Computer Communications 000 (2016) 1–17

Contents lists available at ScienceDirect

Computer Communications

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

Evaluation of link layer mobility in Ethernet networks

Nuutti Varis*, Jukka Manner

Aalto University School of Electrical Engineering, Department of Communications and Networking, P.O. Box 13000, 00076 Aalto, Finland

ARTICLE INFO

Article history:
Received 3 March 2015
Revised 6 June 2016
Accepted 9 July 2016
Available online xxx

Keywords:
Ethernet
Distributed Hash Table
Mobility
Experimental evaluation
System design

ABSTRACT

Ethernet is the prevalent link layer mechanism in data communication networks. Today, it is used in enterprise and home networks, data center networks, telecommunication networks, and in various industrial deployments. In addition, more and more hosts are becoming mobile, connecting to Ethernet networks via Wi-Fi, or through virtualization solutions in data centers. Unfortunately, the Ethernet protocol suite itself has no generic support for host mobility.

In this paper, we evaluate the effects of host mobility in Ethernet networks using real-time emulation. We compare routing bridges, an IETF-driven Ethernet frame forwarding protocol, with our DBridges design. DBridges is an evolution of the routing bridges standard, integrating a one-hop Distributed Hash Tables (DHT) scheme into the protocol. Our solution offers improved scalability characteristics in Ethernet networks, as well as enhanced support for host mobility.

Our evaluation shows that while host mobility without any explicit signaling will remain a best-effort service in Ethernet networks, we can significantly improve its efficiency and reliability in certain use cases, while providing improved scalability and safety properties. We also show that while the host mobility support in DBridges suffers from transient forwarding loops, the problem is exceedingly rare in real networks and is mitigated by the base functionality in routing bridges.

1. Introduction

Ethernet has found its way into most network deployments in the world. It is the de facto link layer protocol for large-scale data centers, in home networks to connect various devices to the Internet through a router, and increasingly in access and aggregation segments of telecommunication networks. Arguably, the most prominent feature of Ethernet that has driven its adoption is simplicity. Each Ethernet device has a unique address that is preconfigured by the manufacturer of the device. This allows the basic Ethernet network to function without configuration; each device can receive messages via the unique Ethernet address, and host locations are discovered from frames arriving on an interface (e.g., a switch port).

Data center networking has emerged as a significant driver for wired Ethernet development in the recent years. Inside the data center, virtualization is commonly used to provide flexible provision of services and applications to customers. Among other benefits, virtualization as a mechanism allows data center operators to dynamically adjust the resource allocation inside the data center, leading to potential energy savings and resource efficiency [1,2]. A large part of the flexibility of virtualization stems from the possibility to dynamically migrate virtual machines to different parts (e.g., racks) inside the data center. From the perspective of the network, the *live migration* [3] of virtual machines often presents itself as a host mobility event.

Conventional 802.3-based Ethernet does not mandate any explicit signaling when a host changes its location; i.e., it allows hosts to be silent. Without an explicit handover procedure, the network will not be able to deliver frames to the mobile host until it has been active in the new location. While standard Ethernet does not have any explicit signaling when a host attaches or detaches from a network, supplemental protocols, such as 802.1X [4] or in many cases the 802.11 wireless Ethernet protocol family can be used to monitor the connection state of Ethernet devices. In the case of virtualization, the live migration process typically ends with an Address Resolution Protocol (ARP) announcement sent to the network to update the location information (i.e., MAC learning tables) in the Ethernet switches.

This paper evaluates and compares two different methods for supporting host mobility in Ethernet networks directly on the link layer. Our efforts concentrate on evaluating link-layer behavior that in the Ethernet family of protocols as a whole, scoping out the mobility management protocols typically in use with wireless de-

* Corresponding author.
E-mail addresses: nuutti.varis@aalto.fi (N. Varis), jukka.manner@aalto.fi (J. Manner).

http://dx.doi.org/10.1016/j.comcom.2016.07.005
0140-3664/© 2016 Published by Elsevier B.V.

vices (e.g., MobileIP [5,6], Inter-Access Point Protocol (IAPP) [7], and many IEEE ratified wireless standards [8,9]), or protocols designed for media-independent handovers [10]. Consequently, host mobility support in wired Ethernet networks can be described on a high level as a method for updating the location of the mobile host (e.g., host MAC address) on switches in the network. To emulate host link-layer behavior consistent with real-world scenarios, we use data collected from wireless devices to implement a mobility model for the evaluation.

Conventional Ethernet switching uses passive location information updates to support host mobility in the network, i.e., each Ethernet frame emitted in the network updates the location of the host in switches without any explicit signaling. The passive updating scheme is directly integrated into the base protocol operation of all Ethernet switches and offers flexible updates of location information in the network. However, as pointed out in the literature [11–14] and practical use cases, conventional Ethernet networks have significant scaling issues that limit the number of active hosts and the overall throughput of the network.

