Computer Communications 73 (2016) 118–131

Contents lists available at ScienceDirect

Computer Communications

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

On search and content availability in opportunistic networks[☆]

Esa Hyytiä[a,*], Suzan Bayhan[b], Jörg Ott[a], Jussi Kangasharju[b]

[a] Department of Communications and Networking, Aalto University, Finland
[b] Department of Computer Science, University of Helsinki, Finland

ARTICLE INFO

Article history:
Available online 2 November 2015

Keywords:
Mobile opportunistic networks
Mobile search
Mobile cloud computing
Dynamic programming
Availability estimation

ABSTRACT

Searching content in mobile opportunistic networks is a difficult problem due to the dynamically changing topology and intermittent connections. Moreover, due to the lack of global view of the network, it is arduous to determine whether the best response is discovered or search should be spread to other nodes. A node that has received a search query has to take two decisions: (i) whether to continue the search further or stop it at the current node (current search *depth*) and, independently of that, (ii) whether to send a response back or not. As each transmission and extra hop costs in terms of energy, bandwidth and time, a balance between the expected value of the response and the costs incurred must be sought. In order to better understand this inherent trade-off, we consider a model where both the query and response follow the same or similar path. We formulate the problem of optimal search for two cases: a node holds (i) exactly matching content with some probability, and (ii) some content partially matching the query. We design *static search* in which the search depth is set at query initiation, *dynamic search* in which search depth is determined locally during query forwarding, and *learning dynamic search* which leverages the observations to estimate suitability of content for the query. Additionally, we show how unreliable response paths affect the optimal search depth and the corresponding search performance. Moreover, we study different methods to a priori learn the availability of the content in the network based on passive observations (e.g., using regression and maximum-likelihood based estimates). Such information is highly valuable when defining the optimal search parameters. Finally, we investigate the principal factors affecting the optimal search strategy.

1. Introduction

Mobile users rely on cloud-based third party services for information sharing and messaging even if they are close by so that their respective devices could exchange information directly, without taking a long distance detour across the Internet and potentially half-way around the world. This holds for a broad spectrum of services from email and chat (e.g., Jabber) to online social networks (e.g., Facebook, Twitter) to data and file sharing (e.g., Dropbox). Important features of such services include that they are well-known and well-managed (including backups), support a global community, and are usually instantly accessible, so that users have them always conveniently at their disposal—as long as they are connected to the Internet.

However, this convenience creates dependencies, at the very least on Internet access and the availability and reachability of the respective third-party service. It also creates cost: the data needs to be moved (repeatedly) through parts of the Internet, consuming network and energy resources, and needs to be processed and stored in the cloud—including backups, as even short-lived data is rarely deleted (immediately or at all) after sharing, as the authors observe in their and others use of, e.g., Dropbox.

Instead of using infrastructure services, nodes can directly exchange content via short-range communication interfaces, e.g., Wi-Fi, Bluetooth, such that only the peers in wireless contact are involved. In this manner, peers can build a network operating in an ad hoc mode and facilitating the communication between non-adjacent nodes via hop-by-hop data forwarding. With the sheer growth of data traffic and fierce competition for the bandwidth, operators benefit from this approach for decreasing the path between the content provider and the consumer. On the other hand, due to the instability of the direct links, this scheme may not guarantee certain delay bounds, limiting its applicability to only delay-tolerant applications. *Delay-tolerant networking* (DTN) [8,9] defines such a networking paradigm facilitating communication without an infrastructure support for a variety of application scenarios including inter-planetary, vehicular, underwater, and opportunistic networks.

Mobile opportunistic networks, also dubbed *Pocket switched networks* [11,19], are of particular interest with the increasing diffusion of

---

[☆] An earlier version of this paper was presented at ACM MSWiM 2014 [13].

[*] Corresponding author. Tel.: +358 50 4354971; fax: +358 9 4512474.
E-mail addresses: esa.hyytia@aalto.fi, esa.hyytia@iki.fi, esa@netlab.tkk.fi (E. Hyytiä), bayhan@hiit.fi (S. Bayhan), jo@netlab.tkk.fi (J. Ott), jakangas@cs.helsinki.fi (J. Kangasharju).

http://dx.doi.org/10.1016/j.comcom.2015.09.011
0140-3664/© 2015 Elsevier B.V. All rights reserved.

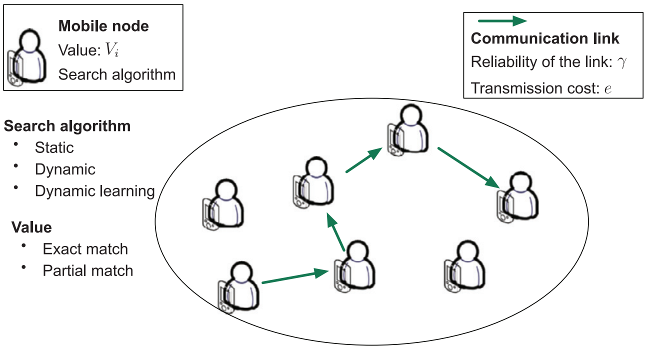

Fig. 1. Search query travels from a node to another and the path forms a linear trajectory in space. The response is assumed to follow the same (or similar) path backwards.

powerful mobile wireless devices (such as tablets and smartphones). Mobile devices carried by humans can exchange information when they come in transmission range of each other and physically carry the content on their way. Certainly, this operation mode is vital for cases where the network infrastructure fails (e.g. after natural disasters), does not exist, or access to infrastructure services or even the Internet at large is blocked [7].

