Performance Aspects of Some Switching Networks

Zbigniew Hulicki
Department of Telecommunication, AGH University of Science and Technology, ul. Czarnowiejska 78, Krakow, Poland

Keywords: Switching Fabric, Multistage Interconnection Network (MIN), Scalable High-Performance Router.

Abstract: The intention of this paper is to examine the performance and reliability aspects of switching fabrics which are used for scalable high-performance routers. Topology and capabilities of switching fabrics are discussed, followed by an examination of the performance vs.– reliability trade-off in diverse scenarios of possible failures. It has been shown that the switching fabric based on the PM2I type multistage interconnection network (MIN) outperforms those based on the cube type MIN under any failure scenario. The simulation results presented should be helpful in predicting the performance vs.– reliability trade-off before actual fabrication of the switching fabric.

1 INTRODUCTION

Routers and switches are viewed as the most critical parts of the current communication infrastructure and will be also needed to provide fast and efficient communication in next-generation networks. Recently there have been various efforts to design an efficient optical switch for use in a high-speed router (Hossfeld et al., 2009). Besides, there is an immense interest in designing a simple and high performance switch which would satisfy the demands of an entirely new scenario for emerging broadband services.

In the past few decades, a number of switch fabrics have been proposed in literature and used in practical implementations to interconnect key components in routers, such as routing engines and line cards (Lien et al., 2010; Rongsen and Delgado-Frias, 2007). Many of the proposed switch designs have been based on multistage interconnection networks (MINs). Various solutions, i.e. single or multiple panel (replicated), can be used in design of MIN based switches (Aulakh, 2006).

The purpose of this paper is to examine the issues dealing with the performance and reliability of single plane MINs used as switch fabrics for scalable high-performance routers and/or OXCs. Topology and capabilities of switching fabric will be discussed first, followed by a specification of the switch model and its performance measures. Then, the performance and dependability of MIN based switch fabrics will be examined, taking into account diverse scenarios of possible faults. Lastly, there are some conclusions and remarks regarding the trade-offs between switch characteristics

2 TOPOLOGY OF SWITCHING FABRIC

In selecting the architecture of switching fabric, four design decisions can be identified: operation mode, control strategy, switching method, and topology of MIN. The topology of a switching network is a key factor in determining a suitable fabric architecture (Skalis and Mhamd, 2010). MINs proposed for the switch fabrics are usually constructed using small crossbar switches organized in stages and may have a uniform or non-uniform connection pattern between stages. SEs can be buffered or non-buffered, but the use of self-routing MINs are favored because small delays can be achieved. Generally, the topologies of MINs tend to be regular and can be grouped into two categories: static or dynamic. The static MINs are simple to build and expand, but fail even in the presence of a single fault. On the other hand, a dynamic MIN is able to reroute the data through a fault free path if the regular path is faulty or busy. Because the reliable operation of a MIN is an important factor in overall router performance, it is important to design a switching network that combines full connection

Hulicki Z..
Performance Aspects of Some Switching Networks.
DOI: 10.5220/0004428200900093
In *Proceedings of the 3rd International Conference on Simulation and Modeling Methodologies, Technologies and Applications* (SIMULTECH-2013), pages 90-93
ISBN: 978-989-8565-69-3

capability – in spite of faults – with a slightly lower performance.

A number of diverse fabric solutions which offer redundant paths between the input and output ports have been proposed to provide fast, efficient and reliable communications at a reasonable cost. However, no single solution network is generally preferred because the cost-effectiveness of a particular design varies with such factors as its application, the required speed of data transfer, the actual hardware implementation of the switching network, the number of input/output ports, and the construction cost.

As far as MINs are concerned, they can be built using a variety of structures. Two significant examples of topological structures based on the single plane design are switching networks based on the $cube_i$ and PM2I interconnection functions (Lien et al., 2010). Therefore, there is an open question about the performance vs.-reliability trade-off that a multiple-path structure of the PM2I network might offer, for various distributed control algorithms and decision rules, when compared to a cube-type network.

3 THE SWITCH MODEL

To evaluate the properties of a fabric based on a given switching network (MIN), the following assumptions are made about the operation and the environment of the interconnection network: the single plane design of switching fabric is taken into account; the slotted traffic source model is used with a uniform random distribution of packet (cell) destinations (because, in this contribution, we are focused on the performance vs. reliability trade-off); the switching network is operated synchronously, meaning that the packets are transmitted only at the beginning of a time slot given by the packet clock and each input link is offered the same traffic load; the buffering is external to the switch fabric, i.e. non-buffered switch fabrics formed from non-buffered SEs are under consideration. Therefore, the queuing effects are unaccounted in this model. All output ports perform the function of a perfect sync and a conflict is said to occur if more than one packet arrives at the same output link at the same time.

