Ad Hoc Networks 25 (2015) 595–605

Contents lists available at ScienceDirect

Ad Hoc Networks

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

Designing energy-efficient link aggregation groups

J. Galán-Jiménez [a,*], A. Gazo-Cervero [b]

[a] Computer Science and Communications Engineering Department, Universidad de Extremadura, Spain
[b] Department of Computer Science, University of Oxford, United Kingdom

ARTICLE INFO

Article history:
Received 6 February 2014
Received in revised form 22 September 2014
Accepted 3 November 2014
Available online 10 November 2014

Keywords:
Energy efficiency
Link aggregation groups
Sleeping
Rate adaptation
Equivalent energy function

ABSTRACT

Link high availability and traffic load balancing features are key requirements in today's data centers networks. Network link aggregation techniques are quite commonly used to provide those features. However, these techniques often come with a cost in terms of energy efficiency, when network links face low traffic situations or even active/passive configurations, because they still consume almost the same amount of energy as if they were operating at their full rate. Most of the network equipment currently deployed is not yet able to take advantage of saving energy in scenarios where network links are not used at their full capacity. By exploiting the use of existing energy efficient techniques for single links, this work presents a study in which an impact is quantified in terms of energy saving opportunities when link aggregation techniques are in use. Optimal link aggregation configurations are determined and energy savings are assessed using real network traffic datasets. Results show that adding links capable of using sleep mode to link aggregation groups is a more energy-efficient technique compared to using less links in link aggregation groups, even if they had more advanced energy-saving features like rate adaptation capabilities.

1. Introduction

It is well known that reduction of energy consumption of network equipment is an open issue in wired networking research [1] due to the environmental impact of producing energy associated with the operation of current ICT equipment [2]. Moreover, provisioning of high availability (active/passive configuration) and load balancing (active/active configuration) is nowadays crucial to corporate and data center networks [3]. To this extent, link aggregation techniques are used to provide active load balancing and redundancy between two nodes through the use of a discrete number of single links.

However, these methods can be considered as inefficient from the energy efficiency point of view due to the amount of energy consumed by the active links in the LAGs (Link Aggregation Groups) when traffic load is low or in absence of failure.

In the last few years, several works have been proposed to improve energy efficiency in single links [1]. By exploiting the use of these techniques, LAGs could incorporate energy efficiency extensions whilst providing both high availability and load balancing features.

One approach is to put the links of the LAG to sleep or wake them up individually (sleeping capabilities) whilst another one is to adapt the operation rate of the links in the LAG as a function of each link traffic load (rate adaptation capabilities). In both cases, the energy consumption derived from the link aggregation could be reduced.

For the former case, the IEEE 802.3az standard (EEE, Energy Efficient Ethernet) could be applied [4]. This

* Corresponding author.
E-mail addresses: jaime@unex.es (J. Galán-Jiménez), alfonso.gazo@cs.ox.ac.uk (A. Gazo-Cervero).

http://dx.doi.org/10.1016/j.adhoc.2014.11.005
1570-8705/

standard is based on the action over sending and idle intervals to be able to connect or disconnect some of the components of a single network interface. For the latter case, the hardware-based Rate Adaptation method could be used [5]. With this technique, each single link in the LAG can be configured with a particular discrete number of energy levels [6].

From the three areas defined in [7] to categorize methods that improve energy efficiency in wired communication networks (network design, system design and protocols design) we focus on the network design process. We aim to evaluate, at this stage, the impact of having a specific number of links in a LAG and if the technology implemented by those links is relevant in terms of energy savings. This evaluation has been carried out through an experimental method based on problems formulation and optimization.

2. Related work

The technique based on adding more line cards to network devices and using link aggregation is traditionally used to achieve a bandwidth increase and to provide network high availability. If a particular link belonging to the link aggregation group fails, traffic load is distributed among the remaining set of links.

