Journal of Systems Architecture 72 (2017) 108–119

Contents lists available at ScienceDirect

Journal of Systems Architecture

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

Interest spaces: A unified interest-based dissemination framework for opportunistic networks

Radu-Ioan Ciobanu, Radu-Corneliu Marin, Ciprian Dobre*, Florin Pop

Faculty of Automatic Control and Computers, University Politehnica of Bucharest, 313 Splaiul Independentei, Bucharest, Romania

ARTICLE INFO

Article history:
Received 10 January 2016
Revised 4 April 2016
Accepted 7 June 2016
Available online 8 June 2016

Keywords:
Opportunistic
Mobile
Dissemination
Framework
Interests
Social
Context-aware

ABSTRACT

Embedded ubiquitous computing systems (EUCS) promise to grow remarkably in the near future. For such systems, new communication technologies are constructed to cope with the increasing functional and temporal demans. Because of the ubiquitousness of mobile devices nowadays, for example, traditional publish/subscribe is no longer an adequate model for data dissemination in mobile networks. Since any node can publish content at any time, the network can get congested easily, so a dissemination paradigm where mobile nodes contribute with a fraction of their resources is needed, through the use of opportunistic networks. Furthermore, a suitable organization for data dissemination in mobile networks should be centered around interests. Thus, we propose a unified interest-based dissemination framework for opportunistic networks entitled Interest Spaces, which simplifies dissemination by just allowing applications to mark data items with certain tags, letting the framework handle the caching, routing, forwarding, and disseminating. Similarly, applications that need to subscribe to channels simply have to specify the tags they are interested in, and the framework does the rest. However, the Interest Spaces framework also allows applications to have more control over the dissemination process if they wish, by specifying various criteria and dissemination rules. In this article, we present the architecture and components of Interest Spaces. We focus on data aggregation at the context layer, showing through simulations the benefits it brings in opportunistic dissemination. We also present a real-life use case for Interest Spaces through Chatty, an opportunistic chat application.

1. Introduction

In traditional client/server and peer-to-peer systems, the publish/subscribe paradigm has a relatively straightforward implementation. Some nodes are publishers, which means they are able to generate data on some channels, while others are subscribers, subscribing to various channels, and thus signaling their interest in the data published on those channels. In between, a catalog or a similar construct acts as intermediary between nodes registering to receive data, and the ones producing it. Generally, such systems assume a limited number of publishers, which are well-known and thoroughly advertised. However, with the advent of Web 2.0, and (even more importantly) due to the exponential increase in the number of computationally-capable mobile devices on the market, the publish/subscribe paradigm has lately suffered significant changes. For example, adapted publish/subscribe paradigms assume that any node is able to publish content, so the number of publishers can easily match the number of subscribers. In this situation, congestion can easily be reached when many nodes publish data on the same channel and the number of subscribers is high. The high number of mobile devices also makes classic publish/subscribe impracticable, since mobile networks require decentralization. Moreover, as stated in [1], conventional peer-to-peer approaches are beginning to not be enough, exactly because of the presence of mobile nodes, as well as because of challenged networking conditions. Therefore, a data dissemination paradigm where mobile nodes contribute with a fraction of their resources is needed, which leads to the situation of opportunistic networks (ONs) addressed by this article. Moreover, the availability of data plays an important role in opportunistic dissemination, as specified in [2].

Imagine a football stadium of 50,000 seats, where 75% of the fans (i.e. 37,500) have mobile devices. If only half of them need to use 3G/4G or WiFi communication at the same time (to find out the score at a different match, for example), they will have problems connecting. This is where opportunistic networks can be very useful. In this situation, instead of connecting to crowded access points, nodes can connect to each other, until they reach a node

* Corresponding author.
E-mail addresses: radu.ciobanu@cti.pub.ro (R.-I. Ciobanu), radu.marin@cti.pub.ro (R.-C. Marin), ciprian.dobre@cs.pub.ro (C. Dobre), florin.pop@cs.pub.ro (F. Pop).

http://dx.doi.org/10.1016/j.sysarc.2016.06.004

that has the required information. A node may possess this information either because it has downloaded it through 3G/4G/WiFi, because it has generated it itself, or because it has received it from a different node that already had it. This is opportunistic communication in a nutshell, and we believe it is the future of mobile communication, and of the publish/subscribe paradigm.

Whereas, in classic publish/subscribe, data is organized into channels to which nodes subscribe, we believe that a more suitable organization for mobile networks should be done according to interests. Thus, instead of nodes specifying a channel to which they subscribe, they can simply state that they are interested in a certain topic (such as "football" in the example above). Conversely, nodes that generate information regarding football simply mark that information with a "football" tag. This abstraction simplifies publishing and subscribing greatly, since nodes do not need to know who the publishers are or how channels are identified.

