Optical Switching and Networking 25 (2017) 109–116

Contents lists available at ScienceDirect

Optical Switching and Networking

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

Nonblocking conditions for (f_1, f_2)-cast Clos networks under balanced traffic

Lisheng Ma[a,b], Bin Wu[c], Xiaohong Jiang[a], Achille Pattavina[d,*]

[a] School of Systems Information Science, Future University Hakodate, 116-2, Kameda Nakano-Cho, Hakodate, Hokkaido 041-8655, Japan
[b] School of Computer and Information Engineering, Chuzhou University, Anhui 239000, PR China
[c] School of Computer Science and Technology, Tianjin University, Tianjin 300072, PR China
[d] Department of Electronics, Information and Bioengineering, Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133 Milano, Italy

ARTICLE INFO

Keywords:
Clos networks
(f_1, f_2)-cast
Multicast
Nonblocking condition

ABSTRACT

Cost efficient design of nonblocking multicast Clos switching networks has been a challenging issue due to the high cost involved in such networks. This paper focuses on the study of wide-sense nonblocking and strictly nonblocking Clos networks under the (f_1, f_2)-cast traffic, where the minimum fanout and the maximum fanout of each connection request are restricted to the values of f_1 and f_2, respectively. The (f_1, f_2)-cast is general in the sense it covers the unicast and multicast as special cases. Our study addresses the specific case in which the (f_1, f_2)-cast traffic loads evenly the input stage switches based on f_1. We conduct analysis on the sufficient conditions for wide-sense nonblocking as well as the necessary and sufficient conditions for strictly nonblocking in such networks. Our results indicate that under (f_1, f_2)-cast traffic, the number of middle stage switches required in the special Clos networks can be reduced. It is expected that this study can provide a fundamental guideline to the cost efficient design of nonblocking multicast Clos networks for different applications, including, e.g., optical switching fabrics for data center environments.

1. Introduction

With the rapid development of information technology, there is an increasing demand for high bandwidth and high capacity switching networks. Multistage interconnection networks are promising to construct large-scale switching networks with tens or hundreds of thousands of servers, because the switching capacity of multistage switch architectures has been increased to a few hundred terabits or even petabits per second [1]. The Clos network architecture [2] can be used to construct large-scale switching networks based on many low-end commodity switches (switching modules), and thus can provide significant advantages in terms of cost and scalability. As a result, Clos networks and their variations (like Fat-trees) have been widely adopted in large-scale data center networks [3] to meet the rapidly growing demands of web search, gaming, e-commerce and multimedia applications. Recently, Clos networks based on arrayed-waveguide gratings and tunable wavelength converters are used to construct optical switch fabric for data center networks [4,5], and nonblocking performance of such optical Clos networks has also been investigated [6].

On the other hand, due to the emerging applications of IPTV, web search and cloud computing [7], etc., it is necessary to support one-to-many connections in Clos networks, where one input is required to be connected to multiple outputs (*multicast*) or even to all outputs (*broadcast*), thus extending the concept of one-to-one (*unicast*) connections. However, supporting multicast in Clos networks is very challenging due to the high cost of the large number of switching modules required. Since Charles Clos introduced the nonblocking three-stage connecting networks in 1953 [2], extensive research efforts have been devoted to the design of nonblocking multicast multistage connection networks (see Section 2 for related literature). This paper focuses on the study of wide-sense and strictly nonblocking multicast Clos networks, where we assume that every switch in the Clos networks has multicast capacity. Note that a multicast traffic can be either open-end or closed-end depending on whether or not a new connection can be established from a busy input by adding extra output ports. In this paper, we consider closed-end multicast traffic where extra output ports cannot be added to existing connections.

Known works on the study of nonblocking multicast Clos networks generally assume that there are simultaneous connection requests from all the input ports of a single input stage switch and then apply such assumption to determine the nonblocking conditions (i.e., the required number of switching modules in the middle stage) for a nonblocking multicast Clos network. It is notable, however, that for multicast in a Clos network, where a connection from one input may request multiple outputs, the number of active input ports is usually smaller than the

* Corresponding author.
E-mail addresses: mls@chzu.edu.cn (L. Ma), binwu.tju@gmail.com (B. Wu), jiang@fun.ac.jp (X. Jiang), pattavina@elet.polimi.it (A. Pattavina).

http://dx.doi.org/10.1016/j.osn.2017.03.004
Received 23 November 2016; Received in revised form 27 March 2017; Accepted 27 March 2017
Available online 01 April 2017
1573-4277/

number of active output ports. Thus, in a multicast Clos network, on average the number of active input ports in each input stage switch is usually smaller than the number of input ports of the switch. To the best of our knowledge, no study considers such observation of reduced number of average simultaneous connection requests from each input stage switch in the nonblocking condition analysis of multicast Clos network.

