Journal of Network and Computer Applications 32 (2009) 1194–1209

Contents lists available at ScienceDirect

Journal of Network and Computer Applications

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

KEPPAN: Knowledge exploitation for proactively-planned ad-hoc networks

Martín López-Nores *, Jorge García-Duque, José J. Pazos-Arias, Yolanda Blanco-Fernández, Manuel Ramos-Cabrer, Alberto Gil-Solla, Rebeca P. Díaz-Redondo, Ana Fernández-Vilas

Telematics Engineering Department, University of Vigo, ETSE Telecomunicación, Campus Universitario s/n, 36310 Vigo, Spain

ARTICLE INFO

Article history:
Received 20 November 2007
Received in revised form
19 February 2009
Accepted 4 May 2009

Keywords:
Mobile ad-hoc networks
Knowledge exploitation
Communication requirements
Proactive-planning

ABSTRACT

This paper explores the idea of managing mobile ad-hoc networks (MANETs) by the communication needs of their nodes, as a means to facilitate the operation of distributed applications. Specifically, we present a middleware layer that enables reasoning about the multiple possibilities there may exist to ensure satisfiability of certain communication needs. This middleware has been explicitly devised to handle partial and changeable knowledge about the networks, and to guide the search for missing information whenever it cannot conclude whether it will be possible to satisfy some needs. These features provide the basis to implement policies with which to coordinate activities in a MANET, in quest for the configuration that best satisfies the communication needs of its nodes. We provide simulation results to show the comparative advantages of our solution, plus a report of experiments to assess its practicality and usability.

1. Introduction

Computing is witnessing a shift from desktop machines reliant on fixed, wired networks to wireless networks of mobile devices that communicate without infrastructure support. The paradigm of mobile ad-hoc networks (MANETs) is becoming increasingly popular, with applications in diverse areas like military surveillance (Rubin and Zhang, 2003), emergency situations (Jiang et al., 2004), vehicular networks (Blum et al., 2004), sensor–actor networks (Melodia et al., 2007) and many others.

The problem we address in this paper is that the physical mobility of the nodes of a MANET frequently results in sporadic and transient connectivity, which may cause an application to malfunction if it cannot communicate with the nodes that provide the services it needs. To prevent that, many authors (Basu and Redi, 2004; Chang et al., 2004; Goldenberg et al., 2004) considered scenarios in which the nodes move purposefully to preserve static *connectivity requirements*, such as maintaining a biconnected network topology. Unfortunately, these *tethered mobility* (TM) approaches tend to over-restrict the movements of the nodes, which may be very harmful for some applications. As an alternative, the authors of Su et al. (2001), de Rosa et al. (2005) and Härri et al. (2005) opted not to ensure that disconnections do never occur, but to anticipate them using *movement prediction* (MP) techniques, and then warn the applications or their users about the foreseeable problems. This approach can prevent breaks between pairs of nodes that are exchanging information at a given moment; however, it cannot provide those nodes with information on how to maintain the communications, e.g. by establishing multihop routes through other nodes, migrating services onto nodes that stay within range, etc.

Knowing the limitations of previous approaches, some authors have wondered about enabling proactive-planning of MANETs through knowledge exploitation (KP) (Sen et al., 2005), driven by the *communication requirements* of the nodes, i.e. by indications that it should be possible to communicate with certain services at specific times and places. The idea starts out by having the nodes exchange information, including whatever they know about:

- The service provision plans of the different nodes: which services they are going to provide, during what time intervals, etc.
- The elasticity properties of the services, i.e. whether they can migrate from one node to another, be cloned or leased for a period of time, etc.
- The motion profiles of the nodes, i.e. their intended/expected movements.
- Miscellaneous facts like the impossibility to make some moves (e.g. due to the presence of walls), the fact that a given service cannot be provided by certain nodes, etc.

Gathering this information, any node can guess a global picture of the MANET to automatically check the satisfiability of its communication requirements, and derive possible reactions if they cannot be fulfilled. The reactions may imply moving oneself to a specific location, telling others to move, accomplishing

* Corresponding author. Tel.: +34 986813967; fax: +34 986812116.
E-mail address: mlnores@det.uvigo.es (M. López-Nores).

1084-8045/$ - see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.jnca.2009.05.003

service relocations, reconsidering the communication requirements themselves, etc. In any case, it is up to the applications or the users of the different nodes to agree on which reactions to apply and enforce the corresponding rearrangements.

Research into knowledge exploitation in MANETs is still incipient. Even though we can trace back its origins to the wealthy literature of autonomous systems (Dobson et al., 2006), the idea was fully described in Handorean et al. (2004), where the authors sketched the formalisms needed to represent pieces of information about a network and to reason about the knowledge they convey. Later, Mecella et al. (2006) presented an architecture to support collaborative applications in MANETs, focusing on the coordination procedures with which to handle problems with the satisfiability of communication requirements (specifically, in emergency scenarios). Murphy et al. (2001) proposed a knowledge exchange and coordination scheme based on *tuple spaces*, which they proved particularly suitable to furnish the view of a shared memory for distributed applications. Also with tuple spaces, Sen et al. (2007) refined the formalisms put forward in Handorean et al. (2004) and built a database of motion profiles to deal with connectivity problems.

