can be controlled with the hop-limit (which we set as m times of the group size), we see that the delivery ratio is lower than that of the *similarity-based* protocol with comparable overhead (comparing *similarity* 0.6 with *RTx* $m = 9$, the former has a 30% higher delivery ratio than the latter) because in many cases the message is transmitted to some node out of the desired group and there is no knowledge to direct its propagation. Further more, the average delay for the delivered messages is much longer than in the other protocols where multiple copies of the message propagate in the network. In addition, we try the *RTx* protocol with various p and m values and find it is not as flexible as the *similarity-based* protocol in which the parameters can be tuned to trade overhead for better delivery ratio.

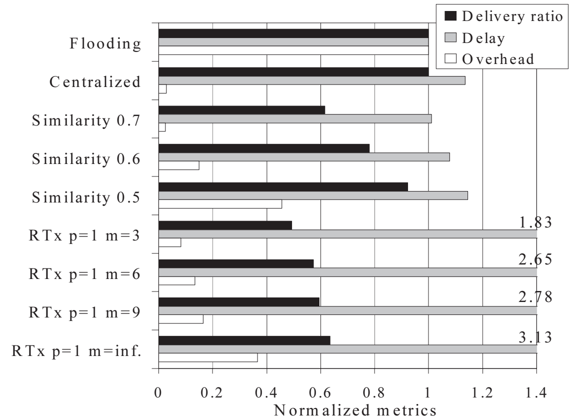

Figure 1: Relative performance metrics of the compared schemes normalized to the performance of *flooding*.

V. Conclusion

In this paper, we propose *profile-casting* as a new service paradigm. We demonstrate that similarity in user mobility can be utilized for scoped message dissemination in DTNs and show improved performance over other candidates (i.e. the epidemic routing or random transmission). This points out the insight from a detailed study of user behavior might provide new directions to improve services and protocols, especially as services become highly personalized in the future mobile networks. Our future directions include extending the concept to the design of other services, not only for mobility profiles, but also other social attributes, and experimenting with various traces or testbeds.

References

[1] W. Hsu, D. Dutta, and A. Helmy, ”Extended Abstract: Mining Behavioral Groups in Large Wireless LANs” In Proceedings of ACM MOBICOM, Sep. 2007. Longer technical report available at http://arxiv.org/abs/cs/0606002

[2] K. Fall, ”A Delay-Tolerant Network Architecture for Challenged Internets,” In Proceedings of ACM SIGCOMM, August 2003.

[3] J. Leguay, T. Friedman, and V. Conan, ”Evaluating Mobility Pattern Space Routing for DTNs,” in Proceedings of IEEE INFOCOM, April, 2006.

[4] A. Vahdat and D. Becker, ”Epidemic Routing for Partially Connected Ad Hoc Networks,” Technical Report CS-200006, Duke University, April 2000.

[5] W. Hsu and A. Helmy, MobiLib USC WLAN trace data set. Downloaded from http://nile.cise.ufl.edu/MobiLib/USC_trace/

[6] D. Kotz, T. Henderson and I. Abyzov, CRAWDAD data set dartmouth/campus/ movement/01_04 (v. 2005-03-08). Downloaded from http://crawdad.cs.dartmouth.edu/dartmouth/campus/movement/01_04