We focus on wide-sense and strictly nonblocking conditions for (f_1, f_2)-cast Clos networks by assuming multicast capacity for each switch. Nevertheless, this assumption makes nonblocking analysis for the rearrangeable case extremely complex, and thus we leave the rearrangeable case for future study. We propose traffic balancing mechanisms to evenly distribute the maximum number of possible incoming connection requests across the input stage switches according to the minimum fanout f_1. More specifically, we determine the maximum allowed number of simultaneous connection requests for each input stage switch based on f_1, and "traffic balancing" in our paper means that this maximum allowed number is the same for each and all input stage switches. Although the actual numbers of simultaneous connection requests may indeed be different among the input stage switches, traffic is still taken as "balanced" across the input stage switches as long as each of those numbers does not exceed the same maximum allowed limit determined by f_1. Such a traffic balancing scenario can significantly reduce the required number of middle stage switches. The main contributions of our work are summarized as follows:

- We extend the literatures of wide-sense and strictly nonblocking Clos networks to the more general (f_1, f_2)-cast traffic scenario, which covers both unicast and multicast as special cases.
- We formulate the relationship between the minimum fanout f_1 and the input traffic pattern for the generalized (f_1, f_2)-cast Clos networks, where similar relationships have not been explored in existing literature.
- We derive wide-sense and strictly nonblocking conditions for the proposed architectures, showing how these conditions include known results about unicast and multicast cases.
- We show that traffic balancing at input stage switches can reduce nonblocking network cost by decreasing the number of middle stage switches, and propose traffic balancing mechanisms by adjusting either network topology or traffic pattern offered to the network.

The rest of the paper is organized as follows. After showing related literature in Section 2, Section 3 provides preliminaries that will facilitate our study. Section 4 derives the conditions for nonblocking (f_1, f_2)-cast Clos networks, with numerical results provided in Section 5. Section 6 describes how traffic balancing can be attained and, finally, Section 7 draws the conclusions.

2. Related literature

In recent years, multicast in multistage switching networks has attracted extensive research attentions. Existing works on nonblocking multicast can be roughly divided into three categories: nonblocking conditions, performance analysis, and routing techniques.

A nonblocking multicast network can be either strictly nonblocking (SNB), or wide-sense nonblocking (WNB), or rearrangeable nonblocking (RNB). In strictly nonblocking, we can establish a connection from an idle input port to a group of idle output ports regardless of how other connections are established. The strictly nonblocking condition under multicast is firstly provided in [8]. The sufficient conditions for a multicast strictly nonblocking three-stage network are derived in [9]. The necessary and sufficient conditions for strictly nonblocking three-stage Clos network are investigated in [10], in which two types of multicast traffic are considered. In [11], the necessary and sufficient conditions are given for a strictly nonblocking three-stage switching network when multicast is offered, where the conditions are given as a function of both the minimum and the maximum connection fanout. In [12], a unified approach of the sufficient and necessary conditions for strictly nonblocking multicast Clos networks are proposed. The necessary and sufficient conditions are determined for f-cast photonic networks in [13]. In [14], the nonblocking conditions for multicast three-stage switching networks are provided, and the network complexity is discussed depending on main network and traffic parameters.

In a wide-sense nonblocking multicast network, we can establish a connection from an idle input port to a group of idle output ports without disturbing the existing connections, if routes for new connections are always chosen suitably. In the literatures, there are only few works on wide-sense nonblocking multicast Clos networks. In [15], the sufficient conditions based on a routing algorithm are presented. Recently, other sufficient conditions are developed in [16], which explores the server redundant.

In a rearrangeable nonblocking multicast network, we can establish a connection from an idle input port to a group of idle output ports by rerouting one or more existing connections. The sufficient conditions on rearrangeable nonblocking multicast Clos networks are discussed in [17]. A summary of nonblocking conditions is presented in [18].

Some efforts have been focused on understanding multistage switching networks by analyzing performance metrics [19–22]. Besides, various routing techniques have been proposed in [3,23–25], as summarized in [26].

3. Preliminaries

The general scheme of a three-stage $N_1 \times N_2$ Clos network is illustrated in Fig. 1; it includes r_1 switches at input stage (IS), labelled $IS_1, \ldots, IS_{r_1}$, r_2 switches at output stage (OS), labelled $OS_1, \ldots, OS_{r_2}$, and m middle stage (MS) switches, labelled $MS_1, \ldots, MS_m$. Each IS switch has n_1 input ports and a link connection to every MS switch, and each MS switch has a link connection to every OS switch, each with n_2 output ports. Therefore $r_1 = N_1/n_1$ and $r_2 = N_2/n_2$. Such three-stage Clos network will be denoted as $C(n_1, r_1, n_2, r_2, m)$. When $N_1 = N_2 = N$ and $n_1 = n_2 = n$, we obtain a symmetrical Clos network, denoted as $C(n, r, m)$, where the number of switches at input and output stage is $r=N/n$.

If a connection request from an input port requests output ports associated with f output switches, we say that this connection has fanout f. In this paper, (f_1, f_2)-cast means that the fanout of each connection request is restricted to the value between f_1 and f_2. We also refer to f_1 and f_2 as *minimum fanout* and *maximum fanout* of the connection, respectively.

For a $C(n, r, m)$ network, if we assume that each input or output port in the Clos network supports at most one connection at a time, there are at most n connections from the MS switches entering each OS switch, and the total number of connections in all OS switches is at most nr. Since each connection request has a fanout of at least f_1, the number of simultaneous connection requests in the IS switches is at

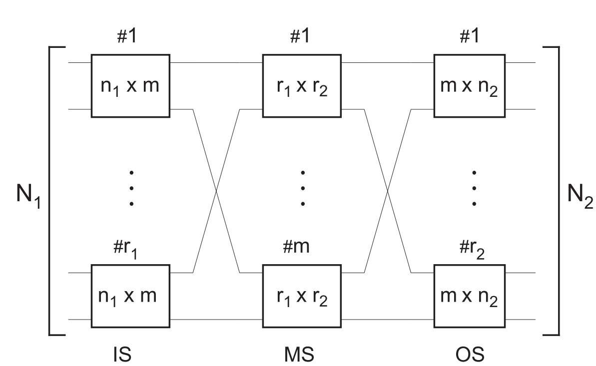

Fig. 1. Configuration of three-stage Clos network.