Computer Communications 48 (2014) 56–70

Contents lists available at ScienceDirect

Computer Communications

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

Performance modelling of opportunistic forwarding under heterogenous mobility

Chiara Boldrini *, Marco Conti, Andrea Passarella

Institute for Informatics and Telematics, National Research Council, Via G. Moruzzi 1, 56124 Pisa, Italy

ARTICLE INFO

Article history:
Available online 12 April 2014

Keywords:
Opportunistic networks
Forwarding protocols
Performance modelling

ABSTRACT

The delay tolerant networking paradigm aims to enable communications in disconnected environments where traditional protocols would fail. Opportunistic networks are delay tolerant networks whose nodes are typically the users' personal mobile devices. Communications in an opportunistic network rely on the mobility of users: each message is forwarded from node to node, according to a hop-by-hop decision process that selects the node that is better suited for bringing the message closer to its destination. Despite the variety of forwarding protocols that have been proposed in the recent years, there is no reference framework for the performance modelling of opportunistic forwarding. In this paper we start to fill this gap by proposing an analytical model for the first two moments of the delay and the number of hops experienced by messages when delivered in an opportunistic fashion. This model seamlessly integrates both social-aware and social-oblivious single-copy forwarding protocols, as well as different hypotheses for user contact dynamics. More specifically, the parameters of model can be solved in a closed form in the case of exponential and Pareto inter-meeting times, two popular cases emerged from the literature on human mobility analysis. In order to exemplify how the proposed framework can be used, we discuss its application to two case studies with different mobility settings. Then, we discuss how the framework can be also extended to accommodate inter-meeting times following a hyper-exponential distribution. This case is particularly relevant as hyper-exponential distributions are able to approximate the large class of high-variance distributions (distributions with coefficient of variation greater than one), which are those more challenging, e.g., from the delay standpoint. Finally, we provide a validation for the framework with both ideal contacts (i.e., exactly following a given distribution) and contacts extracted from a real mobility trace. This evaluation highlights the strength of the framework in terms of its ability both to provide very accurate predictions under ideal mobility and to effectively approximate the behaviour of the delay moments under real mobility.

1. Introduction

With the advent of powerful and lightweight mobile devices, such as smartphones and tablets, the ubiquitous networking vision is quickly becoming a reality. A further step in the direction of communicating anytime anywhere is represented by the Delay Tolerant Networking paradigm, which enables communications also in disconnected environments. In such conditions, the main requirement of protocols for legacy Mobile Ad Hoc NETworks (MANET), i.e., the presence of an end-to-end path connecting the source and the destination of a message, can hardly be satisfied. Typical delay tolerant networks are, e.g., networks made up of subnetworks connected only by satellite links [1], or networks whose nodes are people moving around with their hand-held devices [2]. The latter case is the scenario considered in this paper. In the literature, such networks have been named Pocket Switched Networks (PSN [3]) or simply *opportunistic networks*, because they opportunistically exploit contacts between users. Opportunistic networks are particularly suitable for content generation, sharing, and access, along the paradigm of content-centric networking [4].

In opportunistic networks, messages are dynamically handed over from node to node upon contact, according to the store-carry-and-forward paradigm. Nodes carry messages with them while they move across the network and with their movements they create transmission opportunities that enable communications. Thus, in opportunistic networks the delay accumulated by the messages along the forwarding path critically depends on the way users move. The simplest exploitation of contact opportunities

* Corresponding author. Tel.: +39 050 3153504; fax: +39 050 3152593.
E-mail addresses: chiara.boldrini@iit.cnr.it (C. Boldrini), marco.conti@iit.cnr.it (M. Conti), andrea.passarella@iit.cnr.it (A. Passarella).

http://dx.doi.org/10.1016/j.comcom.2014.03.028
0140-3664/

in order to forward messages is represented by Epidemic forwarding [5], which generates and hands over a new copy of the message for each new encounter. The rationale behind this approach is to leverage as many routes to the destination as possible. Unfortunately, this greedy approach suffers from severe resource consumption and tends to overload the network [6]. Smarter strategies as to who to forward and how many copies should be generated have been devised since then. According to the type of information used when making forwarding decisions, these strategies can be classified as partially social-aware [7,8] and fully social-aware [9–11]. They leverage information about the users, their contact dynamics, the environment they operate in, the social relationships they share, in order to select one (or a bunch of) *best* next hop. Depending on the number of copies generated for the same message, forwarding protocols can also be classified into single-copy or multi-copy schemes. In the first case, at any time, in the network there is just one copy of the message to be delivered, while in the second case more copies are generated, hoping that at least one of them will eventually reach the destination. While multi-copy strategies have been shown to improve the reliability of delivery, they are typically resource consuming.

