Computer Networks 56 (2012) 2949–2962

Contents lists available at SciVerse ScienceDirect

Computer Networks

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

ALPS – Adaptive Location-based Publish/Subscribe

Adrian Holzer [a,*], Patrick Eugster [b], Benoît Garbinato [c]

[a] LARIM, Polytechnique Montréal, Montréal, QC, Canada H3T 1J4
[b] Purdue University, West Lafayette, IN 47907, USA
[c] Distributed Object Programming Lab., University of Lausanne, CH-1015 Lausanne, Switzerland

A R T I C L E I N F O

Article history:
Received 19 July 2011
Received in revised form 2 March 2012
Accepted 18 May 2012
Available online 26 May 2012

Keywords:
Mobile infrastructureless networks
Publish/subscribe abstractions
Context-aware approaches
Adaptive systems

A B S T R A C T

Location-based publish/subscribe – LPS for short – is an important building block for context-aware applications in mobile ad hoc networks (MANETs). In LPS, published messages are routed based on their content as well as on the location of publishers and subscribers.

Existing LPS algorithms can be coarsely classified as follows: (1) *message-centric* approaches consist in broadcasting published messages, (2) *query-centric* approaches broadcast subscriber queries for subsequently routing messages, and (3) *hybrid* approaches broadcast queries and messages each within restricted scopes. Each approach is clearly superior to others for particular communication patterns, e.g., for certain ratios between the number of queries and the number of messages in the network. This paper presents an *adaptive location-based publish/subscribe* (ALPS) algorithm for settings with multiple, unknown, or varying communication patterns. ALPS can be viewed as a parameterized hybrid LPS algorithm that can seamlessly move between message- and query-centricity based on estimations of the current communication pattern.

We evaluate ALPS on two benchmark applications namely in the context of mobile social networking and robot swarms. Our results indicate that ALPS reduces the message complexity by up to a factor 3× compared to the best respective alternative, while performing comparably to the respective optimal solutions with static communication patterns, making ALPS appealing as a *one-size-fits-all* solution.

1. Introduction

Publish/subscribe has become an important building block for many networked applications. The participant decoupling it achieves is especially beneficial to deal with the dynamic nature of mobile ad hoc networks (MANETs), such as mobile sensor networks and mobile ad hoc networks.

1.1. Location-based publish/subscribe

To address the *geographic dependency* of interactions in MANETs (node motion), corresponding publish/subscribe abstractions often include a notion of *location context* complementing content-based queries. This enables the expression of proximity semantics for publishers and/or subscribers [10,17,23], leading to *location-based publish/subscribe (LPS)*. A second key trait of MANETs is the increased *temporal dependency* of interaction, resulting from the ad hoc and transient nature of communication. This trait is captured by associating a notion of *lifetime* with published messages [10,23].

1.2. Implementing LPS

Implementing LPS efficiently and reliably is a daunting task not only due to the inherent dynamism of MANETs [27], but also because of the substantial differences observed in loads induced by distinct networked

* Corresponding author.
E-mail addresses: adrian.holzer@gmail.com (A. Holzer), peugster@cs.purdue.edu (P. Eugster), benoit.garbinato@unil.ch (B. Garbinato).
URL: http://www.hec.unil.ch/dop (B. Garbinato).

1389-1286/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.comnet.2012.05.007

applications. Early LPS algorithms can be coarsely divided into two categories motivated each by specific *communication patterns*, i.e., specific ratios between the number of subscriptions (“queries”) and the number of publications (“messages”) occurring at execution. First, *message-centric* algorithms (e.g., [17]) disseminate messages along with contextual information and perform matching on the subscriber side. Second, *query-centric* algorithms, spearheaded by Yoneki and Bacon [28] and Baldoni et al. [1], rely on the propagation of content queries and location criteria of subscribers to allow publishers and intermediate nodes to perform the matching and routing of messages. In very rough terms, the distinction between the two is similar to that between reactive and pro-active routing.

Message-centric approaches perform best when most messages are relevant to a large fraction of the nodes – essentially a broadcast pattern. Conversely, with few messages or selective content-based dissemination, one can expect query-centric algorithms to prevail. The *difference* in terms of number of messages produced for certain communication patterns can easily constitute up to two orders of magnitude. To strike a better balance, a third approach has been proposed with *hybrid* strategies [2,30]. These disseminate queries to a certain degree. Yet, to be efficient, devices used in different communication patterns need to be equipped with several protocols.

1.3. Contribution

System adaptation and configurability is a core requirement for pervasive computing [16], and several authors have elaborated in seminal work on different faces of adaptation. This paper addresses the problem of adaptation to different LPS *communication patterns* expressed as the ratio between queries and messages by proposing ALPS – an *adaptive LPS* algorithm.

More precisely, we make the following contributions to efficient and reliable LPS in MANETs:

- we lay the foundation for adaptation to different communication patterns by introducing a novel LPS algorithm. This algorithm has two modules, a message dissemination module in charge of relaying messages from publishers to subscribers and a decision module in charge of adaptation. We first present the message dissemination module, which can be seen as a parametrized hybrid protocol, then we present the decision module used to control the parameters of our message dissemination module.
- we analyze and evaluate our algorithm different communication patterns and decision policies using two case studies, namely a robotic swarm and a mobile social network. Our goal is to minimize the message load (i.e., number of single hop message sends) whilst retaining a high delivery rate (i.e., number of nodes delivering a message which is of interest to them).
- we introduce a simple optimization of our decision-making algorithm, showing that we can reduce the cost of the algorithm.

Our evaluation shows that ALPS outperforms static (e.g., message- or query-centric) solutions under communication pattern variations, while performing comparably without variations, making it an appealing *one-size-fits-all* LPS solution.

1.4. Roadmap

Section 2 provides background information on LPS. Section 3 presents ALPS, and Section 4 describes the setup for our evaluation whose results are presented in Section 5. Section 6 presents related work and Section 7 concludes with a perspective on future research.

2. Location-based publish/subscribe

In *location-based publish/subscribe* (LPS), the delivery of a published message by a subscriber is subject to three *conditions*: besides the usual (1) *content match* between the message and the query of the subscriber, the (2) *location match* between the publisher and the subscriber and the (3) *lifetime* of the message as defined by the publisher condition message delivery.

2.1. When does a location-match occur?

A location match hinges on two notions: (a) *message space* and (b) *query space* (see Fig. 1a). A message space is associated with each published message and defines a geographical range centered around the publisher of that message. This space *moves together* with the publisher. Symmetrically, each subscriber issuing a query has an associated query space moving with it.[1]

For a location match to occur, the subscriber and the publisher must both be located in the intersection of these two spaces, as illustrated in Fig. 1.b. LPS requires that all devices in the network have access to some location service, e.g., GPS, triangulation.

2.2. Persistent messages

The *message persistence* is defined as the lifetime – specified by the application – within which a message is relevant. Once this lifetime has elapsed, the message is considered obsolete and will not be delivered anymore. Persistent messages allow systems to capture real life *events* that are indeed not instantaneous, such as traffic jams, and concerts. So it is important to distinguish between the duration of an event, which could be quite long, and the duration of its *delivery* to a given subscriber, which can be assumed to be instantaneous. In this sense, the notion of persistent message makes it possible to capture the semantics of MANETs-based applications that would otherwise be difficult to express.

2.3. Model

In our model, we define two fundamental concepts used in location-based publish/subscribe namely *messages* and

[1] As suggested by Fig. 1a, the message space and the query space are assumed to be significantly larger than the *transmission range*.