Computer Networks 127 (2017) 94–108

Contents lists available at ScienceDirect

Computer Networks

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

GROUPS-NET: Group meetings aware routing in multi-hop D2D networks

Ivan O. Nunes*, Clayson Celes, Pedro O.S. Vaz de Melo, Antonio A.F. Loureiro

The Department of Computer Science from the Federal University of Minas Gerais, Belo Horizonte, MG, 31270-901, Brazil

ARTICLE INFO

Article history:
Received 3 March 2017
Revised 5 July 2017
Accepted 14 August 2017
Available online 16 August 2017

Keywords:
Group mobility
Opportunistic routing
Device-to-device communication
Mobile social networking
Social awareness

ABSTRACT

Device-to-device (D2D) communication will allow direct transmission between nearby mobile devices in the next generation cellular networks. A fundamental problem in multi-hop D2D networks is the design of forwarding algorithms that achieve, at the same time, high delivery ratio and low network overhead. In this work, we study group meetings' properties by looking at their structure and regularity with the goal of applying such knowledge in the design of a forwarding algorithm for D2D multi-hop networks. We introduce a forwarding protocol, namely GROUPS-NET, which is aware of social group meetings and their evolution over time. Our algorithm is parameter-calibration free and does not require any knowledge about the social network structure of the system. In particular, different from the state of the art algorithms, GROUPS-NET does not need communities' detection, which is a complex and expensive task. We validate our algorithm by using different publicly available data sources. In real-world large scale scenarios, our algorithm achieves approximately the same delivery ratio of the state-of-the-art solution with 40% less network overhead.

1. Introduction

In recent years, high data rate applications such as videos, songs, games, and social media have become increasingly popular in cellular networks. Device-to-Device (D2D) communication has been proposed to facilitate high data rate transmissions among near users offering higher throughput, efficient spectral usage, extended network coverage and improved energy efficiency.

D2D refers to the direct transmission of content between devices in close proximity, without need for all data to go through the base station. To enable direct D2D *ad hoc* links in the current cellular network architecture, the 3GPP consortium proposed Device-to-Device Proximity Service (D2D ProSe) [1]. D2D ProSe reuses LTE radio interface for direct transmission between devices, allowing the base station to participate and control D2D connection establishment and creation of routes between devices. Several studies in the context of 3GPP-LTE [2] have shown D2D's great potential for alleviating data traffic demands. Nowadays, D2D spans way further, being considered not only as a 3GPP technology, but as one of the key enabling technologies in 5G cellular networks [3].

D2D modes can be classified into two basic types: **1-hop transmission**, in which a message goes directly from the source to the destination if they are close enough to each other; and **multi-hop transmission**, where the message must be opportunistically routed, device by device, from the source to the destination. This last strategy is more complex, since it depends on the intermittent communication structure of a mobile network and is suited for communications that tolerate larger delivery times. This concept was initially introduced in the context of Delay Tolerant Networks (DTNs), but it has many applications to D2D networks as well [25]. A fundamental difference from D2D networks to pure DTNs is the possibility of a centralized control plane with a distributed data plan, as depicted in Fig. 1. In this case, the base station controls the forwarding policy while the data is transmitted from device-to-device until it reaches the destination [32].

Forwarding algorithms in multi-hop D2D networks aim to achieve cost-effective delivery, i.e., the highest possible delivery ratio with the lowest possible network overhead. In this case, the delivery ratio is measured as the percentage of messages routed opportunistically that are successfully delivered to the destination. Successfully delivered messages are the ones that the base station will not need to deliver itself, enabling bandwidth offload. The network overhead is measured by the average number of times that the content needs to be D2D-transmitted for the message to reach its destination. A high number of transmissions may negatively im-

* Corresponding author.
E-mail addresses: ivanolive@dcc.ufmg.br (I.O. Nunes), claysonceles@dcc.ufmg.br (C. Celes), olmo@dcc.ufmg.br (P.O.S. Vaz de Melo), loureiro@dcc.ufmg.br (A.A.F. Loureiro).

http://dx.doi.org/10.1016/j.comnet.2017.08.012
1389-1286/

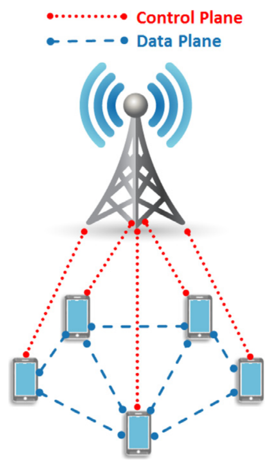

Fig. 1. D2D hybrid architecture: centralized control and distributed data planes.

pact the users' experience by, for example, increasing the devices' energy expenditure.