Active location information update mechanisms have been proposed in the past by various research papers and industry efforts. Active location updates, signaled through an external mechanism, are generally used to either eliminate or minimize the flooding behavior in Ethernet networks used by frame forwarding and higher layer address resolution (e.g., ARP). The active location information update mechanism can also be used to support mobility in the network by signaling location updates to switches in the network instead of learning the location of hosts passively from forwarded frames.

Our previous work [15] introduced DBridges, a design that merges the one-hop Distributed Hash Table (DHT) scheme from SEATTLE [16] with the IETF routing bridges standard [17], and evaluated our design with static hosts. In this paper, we identify problem areas in our original design when the network contains a significant amount of host mobility. We also present an evolution of the original DBridges design that significantly improves the behavior of the system in networks with high degree of host mobility. Finally, we also evaluate our new system design in relation to the signaling characteristics of host mobility on the Ethernet link-layer by analyzing network behavior from several perspectives:

- The continuous signaling load in the network during mobility events,
- the effect and overhead of host mobility on the signaling,
- the effect of host mobility on overall path length of frames,
- the length of convergence periods in the network, and
- the connection interruptions on the host application layer.

The evaluation is performed on a network topology that resembles a wireless access network. Our evaluation results show that in many cases, DBridges significantly improve the efficiency and reliability of the signaling in Ethernet networks with mobile hosts. Furthermore, we show that in worst case situations, our solution performs similarly with conventional Ethernet networking while simultaneously reducing the signaling traffic in the network. We also describe a forwarding loop problem in the SEATTLE DHT scheme and show how the problem is exceedingly rare in real networks and mitigated by our choice of routing bridges as the base system.

The rest of the paper is structured as follows. Section 2 covers some of the work done on Ethernet mobility and gives a brief overview of the various designs that support Ethernet mobility. Section 3 covers the functionality of Ethernet networks in relation to host mobility and introduces the passive and active location information update mechanisms in more detail. Next, Section 4 describes routing bridges, the Ethernet frame forwarding protocol used as a foundation for our design. We will also discuss the functionality and issues with host mobility in terms of routing bridges. Section 5 introduces host mobility support in DBridges and gives a high-level overview of the DHT scheme. We also discuss some of the inherent problems with the active location information update mechanism. Sections 6 and 7 include our evaluation, starting with an overview of our evaluation setup and environment, followed by a discussion of our results. Finally, Section 8 concludes the paper.

2. Related work

There have been several proposals for solving scalability in Ethernet networks where a part of the problem set is host mobility. Most of the other work includes host mobility as one of the requirements for the system. The primary motivator for the requirement is the prevalent use of virtualization, and as a consequence, the live migration of virtual machines. However, the majority of the work concentrates on evaluating other issues with Ethernet, such as bandwidth, fault tolerance, security, cost-efficiency, and segment size (both in terms of number of active hosts, and number of switches). The following section discusses different domains, where Ethernet link-layer mechanisms are typically involved in the host mobility solution of the network.

In addition, the increase in cloud computing has brought Software Defined Networking (SDN)-based network architectures to the academic forefront, and shifted focus away from a traditional distributed network architectures such as various link-state protocol-based approaches (including DBridges). Recently, Network Function Virtualization (NFV) [18] has gained traction due to its promise of increased elasticity in the network. As a consequence a significant body of recent work has been published on the research problems surrounding NFV. In this context, mobility is often treated as a higher level construct (e.g., migration of virtual network functions) than what is discussed in this paper (e.g., the link-layer signaling effects of host mobility events).

In contrast, we concentrate on evaluating the effects of host mobility on our solution and compare it to a conventional Ethernet solution (i.e., routing bridges). To our knowledge, our evaluation represents one of the first in-depth studies of the network effect of mobility in Ethernet networks. Consequently, our related work goes over proposals that deal with host mobility in Ethernet networks, however relatively few actually evaluate the effect of mobility. The majority of the mobility management research that considers experimental evaluation is done in the domain of wireless Ethernet, and as such is not directly comparable to our work.

Using Ethernet as part of the mobility solution in the context of telecommunications networks has been proposed in the past. Mobile Ethernet [19] offers either a centralized or a fully distributed host location information storage in a hierarchical ring-like topology. The hierarchical constraint is used to reduce the amount of conventional MAC learning done in the network and to segment the network around a ring-like core network. Mobile Ethernet evaluates the overhead of two different frame forwarding schemes in terms of generated traffic concentrated on the core network switches. However no evaluation of other scaling characteristics is done.

Recently, there has been significant academic activity surrounding the forthcoming 5G family of cellular data communication standards. One of the key research topics in 5G standardization is to improve the elasticity and efficiency of the telecommunications network architecture. The presented works typically offer benefits through NFV-based solutions, or by simplifying the user data and control planes. Ameigeiras et al. [20] propose a two-layer approach to solve the mobility-related inefficiencies in the evolved packet core. They use link-layer mobility (through encapsulation or frame rewriting on the edges of the network) on the access network, and