



Opportunistic content diffusion in mobile ad hoc networks



Bryce Thomas^{a,*}, Raja Jurdak^b, Ian Atkinson^c

^aDepartment of Information Technology, James Cook University, Townsville, Queensland, Australia

^bCSIRO, Brisbane, Queensland, Australia

^ceResearch, James Cook University, Townsville, Queensland, Australia

ARTICLE INFO

Article history:

Received 17 January 2015

Revised 20 February 2016

Accepted 29 February 2016

Available online 5 March 2016

Keywords:

Manets

Content-diffusion

Contact-networks

ABSTRACT

Opportunistic wireless content sharing via Mobile Ad hoc NETWORKS (MANETs) can increase throughput, lower latency, extend network coverage and reduce load on infrastructure. While the benefits of content diffusion clearly depend on the underlying movement dynamics and content demand, the impact of these factors on diffusion remains largely unexplored. We analyze content sharing potential based on device encounters inferred from a large multi-site wireless LAN trace. We explore the impact of time, location, and number of sources on diffusion, finding that contexts with higher activity generally promote faster diffusion, while additional content sources improve diffusion mainly in the short-term. We then apply real-world demand patterns from a popular campus maps application to content diffusion simulations. We find that up to 70% of map requests could theoretically be served from the peer network over the first 12 h. Finally, our analysis of the impact of trace uncertainties and individual device variation on diffusion potential reveals large differences based on the selected assumption and chosen source devices. We discuss these results and their implications for content-diffusion in MANETs.

Crown Copyright © 2016 Published by Elsevier B.V. All rights reserved.

1. Introduction

Enabling wireless user devices to directly share common-interest content is a conceptually attractive approach to enhancing wireless networks. Each user device caches content retrieved from the infrastructure and makes it transparently available to co-located peers, either pre-emptively or on demand. Devices' content demands are preferentially served from a nearby peer with the infrastructure serving as a fallback when a cached copy is unavailable. The potential benefits of such a scheme include higher throughput, lower latency, greater spectrum reuse, extended network coverage and reduced load on infrastructure.

1.1. Motivating example

We present a mobile map sharing application as a motivating example. Suppose User A is using their mobile device to navigate a geographic region after having downloaded the region's map from the infrastructure (e.g. a cell tower or wireless access point). Now suppose User B enters the same region and encounters User A. User A proceeds to pre-emptively share the map data with User B. Shortly afterwards, User B would also like to view a map of the

region. Rather than having to retrieve the mapping data from the infrastructure, User B already has a local copy available received earlier from User A. We highlight several potential benefits of this peer sharing:

- Being in close geographic proximity allows the devices to transmit at lower power, reducing battery consumption and increasing opportunities for spectrum reuse in adjacent areas.
- User A and B can establish a short-range dedicated connection, increasing throughput. This is particularly important if User B were to retrieve the map on demand, rather than receiving it pre-emptively.
- The devices can communicate with very low latency as a result of the short-range nature of the connection and because the devices are not contending with other devices for access to the infrastructure. Again, this is important for on-demand retrieval.
- If User B is not in range of the infrastructure, User A effectively extends User B's coverage by making otherwise unreachable content available.
- Finally and in many cases most importantly, load has been taken off the fixed wireless infrastructure. Wireless infrastructure and cellular data infrastructure in particular is often viewed as being in a perpetual state of underprovision. Partially offloading content delivery from the infrastructure onto a Mobile Ad hoc NETWORK (MANET) may prove a useful strategy for

* Corresponding author. Tel.: +1 206 455 3104.

E-mail addresses: bryce.m.thomas@gmail.com (B. Thomas), raja.jurdak@csiro.au (R. Jurdak), ian.atkinson@jcu.edu.au (I. Atkinson).

reducing the necessary cost or frequency of infrastructure upgrades.

Continuing the maps example, assume that some time later User A transitions to a new geographic region. As a result of A's mobility, maps of the prior region are now available to devices in the new region. This is an example of how content may spread with the aid of device mobility.

We have presented mapping as just one motivating example of MANET-based content sharing via diffusion. The use cases of content diffusion however generalize to any application premised on or enhanced by the ability to move content quickly and efficiently. Content diffusion may prove particularly useful for other applications which like maps exhibit locality of reference [1] in content interests, i.e. content interests tend to be spatially and/or temporally correlated. This includes web content, app content and even personal area networks (PANs) where a single user carries multiple cloud-connected devices synchronizing identical data.

