

Real-Time Performance Analysis of Adaptive Link Rate

Baoke Zhang, Karthikeyan Sabhanatarajan, Ann Gordon-Ross*, Alan George*

HCS Research Lab, ECE Department, University of Florida
{zhang,sabhanatarajan,ann.george}@hcs.ufl.edu

* Also with the NSF Center For High-Performance Reconfigurable Computing at the University of Florida

Abstract—High speed links are widely deployed in modern day computer networks to meet the ever growing needs for increasing data bandwidth. However, with the increase in the link rate, the power consumption of the network interfaces increases exponentially, compounding growing concerns about network power consumption. Fortunately, network traffic characteristics show that rapid link rates are not always required. During times of reduced network traffic, the Adaptive Link Rate (ALR) mechanism allows link rates to be reduced with little impact on network performance. Current research has focused on policies to control when and how to change link rates, and have shown promising energy savings. However, these works have been largely simulative, and have not addressed many of the challenges involved in implementation. In this paper, we develop a hardware prototype ALR system and address real-time challenges involved in realizing such an implementation. We also identify new considerations for control policy development given current technology capabilities as well as future projections.

Keywords—Adaptive link rate (ALR), local area networks, energy efficient Ethernet, Ethernet, hardware prototyping

I. INTRODUCTION

The aggregate power consumption of computer networks has been increasing at a rapid rate [4] due to the growing number of connected devices such as PCs, switches, and routers. To support the network traffic introduced by these devices, link rates of the connecting infrastructure have also been increasing. We find that as the link rate increases, the power consumption of the links increases exponentially. (Figure 4)

To address this rapid increase in energy use, recent research has focused on reducing power consumption of both network devices and their infrastructure. One potential method for reducing power consumption is through exploiting the characteristics of network traffic, or the link utilization. Research shows that network packets are typically transmitted in bursts [16], with variable length periods of idleness or reduced link utilization between bursts. During these periods, energy is wasted by operating the network interfaces and links at the highest rates available. Devices may adaptively shut down and/or vary link rates in response to link utilization to reduce power consumption [1][4][9]. The ability to vary the link rate is a technique known as adaptive link rate (ALR) [9].

Determining when to change link rates is a challenging problem due to the unpredictability and bursty nature of

network traffic. Control policies determine the appropriate time to change link rates and the proper link rate to change to, taking into consideration several factors such as mean packet delay and packet loss. Mean packet delay is the total transit time for a packet, and if increased by too much, could result in human perceivable delay and reduced quality of service. Since the process of switching the link rate can introduce non-negligible delay, control policies should consider this link rate switching time. Furthermore, these topics are currently a focus of the IEEE 802.3az study group [15].

A large switching time is one of the biggest challenges in realizing ALR, and can increase mean packet delay very rapidly. Additionally, significant switching times can lead to buffer overflows and subsequent packet loss. However, in order to achieve fast link switching times, the physical implementation currently faces several challenges at multiple levels. These challenges are device synchronization at the MAC and PHY layers.

Most previous research focused heavily on addressing control policy challenges through simulation models [1][8][9][10]. Due to the absence of a real-time ALR capable system, these efforts did not consider the implementation challenges for MAC and PHY synchronization. In this paper, we develop a hardware prototype ALR system to address these challenges. To the best of our knowledge, we are the first to prototype ALR in an FPGA system. We measure power consumption and link rate switching times using our hardware prototype, and discuss challenges and solutions involved in a direct implementation. Our measurements indicate that the switching time between link rates is on the order of milliseconds, which is at least 70 times larger than that assumed in previous works. Using our real-time measurements of link switching times, power, and energy consumption, we identify new considerations for future control policy development.

II. BACKGROUND AND MOTIVATION

A. Background

Reducing energy consumption of network devices is a major research focus. Gupta et al. [12] proposed a method to reduce energy consumption in Ethernet switches by dynamically shutting down the transceiver completely depending on traffic arrivals, buffer occupancy, and a bounded maximum packet delay. Hays [13] proposed the active/idle