Since opportunistic networks are totally decentralized, nodes need to cooperate in order to successfully create a publish/subscribe environment that benefits all the nodes in the network. For example, let us assume that two friends carrying mobile devices meet in a bar for a drink. One of them is a football fan, and will go to a match after the two friends separate, while the other one likes rock music and is preparing to go to a concert of his favorite band. In the bar they meet in, there are many other people with mobile devices, each of them having different tastes. In order to have an efficient data dissemination using only the mobile devices present in the bar, each of them must be able to correctly detect the context of its user (such as preferences, behavior patterns, position, etc.), to infer the tags for which that user is a good forwarder. Consequently, the mobile device belonging to the friend that likes football should start collecting all football-related information from the devices nearby (belonging to the other people in the bar). Similarly, the rock music fan's phone should gather all rock-related data. When the two friends separate, the football fan goes to the match, where all the information his device has collected will be useful for a large number of people (since probably all the people in the stadium are football fans). Something similar should happen for the rock fan, and for all the other people present in the bar.

As we have shown, data dissemination in opportunistic networks can only be accomplished successfully through the collaboration of users performing a modified version of publish/subscribe based on interests. For this reason, in this article we propose an interest-based data dissemination framework for opportunistic networks, entitled Interest Spaces. It will be able to disseminate data to interested nodes, by taking advantage of their context information (such as location, interests, social connections, encounter history, etc.). The advantage of the proposed framework is that it offers a unified interface for data dissemination in various situations. Thus, the contributions of this paper are as follows:

- Propose a unified interest-based dissemination framework for opportunistic networks entitled Interest Spaces, which simplifies dissemination by just allowing applications to mark data items with certain tags, letting the framework handle the caching, routing, forwarding, and disseminating
- Present the architecture and components of Interest Spaces
- Highlight the benefits of data aggregation in Interest Spaces dissemination through experimental testing on a mobility trace
- Present Chatty, an Android opportunistic chat application which is a real-life use-case for Interest Spaces.

The rest of this article is structured as follows. Section 2 presents various data dissemination frameworks based on context and interests, showing what their issues are, and how the Interest Spaces framework means to address them. In Section 3, we propose Interest Spaces, offering details about its architecture, as well as its composing layers. In Section 4 we put together the functionality while focusing on data aggregation, followed by the presentation of a case study in Section 5. Finally, in Section 6, we present our conclusions and future work.

2. Related work

Several context-based models and frameworks for mobile networks have been proposed. For example, Context Spaces [3] is a framework that offers a model for representing context as a constrained view of the world, which can be analyzed in order to decide if certain situations occur. Interest Spaces has a similar way of representing context, but using an ECA (event condition action) paradigm. Thus, situations are not necessarily events that are observed to occur, but rather actions that should be taken if certain conditions are met.

Another example is ContextCast [4], a framework that offers a protocol for maintaining a self-organizing routing backbone which permits geographical addressing and resource discovery. However, one of the main caveats of ContextCast is that it assumes that the nodes have Internet connectivity through TCP at all times. Furthermore, ContextCast groups nodes based on geographical location, but does not take advantage of a mobile device's capabilities of detecting which nodes are within wireless communication range. The Interest Spaces framework is able to group nodes not only geographically, but also based on common interests. It is also specifically designed for ONs, so it assumes that nodes can communicate directly when they are in range, or indirectly through other nodes.

SICC [5] is a framework for MANETs that can be used to describe the context of an application by extending it to encompass a neighborhood within the network. This allows applications to define their area of effect, giving them control over the dissemination process. However, MANETs need routing information to function correctly. In a dense network with thousands of devices, the amount of storage space required to store all the routes of a node becomes extremely high. The advantage of the Interest Spaces framework is that it is built for ONs, where disconnections are the rule, rather than an undesired effect, and routes are opportunistically built depending on a node's encounters. Moreover, while SICC allows nodes to define a context composed of peers in their geographic vicinity, applications using our proposed framework define interest spaces, which can contain nodes located anywhere, but sharing common interests.

Another framework that offers services similar to Interest Spaces is Floating Content [6]. It allows applications to define anchor zones, which are geographical areas where nodes enter, spend a certain amount of time, and leave. In this case, the nodes are mobile devices belonging to humans, and the anchor zones are relatively small-sized areas where many people congregate, which represent the boundaries of an ad hoc ON. While inside the anchor zone, nodes may copy the data either if they need it for themselves, or if they transport it for the benefit of others. If a node is interested in a certain data item from the anchor zone, it replicates it. Data dissemination using Floating Content is done to all nodes in an anchor zone, so a node receives messages regardless of whether they are required or not, which may lead to congestion and unnecessary transfers. On the contrary, Interest Spaces only delivers data to nodes that have declared their interest in it, or to cache nodes, which are able to speed the dissemination process for other nodes.

This shows that there are several existing frameworks that deal with data dissemination in mobile networks. However, some of them are built for MANETs, which require storing a lot of data for the routing tables, as well as many computations that must be done to keep the routing tables updated. We believe that some solutions shown above are limited because they group nodes ge-