Two classes of packet switched MINs, one with redundant interconnection paths and one without are under consideration (i.e. the cube-type class of MINs known as unique-path networks and the PM2I class of MINs which have multiple paths between a given network input-output pair). It is clear that faults in the switching network often result in degraded performance of a fabric. Hence, in order to capture effects of physical failures on the operation of switching fabric, generally three models are used: the stuck at fault model, the link fault model, and the SE fault model (Aulakh, 2006). The link and SE fault models will be used in this contribution. In the former model, physical failures result in a faulty link whereas, in the latter model in a faulty SE. Moreover, in both cases the single and multiple fault model versions will be considered. Because failures occur at random, and because the combinations of the type, number and location of faults can widely vary, the influence of failures on the performance of the switching fabric will be examined by simulation.

4 PERFORMANCE EVALUATION

The performance of switching fabric, based on both of the aforementioned classes of packet switched MINs, has been studied using simulated experiments. A discrete time event-driven simulator has been designed to carry out simulations. The simulator uses the switch model and performance measures described in the previous section. It enables performance evaluation of the fabric under different traffic and/or interconnection patterns, including both uniform and non-uniform traffic patterns, as well as an evaluation of the fabric's dependability under diverse fault models, including combinations of the type, number and location of faults. Moreover, in order to stop packets from entering the switching fabric once its resources are exhausted, the switching network of a router often employs a backpressure mechanism (Lien et al., 2010). Therefore, except for non-buffered MINs, the simulator also allows for a modification to the switching fabric, i.e. an implementation of buffering (introduced at the input/output of the switching network or to the SEs) or using different types of SEs.

This section presents only selected results from the simulation experiments because the results point to a similar trend and the packet loss (or packet drop rate) can serve as a good indicator of the performance vs. –reliability trade-off. The simulation results have been shown and compared in a few figures (Figs. 1, 2, 3, 4, and 5). Each data value given in these figures is the result after 100,000 clock cycles in the simulation, where this number of cycles is found to yield steady-state outcomes (95% conf. lev.).

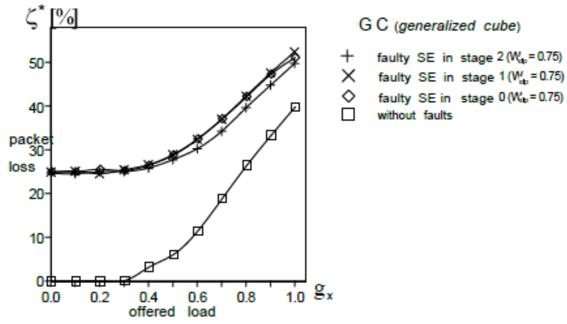

Figure 1: Comparison of packet loss in a switch with faulty SE.

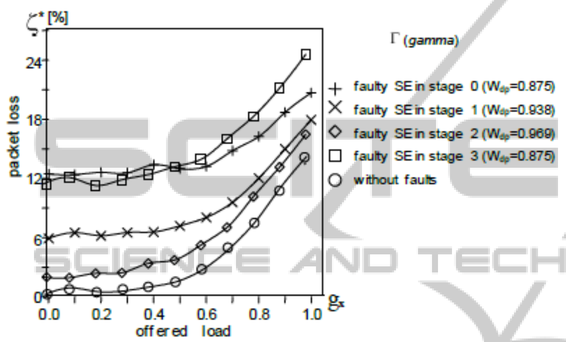

Figure 2: Comparison of packet loss in a switch with faulty SE.

The packet loss versus offered load under a single SE fault for both the cube type (GC) and the PM2I type (Γ) MIN is demonstrated in Fig. 1 and Fig. 2 respectively. One can observe a qualitatively similar nature of the performance of both switching networks. Besides, it is clear (cf. Fig. 1 and 2) that faultiness of a SE has crucial impact on the switch performance. In both cases, the faulty SE determines the degradation to the switch performance. However, in the PM2I type MINs this effect will depend on the stage location of the faulty SE. The packet loss is higher if the faulty SE is located in the output or input stage of the switching network (Γ), and the loss is lesser if the faults of SEs occur in the inner stages of the MIN. In the cube-type MINs, the faulty SE exhibits substantial impact to the performance regardless of the fault locations.

Since faults have a strong tendency to happen in continuous (or nearby) areas (Rongsen and Delgado-Frias, 2007), the next tests had to consider the multiple fault model scenario. The packet loss versus offered load under a double SE fault for the modified cube type (ESC) MIN is shown in Fig. 3. It is observed that the location of the faulty SEs determines the degradation of switching fabric performance. A double SE failure has a more detrimental effect on the performance (packet loss) if defects of SEs happen in the inner stages of the MIN, but the worst case is when the faults of SEs happen in nearby areas.

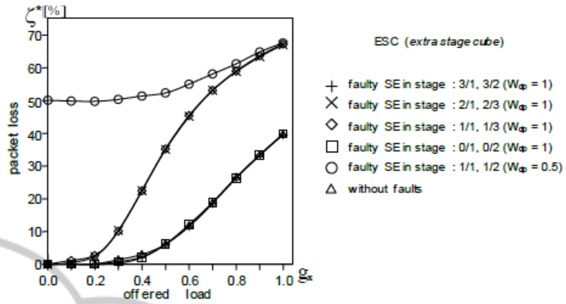

Figure 3: Comparison of packet loss in a switch with multiple faults.

