Theoretical Computer Science 562 (2015) 75–89

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Design of fault-tolerant on-board networks with variable switch sizes

O. Delmas [a,b], F. Havet [c,*,1], M. Montassier [d,e], S. Pérennes [c,1]

[a] Univ. Bordeaux, LaBRI, UMR 5800, F-33400 Talence, France
[b] CNRS, LaBRI, UMR 5800, F-33400 Talence, France
[c] COATI Project, I3S (CNRS/UNSA) & INRIA, 2004 route des Lucioles BP 93, 06902 Sophia-Antipolis Cedex, France
[d] Univ. Montpellier 2, France
[e] CNRS – LIRMM, 161 rue Ada, 34095 Montpellier Cedex 5, France

A R T I C L E I N F O

Article history:
Received 10 May 2013
Received in revised form 5 September 2014
Accepted 13 September 2014
Available online 18 September 2014
Communicated by G. Ausiello

Keywords:
Fault tolerance
Switching networks
Flow networks
Vulnerability

A B S T R A C T

An (n,k,r)-*network* is a triple $N = (G, \mathrm{in}, \mathrm{out})$ where $G = (V, E)$ is a graph and in, out are non-negative integer functions defined on V called the *input* and *output* functions, such that for any $v \in V$, $\mathrm{in}(v) + \mathrm{out}(v) + \deg(v) \leq 2r$ where $\deg(v)$ is the degree of v in the graph G. The total number of inputs is $\mathrm{in}(V) = \sum_{v \in V} \mathrm{in}(v) = n$, and the total number of outputs is $\mathrm{out}(V) = \sum_{v \in V} \mathrm{out}(v) = n + k$.

An (n,k,r)-network is *valid*, if for any *faulty* output function out' (that is such that $0 \leq \mathrm{out}'(v) \leq \mathrm{out}(v)$ for any $v \in V$, and $\mathrm{out}'(V) = n$), there are n edge-disjoint paths in G such that each vertex $v \in V$ is the initial vertex of $\mathrm{in}(v)$ paths and the terminal vertex of $\mathrm{out}'(v)$ paths.

We investigate the design problem of determining the minimum number $\mathcal{N}(n,k,r)$ of vertices in a valid (n,k,r)-network and of constructing minimum (n,k,r)-networks, or at least valid (n,k,r)-networks with a number of vertices close to the optimal value.

We first give some upper bounds on $\mathcal{N}(n,k,r)$. We show $\mathcal{N}(n,k,r) \leq \lceil \frac{k+2}{2r-2} \rceil \lceil \frac{n}{2} \rceil$. When $r \geq k/2$, we prove a better upper bound: $\mathcal{N}(n,k,r) \leq \frac{r-2+k/2}{r^2-2r+k/2} n + O(1)$.

Next, we establish some lower bounds. We show that if $k \geq r$, then $\mathcal{N}(n,k,r) \geq \frac{3n+k}{2r}$. We improve this bound when $k \geq 2r$: $\mathcal{N}(n,k,r) \geq \frac{3n+2k/3-r/2}{2r-2+\frac{3r}{\lfloor \frac{k}{r} \rfloor}}$.

Finally, we determine $\mathcal{N}(n,k,r)$ up to additive constants for $k \leq 6$.

1. Introduction

The design of modern telecommunication satellites is complex, and an important industrial priority is to provide robustness at the lowest possible cost. Alcatel Space Industries is a major manufacturer of telecommunication satellites. A key component of their satellites is an interconnection network which redirects signals received by the satellite to a set of amplifiers implanted in the satellite, from where the signals are then retransmitted (a detailed overview of the model and its motivations can be found in [5,2]). Because of its reliability, wave guide technology was chosen by Alcatel Space Industries to build these on-board networks (for background information see [7,11]). Such interconnection networks consist of expensive four-port switches, of wave guides linking these switches, of inputs (where the signals enter the network) and of outputs

* Corresponding author.
[1] Partially supported by the ANR project Gratel (grant ANR-09-BLAN-0373).

http://dx.doi.org/10.1016/j.tcs.2014.09.034

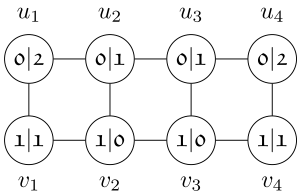

Fig. 1. A first solution: the network N_1.

(where the signals leave the network). Before being retransmitted, the signals must be amplified, so the outputs are amplifiers based on Travelling Wave Tube Amplifier technology [7,11]. However, amplifiers are prone to failure. While switches are also prone to failure, wave guide technology ensures that the probability of a switch developing a fault is much smaller than the probability of a fault arising in an amplifier. For this reason, only faults in amplifiers are considered in models [2]. Many techniques have been proposed to increase the reliability and fault-tolerance of multistage interconnection networks or switching networks (see [1,10,6]). These techniques only consider networks whose switches (or whose links between switches) are subject to failure. They do not consider faulty outputs. These works focused on aspects such as deadlock and adaptive routing schemes, aspects that are not relevant to our problem.

