Applied Soft Computing 40 (2016) 416–426

Contents lists available at ScienceDirect

Applied Soft Computing

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

Optimizing groups of colluding strong attackers in mobile urban communication networks with evolutionary algorithms☆

Doina Bucur[a], Giovanni Iacca[b,*], Marco Gaudesi[c], Giovanni Squillero[c], Alberto Tonda[d]

[a] Johann Bernoulli Institute, University of Groningen, Nijenborgh 9, 9747 AG Groningen, The Netherlands
[b] INCAS[3], Dr. Nassaulaan 9, 9401 HJ Assen, The Netherlands
[c] Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129 Torino, Italy
[d] INRA UMR 782 GMPA, 1 Avenue Lucien Brétignières, 78850 Thiverval-Grignon, France

ARTICLE INFO

Article history:
Received 7 July 2015
Received in revised form 9 September 2015
Accepted 9 November 2015
Available online 17 December 2015

Keywords:
Cooperative co-evolution
Delay-Tolerant Network
Evolutionary algorithms
Network security
Routing

ABSTRACT

In novel forms of the Social Internet of Things, any mobile user within communication range may help routing messages for another user in the network. The resulting message delivery rate depends both on the users' mobility patterns and the message load in the network. This new type of configuration, however, poses new challenges to security, amongst them, assessing the effect that a group of colluding malicious participants can have on the global message delivery rate in such a network is far from trivial. In this work, after modeling such a question as an optimization problem, we are able to find quite interesting results by coupling a network simulator with an evolutionary algorithm. The chosen algorithm is specifically designed to solve problems whose solutions can be decomposed into parts sharing the same structure. We demonstrate the effectiveness of the proposed approach on two medium-sized Delay-Tolerant Networks, realistically simulated in the urban contexts of two cities with very different route topology: Venice and San Francisco. In all experiments, our methodology produces attack patterns that greatly lower network performance with respect to previous studies on the subject, as the evolutionary core is able to exploit the specific weaknesses of each target configuration.

1. Introduction

The so-called *Social Internet of Things* calls for nearly ubiquitous communicating devices. There is today a need to integrate low-cost, low-power devices to support networking services in more effective and efficient ways. In such a scenario, new solutions are continuously developed and deployed, while approaches that just a few decades ago were used only in highly complex, niche applications are now literally brought down to earth—Delay-Tolerant Networks (DTNs) are a technology originally developed for space communications that, over the years, made its way down to quite mundane applications [1]. Emerging technologies and applications are posing serious problems to designers. In most cases there is not enough time to thoroughly validate them, or even to simply analyze their possible failures and problems. Engineers are forced to resort to their experience to choose heuristics that look reasonable, and then observe the actual outcome from real applications. Security in DTNs is a paradigmatic case: such networks need to remain open to all willing participants, and few malicious participants may try to disrupt communications, for instance, routing no messages to other nodes or injecting large number of messages into the network. While such a risk is plausible, precisely assessing DTNs' vulnerabilities is hard.

This paper focuses precisely on evaluating the amount of damage that can be caused to a DTN by a group of synchronized attackers with deep knowledge about the network. Given a scenario, we propose to optimize attackers for minimizing the performances of the network using a heuristic methodology. It is important to note that the adoption of such methodology is more a necessity than a choice: determining the most effective attack for a given network was proven to be NP-hard [2], the complexity and number of variables involved in the problem preclude the use of formal techniques and the size of the scenarios prevent exhaustive analyses.

☆ This paper is an extended, improved version of the paper *Black Holes and Revelations: Using Evolutionary Algorithms to Uncover Vulnerabilities in Disruption-Tolerant Networks* presented at EvoComNet2015 and published in: Applications of Evolutionary Computing, Proceedings of 18th European Conference, EvoApplications 2015, Copenhagen, Denmark, April 8–10, 2015, LNCS 9028, pp. 29–41, Springer, 2015.

* Corresponding author.
E-mail addresses: d.bucur@rug.nl (D. Bucur), giovanniiacca@incas3.eu (G. Iacca), marco.gaudesi@polito.it (M. Gaudesi), giovanni.squillero@polito.it (G. Squillero), alberto.tonda@grignon.inra.fr (A. Tonda).

http://dx.doi.org/10.1016/j.asoc.2015.11.024
1568-4946/

The idea of using heuristic methods to disprove a property of a system when formally proving it is not possible, is not a novelty in itself. The simplest approach, namely random sampling of the parameter space, is often used under the assumption that the effort employed failing to find a counter example may be sensibly linked to the degree of confidence that a counter example does not actually exist.