The wealth of data produced or downloaded by the mobile devices requires efficient search algorithms that can locate the relevant content quickly and cost-effectively rather than naïve enquiry of each node upon a contact. Searching content in mobile opportunistic networks is a difficult problem due to the dynamically changing topology and intermittent connections. A question arising in this context is what are the fundamental determinants of search in mobile opportunistic networks. In this work, we aim to provide insights on this question by designing static and dynamic search schemes. We focus on a single query that visits a node after another along some (natural) path as illustrated in Fig. 1 (i.e., the query is not replicated). The response follows the same (or similar) path backwards. That is, the response path is assumed to equally long, and moreover, it can be unreliable, e.g., due to mobility during the search. More specifically, we assume that searches terminate relatively quickly (say, order of ten seconds) and a link backwards exists if the response can be transmitted shortly (but not necessarily immediately). In other words, we do not require persistent end-to-end paths.

In practice, the search path can form naturally based on some path selection criteria such as a similarity metric for nodes, which reflects positively to the probability of finding relevant information in the node. Similarly, the actual search can consist of multiple (independent) linear paths.

The main contributions of this paper are as follows:

1. We provide theoretical modeling of several search strategies: static, dynamic, and learning dynamic search. In *static search*, the search is extended to a predefined number of nodes n whereas in *dynamic strategies* a node may stop the search before depth n depending on what has been found so far (and whether some response has already been sent back). Both of these assume that each node knows the distribution of information in the nodes (value of response to given search). Our final search strategy, referred to as the *learning strategy*, is more robust and estimates the value distribution dynamically as the search progresses from a node to another.
2. To assist the search, we also study different methods to a priori estimate the content availability based on other similar queries nodes have observed in the past. As these methods rely on passive observations, the only additional cost is a negligible increase in the computational effort, while the benefits can be considerable when each query has a good understanding of the operating environment from the beginning. To the best of our knowledge, our paper is the first paper on content availability estimation in an opportunistic network.
3. For both the static and dynamic search, we model the search process under both *exact matching* and *partial matching* content items. The former corresponds to search of the specific content (yes or no), whereas in the latter a multitude of answers is possible each with a different value.
4. Although we assume that response follows the same or equally long path as the query, we model the unreliability of the link between two nodes on the response path and analyze how it affects the optimal search (cf. mobility).

Rest of this paper is organized as follows. First, in Section 2 we briefly review the related work. Then, Section 3 introduces our model and notation. Section 4 presents the analysis of different search strategies, whereas Section 5 discusses four methods to a priori estimate the average content availability in the opportunistic network. These are followed by a performance evaluation in Section 6. Section 7 concludes the paper.

2. Related work

In a broad context, we can consider every forwarding algorithm in a DTN as a search scheme for a specific target node. We exclude broadcast algorithms as they aim to reach each and every node. In content search, first the sought content is mapped to some node(s) that have a high likelihood of holding this particular content. Next, nodes upon encounters forward the query with the aim of reaching the specified destination(s) that matches the mapping between content and the node profile. For example, *seeker-assisted search* (SAS) [3] vaguely maps a content to the nodes of a particular community which is a group of nodes sharing common interests. Hui et al. [11], design Haggle – a content sharing scheme, by leveraging the node's self-declared interests to locate the contents that might fall in the interest of the node. In Haggle, each content and node have some attributes that are manually defined. These attributes provide the basis of mapping between a content and its target nodes. See also the *Bubble* forwarding algorithm [12], which tries to exploit the social structures when making the forwarding decisions.

Rational search schemes should direct the search towards the nodes that have a higher likelihood of having the sought item. On the way to these "potential content providers", nodes with good relaying capabilities can be employed as intermediate carriers. The forwarding decision can exploit various characteristics of the network, e.g., centrality of the nodes, (sub-)groups in the network, content and node relevance. For example, SAS [3] exploits the *homophily principle*, tendency to associate and interact with similar others, and directs the search towards the nodes of the same community as the content might have been sought and be readily available at a node in this community. In order to avoid searching only a specific part of the network, SAS expands the search also to elsewhere, although the probability of finding the content might be lower. Likewise, DelQue [10] defines *geo-community* concept to associate the interests with the locations (e.g., people interested in basketball contact each other in gyms). In Haggle, nodes exchange contents at each encounter so that contents are constantly *pushed* towards the nodes with some interests for this content rather than an explicit search. In this paper, similar to SAS, we consider a *pull-based* search scheme in which nodes issue queries for finding specific contents.

Deciding when to stop the forwarding of the search query is another challenge as nodes operate in distributed fashion relying on their local knowledge. Although a query might have reached the content provider and a response message may be already on the way, this may not be signaled immediately to other parts of the network. Hence, each node should decide on forwarding or terminating the spread. An early termination may result in search getting no