Computer Communications 73 (2016) 66–79

Contents lists available at ScienceDirect

Computer Communications

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

Stalk and lie—The cost of Sybil attacks in opportunistic networks

Sacha Trifunovic [a,*], Andreea Hossmann-Picu [b]

[a] Computer Engineering and Networks Laboratory, ETH Zurich, Switzerland
[b] Computer Networks and Distributed Systems Group, University of Bern, Switzerland

ARTICLE INFO

Article history:
Available online 19 May 2015

Keywords:
Sybil attack
Opportunistic networking
Trust
Community structure
Social network

ABSTRACT

Opportunistic Networks are envisioned to complement infrastructure-based communication in overloaded cellular settings, in remote areas, or during and immediately after large scale disasters. On account of their highly distributed and dynamic nature, as well as of their dependence on the honest cooperation of nodes, Opportunistic Networks are particularly vulnerable to Sybil attacks. In a Sybil attack, a node assumes multiple identities and attempts to form many links to the rest of the network, with the aim of gaining access to resources, influencing the network, circumventing detection of misbehavior ("spread the blame"), etc.

Sybil attacks have been studied extensively in the context of distributed systems and online social networks and many defense mechanisms have been proposed based on the graph structure of these systems. However, the Opportunistic Networking setting brings new challenges, specific to the network conditions: forming links may require significant resources from the attacker (e.g. time, speed, multiple devices, etc.), and each link is ephemeral. It also brings new opportunities for the attacker such as the possibility to manipulate the social graph. In this paper, we study the types and effectiveness of Sybil attacks that are possible in Opportunistic Networks, under various resource constraints on the one hand, and attack boosting graph faking attempts on the other hand. We use four state of the art Sybil defense algorithms to evaluate each attack and graph faking attempt based on the influence the attacker can gain through it. We then introduce three defensive measures against these graph manipulation attempts and evaluate their effectiveness. Finally, we quantify the resources required by an attacker to influence one or several communities. We find that Sybil attacks, even with relatively unconstrained resources, are much harder to implement in the Opportunistic Networking setting, due to the link establishment mechanisms using mobility. An attacker needs to invest several hours per day to successfully infiltrate a community. Additionally, this naturally limits the amount of communities an attacker can infiltrate. Instead of disregarding the decentralized nature of opportunistic networks as a curse we can rely on its underlying mobility to naturally defend us against Sybil attacks in such networks.

1. Introduction

In Opportunistic Networks (OppNets), mobile phone users may cooperate to complement existing wireless communication services (cellular, Wi-Fi) and to enable communication in case of failure or lack of infrastructure (disaster, censorship, remote areas). Wireless peers communicate when they are in proximity (in *contact*), forming an impromptu network, whose connectivity graph is highly dynamic and only partly connected. Using redundancy (e.g., coding, replication) and smart mobility prediction schemes, data can be transported over a sequence of such contacts, despite the lack of end-to-end paths [1,2].

The feasibility of communication over an OppNet highly depends on the *honest* cooperation of the nodes, by actively contributing resources to run the network instead of only receiving the communication service passively from the system. The issues of fairness [3–5], benign selfishness (non-cooperation) and of motivating users to cooperate have already received a fair amount of attention from researchers: there exist many studies on the effects of selfishness [6,7], as well as a number of incentive systems to ensure participation in the network [8,9]. However, the possibility of more pro-actively malicious users has hardly been considered so far.

One of the most powerful and most versatile ways to disrupt the incentive and/or security mechanisms of cooperative systems is the *Sybil attack*, in which the adversary creates many fake identities (Sybils) and uses them to undermine the system's normal operation. Sybils can be used to obtain a large share of resources from resource allocation algorithms, to bias recommendation or rating systems, to

* Corresponding author. Tel.: +41 446320220.
E-mail address: sascha.trifunovic@gmail.com, trifunovic@tik.ee.ethz.ch (S. Trifunovic).

http://dx.doi.org/10.1016/j.comcom.2015.04.007
0140-3664/

intercept seemingly disjoint routing paths, to circumvent the detection of misbehavior by "spreading the blame" [10], etc.

Sybil attacks have been studied extensively in the context of distributed systems and online social networks [11–15]. However, the OppNet setting brings both new possibilities as well as new challenges, specific to the network conditions. On the one hand, the highly distributed and dynamic nature of OppNets makes them easy targets for Sybil attacks (as it is practically impossible to rely on a single centralized authority to certify legitimate users). On the other hand, a Sybil attack requires the establishment of a number of connections (known as "attack edges") between the adversary's identities and the rest of the network; in OppNets, where the network structure is inherently social as the links are by-products of node mobility, forming intentional links may require significant effort, as the two connected peers must be physically co-located for a significant amount of time. We can use this property of OppNets, i.e., its underlaying social network, to perform state of the art social Sybil defense on the network structure itself instead of some overlaid friendship graph.[1]

Our goal in this paper is to provide a thorough exploration and dissection of the *procedure of carrying out Sybil attacks in OppNets*. We first explain how the social contact graph of OppNets, the basis of social Sybil defense, is built and how an attacker can manipulate this graph and fabricate Sybil identities. We then analyze the impact of these graph manipulation attempts on the strength of the Sybil attack. To counter graph faking attempts, we propose three different defense mechanisms to limit their benefit to an attacker. Finally, we quantify the amount of effort and resources an adversary needs to spend to successfully infiltrate Sybil identities by using four state of the art social Sybil defense algorithms. For our analysis we use a community structure based mobility model as well as four real world mobility traces representing distinct settings and collected using distinct technologies. Consistently among all traces, we show that, despite the mostly disconnected state of an OppNet, Sybil attacks are much harder to implement here, even under generous resource allowances. An attack requires an investment of several hours a day and is naturally limited to one or a few communities. This is mainly due to the link formation mechanisms via mobility, which demand continuous effort from an attacker.