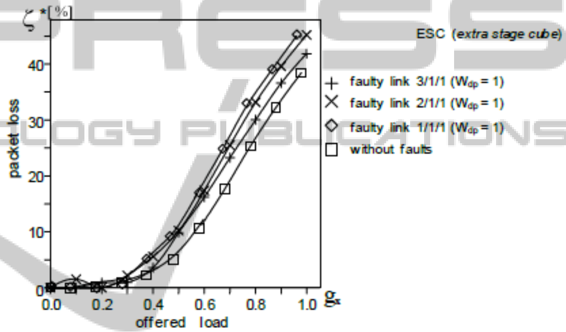

Figure 4: Packet loss vs. traffic load in the MIN based switch with faulty link.

On the other hand, there are SE and link failures inside the switching fabrics. Therefore, it was interesting to know what the effects would be of link failures on switching fabric performance. In consecutive tests, the impact of single faults has been evaluated by randomly designating one link to be faulty. The packet loss versus offered load under a single link failure in different locations is depicted in Fig. 4, and the results under a fault-free situation are also included for comparison. It is observed that a single link failure also leads to the performance degradation; however, its impact is not as detrimental as it was for an SE failure. The location of the faulty link has (practically) small effect on the performance degradation as well. Moreover, the effects of link failures in the switching fabric based on the PM2I type (Γ) MIN are different. One can observe (cf. Fig. 5) more detrimental effects on switching fabric performance when a straight link is faulty and a failure occurs between the input and output stage of switching network. Irrespective of the location, other link failures have a negligible

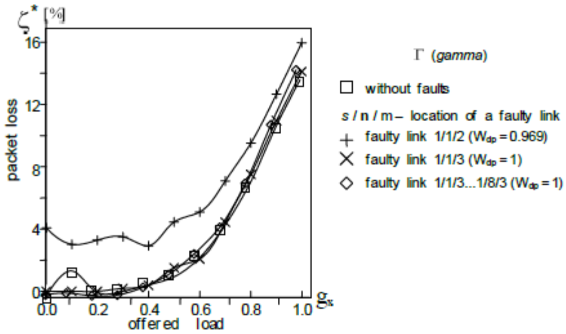

Figure 5: Packet loss vs. traffic load in a switch with faulty link in different location.

effect on performance degradation. Such capabilities of the (Γ) switching network result from its architecture which already has built-in redundancy (multiple paths between a given network input-output pair in the PM2I type MINs) that provides fault tolerance. Furthermore, it is also clear (cf. Fig. 3 and Fig. 4 and/or Fig. 2 and Fig. 5) that an SE failure has more detrimental effects on switching fabric performance than a link failure. One could expect this because SE hardware is much more complex and, therefore, more prone to faults than the internal link connections.

4 CONCLUSIONS

This paper presents the analysis of the performance and reliability of single panel multistage switching fabrics. Two classes of packet switched MINs, those with redundant interconnection paths (the PM2I type MINs) and those without (the cube type MINs), have been examined by simulating performance and reliability conditions. Simulations under different fault model scenarios have shown that the performance of multistage switching fabrics deteriorates after an increase in load. However, the switch based on the PM2I type MIN outperforms that based on the cube type MIN under any failure scenario. Moreover, the simulation revealed that PM2I type multistage switching fabrics are highly fault tolerant against internal link or SE failures. This property can be preferable for scalable core routers or backbone OXCs. Such single panel multistage switching fabrics seem to be an attractive alternative to the multiple panel architecture of switching networks. Furthermore, presented results of the simulation analysis should be also helpful in predicting the performance vs. reliability trade-off before actual fabrication of the switching fabric.

REFERENCES

Aulakh, N. S., 2006. ‘Reliability analysis of mux-demux replicated multistage interconnection networks’, *Experimental Techniques*, July/Aug. 2006, vol. 30, (4), pp. 19–22.

Hossfeld, T., Leibnitz, K., Nakao, A., 2009. ‘Modeling of modern router architectures supporting network virtualization’, *GLOBECOM Workshops*, 2009 IEEE, pp. 1–6.

Lien, Ch.-M., Chang, Ch.-S., Cheng, J., Lee D.-S., Liao, J.-T., 2010. ‘Using banyan networks for load-balanced switches with incremental update’, *GLOBECOM Workshops*, 2010 IEEE, pp. 1–6.

Rongsen, H., Delgado-Frias, J.G., 2007. ‘Fault tolerant interleaved switching fabrics for scalable high-performance routers’, *IEEE Trans. Paral. Distrib. Syst.,* vol. 18, no. 12, Dec. 2007, pp. 1727-1738.

Skalis, N., Mhamd, L., 2010. ‘Performance guarantees in partially buffered crossbar switches’, *GLOBECOM Workshops*, 2010 IEEE, pp. 1–6.