Link aggregation provides a way to increase bandwidth and add resilience between two network devices. One of the link aggregation techniques lies in gathering a set of links which connect two devices to form an single logical link known as Link Aggregation Group (LAG), bond, trunk, bundle or EtherChannel.

In order to do this, both proprietary and standards-based protocols have been developed. PAgP (Port Aggregation Control Protocol) [8] stands out among the proprietary ones, which is used to connect Cisco devices. LACP (Link Aggregation Control Protocol), in turn, is the link aggregation protocol defined in IEEE 802.3ad standard [9].

However, as there are only two individual devices (switches or routers) at the end of the aggregated link, a complete loss of connectivity would be produced if one of these two devices fails. To prevent this possible situation, different link aggregation techniques focused on device-level redundancy are also used. Several protocols, such as VRRP (Virtual Router Redundancy Protocol) [10] or STP (Spanning Tree Protocol), defined in 802.1D standard [11] have been proposed. However, network equipment vendors have recently shown interest in incorporating alternative mechanisms to solve some of the problems identified when using these techniques, such as the problem of potential reduction of available bandwidth derived from blocked ports by STP. IETF TRILL protocol (Transparent Interconnection of Lots of Links) [12] and SPB (Shortest Path Bridging) technology, specified in IEEE 802.1aq standard [13] are two examples which have already been deployed by most of the network equipment vendors despite their youth.

Fig. 1 shows the relationship between the different features provided by link aggregation techniques and the protocols developed in this scope. Four features have been identified: link-level redundancy, device-level redundancy, capacity aggregation and high availability.

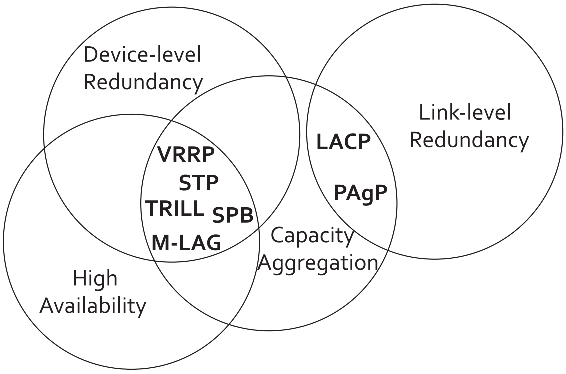

Fig. 1. Link aggregation techniques and technologies.

Some works address the problem of energy efficiency in link aggregation, focusing on reducing the wasted energy when all the links in the LAG are active and the global link load is low.

In this way, several methods have proposed to save energy in aggregated links through the possibility of switching off single links in the LAG during low utilization periods. In [14], the authors propose that the number of active links in a LAG must be adapted to the current traffic load. In situations with low traffic load, only some of the links in the LAG would remain active whilst the rest would enter a low energy mode (sleeping or standby).

This idea is further studied in [15,16], where a set of experimental results is presented. Transitions between active and sleeping modes need a global coordination among the links using LACP, which results in a necessary time interval (hundreds of milliseconds) to be carried out [17].

The authors of [18] assume the use of a network management system which is able to choose the links in the LAG depending on the traffic profile. Thus, given a network topology and a traffic demand matrix as system inputs, an optimization problem is solved to maximize the number of links to put to sleep, assuring enough provisioning to distribute traffic every instant. Three different methods are proposed to reach the optimal solution. As energy savings achieved by the three methods are similar, authors encourage network operators to use the simplest one.

Differently from [18], the work of [19] propose a method to decide how to act individually over a particular link depending on its utilization, i.e. by putting it to sleep or waking it up. A fixed link utilization threshold (90% during simulations) is used to determine the link to activate or switch off at every moment (one single link each time). With data from real traces collected during one month of observation, simulation results achieve up to 86% of energy savings for a threshold of 90%, although there is a low risk of traffic overflow.

In order to improve their method and reduce the risk of traffic overflow, the authors propose another approach, which considers both the link utilization threshold and the number of active links in the LAG (1–3) [20]. In this