All of the aforementioned studies assume the knowledge available about the networks to be perfect, which is rarely true in real environments. Rather the opposite, the knowledge is most often partial (i.e. incomplete) and changeable. When a node appears in a MANET, it only knows about its own location, services and movements, which yields a very limited view of the network to reason about communication requirements. That view can be augmented, progressively, with information retrieved from other nodes, but there are seldom guarantees about its completeness and freshness. Indeed, it usually happens that a node cannot expose complete information about itself; for example, although there are many cases that a node follows a fixed or predictable trajectory—e.g. nodes attached to rails, unmanned vehicles in military applications (Rubin and Zhang, 2003) or *virtual nodes* whose movement can be controlled at will (Dolev et al., 2004)—complete motion profiles are generally unavailable. This problem was already pointed out in Biskupski et al. (2007), where the authors proposed an abstract agent-based model of self-organizing MANETs, with notions of partiality and changeability in the knowledge available about a network. Another conceptual model was given in Holzer et al. (2007), focusing on the formalization of autonomous and emergent behavior.

The preceding discussion reveals that research into knowledge exploitation remains at a rather theoretical level, since there are no working solutions to the problem of handling imperfect knowledge. Here, we tackle this question by introducing a middleware layer that provides mechanisms for the nodes of a MANET to collaborate in making up the network that best satisfies their communication requirements. This proposal makes contributions related to the following aspects:

- Analyzing what can be concluded about the possibility to satisfy communication requirements in the face of incomplete knowledge about the MANETs.
- Characterizing how the conclusions are affected by changes in the knowledge bases, either additions or modifications.
- Exposing the multiple possibilities there may exist to ensure satisfiability of certain requirements.
- Assessing practicality of the algorithms that process the knowledge and the amount of information that a mobile device has to handle.

Hereafter, we shall focus on the tasks related to knowledge exploitation, i.e. the modeling of the networks and the reasoning enabled about communication requirements. The various means that may be used to gather location information in MANETs and to ensure its trustworthiness have been surveyed in Friedman and Kliot (2006), Wu and Wu (2006) and Marias et al. (2006), respectively. Likewise, mechanisms to accomplish the dissemination of information among different nodes can be sought in Chen and Wu (2003) and Oyabu et al. (2005) or in the literature of tuple spaces (Gelernter, 1985; Freeman et al., 1999; Xu et al., 2006).

Our proposal, called KEPPAN (*knowledge exploitation for proactively-planned ad-hoc networks*), is described in Section 2 along with toy examples illustrating all of its features. Later on, Section 3 includes simulation results to show the comparative advantages of this scheme, plus a report of experiments to assess its practicality and usability. Finally, Section 4 provides a summary of conclusions and the motivation of future work.

2. The KEPPAN middleware

The core of our approach to knowledge exploitation is the scheme shown in Fig. 1, which represents the internal behavior of every node in a MANET. To begin with, through whichever means, a node may receive information about the network at any time. That information serves as the input for a *synthesis* procedure to automatically generate and update discrete models, that capture the envisaged availability of the different services against time and space. Those models are used by an *analysis* procedure whenever it is necessary to check whether it will be possible to satisfy the node's communication requirements. The synthesis and analysis algorithms run independently of one another, in such a way that the latter always work with the last stable update of the discrete models.

The main point of Fig. 1 is to note that the analysis of communication requirements can return three different results:

- [*positive*]: the requested services will be available when and where indicated;
- [*negative*]: the requirements cannot be satisfied in the envisaged network;
- [*uncertain*]: the knowledge available does not suffice to conclude whether it will be possible to satisfy the requirements.

According to these outcomes, it would be odd to reason about MANETs in terms of the classical Boolean logic, because it cannot reflect uncertainty: everything is either true or false, and this can lead to erroneous conclusions about the satisfiability of the communication requirements. To avoid taking for false what is indeed unknown, we advocate the use of formalisms based on Kleene's (1952) three-valued semantics, which enables an explicit differentiation of what is known to be true ($\mathscr{T}$, meaning *allowed*, *possible*, *reachable* or *available*), what is known to be false ($\mathscr{F}$, meaning the opposite), and what is still unknown ($\perp$) (Fig. 2).

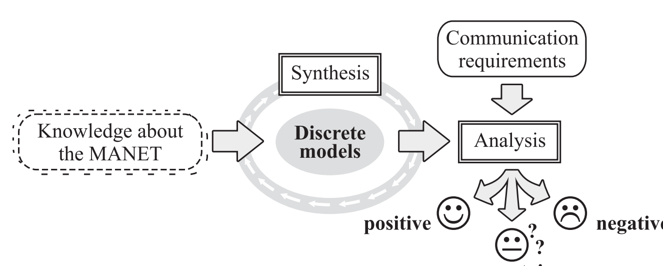

Fig. 1. Synthesis of models and analysis of communication requirements.