This work extend our previous analysis of the cost of Sybil attacks in OppNets [16] by (i) considering three additional Sybil defense algorithms, (ii) analyzing the effect of an attacker's attempt to fake the social contact graph, (iii) proposing and analyzing three different defense mechanisms against graph manipulation attempts, and (iv) by verifying our results with two additional contact traces with complementary properties.

The rest of the paper is organized as follows: In Section 2, we present the state of the art in Sybil defense mechanisms and discuss how and whether they apply to OppNets. Next, in Section 3, we review the possible variations on how to implement each of the three main elements (ID fabrication, link creation, and link manipulation) of a Sybil attack in the OppNet context. Then, we show our methodology for assessing and comparing the effect of each of these variations on the attack strength in Section 4. To begin our evaluation, we analyze the effect of link manipulation in Section 5 using a simple test scenario as well as four real world traces and a mobility model. In Section 6 we propose and analyze three defense mechanisms to cope with such link manipulation attempts. Finally, in Section 7, we use the four real traces and the mobility model to quantify the strength of Sybil attacks on OppNets, depending on used resources and on how the attack is implemented. Finally, we briefly discuss our results in Section 8 and conclude in Section 9.

[1] An overlaid friendship, trust, or reputation graph may still be used in addition to improve the overall Sybil defense.

2. Background and related work

Carrying out a (simple) Sybil attack is generally quite straightforward in traditional distributed systems as opposed to OppNets. Therefore, research on the Sybil attack focuses on devising effective detection/defense mechanisms, with little attention on how to best implement the attack. The goal of Sybil detection is to accurately identify Sybil identities (i.e. accept all legitimate identities, but no counterfeit ones). For our purpose of dissecting and comparing Sybil attacks and their effectiveness in OppNets, algorithms for Sybil detection can be very useful in assessing the cost-benefit tradeoff of each element of an OppNet Sybil attack.

The state of the art in Sybil detection consists in leveraging the social network underlying the distributed system under consideration. Assuming the Sybil identities can create only a limited number of connections to honest nodes (i.e. attack edges), the resulting graph will then consist of two regions, loosely connected to each other: the honest nodes and the Sybils. Depending on the structure of the honest region, defense mechanisms do one of the following:

(i) Identify and exclude all Sybils (universal Sybil defense). If the honest region of the graph is relatively well, but flatly connected internally and fast mixing, then the Sybils are easily detectable via community detection or similar algorithms. Many Sybil defense solutions are based on this idea [11–15].
(ii) Enable honest nodes to white-list a set of nodes of any given size, ranked according to their trustworthiness. Since it has been shown that, in practice, the honest region often has internal structure (e.g. "communities" of tightly-knit nodes, relatively loosely connected with one another), the goal of Sybil defense has shifted accordingly [17,18].

OppNets have been shown to have a very strong social component already at the network layer [1,2] (due to the nodes/phones being near perfect proxies for the human users), therefore the abovementioned Sybil defense solutions should be directly applicable to the OppNet setting. However, this requires the representation of the network as a static graph. In OppNets, communication occurs between two wireless peers whenever their mobility brings them within radio range of each other. Deriving a graph from such contacts is not straightforward: the contacts are short-lived and the timing information must be stored for each edge. This makes the resulting time-varying graph very cumbersome and unintuitive. A more practical and widely accepted representation can be obtained by aggregating contacts into a (static) weighted graph, with weights derived from contact statistics (e.g. frequency, duration, age of last contact, or combinations thereof). It is *this* graph that was shown to also reflect to high extent the social relationships among the users [19], and on which existing Sybil defense schemes are readily applicable.

Some literature on Sybil detection exists also for network environments similar to OppNets. For example, adaptations of the above algorithms have been proposed for small mobile networks with very sparse social ties, as resulting from secure pairing [20]. However, the feasibility of such a pairing-based network is questionable, as it requires incentives for the users to actually pair their devices. Using the OppNet structure to establish trust has been proposed before, but the ranking is performed via an algorithm based on node similarity [21]. In the context of mobile ad hoc networks (MANETs), Sybil detection schemes are mostly based on peculiarities of this environment (some of which are also exhibited by OppNets). However, such solutions usually require specialized hardware [22–24] or are based on strong, unrealistic assumptions [25].

To assess the effectiveness of OppNet Sybil attacks, we will apply the latest Sybil defense algorithms from the second category as well as two similarity based rankings to the weighted OppNet graph. We provide a more detailed description of how these algorithms operate