



Contents lists available at ScienceDirect

Pervasive and Mobile Computing

journal homepage: www.elsevier.com/locate/pmc

Fast track article

Exploiting user interest similarity and social links for micro-blog forwarding in mobile opportunistic networks

S.M. Allen^a, M.J. Chorley^{a,*}, G.B. Colombo^a, E. Jaho^b, M. Karaliopoulos^b, I. Stavrakakis^b, R.M. Whitaker^a^a School of Computer Science & Informatics, Cardiff University, Cardiff, UK^b Department of Informatics and Telecommunications, National & Kapodistrian University of Athens, Athens, Greece

ARTICLE INFO

Article history:

Received 9 March 2011

Received in revised form 13 December 2011

Accepted 20 December 2011

Available online 27 December 2011

Keywords:

Micro-blogging

Social networks

Pervasive

Mobile

Content dissemination

ABSTRACT

Micro-blogging services have recently been experiencing increasing success among Web users. Different to traditional online social applications, micro-blogs are lightweight, require small cognitive effort and help share real-time information about personal activities and interests. In this article, we explore scalable pushing protocols that are particularly suited for the delivery of this type of service in a mobile pervasive environment. Here, micro-blog updates are generated and carried by mobile (smart-phone type) devices and are exchanged through opportunistic encounters. We enhance primitive push mechanisms using social information concerning the interests of network nodes as well as the frequency of encounters with them. This information is collected and shared dynamically, as nodes initially encounter each other and exchange their preferences, and directs the forwarding of micro-blog updates across the network. Also incorporated is the spatiotemporal scope of the updates, which is only partially considered in current Internet services.

We introduce several new protocol variants that differentiate the forwarding strategy towards interest-similar and frequently encountered nodes, as well as the amount of updates forwarded upon each encounter. In all cases, the proposed scheme outperforms the basic flooding dissemination mechanism in delivering high numbers of micro-blog updates to the nodes interested in them. Our extensive evaluation highlights how use can be made of different amounts of social information to trade performance with complexity and computational effort. However, hard performance bounds appear to be set by the level of coincidence between interest-similar node communities and meeting groups emerging due to the mobility patterns of the nodes.

© 2011 Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/3.0/>).

1. Introduction

Online micro-blogging services have become very popular in recent years. The basic idea behind them is to allow users to post short messages and automatically receive updates from other specific users who they decide to 'follow'. Although the 'follower' relationship is reminiscent of traditional online social networks (OSNs), it is also substantially dissimilar from a typical online friendship in that links are essentially uni-directional and may not be reciprocated. Also, users do not necessarily receive any kind of information about their followers and their interests, while followers can be blocked from receiving updates if so desired.

* Corresponding author.

E-mail addresses: stuart.m.allen@cs.cardiff.ac.uk (S.M. Allen), m.j.chorley@cs.cardiff.ac.uk (M.J. Chorley), g.colombo@cs.cardiff.ac.uk (G.B. Colombo), ejaho@di.uoa.gr (E. Jaho), mkaralio@di.uoa.gr (M. Karaliopoulos), ioannis@di.uoa.gr (I. Stavrakakis), r.m.whitaker@cs.cardiff.ac.uk (R.M. Whitaker).

<http://dx.doi.org/10.1016/j.pmcj.2011.12.003>

1574-1192/© 2011 Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/3.0/>).

Thanks to user-originated features such as ‘retweeting’ (the forwarding of a received tweet, allowing its spread far beyond the set of followers of its original source), micro-blogging has become an effective tool for information diffusion, similar to news media services [1]. Additional user-originated features include a form of tagging (hash-tag) that allows categorisation of updates by topic. Micro-blogging has found applicability in emergency scenarios (updates during riots in Kenya, Egypt, Iran and Libya), and facilitated information dissemination at institutional level or in critical situations (help during large-scale emergencies, live updates to track traffic delays), thus serving as a powerful instrument of cooperation [2,3]. It has also been recommended as an effective alternative for reducing overload in working environments [4].

Capabilities such as targeting a specific user in a post (reply), or sending direct messages to users suggest a (weak) definition of ‘friendship’ between users that have participated in a given number of these more direct interactions. Alternative interpretations consider ‘friends’ those users for whom the follower relation is reciprocated [1]. According to an exploratory study of Twitter usage [5], it is the compulsory brevity of the updates that further allows the reader to effectively filter large numbers of messages. This feature reduces the cognitive threshold for the writer to decide to share and the burden for the reader to process all updates. Because of their particular requirements in terms of message size and their purpose to inform, warn, share and offer opinions, micro-blogs have been often compared to the concept of ‘utterance’ in linguistics [6,7,3,8].