Despite the variety of practical forwarding solutions based on different heuristics (such as encounter frequency and sociality metrics) no general framework has been introduced so far for the analysis of opportunistic forwarding protocols in a structured way. Some models exist in the literature (e.g., [12,13,8,14,15]), but they are specific to the protocols being studied and can hardly be re-used when the protocols are changed. The situation is even worse for social-aware schemes, which, despite their popularity, are typically difficult to model analytically. Moreover, the absence of a general consensus on some fundamental properties of user movement patterns (e.g., the distribution of the inter-meeting times) makes it even more complex to found a model on a solid basis. In fact, the performance of message forwarding closely depends on the users' contact dynamics [16]. From the analysis of real movement traces many hypotheses (e.g., [16–22]) have been made as to which distribution better describes significant quantities such as the time between consecutive contacts, or the duration of a contact, but without ultimate consensus.

The contribution of this paper is twofold. First, a general framework for the analysis of single-copy forwarding schemes is introduced. This model, based on Markov chains, allows us to compute significant quantities, such as the first and second moments of the number of hops and delay, which characterise the forwarding performance. These moments can then be used to approximate the full distribution of the delay and number of hops. Clearly, the full distribution of, e.g., the delay is more informative than just its expectation, as it allows us to analyze, for example, the dependency of the delay on the TTL. This general framework also takes into account social-awareness, which can be incorporated seamlessly into the model. In addition, our framework is independent of specific mobility assumptions, thus it would remain usable even if new insights on the way users move were provided.

The second contribution is the instantiation of the framework in three different mobility scenarios. More specifically, we solve the parameters of the framework *in closed form* in the case of exponential and power law inter-meeting times, which are popular assumptions for inter-meeting times emerged in the literature [19,16,23,24]. In addition, we also provide a complete solution to the framework in the case of hyper-exponentially distributed inter-meeting times. To the best of our knowledge, this is the only modelling framework that accommodate heterogeneity not only at the level of the parameters of the distribution (e.g., different pairwise rates for exponential intermeeting times) but also at the level of the distribution itself.The result about the hyper-exponential case is particularly significant, since the hyper-exponential distribution can approximate the behaviour of a large class of distributions, those having a coefficient of variation greater than 1. The coefficient of variation [25] is defined as the ratio between the standard deviation and the mean, and measures the dispersion of a probability distribution. The higher the coefficient of variation, the more distant a sampled value can be from the mean. High-variance distributions are extremely important in opportunistic networks for two reasons. First, they have often emerged as a plausible hypothesis for inter-meeting times (apart from the power law hypothesis, recently the LogNormal one has also gained a lot of popularity [24]). Second, high-variance distributions can drastically affect the delay experienced by messages, causing the expectation of the delay to diverge in extreme cases [26,16].

The characteristics of single-copy schemes have been analytically studied in the literature for what concerns social-oblivious strategies [8,16], but, to the best of our knowledge, the one proposed in this paper is the first general framework that takes into account the social-awareness of the forwarding process. Moreover, results obtained for single-copy schemes are important to multi-copy schemes as well. Consider for example multi-copy schemes in which replication can occur only at the source node. Each copy travels along a path independently of the others. While the delivery from the source node to the first relays is significantly different from a single-copy delivery due to the multi-copy generation, from the first relay to the destination the delay can be approximated using single-copy results. The extension of the framework to the multi-copy case is currently under study.

The paper is structured as follows. In Section 2 we review the state of the art on forwarding protocols and performance modelling for opportunistic networks. In Section 3 we describe the scenario we consider and the assumptions we make, based on which, in Section 4, we define our general modelling framework. After defining in Section 5 our reference forwarding schemes, in Section 6 the general framework is specialised under the assumptions of exponential and power law inter-meeting times. In order to exemplify how the proposed model can be used, we discuss its application to two case studies with different mobility settings in Section 7. After that, in Section 8 we extend the model to include the case of hyper-exponential inter-meeting times, and we discuss how this case can be used to solve the framework approximately in the general case of high-variance inter-meeting times. Finally, in Section 9 we evaluate the quality of the predictions provided by the framework. More specifically, we first compare the predicted first two moments of the expected delay to their values measured from simulations assuming that intermeeting times are ideally Pareto and hyper-exponentially distributed, and then we compare predicted and measured delays when simulations are performed using a real contact trace as input.

2. Related work

2.1. Opportunistic forwarding

According to the type of information that they exploit when making forwarding decisions, forwarding protocols can be classified into social-oblivious, partially social-aware and fully social-aware protocols [9]. In the following we overview some of the most significant protocols for each of these categories. For a more detailed survey, we refer the reader to Al Hanbali et al. [27]. Social-oblivious protocols do not use at all information on the way nodes meet or relate with each other. This is the case of the Epidemic protocol [5], whose strategy is to generate and hand over a new copy of the message to each node encountered, and of the Direct Transmission protocol [28], in which messages can only be delivered to the destination when encountered directly. Their performance is