

Design of fault-tolerant networks for satellites (TWTA* redundancy)

Jean-Claude Bermond, Eric Darrot[†]

MASCOTTE Join Project CNRS/INRIA/University of Nice – Sophia Antipolis
2004 Route des Lucioles, B.P. 93, F-06902 Sophia Antipolis, France
Jean-Claude.Bermond@sophia.inria.fr, eric.darrot@laposte.net

Olivier Delmas[‡]

Université Bordeaux 1
Laboratoire Bordelais de Recherche en Informatique (UMR 5800)
351 cours de la Libération, 33405 Talence Cedex, France.
Olivier.Delmas@labri.u-bordeaux.fr

Abstract

This article deals with the design of networks to be placed on satellites. These networks should connect inputs (corresponding to signals arriving at the satellite) to outputs (corresponding to amplifiers), even in case of failures of amplifiers. The networks are made of links and expensive switches, hence we want to minimize the number of switches subject to the following conditions: each input and each output is connected to exactly one switch; each switch is adjacent to exactly four links; there are n inputs and $n+k$ outputs; among the $n+k$ outputs, k can fail permanently; and finally all the input signals should be sent to valid amplifiers, i.e., outputs, via disjoint paths. So, the aim is to design networks having as few switches as possible and satisfying the following property: *there exist n edge-disjoint paths from the n inputs to any set of n outputs chosen from the $n+k$ total number of outputs*. We call such networks *valid k -fault tolerant networks*. Let $\mathcal{N}(n, k)$ denote the minimum number of switches of a valid network with n inputs, $n+k$ outputs and k output failures.

In this article, we give some details on the problem, and some preliminary results such as the fact that $\mathcal{N}(n, k) \geq n$. We also propose a general construction which yields (under some conditions) $\mathcal{N}(n+n', k) \leq \mathcal{N}(n, k) + \mathcal{N}(n', k)$.

Keywords: fault-tolerant networks, switching networks, routing, connectivity, satellite, TWTA redundancy.

1 Introduction

ALCATEL SPACE INDUSTRIES asked us to consider the following problem. Signals arriving at a telecommunication satellite (via input links) have to be routed through a network to amplifiers

*Outputs of such networks use “*Travelling Wave Tube Amplifiers*” (TWTA) technology.

[†]This work has been supported by Alcatel Space Industries (Toulouse, France) and the french CNRS Telecom program.

[‡]Additional support from the Aquitaine Region Project #98024002

(outputs, also called *Travelling Wave Tube Amplifiers*, i.e., TWTA [3]). The links of the networks are made of wave guides (for background information, see [3]).

The problem comes from the fact that the amplifiers may fail during the satellite's lifetime and cannot be repaired. So one needs to have more amplifiers than the number of signals. One wants to be able to route the arriving signals to valid amplifiers, that is to find disjoint paths between the inputs and valid outputs in the interconnection network. The network consists of switches with 4 ports and which can realize the connections displayed in Figure 1.




Figure 1: A switch can use 4 different states.

Due to technological constraints, two signals cannot cross in a switch. A switch can be joined to an input or an output or another switch or a dead end. If the network is able to route signals to amplifiers in the presence of k faulty amplifiers, we will say that the network is a *valid k -fault tolerant network*.

There are variants of the problem according to whether each signal needs a specific amplifier (see the articles [1]) or can be sent to any of the amplifiers (that is all the amplifiers have the same function). Here, we suppose we are in this latter case of identical amplifiers. Then, we do not worry about the impossibility of realizing two opposite connections in a switch. Indeed, as shown by Figure 2, suppose that the routes from input i to output o and that of input i' to output o' cross in a switch (Figure 2(a)). Then it suffices to send input i to output o' and input i' to output o (Figure 2(b)) to avoid the crossing and so we can use the switch in a valid manner.

In all our figures we use the following notation: an **arrow** represents an *input*, a **circle** represents a *switch* and a **rectangle or square** represents an *amplifier*, i.e., an *output*.




Figure 2: We can avoid a forbidden state (a) by applying the transformation (b).

Example 1 Figure 3 represents a 4-fault tolerant valid network with 4 inputs and 8 amplifiers (among them 4 can fail). This example shows how to route the signals in the inputs i_1 , i_2 , i_3 and i_4 to the 4 remaining valid outputs (the failed ones are indicated with a cross inside).




Figure 3: A 4-fault tolerant valid network with 4 inputs, 8 amplifiers and 4 failures represented by crosses inside amplifiers.

Realizing a valid k -fault tolerant network is easy, but it is difficult to optimize it. For many reasons (layout properties, reliability, energy saving, etc), but mainly to decrease launch costs, it is crucial to minimize the network's physical weight, i.e., to minimize the number of switches, the number of links (waveguides) and their length. Remember, on-board networks are not simple small electronic devices. The technology of such networks is based on wave guides with specific switches and TWTA amplifiers; the weight of such components is not negligible (cf. [3]). As launch and switch costs are dramatically high, it is worth saving even one switch. Other optimization criteria can be considered, such as minimizing the length of the routes (either in terms of the number of switches crossed or the physical length of the links) to decrease power loss, the number of routes to change or the number of elementary switch rotations in case of a failure (to increase energy savings), etc.

Here, we restrict ourselves to find a *valid k -fault tolerant network* with

- identical amplifiers;
- a number of signals equal to the number of inputs;
- a minimum number of switches.

In Section 2 we formalize the problem. Then, in Section 3 we give some preliminary results and show that in the presence of 1 or 2 faults an optimal design needs as many switches as inputs. In Section 4 we prove a composition theorem which enables us to construct large networks from smaller ones; this construction happens to be optimal for 3 or 4 faults. Finally, in the conclusion we indicate recent results obtained on the problem and new perspectives.

2 Formalization of the problem

The problem formulated in the introduction can be restated more formally as follows.

Definition 2 An (n, k) -network $G = (\{I, O, S\}, E)$ is a graph (V, E) where V is partitioned into the 3 subsets (I, O, S) respectively called the *inputs* the *outputs* and the *switches*, satisfying the following constraints:

- there are n inputs and $n + k$ outputs;
- each input and each output is connected to exactly one switch;
- switches have degree at most 4.

Definition 3 An (n, k) -network $G = (\{I, O, S\}, E)$ is **valid** if for any subset $O' \subset O$ with $|O'| \geq n$ there exists in G n edge-disjoint paths from I to O' .

Definition 4 We will denote by $\mathcal{N}(n, k)$ the minimum number of switches (i.e., the cardinality of S) of any valid (n, k) -network. A valid (n, k) -network with $\mathcal{N}(n, k)$ switches will be called **minimum**.

Our design problem consists of determining $\mathcal{N}(n, k)$ and of constructing minimum valid networks or at least valid networks having a number of switches close to this optimal value. We note that testing the validity of a network can be done by solving $\binom{n}{k}$ flow problems and therefore we can use network flow theory. However constructions in the literature of n -connected graphs cannot be used; indeed the reader has to be careful that in our networks we have 3 classes of vertices: two of them (inputs and outputs) have a given cardinality and we want to minimize the cardinality of the third class (switches). Furthermore we are interested in paths only between the inputs and the outputs (and not between any pair of vertices).

3 Preliminary results and the case of $k \leq 2$ failures

The design problem is of interest only if $k \geq 1$. Otherwise, if there is no failure, we can connect directly each input to an output link and so there is no need of a network.

Lemma 5 *In a minimum valid (n, k) -network, with $k \geq 1$ failures, there is no switch connected to 2 or more inputs.*

Proof. If there was a switch connected to 3 (or *a fortiori* 4) inputs, it will be impossible to route the input signals through the only remaining link. Now, suppose that a switch s_0 is connected to 2 inputs. If s_0 is also connected to an output, the failure of this output leaves only one available link for two input signals. Otherwise, s_0 is connected to two other switches s_1 and s_2 . In this case we can remove s_0 and connect directly the two inputs respectively to s_1 and s_2 (as amplifiers are identical). Therefore, we obtain a new valid network with one fewer switch which is in contradiction with the minimality hypothesis. \square

Corollary 6 *If $k \geq 1$, then $\mathcal{N}(n, k) \geq n$.*

Proposition 7 $\mathcal{N}(n, 1) = \mathcal{N}(n, 2) = n$.

Proof. By Corollary 6, it suffices to construct a valid $(n, 2)$ -network with n switches. Such a network is given in Figure 4; it contains n switches s_j , $1 \leq j \leq n$, each being connected to the

switch s_{j+1} for $1 \leq j \leq n-1$ and to the switch s_{j-1} for $2 \leq j \leq n$. Each switch s_j is connected to an input i_j and an output o_j . Furthermore, s_1 is connected to output o_0 and s_n to output o_{n+1} .




Figure 4: A minimum valid 2-fault tolerant network.

Now, suppose the two failing outputs are o_{j_1} and o_{j_2} with $j_1 < j_2$.

- If $j_1 = 0$ we route the input i_j as follows: for $1 \leq j < j_2$ to the output o_j via s_j ; for $j_2 \leq j < n$ to output o_{j+1} via s_j and s_{j+1} ; for $j = n$ to output o_{n+1} via s_n .
- If $j_2 = n + 1$ then we route the input i_j as follows: for $j = 1$ to o_0 via s_1 ; for $1 < j \leq j_1$ to o_{j-1} via s_j and s_{j-1} ; for $j_1 < j \leq n$ to o_j via s_j .
- If $1 \leq j_1 < j_2 \leq n$ then we route the input i_j as follows: for $j = 1$ to o_0 via s_1 ; for $2 \leq j \leq j_1$ to o_{j-1} via s_j and s_{j-1} ; for $j_1 < j < j_2$ to o_j via s_j ; for $j_2 \leq j \leq n-1$ to o_{j+1} via s_j and s_{j+1} ; for $j = n$ to o_{n+1} via s_n .

So, the network is valid. □

4 A recursive construction

We have to warn the reader that if you have a valid (n, k) -network with N switches and a valid (n', k) -network with N' switches, then the trivial union of them is unfortunately not a desired network as it contains $n + n' + 2k$ outputs (i.e., amplifiers) and can tolerate only k failures. However is also possible to give a recursive construction which enables us to double the number of failures. Consider two valid (n, k) -networks G_1 and G_2 . Replace $(n + k - 1)$ outputs in G_1 and $(n + k - 1)$ outputs in G_2 by $2(n + k - 1)$ links connected pairwise (one from G_1 paired with one from G_2) to $n + k - 1$ switches with 2 outputs. One can check that we obtain a valid $(2n, 2k)$ -network. This construction can be interesting due to its symmetry and simplicity. Unfortunately, it gives poor asymptotic bounds and we found a more efficient one (see the next theorem).

Definition 8 An I_2 -switch is a switch connected to one input, two outputs and one other switch.

In the example of Figure 3 the 4 peripheral switches are I_2 -switches.

Theorem 9 Let k be even. If there exists a valid (n, k) -network with N switches and a valid (n', k) -network with N' switches both containing at least $\frac{k}{2}$ I_2 -switches, then there exists a valid $(n + n', k)$ -network with $N + N'$ switches.

Proof. Let G (resp. G') be a valid (n, k) -network (resp. (n', k) -network) and let $s_1, s_2, \dots, s_{\frac{k}{2}}$ (resp. $s'_1, s'_2, \dots, s'_{\frac{k}{2}}$) be $\frac{k}{2}$ I_2 -switches in G (resp. G'). For $1 \leq i \leq \frac{k}{2}$, let γ_i and ω_i (resp. γ'_i and ω'_i) be the two outputs connected to s_i (resp. s'_i).

Now, we construct a network H as follows. Consider a copy of G and a copy of G' . For $1 \leq i \leq \frac{k}{2}$ remove the outputs ω_i and ω'_i and join s_i and s'_i (see Figures 5 and 6). H has $n + n'$ inputs and $(n + k) - \frac{k}{2} + (n' + k) - \frac{k}{2} = n + n' + k$ outputs. So, H is an $(n + n', k)$ -network. Furthermore, H has $N + N'$ switches. To finish the proof it remains to show that H is valid.

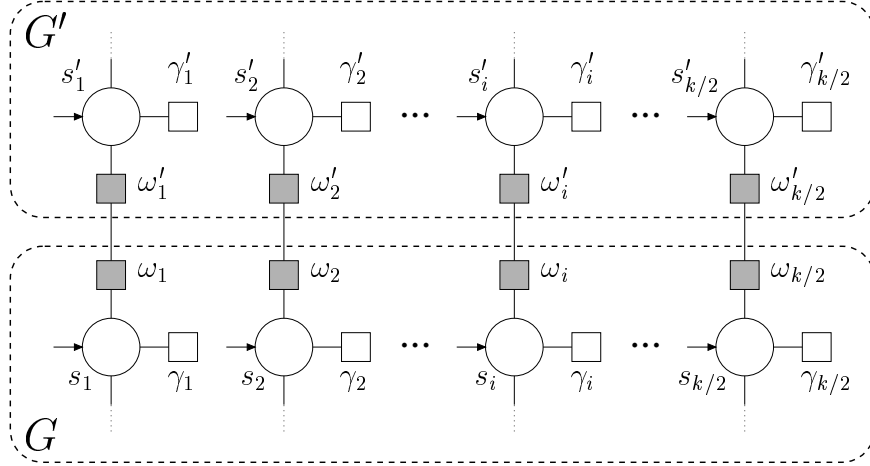


Figure 5: A valid $(n + n', k)$ -network H (k even) can be obtained by the composition of a valid (n, k) -network G with a valid (n', k) -network G' each with $\frac{k}{2}$ I_2 -switches.

Let K be a set of k failed outputs in H and let k_G (resp. $k_{G'}$) be the number of failed outputs in the copy of G (resp. G') in H . We have $k_G + k_{G'} = k$. Without loss of generality, let us suppose that $k_G \geq k_{G'}$ (the reverse case can be dealt with by exchanging the roles of G and G'); therefore $k_G \geq \frac{k}{2}$. Among the $k_{G'}$ failures, we distinguish those on outputs of type γ'_i and we denote by $J' = \{j_1, \dots, j_h\}$ the indices of the failed γ'_i , i.e., $\gamma'_{j_1}, \dots, \gamma'_{j_h}$.

Now, let us route the n inputs in the copy G . In G , consider the following set of k failures: the k_G failures of H , plus $k_{G'}$ outputs of type ω_j , $j \in J$ with $J \supseteq J'$ (that is we necessarily choose as failed the h outputs $\omega_{j_1}, \dots, \omega_{j_h}$; the others are arbitrarily chosen among the ω_j). As G is a valid (n, k) -network we can route in G the n inputs of G to n outputs that are not failed. In H , we route an input of the copy of G using the above routing obtained in G except if it was routed to an output ω_j in which case we route it to γ'_j using at the end the switches s_j and s'_j . Note that routing to γ'_j can happen only if $j \notin J$; therefore $j \notin J'$ and so γ'_j is not failed in H .

Now, it remains to route the n' inputs in the copy of G' . For that, we can use an output in the copy of G' except if it is failed in H (there are $k_{G'}$ such failures) or used by the inputs of G that is of the form γ'_j with $j \notin J$ (there are $k - k_{G'}$ such amplifiers). So, there remain n' valid outputs in G' and we use the routing in G' as it is a valid (n', k) -network. \square

Notice that if G (resp. G') contains m (resp. m') I_2 -switches, with $m \geq \frac{k}{2}$ and $m' \geq \frac{k}{2}$, then H contains $(m + m' - k)$ I_2 -switches. In particular, if G and G' contain at least k I_2 -switches, H also contains at least k I_2 -switches and we can apply recursively the theorem.

Corollary 10 *If there exists a valid (n, k) -network with N switches and at least k I_2 -switches, then for any integer p there exists a valid (pn, k) -network with pN switches and k I_2 -switches.*

Proof. By induction on p . The property is true for $p = 1$ as G is a valid (n, k) -network. Suppose it is true for p and let G' be the valid (pn, k) -network with k I_2 -switches. By applying Theorem 9 to G and G' , the graph H obtained is a valid $((p + 1)n, k)$ -network with still k I_2 -switches. \square

Proposition 11 $\mathcal{N}(4p, 4) \leq 5p$.

Proof. For $k = 4$, Figure 3 displays an example of a valid $(4, 4)$ -network with 5 switches, 4 of them being I_2 -switches. So, by Corollary 10 we can construct a valid $(4p, 4)$ -network with $5p$ switches. \square

The cases $p = 2$ and 3, that is $n = 8$ and 12, are illustrated in Figure 6.




Figure 6: Examples of recursive composition. (a) displays a valid $(4, 4)$ -network. With this network, we obtain by composition a valid $(8, 4)$ -network (b) and a valid $(12, 4)$ -network (c).

In a preliminary version [2] of this paper, we conjectured that $\mathcal{N}(4p, 4) = 5p$. This result has been proved by S. Pérennes and D. Tóth [4].

Theorem 12 (S. Pérennes, D. Tóth) $\mathcal{N}(n, 3) = \mathcal{N}(n, 4) = \lceil \frac{5n}{4} \rceil$.

We similarly obtain constructions for $k = 6, 8, \dots$ starting with small examples [2]. For $k = 6$ a valid $(6, 6)$ -network is given in Figure 7. One can prove that it is a minimum network [2, 4]. It has 9 switches, among them 6 are I_2 -switches. So, we have $\mathcal{N}(6p, 6) \leq 9p$.




Figure 7: A minimum $(6, 6)$ -network with 9 switches, including 6 I_2 -switches.

Unfortunately, for general p , these networks are not minimum. For example in Figure 8, we have illustrated a valid $(12, 6)$ -network with only 17 switches to be compared with the 18 switches obtained by Theorem 9. It is also a minimum network (see [4]).




Figure 8: A minimum $(12, 6)$ -network with only 17 switches.

5 Conclusion

In this paper we have given preliminary results and a recursive construction of valid (n, k) -networks. Using this construction and other variants we have been able to construct small networks of practical importance; for example for the satellite ASTRA1K we have reduced the number of switches used by 50 (from a total of 249).

Recently, S. Pérennes and D. Tóth [4] have found a general methodology which enables them both to obtain lower bounds on $\mathcal{N}(n, k)$ and to derive constructions of valid (n, k) -network asymptotically matching the lower bounds. For example they have shown $\mathcal{N}(n, 6) = \frac{5n}{4} + \sqrt{\frac{n}{8}} + O(1)$, $\mathcal{N}(n, 8) = \frac{4n}{3} + \frac{2}{3}\sqrt{\frac{n}{3}} + O\left(n^{\frac{1}{4}}\right)$, $\mathcal{N}(n, 10) = \frac{11n}{8} + \Theta(\sqrt{n})$ and for larger k they show that $\mathcal{N}(n, k) \leq \frac{3n}{2} + \frac{k}{2}$; and this bound is asymptotically tight. In all the example we have found, the values of $\mathcal{N}(n, k)$ are the same (or differ by at most one) for $k = 2p$ or $2p + 1$; perhaps this property is always true.

Many variants of this problem are of practical interest and are considered in forthcoming papers. An important case appears to be one in which, without failures, an input signal can be routed directly to some amplifier. That means that a switch connected to an input should also be connected to at least one output. That is the case for our optimal solutions for $k = 4$ but not for the optimal solutions for greater values of k . Partial answers are given in [4]; for example, with 6 failures one needs $n + \frac{1}{2}n + O(1)$ switches. D. Tóth also considered the case where the network should be planar which is interesting for the practical realization. As stated in the introduction other optimization criteria can be considered.

Furthermore in practice the inputs are not all used simultaneously and we have to route p signals arriving in p entries (chosen among n) to p amplifiers (chosen among $p + k$). We define a (p, n, k) -network with $p \leq n$ as having n inputs and $p + k$ outputs. It is valid if we can route any set of p inputs to any set of p outputs. We could define $\mathcal{N}(p, n, k)$ as the minimum number of switches of a valid (p, n, k) -network. The determination of $\mathcal{N}(p, n, k)$ appears very difficult.

A completely different class of problems appears if an input signal needs to be routed to a specific kind of amplifier. An extremal case appears when we have to realize a permutation of the inputs on the valid outputs. When there is no failure this is an old problem studied in multistage interconnection networks. For details on this problem we refer the reader to [1].

Acknowledgements

We thank very much Bruno Vaillant (ALCATEL SPACE INDUSTRIES) for introducing us to the problem and for giving us some of the basic ideas of the paper. We thank S. Pérennes and C. Peyrat for their helpful remarks. We would specially like to thank the referees for their relevant comments.

References

- [1] B. Beauquier and E. Darrot. On arbitrary size Waksman networks and their vulnerability. Research Report 3788 — Thème 1, INRIA Sophia Antipolis, France, October 1999. To appear in *Parallel Processing Letters*.
- [2] E. Darrot. *Diffusion de messages longs et propriétés structurelles dans les réseaux d'interconnexion*. PhD thesis, Université de Nice - Sophia Antipolis, France, June 1997.
- [3] K. Y. Eng and M. Hecht. Switch matrix for TWTA redundancy on communication satellites. In *Record of the IEEE 1978 National Telecommunications Conference*, pages 185–191, Birmingham, AL, USA, December 1978.
- [4] S. Pérennes and D. Tóth. Design of fault tolerant flow networks. In preparation.