Repeated random sampling has also be used as a means to estimate numerical quantities when complexity and dimensionality of a problem impedes the application of analytic analyses. In the specific case of DTNs performance, it has been considered in [3], although limited only to small networks and attackers with no information about the environment. However, random sampling is unlikely to provide any interesting result when the goal is to detect a very specific corner-case scenario, such as the damage caused by specialized attackers that are fully aware of the network characteristics. Finally, when the search space is too vast, even the effort required to get a significant sampling could be excessive.

In this work, we move forward from random sampling by using an evolutionary algorithm (EA) to *optimize the attackers' parameters* in order to inflict the maximum possible damage to the network. We overcome the limitations of random sampling by using the capability of the EA to drive random search towards specific regions of large search spaces. Furthermore, we extend the features of a classical evolutionary algorithm to enable it to find a *team of colluding attackers*. As the members of such a team cooperate in order to maximize the cumulative damage, even at the expense of the damage caused by each single attacker, the approach is a form of *cooperative co-evolution*, an open area of research for which very few successful strategies have been found so far.

We tested the proposed methodology on medium-sized networks describing urban scenarios with different topologies, where a number of agents, i.e., the network nodes, move realistically. The results clearly demonstrate the efficacy of the approach: we found scenarios where even few (up to 10% of the total network size), highly optimized attackers can reduce the global data delivery in the network by over 90%, when compared to the network with no attackers. We also observed that the composition of the attacker team obtained by evolution changed when cooperative co-evolution is used, demonstrating that such scheme leads to synergistic solutions not found by classical evolutionary algorithm.

The rest of the paper is organized as follows: the next section summarizes the research background; Section 4 details the proposed methodology; Section 5 reports the experimental evaluation; Section 3 surveys the related work; finally, Section 6 concludes the paper.

2. Background

This section first gives an overview of the application domain of Delay-Tolerant Networks. Sections 2.1 and 2.2 describe the paradigm of routing in DTNs, and First Contact, the DTN protocol under study. Section 2.3 summarizes the mobility model that an urban DTN node follows, from the literature. Section 2.4 describes the two main types of security attacks relevant: black hole and flooding attacks. Finally, Section 2.5 gives an overview of the EA field.

2.1. Delay-Tolerant Networks: performance objectives

Delay-Tolerant Networking was designed to cater for *message routing* in practical applications with *heavy node mobility*. In such applications, the connectivity pattern between nodes in the network can be either predictable or unpredictable with time. An example DTN with predictable connectivity is that of a mixed terrestrial-and-space network where some of the nodes are Low-Earth Orbiting Satellites, and the rest are ground users; this was the application for which DTN-specific routing protocols were originally designed [4]. More recently, DTNs have also been proposed in scenarios with nearly unpredictable connectivity. This is the case of animal-tracking applications [5] and opportunistic urban networks. An example of the latter is the 30-bus experimental DieselNet [2], in which urban vehicles constrained to city roads act as mobile message routers. It is urban scenarios with unpredictable connectivity that we study in this paper.

Given an application scenario, the main *performance factors* for a DTN message-routing protocol quantify the protocol's ability to route messages in that scenario, and the timeliness of the routing:

Delivery rate The percentage of messages injected in the network by nodes which were successfully delivered to their destination nodes.

Message delay The average time interval between a message injection in the network until its delivery.

2.2. DTN routing: the First Contact protocol

A DTN routing protocol essentially implements a logic to achieve message routing in the mobile network, end-to-end from the source of a message to its destination, over a connectivity graph which varies in time and is by nature disconnected. Given these scenarios, the protocol logic cannot be based on standard distributed algorithms for computing shortest paths end-to-end in a graph: the routing cannot converge on correct routes when the network graph is highly dynamic. Instead, DTN message communication on a path between source and destination include long-term *storage* of the message in the nodes' (finite) node buffers, until an opportunity for further delivery of the message arises. This ability to safely *delay* the forwarding of a message is typical of DTN routing protocols.

DTN protocol design follows a simple taxonomy based on the following features:

Network knowledge A protocol aiming to compute optimal paths at a node would be helped if the node is able to predict the future network conditions: the pattern of contact with other nodes, the set of nodes with congested buffers, and the pattern of traffic demands. While network knowledge may be acquired in practice by an attacker via monitoring the network, many protocols cannot assume any (i.e., are *zero-knowledge*).

Message replication *Forwarding protocols* simply route the original message through the network. *Replicative protocols* introduce into the network a number of copies of each original message, each of which is then forwarded independently with the aim that at least one copy reaches the destination.

We study here one of the simplest and most common DTN routing protocols, namely *First Contact* (FC) [4]. FC is zero-knowledge and forwarding; it routes messages opportunistically using any available contacts with other nodes. A single copy of each message in the network exists at a time, and it is forwarded to the first available contact (if more contacts are available, one is chosen randomly among all the current contacts). On simple network topologies, FC was shown to have performance comparable to partial-knowledge protocols; this degrades in complex topologies to varying degrees, depending on the network load.