Considering these metrics, the most successful strategy for opportunistic cost-effective forwarding in D2D Networks, Bubble Rap [21,25], relies on two aspects: information about the nodes' centrality and the static social communities they are part of. The first element can be approximated by the node popularity within the mobile network. The latter, however, presents some significant problems. First, communities are computationally expensive to identify [30]. Second, they are hard to detect in a distributed way, since individual nodes do not have information about the contact graph of the whole network. Existing distributed community detection algorithms achieve at most 85% precision in the detected communities [22]. Another problem of community detection algorithms is the parameter calibration, which depends on the parameters to be adjusted for each given scenario [37]. In a real-time application, such as D2D communication, such calibration is not feasible. Moreover, there is no established truth for community detection. In fact, Abrahao et al. [9] showed that, for the same scenarios, different community detection algorithms yield very different results for communities' compositions. Finally, static communities detection does not account for the dynamism in humans' social relationships, i.e., how they change over time.

Given all these issues, in this work we go beyond and propose the use of social groups' meetings instead of communities. A social group meeting is defined as *a group of people who are together, in space and time, for some social reason or common goal*. For example, friends hanging out at a bar share the social motivation of being together to relax and talk to each other. Students in the classroom share the objective of learning the class' subject content. People in a bus are together because they share the same goal of going to the same route. These are examples of social group meetings. As human beings have regular schedules and routines, it is reasonable to expect social group meetings to present some regularity as well.

From the implementation point of view, different from communities, detecting a social group meeting in a distributed fashion is an easy task. A device can detect a group meeting of which it is part of by simply looking at the list of devices that remained nearby for more than a given time, for example, 10 min. Moreover, the group meeting detection method does not change depending on the scenario nor require parameters' calibration for each specific network, as in community detection schemes. In addition to those desirable characteristics, by looking only at recent group meetings or by giving a higher importance to more recent meetings, it is possible to account for the dynamic nature of social relationships. All of these favorable characteristics motivated us to employ group meetings in the design of an opportunistic routing scheme that is better suited for D2D networks than the current social-aware proposals. Generally speaking, *the main contribution of the present work is to introduce social group meetings and their regularity as a better metric for social context, specially for opportunistic D2D communication and ad hoc networks, in which community detection is a hard task*. Specifically, the contributions of this work are fourfold:

- A methodology for detecting and tracking group meetings from pairwise contact traces;
- A characterization and a model of group meetings, which can be used to predict future meetings using the information about the most recent ones;
- An analysis of the state-of-the-art synthetic mobility models, which allowed us to conclude that group meetings' properties are not well captured by the synthetic traces generated from such models;
- An opportunistic forwarding algorithm for D2D networks, aware of group meetings, which does not need to detect communities nor calibrate parameters, and that achieves better cost-effectiveness than the state-of-the-art solution in real large-scale scenarios.

This paper is organized as follows. Section 2 presents the related work and discusses the corresponding contributions with respect to human mobility, cost-effective opportunistic routing, and D2D communication. Section 3 formalizes our methodology to detect and track group meetings from pairwise contact traces. Section 4 describes the main properties of group meetings that make them interesting in the design of a new forwarding strategy. Section 5 introduces GROUPS-NET, a group meetings aware routing protocol. Section 6 presents a comparison of real and synthetic mobility traces regarding the presence of group mobility. Section 7 evaluates GROUPS-NET, contrasting its performance with the state-of-the-art solution, Bubble Rap, in different network scales. Finally, Section 8 presents final remarks and future work.

2. Related work

2.1. Human mobility

Many studies have used diverse data sources to look into human mobility patterns in the perspective of both individual behavior and collective dynamics. Gonzalez et al. [17] show that human trajectories present a high degree of temporal and spatial regularity. More recently, a dichotomy in individual mobility was revealed in [36] and [42]. Those studies suggest that two mobility profiles, called returners and explorers, govern people movements based on preferential returns and explorations of new places. In relation to collective dynamics, Candia et al. [12] analyze large-scale collective behavior from aggregated call detail records. Their work reveal that the spatio-temporal fluctuations of individuals in a city strongly depend on the activity patterns and routines. Yet, Issacman et al. [23] propose an approach for modeling how people move in different metropolitan areas.

All the studies above are concerned with identifying the intrinsic properties of human mobility in order to provide knowledge and models for different applications, such as the impact of large-scale events in urban mobility [11], typical transitions between points of interest in a city [41], and characterization and prediction of traffic conditions [10,27,43]. Some studies have investigated the importance of understanding the properties of human mobility for designing communication protocols based on opportunistic encounters among people [13,35,39,40]. These studies explore pairwise contacts between users from the individual mobility perspective considering the following metrics: contact rate, inter-contact