In this paper, following [2,4,9], we focus on designing networks that are capable, in the presence of faulty output ports, of rerouting input signals to operational output ports. Since the components of a satellite cannot be repaired, redundant amplifiers are added, and the interconnection network satisfies the following fault tolerance property: the network connects the set of input ports with the set of output ports, and for any set of at most k output port failures, there exists a set of edge-disjoint paths connecting the input ports to the operational output ports. Since each switching device is expensive, these interconnection networks are constructed using the fewest possible switches, or at least a number of switches close to the minimum value. The networks are controlled centrally from Earth. Each time an amplifier in use develops a fault, the controller sends messages to the switches to change their settings, so as to ensure that the inputs remain connected to functioning amplifiers. Variants of the problem have been considered in which there are two kinds of input, the aim being to guarantee a certain quality of service [3,9]. They will not be considered in this paper.

Current switches have four ports. The problem was initially studied for such switches in [2] ($k \leq 4$ failures), and then in [4] (up to 12 failures). For this, the cheapest type of switch, all wave guides are drawn in the plane and due to technological constraints, may not cross. For four-port switches, this was not problematic since there is a 2-dimensional switch which is as powerful as the one realizing all possible matchings of ports (see [2]). However, for a larger number of ports, the types of switches that can be built in the plane under this non-crossing constraint are not very powerful and do not allow the construction of networks with few switches. For this reason, in this paper we seek to design on-board networks with more powerful switches, that is, 3-dimensional switches with more than four ports. In practice, such a switch will be expensive. Hence less powerful but cheaper switches are also envisioned. For the sake of simplicity, we consider here a basic model in which every switch has $2r$ ports and can realize all matchings among them. The aim is to provide elements to determine the number of ports minimizing the cost of the network (this will depend on the cost of construction of $2r$-port switches). Obviously, the larger the number of ports, the more expensive will be the switches, but then fewer will be required. So the cost of such a network involves a trade-off between the total number of switches and their unit cost. In this paper, we give some bounds on the minimum number of $2r$-port switches in interconnection networks with n inputs and $n+k$ outputs.

Generalizing the definition of (n,k)-networks introduced in [2,4], we define (n,k,r)-networks as follows: An *(n,k,r)-network* is a triple $N = (G, \mathrm{in}, \mathrm{out})$ where $G = (V, E)$ is a graph (where each vertex is a switch) and in, out are non-negative integer functions defined on V called *input* and *output* functions, such that for any $v \in V$, its number of ports $\mathrm{por}(v)$ defined by $\mathrm{por}(v) = \mathrm{in}(v) + \mathrm{out}(v) + \deg(v)$ is at most $2r$. ($\deg(v)$ denotes the degree of v in the graph G, that is the number of edges of G incident to v.) Let i and o be two non-negative integers. An *$(i|o)$-switch* is a switch s with i inputs and o outputs, i.e. with $\mathrm{in}(s) = i$ and $\mathrm{out}(s) = o$. The total number of inputs is $\mathrm{in}(V) = \sum_{v \in V} \mathrm{in}(v) = n$ and the total number of outputs is $\mathrm{out}(V) = \sum_{v \in V} \mathrm{out}(v) = n + k$.

Any integer function out' defined on V such that $0 \leq \mathrm{out}'(v) \leq \mathrm{out}(v)$ for any $v \in V$, and $\mathrm{out}'(V) = n$ is called a *faulty output function*. Note that $\mathrm{out}(v) - \mathrm{out}'(v)$ is the number of faults at vertex v. An (n,k,r)-network is *valid*, if for any faulty output function out', there are n edge-disjoint paths in G such that each vertex $v \in V$ is the initial vertex of $\mathrm{in}(v)$ paths and the terminal vertex of $\mathrm{out}'(v)$ paths.

Let us denote the minimum number of vertices in a valid (n,k,r)-network by $\mathcal{N}(n,k,r)$. A valid (n,k,r)-network with exactly $\mathcal{N}(n,k,r)$ vertices is called a *minimum* (n,k,r)-network. The design problem consists of determining $\mathcal{N}(n,k,r)$ and of constructing minimum (n,k,r)-networks, or at least valid (n,k,r)-networks with a number of vertices (i.e. switches) close to the optimal value.

Let us present an example: We would like to construct valid $(4,4,2)$-networks. A first solution is depicted in Fig. 1. The network N_1 is composed of eight switches u_i, v_i for $1 \leq i \leq 4$. The associated graph $G = (V, E)$ is the 4×2 grid. The input and output functions are defined as follows: $\mathrm{in}(v_i) = 1$, $\mathrm{in}(u_i) = 0$ for $1 \leq i \leq 4$, and $\mathrm{out}(v_2) = \mathrm{out}(v_3) = 0$, $\mathrm{out}(v_1) = \mathrm{out}(u_2) = \mathrm{out}(u_3) = \mathrm{out}(v_4) = 1$, $\mathrm{out}(u_1) = \mathrm{out}(u_4) = 2$.

For any faulty output function out', it is easy to see that there are four edge-disjoint paths in G such that each vertex $v \in V$ is the initial vertex of $\mathrm{in}(v)$ paths and the terminal vertex of $\mathrm{out}'(v)$ paths. This implies that this network is valid. It