Available online at www.sciencedirect.com
ScienceDirect
Computer Communications 29 (2006) 3170–3181
computer communications
www.elsevier.com/locate/comcom

Analysis and evaluation of a multiple gateway traffic-distribution scheme for gateway clusters

Pan-Lung Tsai *, Chin-Laung Lei

Department of Electrical Engineering, National Taiwan University, Taipei 106, Taiwan

Received 11 August 2005; received in revised form 20 April 2006; accepted 27 April 2006
Available online 24 May 2006

Abstract

Next-generation Internet gateways are expected to deal with higher volume of network traffic and also perform more sophisticated tasks besides packet forwarding. As the scale-up approach does not escape from the tradeoff between functionality and performance, architectural improvements such as clustering become necessary in the design of future Internet gateways. In this paper, we investigate different clustering architectures for high-performance, feature-rich Internet gateways and formally define the optimization problem behind these architectures as *Multiple Gateway Traffic-Distribution Problem*, both in a discrete and a continuous form. In addition to proposing various algorithms that solve the problem exactly and approximately, we also develop an on-line, self-adjusting scheme based on the solution algorithms. The numerical results of simulation suggest that the proposed approximate solution algorithms are effective and efficient, and the derived adaptive scheme is able to make the best decision on traffic distribution when dealing with the dynamic nature of network traffic in practice.

Keywords: Internet gateway; Approximation algorithm; Self-dispatching; Scalability; Cluster computing

1. Introduction

As the Internet keeps evolving, more and more network applications and services have been invented to provide a wide variety of advanced functions. In addition to deploying dedicated servers for each of these innovative applications and services, researchers and vendors have also spent their efforts on developing novel approaches to seamlessly integrate new functions with existing routing architectures. As a result, next-generation Internet routers, especially those deployed on the edges (i.e., the access routers mentioned in [1]), are expected to accomplish sophisticated tasks like URL filtering, anti-virus, anti-spam, and bandwidth control, rather than just routing.

In this paper, we use a more general term, gateways, to represent such versatile routers. Following this terminology, gateways need to implement some or all of the advanced traffic-processing functions including packet filtering, packet rewriting, packet scheduling, connection splicing, and even pattern matching within the payloads, in addition to the primitive packet forwarding, fragmentation, and reassembly. On the other hand, the advancement of networking technologies and optical components has resulted in a data rate of multiple gigabits per second or even higher in modern computers and communication networks. Consequently, though it is possible to build a standard router that performs basic packet forwarding at very high speed [2,3], the tradeoff between functionality and performance prevents current implementation of gateways from processing network packets in wire speed. Under such circumstances, a more scalable solution for the design and implementation of the gateways is highly desired.

Although hardware-based solutions (e.g., using ASICs) have been long proven in the field to be capable of delivering high throughput for well-defined operations [4], they are also infamous for their inflexibility. In particular, the revision of hardware designs in response to the invention

* Corresponding author. Tel.: +886 2 33663700x6604; fax: +886 2 23638247.
E-mail addresses: charles@fractal.ee.ntu.edu.tw (P.-L. Tsai), lei@cc.ee.ntu.edu.tw (C.-L. Lei).

0140-3664/$ - see front matter
doi:10.1016/j.comcom.2006.04.015

of new protocols is both time-consuming and costly. Conversely, software-based network systems [5–8] are flexible enough but usually fail to achieve satisfactory performance for demanding applications. Since the scale-up approach [9] applying to a single network system is still governed by the tradeoff between functionality and performance, gateway clusters [10,11] that follow the scale-out approach [9] are proposed as the remedy.

In this paper, we investigate different clustering architectures and formally define the problem of traffic distribution over clustered gateways. Following the formal definition of the optimization problem, we propose a number of algorithms for solving the problem exactly and approximately. We also develop an on-line, self-adjusting scheme based on the solution algorithms. The rest of the paper is organized as follows. Section 2 reviews the two broad categories of clustering architectures as well as the models that help to construct the traffic-distribution scheme for gateway clusters. Section 3 formally defines the traffic-distribution problem and presents the solution algorithms, followed by the simulation results of the algorithms in Section 4. Section 5 describes the proposed adaptive scheme based on the solution algorithms, and Section 6 concludes the paper by summarizing our achievements.

2. Related work

Clustering can be viewed as an architectural improvement to the design of network systems. A cluster is inherently scalable under the condition that the input load is elegantly distributed over the clustered units. In this section, we first examine the two broad categories of clustering architectures, one with a dedicated traffic dispatcher and the other without any (as shown in Fig. 1), as well as their corresponding traffic-dispatching techniques. Then we outline the construction of the traffic-distribution scheme and also discuss the literatures relevant to our modeling.

2.1. Gateway clusters with dedicated traffic dispatchers

A straightforward approach to distribute processing load over multiple gateways within a cluster is to deploy a dedicated traffic dispatcher in front of the gateways so

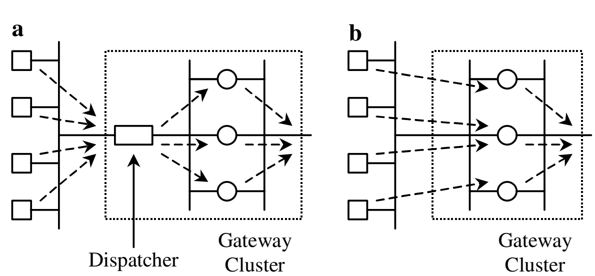

Fig. 1. Two categories of gateway clusters: (a) a gateway cluster with a dedicated traffic dispatcher, and (b) a fully decentralized gateway cluster. The squares represent hosts that generates input load and the circles represent individual gateways. The arrows with dashed lines depict the path of network traffic.

that the dispatcher may redirect individual network packets to different gateways. The decision of redirection can be made according to simple rules such as statically round robin, or it can be as flexible as evaluating complicated criteria like selecting the gateway with least load over the past 10 s. As long as the traffic dispatcher is carefully designed, it is possible to keep the load of the clustered gateways balanced over time and thus to maximize the throughput of the entire cluster.

There are several options for the traffic-dispatching mechanisms adopted in this architecture. The dispatcher may make use of standard routing, network address translation (with network-layer headers being altered) [12], tunneling (e.g., GRE [13,14] and IP-within-IP [15]), etc., or hybrid mechanisms [16,17], depending on the performance requirements, granularity of traffic dispatching, and the complexity of the application running on the gateways behind the dispatcher.

The primary advantages of using a dedicated dispatcher are its simplicity and flexibility. Since the issue of load distribution is isolated and can be taken care of solely by the dispatcher, the gateways in the cluster may focus on application-specific processing and are virtually independent from one another and from the dispatcher. In addition, the dispatcher itself may serve as the representative entity of the entire gateway cluster to other hosts on the network and hence help the cluster maintain the property of transparency without additional mechanisms. The features of simplicity and flexibility together make such dispatcher-based clustering architecture applicable to many classes of network applications and services, ranging from layer 3 to layer 7 in OSI models, and inspire network-equipment manufacturers to create products such as server load balancers [18–20].

2.2. Gateway clusters with self-dispatching mechanisms

Despite the advantages mentioned above, the major drawback of the dispatcher-based clustering architecture is obvious. The centralized design adopted makes the dispatcher the potential performance bottleneck and a single point of failure. In contrast to the dispatcher-based clustering architecture, the second architecture described in this section completely eliminates the need for traffic dispatchers by taking advantage of various self-dispatching techniques, and can be fully decentralized.

To accomplish the task of traffic dispatching and maintain the transparency at the same time, gateway clusters adopting this second architecture usually perform special processing on ARP [21] requests (e.g., proxy ARP [22,23]) and frame filtering. For example, the clustered units in [24] are configured to answer ARP requests either with a nonexistent Ethernet address or with a layer-2 multicast address so that the switch residing between the cluster and other hosts will always flood the frames sent to the cluster. In [25], the clustered units are simply configured with the same layer-2 address to ensure that ARP replies will contain the correct answer. The gateway cluster