Although micro-blogging has essentially been thought of as an online service, the particular structure of the induced followers’ networks [9] makes them an ideal mechanism for the rapid dissemination of information amongst ad-hoc social communities. The application of these services to the mobile domain suggests opportunities for sharing micro-blog posts directly among local devices that can send and receive content while taking into account the local or temporal context.

In this work, we explore scalable decentralised push-based protocols for micro-blogging using mobile devices in pervasive opportunistic environments [10]. According to our scenario, low payload micro-blogs (utterances) are generated by the devices (nodes), stored directly in their memory, and opportunistically exchanged upon their pairwise interactions. Our work extends the basic flooding concepts introduced in [8] by exploring in detail the role of similarity of interest between users. As a consequence of the peculiar nature of micro-blogging services, where forwarding an update does not imply any direct knowledge of the current status of followers, a user forwarding an utterance will have no knowledge about their follower’s stored micro-blogs or preference for particular content. A push-based strategy such as this must therefore follow some form of (pure or controlled) flooding strategy. However, to enhance these strategies’ performance by attempting to reduce the number of irrelevant utterances delivered to users, we allow them to store limited social context about a subset of users and use it to direct the forwarding of micro-blogs to encountered nodes.

The contribution we make is to explore different ways to use friendship, interest similarity and familiarity between nodes to better suit various possible real-life scenarios. This is important because it allows protocols to be adaptive to social content. In the baseline scenario, friendship links form between individual nodes having similar interests. Different interest similarity thresholds can modulate the selectiveness of nodes in forming such relationships. Information about interest profiles is dynamically shared within the resulting social groups of friends and used to prioritise the dissemination of content among them. Note that basing the social group exclusively on similar interests aims to reduce the computational effort for storing, processing, and selecting updates not closely related to a node’s own interests. At a second level, nodes may keep account of nodes they encounter frequently (*familiar* nodes). The relationships with these nodes, which do not necessarily share similar interests, can be seen as a further (weaker) form of friendship that can be exploited to disseminate content on behalf of other ‘familiar’ individuals.

We do not aim to deliver all relevant content to all nodes that may wish to receive it. Instead, we aim to deliver content that is interesting to a node given some spatial or temporal context. We do not guarantee that users will always receive all the interesting content that may be in the system for them, but aim to ensure that the content they do receive is of the most interest. We aim therefore to reduce the dissemination of irrelevant content, thus removing the need for users to filter the content they receive. In order to quantitatively assess the system’s ability to meet these aims we use two metrics describing information retrieval quality, precision and recall.

Although the protocol could be extended to fuzzier classifications, we assume in this work that each utterance is characterised by a well-defined topic (tag). Nodes gain positive utility when receiving an update that matches one of their individual interests (tags). The obtained utility also accounts for the spatial and temporal validity of micro-blogs. The local aspect, in particular, even if it is not originally considered in on-line services, is particularly suited to mobile pervasive scenarios.

We assess different strategies for selection and pushing of utterances and find that the most successful strategies take into account not only the individual interest profiles of friend nodes but also use a ‘community profile’ to push items to nodes that are not friends, but may be familiar due to repeated interactions. A performance tradeoff must be made between using all non-friend nodes and using only those that are familiar to us. Further tests reveal that considering the spatial and temporal validity of content has an important impact on the system performance, and that pushing more than one utterance per encounter may deliver better performance in terms of recall, but worsen the performance in terms of precision. Finally we examine the effect of basing the friend set of a node on both familiarity and similarity rather than just similarity.

The remainder of the document is organised as follows. Section 2 describes the main components of the proposed micro-blog dissemination protocol, including the social information maintained by nodes and the criteria for forwarding micro-blogs to encountered nodes. Section 3 presents the methodology and experimentation scenarios we have devised for evaluating the protocol and revealing its main tradeoffs. Our experimental evaluation is structured into five sets of

Download English Version:

<https://daneshyari.com/en/article/6888875>

Download Persian Version:

<https://daneshyari.com/article/6888875>

[Daneshyari.com](https://daneshyari.com)