1.2. Contributions

Though wireless peer-to-peer (P2P) content sharing is an intellectually attractive approach to improving network efficiency and performance, a lacuna exists in the literature around real-world parameters influencing content diffusion potential. Existing works [2,3] explore some facets of epidemic content diffusion including the resulting network topologies and diffusion potential under various constraints on participation. Our earlier work in [4] provides a preliminary examination of how site, time of day, day of week, number of content sources and empirical patterns of content demand influence content diffusion potential in wireless LANs. In the present paper we build on our prior work by analyzing the impact of *uncertainty and variation* in trace-driven diffusion simulations. We find diffusion potential to be relatively sensitive to the assumptions chosen to compensate for inherent timing uncertainties in wireless LAN traces. We also find a relatively large amount of variability in diffusion potential between individual content source devices. We discuss currently accepted assumptions of the research community as they pertain to inferring device encounters and highlight why verifying the validity and then perhaps improving these assumptions would be beneficial.

1.3. Paper structure

The following section covers related work. Section 3 provides background information on the area of content diffusion and formally defines how device encounters are inferred from wireless LAN traces. Our primary wireless LAN trace from a large university campus is described in Section 4, along with its associated uncertainties in session timestamps. Our first set of simulations analyze *universal diffusion* on the empirical trace, i.e. how quickly an arbitrary piece of content might spread throughout a network. These simulations are described in Section 5 and the results are presented in Section 6. We then focus on a realistic *application-specific* use-case for content diffusion in Section 7—diffusing electronic maps based on the LAN trace and on empirical usage statistics from a university navigation app. Section 8 provides a discussion of our findings regarding the impact of trace uncertainties and presents avenues for future work. Section 9 concludes the paper.

2. Related work

Our work fits broadly into the existing body of research around MANET [5] communications and Delay Tolerant Networking (DTN) [6]. Though present-day device and protocol support for seamless device-to-device communication is somewhat deficient, we are

particularly motivated in our analysis by promising next generation protocols like Content-Centric Networking (CCN) [7]. The pertinent feature of CCN (and similar protocols) is enabling trustworthy content to be retrieved from untrusted hosts.

Most directly related to our work are empirical studies of device mobility and encounters, and the ad hoc content diffusion opportunities these create. Eagle & Pentland [8] recorded 9 months of Bluetooth encounters of 100 mobile devices given to students and faculty at MIT university. Wang et al. [9] recorded 3 days of Bluetooth encounters of 41 “iMote” devices given to participants at the 2005 Infocom conference. Su et al. [3] recorded device encounters of two groups of students given PDAs, each group being around 20 students in size and the two experiments lasting 2.5 and 8 weeks, respectively. Hsu & Helmy [2] analyzed device encounter patterns in traces collected from four university campuses and the Infocom 2005 conference.

Of the aforementioned works, [2] and [3] explicitly analyzed ad hoc multi-hop message dissemination facilitated by device mobility and encounters. Our own work compliments these prior studies by i) analyzing site, time of day, day of week and number of content sources as diffusion parameters; and ii) providing new findings on application-specific diffusion, trace uncertainties and diffusion variation. Furthermore, we perform our simulations using a late 2012 trace, which compared to traces used in past studies is substantially newer (in some instances over a decade), larger in size, and is collected with greater temporal and procedural consistency across sites.

A number of other studies [10–13] have characterized wireless network usage and user behavioral patterns. In addition to these, there have been a multitude of works on mobility models intended to describe the movement of devices in space and time, many of which are reviewed in [14]. Again our work is complementary to these studies, though we focus specifically on information diffusion potential in the context of empirical data, not network characterization or mobility modeling.

3. Background and definitions

3.1. Opportunistic mobile content diffusion

Opportunistic mobile content diffusion refers to the dissemination of content directly between mobile devices during incidental encounters, i.e. where and when opportunities naturally arise. Content may originate directly from a device or have been downloaded from an infrastructure network at an earlier point in time. For example, a sensor reading may originate from a mobile device, while a cached web page originates from an Internet-connected infrastructure network. Once one or more mobile devices possesses a given piece of content, that content can be shared directly with other mobile devices. These other devices may then further propagate the content causing a (time respecting [15]) transitive spread of content throughout the network. Even a device with no interest in a piece of content may act as a data mule [16] that receives, caches and then further propagates the content during subsequent opportunistic encounters.

3.2. Ideal diffusion

We define *ideal diffusion* as a special case of opportunistic content diffusion that takes place *every time* an opportunity arises. Essentially this is a form of flooding—each time two devices encounter, they share with one another their respective contents.

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات