

Performance analysis of large multicast switches with multicast virtual output queues

Min Song^{a,*}, Weiyong Zhu^a, Andrea Francini^b, Mansoor Alam^c

^aDepartment of Electrical and Computer Engineering, Old Dominion University, 231 Kaufman Hall, Norfolk, VA 23529, USA

^bEmbedded Systems Development Department, Bell Labs—Lucent Technologies, 101 Crawfords Corner Road, Holmdel, NJ 07733, USA

^cEECS Department, The University of Toledo, 2801 Bancroft Street, Toledo, OH 43606, USA

Received 13 January 2004; revised 23 August 2004; accepted 14 September 2004

Available online 12 October 2004

Abstract

In multicast switches, the accommodation of multicast traffic in multiple queues per input buffer reduces the throughput degradation caused by head-of-line (HOL) blocking. Such an arrangement, called multicast virtual output queuing (MC-VOQ), is very promising in theory but can only be implemented in practice with heavy approximation. Complete avoidance of the HOL blocking problem would in fact require a distinct queue for each fanout set possible, leading to an exponential growth of the number of queues needed with the switch size. If only a limited number of queues can be used per input buffer, criteria must be identified for setting the number of queues, for associating the fanout sets with the individual queues, and for scheduling the transmission of packets out of the queues.

This paper presents an analytical model for the investigation of saturation throughput and packet delay in MC-VOQ multicast switches. The model relies on the assumption of Poisson-distributed uniform input traffic and random queuing and scheduling policies. Extensive simulation experiments validate the results of the analysis for large switch sizes.

© 2004 Elsevier B.V. All rights reserved.

Keywords: Performance analysis; Multicast; Packet switch; Delay analysis; Input queuing; Queuing theory

1. Introduction

As the Internet has evolved from a convenience to a mission-critical platform for conducting research, education, and business, the spread of multicast applications such as distributed interactive simulations, distance learning, and video-on-demand is rapidly growing. As a result, there is an increasing demand for high-performance packet switches that efficiently support multicast traffic.

The distinguishing feature of a multicast switch is its ability to deliver a multicast packet to multiple output ports. A multicast packet switch must be capable of replicating packets, so that all destinations of a multicast packet can receive their respective copies. We focus on crossbar-based

switch fabrics, because of their intrinsic multicast capabilities [3,7]. We only address crossbars with *speedup* one (in a crossbar, the speedup is the capacity ratio between the fabric ports and the switch ports).

While virtual output queuing is a widely accepted queuing scheme for unicast traffic, it is not practical for multicast switches in its exhaustive form, which offers one queue per possible fanout set [5,10]. One alternative queuing scheme for multicast switches consists in the allocation of a single FIFO (first-in first-out) queue at each input port to handle all multicast traffic that transits through the port [9]. The performance of such multicast switches was studied by Hayes et al. [4] and Hui [5]. While the former focuses on small-sized switches and the latter focuses on large-sized switches, both studies assume one FIFO queue at each input port and the random packet selection policy with fanout splitting. The performance of multicast switches with one input queue per input buffer with no fanout splitting was studied in [1,8]. Although

* Corresponding author. Tel.: +1 757 6835294; fax: +1 757 6833220.

E-mail addresses: msong@odu.edu (M. Song), wzhu@odu.edu (W. Zhu), francini@lucent.com (A. Francini), malam@eeecs.utoledo.edu (M. Alam).

the single-queue scheme can be easily implemented due to its simplicity, head-of-line (HOL) blocking heavily impairs its performance.

To overcome HOL blocking, window-based queuing was introduced [2], allowing packets behind the HOL one to be transmitted prior to the HOL packet. The increased scheduling space results in an approximate 10% throughput increase on average. A major price paid for this enhancement is the increased complexity of the input queues, which must have random access capabilities. Another approach to alleviate HOL blocking consists in the allocation of a plurality of multicast queues at each input port, so that the packet scheduler can choose among multiple HOL packets. Simulation studies for such multicast switches with *non-exhaustive* multicast virtual output queuing (MC-VOQ) were performed by Gupta and Aziz [3], Song et al. [10], and Bianco et al. [11]. Particularly interesting are the results presented in [11], which indicate that the number of input queues, the queuing policy that couples the destination masks with the available input queues, and the scheduling policy that sets the service order over the available HOL packets are all critical to the throughput performance of the switch. All such indications are obtained under *gathered traffic* conditions, i.e. under traffic scenarios where only a small number of input ports supply multicast traffic to the switch. Uniform distribution of the traffic feeds over all input ports makes the performance impact of different design choices almost irrelevant.

The contribution of this paper consists of a first step towards the analytical validation and possible generalization of the simulation results presented in [11]. We provide analytical characterization of the behavior of a large multicast switch when Poisson traffic is uniformly distributed over all input ports. Our results confirm that under non-gathered traffic conditions a small number of multicast queues is sufficient to maximize the saturation throughput, even if low-end queuing and scheduling policies are deployed. We show simulation results that match the analytical ones for reasonably large switch sizes. Our next step, to be presented in a follow-up paper, will be the extension of the model to the case of gathered traffic.

The rest of the paper is organized as follows. Section 2 illustrates the multicast switch architecture used in this paper. Section 3 introduces the initial model and a modified equivalent model. The saturation throughput analysis is given in Section 4. In Section 5, we derive the packet delay and service time. The theoretical and experimental results are jointly presented in Section 6. Finally, conclusions are given in Section 7.

2. Multicast switch architecture

Fig. 1 shows the $N \times N$ multicast packet switch studied in this paper. The switch size N is assumed to be a large number. At each input port, there are M FIFO queues




Fig. 1. The switch architecture with M multicast queues at each input port.

dedicated to multicast traffic. A multicast packet arriving at the input interface is first queued into one of the M multicast queues and then switched from the input port to its target output ports. To preserve the order of packet delivery on the output links, packets of the same flow get queued at the same FIFO. The switch fabric is a bufferless crossbar with speedup one, i.e. at each time slot, each input can transmit no more than one packet and each output can receive no more than one packet. However, it is obviously possible that multiple copies of the same HOL multicast packet are forwarded to different output ports during the same time slot. In essence, the service discipline is based on fanout splitting [9]. The random scheduling policy [11] resolves contention at the outputs. When more than one HOL packet contend for the same output, one of them is selected randomly with equal probability.

We focus on saturation throughput and delay performance. We define throughput as the average number of packets delivered to an output per time slot. Delay is the number of time slots between the arrival and departure times of a packet at its assigned input queue.

3. Modeling

3.1. Initial model

Some assumptions need to be made before any further discussion. The incoming multicast traffic at each input link consists of a mix of multicast Poisson-distributed flows. The traffic arrival rates at the input links are the same for all inputs. We logically organize the input multicast queues in groups: all the k th queues across all input ports belong to group Q^k . Each multicast flow is randomly associated with one of the M queues with equal probability: packets of the same multicast flow always get queued to the same input queue and serviced in the same order they came in.

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات