

Design of Minimal Fault Tolerant On-Board Networks: Practical constructions ^{*}

Jean-Claude Bermond¹, Frédéric Giroire², and Stéphane Pérennes¹

¹ Mascotte Project, CNRS/I3S/INRIA, 2004 route des Lucioles, B.P. 93, _ F-06902 Sophia-Antipolis Cedex, France.

² Project Algorithms, INRIA Rocquencourt, F-78153 Le Chesnay, France

Abstract. The problem we consider originates from the design of efficient on-board networks in satellites (also called Traveling Wave Tube Amplifiers). Signals incoming in the network through ports have to be routed through an on-board network to amplifiers. The network is made of expensive switches with four links and subject to two types of constraints. First, the amplifiers may fail during satellite lifetime and cannot be repaired. Secondly, as the satellite is rotating, all the ports are not well oriented and hence not available. Let us assume that we have $p + \lambda$ ports (inputs) and $p + k$ amplifiers (outputs), then a (p, λ, k) -network is said to be *valid* if, for any choice of p inputs and p outputs, there exist p edge-disjoint paths linking all the chosen inputs to all the chosen outputs. Then, the objective is to design a valid network having the minimum number of switches denoted $N(p, \lambda, k)$. In the special case where $\lambda = 0$, these networks were already studied as *selectors*. Here we present validity certificates from which derive lower bounds for $N(p, \lambda, k)$; we also provide constructions of optimal (or quasi optimal) networks for practical values of λ and k ($1 \leq \lambda \leq k \leq 8$) and a general way to build networks for any k and any λ .

1 Introduction

Motivation. The problem we consider here was asked by Alcatel Space Industries and consists of designing efficient on-board networks in satellites (problem called Traveling Wave Tube Amplifiers). The satellites under consideration are used for TV and video transmission (like for example the Eutelsat or Astra series) as well as for private applications. Signals incoming in a telecommunication satellite through ports have to be routed through an on-board network to amplifiers. A first constraint is that the network is built of switches with four links. But other constraints appear. On the one hand the amplifiers may fail during satellite lifetime and cannot be repaired. On the other hand, as the satellite is rotating, all the ports and amplifiers are not well oriented and hence not available. So more amplifiers and ports are needed than the number of signals which have to be routed. Note that in this context, contrary to classical networks, there

^{*} This work has been partially funded by the European project IST FET AEOLUS,

are no failure of links nor switches. Indeed the switches are very reliable rotating mechanical systems and links are just big wave-tubes. However two different signals cannot use the same tube.

To decrease launch costs, it is crucial to minimize the network physical weight, i.e. for us, to minimize the number of switches. Each switch weighs about 200g and saving a switch implies a gain of more than 20 000 Euros; therefore it is worth saving even one. Space industries are interested in designing such networks for specific values of the parameters. However the general theory is of interest by itself.

Problem. We consider here *networks*, that is graphs connecting *inputs* to *outputs* and where vertices represent the switches. We define a (p, λ, k) -*network* as a network with $p + \lambda$ inputs and $p + k$ outputs. A (p, λ, k) -network is said to be *valid*, if, for any choice of p inputs and of p outputs, there exist p edge-disjoint paths linking all the chosen inputs to all the chosen outputs. For symmetry reason, we may assume in the following that $k \geq \lambda$ and we note $n := p + k$. Note that when chosen the paths become directed from the input to the output. But the disjointness condition is undirected as there cannot be two signals in a wave-tube (edge).

We study the case where the switches of the network have degree four (although the theory can be generalized to any degree) which is of primary interest for the applications. The problem is to find $N(p, \lambda, k)$, the *minimum* number of switches in a valid (p, λ, k) -network and to give *constructions* of such optimal networks (see Figure 4).

Note that finding a minimal network is a challenging problem: the number of possible networks grows exponentially and even testing the validity of a given network is hard. Indeed if we fix the valid inputs and the valid outputs testing the validity reduces to a flow problem but the number of possible choices for inputs and outputs grows exponentially as they are binomial coefficients. Still, the problem is in Co-NP, since one prove that a network is not valid by exhibiting a bad cut. In fact, deciding if a given (p, λ, k) -network is valid is a Co-NP complete problem, see [1].

In the specific case $\lambda = k$, it is interesting to design networks with a particular property: every switch linked to a port is also linked to an amplifier; indeed if there is no failure the incoming signal is routed directly to the amplifier connected to its entering switch. This minimizes the length of signals and avoids the interferences. These networks are called *simplified networks*. Observe that in that case every switch is linked to either two or four switches.

Related Work. When $\lambda = 0$, a valid network is called a *selector* (For a survey on selectors, see [2] or the seminal work of Pippenger [3]). A general theory of selectors can be found in [4] where several results are obtained for small values of k . For example it is proved that $N(p, 0, 4) = \lceil \frac{5p}{4} \rceil$.

In [5] the case of selectors with switches of degree $2k > 4$ is considered. In [6] the authors consider a variant of selectors where some signals have priority and

should be sent to amplifiers offering the best quality of service. In [7] the authors study the case where all the amplifiers are different and where a given input has to be sent to a dedicated output the problem being related to permutation networks.

Results. We first present a simple cut criterion which implies the validity of networks. This criterion will be useful both to prove the validity of the designed networks (giving upper bounds) and also to find lower bounds for the minimal number $N(p, \lambda, k)$ of switches of valid (p, λ, k) -networks.

In Section 3, we present ways to build valid networks close to minimal for small values of λ and k ($1 \leq \lambda \leq k \leq 8$). For instance, for $k \in \{3, 4\}$ and $0 < \lambda \leq 4$

$$N(p, \lambda, k) = \lceil \frac{5n}{4} \rceil$$

For $k \in \{5, 6\}$ and $0 < \lambda \leq 6$

$$N(p, \lambda, k) \leq \lceil \frac{3n}{2} \rceil$$

Examples of $(p, 4, 4)$ and $(p, 6, 6)$ -networks are given in Figure 3. In Section 4, we present a general way to build networks for any k and any λ .

2 Preliminaries

In this section, we define more formally the problem and introduce notations used throughout the paper.

We state a cut criterion (Proposition 1): this criterion is fundamental because it characterizes the validity of (p, λ, k) -networks. It is extensively used to prove that networks are valid. In Section 3 we use the cut criterion to detect forbidden patterns leading to lower bounds for the number of switches of valid networks.

Proofs of lower and upper bounds are simplified by the construction theorems given here (Theorems 1 and 2). We also give one way to build any (p, λ, k) -networks from (p, k) -selectors leading us to a linear bound for the number of switches of (p, λ, k) -networks.

Notations. Let $G = (V, E)$ be a graph and let $W \subseteq V$ be a subset of vertices of G . We denote by $\Delta(W)$ the set of edges connecting W and $\bar{W} = V \setminus W$ and by $\Gamma(W)$ the set of vertices of \bar{W} adjacent to vertices of W .

The cardinality of $\Delta(W)$ is denoted by $\delta(W)$. More generally, the convention used in this paper is that, if a set is denoted by an upper case letter, the corresponding lower case letter denotes its cardinality.

For the sake of simplicity, when a subset has a single element, we note $\delta(v)$ instead of $\delta(\{v\})$.

(p, λ, k) -networks and valid (p, λ, k) -networks. A (p, λ, k) -network is a triple $N = \{(V, E), i, o\}$, where (V, E) is a graph. i, o are positive integral functions defined on V , called input and output functions, such that for any $v \in V$,

$i(v) + o(v) + \delta(v) \leq 4$ (for parity reasons, some switches can be linked to a dead-end. The total number of inputs is $i(V) = \sum_{v \in V} i(v) = p + \lambda$, and the total number of outputs is $o(V) = \sum_{v \in V} o(v) = p + k$. We can see a network as a graph where all vertices but the leaves have degree 4, and where the leaves correspond to inputs or outputs or dead-ends. A *non-faulty output function* is a function o' defined on V such that $o'(v) \leq o(v)$ for any $v \in V$ and $o'(V) = p$. A *used input function* is a function i' defined on V such that $i'(v) \leq i(v)$ for any $v \in V$ and $i'(V) = p$. A (p, λ, k) -network is said *valid* if for any faulty output function o' and any used input function i' , there are p edge-disjoint paths in G such that each vertex $v \in V$ is the initial vertex of $i'(v)$ paths and the terminal vertex of $o'(v)$ paths.

Design Problem. Let $N(p, \lambda, k)$ denotes the minimum number of switches of a valid (p, λ, k) -network. The *Design Problem* consists in determining $N(p, \lambda, k)$ and in constructing a minimum (p, λ, k) -network, or at least a valid (p, λ, k) -network with a number of vertices close to the optimal value.

We introduce a variation of the problem: consider networks with $p + \lambda$ switches with exactly one input and one output (we call such a switch a *doublon*), and with $k - \lambda$ switches with only one output. To find minimum valid network like these is what we call the *Simplified Design Problem*. Networks of this kind are especially good for practical applications, as they simplify the routing process, minimize path lengths and lower interferences between signals. The minimal number of switches of such networks is noted $N'(p, \lambda, k)$.

Excess, Validity and Cut-criterion. We show that, to verify if a network is valid, instead of solving a flow/supply problem for each possible configuration of output failures and of used inputs, it is sufficient to look at an invariant measure of subsets of the network, the *excess*, as expressed in the following proposition.

Definition 1 (Excess $\varepsilon(W)$) Let $\{(V, E), i, o\}$ be a (p, λ, k) -network and $W \subseteq V$ a subset of vertices. The excess in inputs of W is defined as

$$\varepsilon_i(W) := \delta(W) + o(W) - \min(k, o(W)) - \min(i(W), p).$$

The excess in outputs of W is defined as

$$\varepsilon_o(W) := \delta(W) + i(W) - \min(\lambda, i(W)) - \min(o(W), p).$$

Lemma 1 Let $\{(V, E), i, o\}$ be a (p, λ, k) -network. Consider a subset $W \subseteq V$. We have $\varepsilon_o(\overline{W}) = \varepsilon_i(W)$.

Proof. As $\delta(\overline{W}) = \delta(W)$, $o(\overline{W}) = p + k - o(W)$, $i(\overline{W}) = p + \lambda - i(W)$,

$$\begin{aligned} \varepsilon_o(\overline{W}) &= \delta(W) + p + \lambda - i(W) - \min(\lambda, p + \lambda - i(W)) \\ &\quad - \min(p + k - o(W), p) \\ &= \delta(W) + p + \lambda - i(W) - (\lambda - i(W)) \\ &\quad + \min(i(W), p) - (p - o(W) + \min(k, o(W))) \\ &= \delta(W) + o(W) - \min(k, o(W)) - \min(i(W), p) \\ &= \varepsilon_i(W). \end{aligned}$$

So we can mainly restrict our attention to $\varepsilon(W) := \varepsilon_i(W)$.

Proposition 1 (Cut Criterion) *A (p, λ, k) -network is valid if and only if, for any subset of vertices $W \subset V$, the excess of W satisfies $\varepsilon(W) \geq 0$.*

The intuition is that the signals arriving in W (in number at most $\min(i(W), p)$) should be routed either to the valid outputs of W (in number at least $o(W) - \min(k, o(W))$) or to the links going outside (in number $\delta(W)$). The omitted formal proof reduces to a supply/demand flow problem.

Proposition 2 (Cut Criteria and Symmetry) *A (p, λ, k) -network is valid if and only if one of the following proposition is true.*

1. For all W , we have $\varepsilon_i(W) \geq 0$.
2. For all W , we have $\varepsilon_o(W) \geq 0$.
3. For all W , with $o(W) \leq \lceil \frac{p+k}{2} \rceil$ or $i(W) \leq \lceil \frac{p+\lambda}{2} \rceil$, we have $\varepsilon_i(W) \geq 0$ and $\varepsilon_o(W) \geq 0$.

Proof. Direct by Proposition 1 and Lemma 1.

Proposition 3 (Cut Criteria and Connectivity) *For all cut criteria of Proposition 2, it is sufficient to consider only connected subsets W with connected complements \overline{W} .*

This comes from the submodularity of ε (the proof is omitted here). Intuitively, if a W has a negative excess, one of its connected component also has a negative excess. If W has a negative excess, \overline{W} has negative excess in outputs, and so one of its connected component.

Construction Theorems




Fig. 1. Valid symmetric $(2,0,2)$ -network and $(2,2,0)$ -network.

Theorem 1 *The problem is symmetric in inputs and outputs, that is $N(p, \lambda, k) = N(p, k, \lambda)$*

Proof. As soon as we have a valid (p, λ, k) -network, we can immediately build a valid (p, k, l) -network which is the same after we have inverted the inputs and the outputs (see example of Figure 1). The validity is given by Proposition 2.

Theorem 2 A valid (p, λ, k) -network \mathcal{R} can be built from a valid (p, k, k) -network \mathcal{R}' by removing $k - \lambda$ inputs.

Proof. \mathcal{R} and \mathcal{R}' have the same set of switches, V . Let W be a subset of V . $i_{\mathcal{R}}(W)$ (resp. $i_{\mathcal{R}'}(W)$) denotes the number of inputs attached to W in \mathcal{R} (resp. \mathcal{R}'). As $i_{\mathcal{R}}(W) \leq i_{\mathcal{R}'}(W)$, $\varepsilon_{\mathcal{R}}(W) \geq \varepsilon_{\mathcal{R}'}(W)$, finishing the proof.




Fig. 2. A (p, λ, k) -network built with two selectors in series.

Links with (p, k) -networks or (p, k) -selectors. (p, λ, k) -networks where all the inputs are used ($\lambda = 0$) have been studied in [4] and [8] and are called (p, k) -networks or (p, k) -selectors. We have

$$N(p, 0, k) = N(p, k, 0) = N(p, k). \quad (1)$$

Definition 2 (2selectors (p, λ, k) -networks) A 2selectors (p, λ, k) -network is formed by a (p, λ) -selector and a (p, k) -selector in serial, as indicated in the following theorem.

Theorem 3

$$N(p, \lambda, k) \leq N(p, \lambda) + N(p, k)$$

$$N(p, \lambda, k) \leq O(p + k)$$

Proof. The proof is constructive. We can construct a valid (p, λ, k) -network with two valid selectors, a (p, λ) -selector and a (p, k) -selector as shown in Figure 2. The idea is to use the first selector in a symmetric way by replacing its $p + \lambda$ outputs by our inputs, then to link the inputs of the two selectors all together. The outputs of the second selector will be our outputs. Any subset of size p of the inputs can be routed to the p central links by the first selector. The second one can route these links to any subset of our outputs. So the network is valid. We call the networks built this way 2selectors (p, λ, k) -networks (Definition 2).

As it is shown in [4] and [8], that the minimum number of switches of a (p, k) -selector is linear, we have the same result for (p, λ, k) -networks.

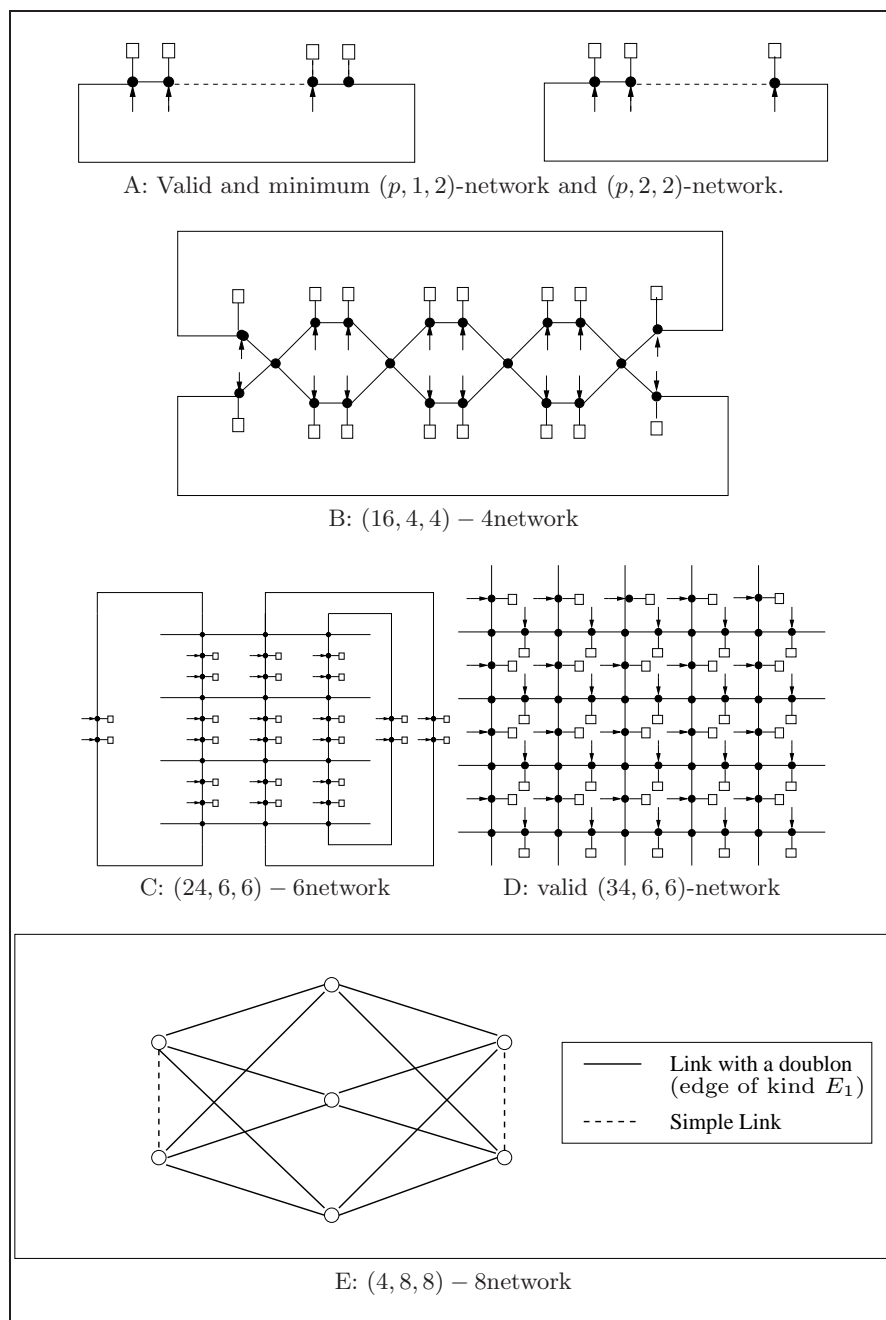


Fig. 3. Constructions for small k

3 Constructions for Small Cases

Recall that we may assume here that $1 \leq \lambda \leq k$ (see Theorem 1 and Equation 1). We first present results for $p = 1, 2$ and $k = 1, 2$, then for $p \geq 3$ and $k \geq 3$. We construct valid networks for $1 \leq k \leq 8$, called respectively 2networks, 4networks, 6 networks and 8networks. The number of switches of these networks are optimal for $\lambda = k$. For $\lambda < k$ we use Theorem 2 to obtain networks close to minimal. The proofs are mostly omitted here.

Case $p = 1, 2$ and any λ, k .

Theorem 4 $N(1, \lambda, k) = \lceil \frac{\lambda+k}{2} \rceil$ and $N(2, \lambda, k) = \lceil \frac{\lambda+k}{2} \rceil + 2$

Proof. When $p = 1$, we build a network consisting of a path where one end vertex has 3 inputs and the other 3 outputs. The internal vertices have either 2 inputs or 2 outputs when λ and k are even as shown in Figure 4. One vertex can have only one input or one output or one input and one output according λ or k or both are odd. When $p = 2$, we build similarly a network consisting of a cycle with a maximal number of switches with 2 inputs or outputs, as shown in Figure 4. The numbers of switches of the above networks attain these bounds.

To prove the validity we will use the cut criterion. When $p = 1$, let W be a connected subset of V (see Proposition 3 for this choice of W). As $\delta(W) \geq 1$, we have $\varepsilon(W) \geq 1 + o(W) - \min(o(W), k) - 1 \geq 0$. The cut criterion finishes the proof. When $p = 2$, we have $\delta(W) \geq 2$. So $\varepsilon(W) \geq 2 + o(W) - \min(o(W), k) - 2 \geq 0$. The networks are valid.

When $p = 1$, the network is minimal because a valid minimal network has to be connected. When $p = 2$, if we construct a network with one vertex less, we have a node v with 2 inputs and 1 output (or 2 outputs and 1 input). $\varepsilon(\{v\}) = 1 + 1 - 1 - 2 = -1$ ($\varepsilon_o(\{v\}) = -1$) and the network would not be valid.

Case $k = 1, 2$ (2networks).

A 2network consists of a cycle of $p + \lambda$ vertices with one input and one output, plus one vertex with one output if $\lambda = 1$ and $k = 2$, connected in a cycle, as shown in Figure 3 A.

Theorem 5 $N(p, \lambda, 2) = p + 2$. $N(p, \lambda, 1) = p + 1$.

Proof. If $W \subseteq V$, we have $\delta(W) = 2$. As $k \leq 2$, $\delta(W) - \min(k, o(W)) \geq 0$. Moreover, as, by construction, $i(W) \leq o(W)$, we have $o(W) - \min(p, i(W)) \geq 0$. So $\varepsilon(W) \geq 0$ and 2networks are valid.

If we construct a network with one vertex less, we obligatory have a vertex with $i(v) = 1$ and $o(v) = 2$. Consider the subset W made of this switch $\varepsilon_o(W) = 1 + 1 - 1 - 2 = -1$. This network would not be valid. So 2networks are minimal. For $k = 1$ we use a connection of the switches via a path.

Case $p \geq 3$ and $k \geq 3$

Let \mathcal{R} be a (p, λ, k) -network. Recall that $n := p + k$. In what follows we suppose $p \geq 3$ and $k \geq 3$.

Definition 3 (Doublons, R-Switches) A doublon of \mathcal{R} is a vertex with $i(v) = o(v) = 1$. An R-switch is a vertex that is not a doublon.

Definition 4 (Edges of kind E_0 , E_1 and E_2) We build a graph G associated to \mathcal{R} . Its vertices are the R-switches of \mathcal{R} . Its edges are of three kinds, respectively E_0 , E_1 and E_2 : the edges of \mathcal{R} between two R-switches, the edges corresponding in \mathcal{R} to a path of length 2 with a doublon in the middle and those corresponding to a path of length 3 with 2 doublons in the middle.

Note that the R-switches and doublons partition \mathcal{R} and that the cut criterion gives immediately that edges of other kinds corresponding to paths of length more than 3 with doublons in the middle are forbidden. Indeed, if we consider the set W consisting of these doublons, we have $\delta(W) = 2$ and $o(W) = i(W) \geq 3$.




Fig. 4. Valid and minimum $(1, \lambda, k)$ -network and $(2, \lambda, k)$ -network

Case $k \leq 3, 4$, (4networks).

A 4network is built from blocks made of one node with no inputs nor outputs and 4 doublons on 2 edges of kind E_2 . A block is connected to two identical blocks in serial as shown in the $(16, 4, 4)$ – 4network of Figure 3 B. Each block has 5 nodes, 4 inputs and 4 outputs except one or two if $(p+k) \bmod 4 \neq 0$ or $(p+\lambda) \bmod 4 \neq 0$. Counting the number of switches and checking that the 4networks are valid gives the following theorem:

Theorem 6 For $k = 3, 4$, $N(p, \lambda, k) \leq n + \frac{n}{4} + c'_4 - c$, $c'_4 = \lceil \frac{n \bmod 4}{4} \rceil$ and $c = \lfloor \frac{k-\lambda}{2} \rfloor$.

On the other hand we can prove the following lower bound:

Theorem 7 For $k \geq 3$, $N(p, \lambda, k) \geq n + \frac{n}{4} - c''_4$, where $c''_4 = \frac{k-\lambda}{2} + \frac{k-\lambda}{8}$.

Notice that the difference between the number of switches of a (p, λ, k) –4network and the lower bound is at most 1 and that in the case $\lambda = k = 4$ we obtain:

Corollary 1 $N(p, 4, 4) = n + \frac{n}{4}$.

Case $k = 5, 6$, (6networks).

A 6network is built from blocks made of 3 switches connected in circle, each of them is connected to two doublons on an edge of kind E_2 . A block is connected

to two identical blocks in serial as shown in the $(24, 6, 6) - 6$ network of Figure 3 C. Each block has 9 nodes, 6 inputs and 6 outputs except one or two if $(p + k) \bmod 4 \neq 0$ or $(p + \lambda) \bmod 4 \neq 0$. Counting the number of switches of a $(p, \lambda, k) - 6$ network and checking the validity of the 6networks we obtain:

Theorem 8 For $k = 5, 6$, $N(p, \lambda, k) \leq n + \frac{n}{2} + c'_6 - c$, where $c'_6 = 3 \lceil \frac{n \bmod 6}{6} \rceil$ and $c = \lfloor \frac{k-\lambda}{2} \rfloor$.

Theorem 9 For $k \geq 5$ and for the Simplified Design Problem

$$N'(p, \lambda, k) \geq n + \frac{n}{2} - c''_6,$$

with $c''_6 = \frac{k-\lambda}{2} + \frac{k-\lambda}{4}$.

Notice that the difference between the number of switches of a $(p, \lambda, k) - 6$ network and the lower bound is at most 4. In the general case we have obtained the following bound:

Theorem 10 For the General Design Problem and for $k \geq 5$

$$N(p, \lambda, k) \geq n + \frac{3n}{8} - c'',$$

with $c'' = \frac{k-\lambda}{2} + \frac{3(k-\lambda)}{16}$.

We found also an other family of valid networks with the same number of switches. Nodes with no inputs nor outputs are connected on a grid on a sphere with edges of kind E_1 , as shown in Figure 3 D.

Case $k = 7, 8$, (8networks).

An 8network is built with n doublons, $\frac{n}{4}$ nodes in N_4 and $\frac{n}{3}$ nodes in N_3 . Nodes in N_4 and N_3 have no inputs nor outputs. A node in N_4 is connected to four nodes in N_3 via an edge of kind E_1 (see Definition 4). A nodes in N_3 is connected to three nodes in N_4 via an edge of kind E_1 and to one node in N_3 via an edge in E_0 . N_3 is divided in four groups with the condition that two nodes linked with an edge in E_0 are in different groups. Each node in N_4 is connected to there four groups as shown in the $(4, 8, 8) - 8$ network with 19 nodes of Figure 3 E. Here again counting the number of switches and checking that the 8networks are valid we obtain:

Theorem 11 For $k = 7, 8$,

$$N(p, \lambda, k) \leq n + \frac{7n}{12} + c'_8 - c.$$

with $c'_8 = 7 \lceil \frac{n \bmod 12}{12} \rceil$ and $c = \lfloor \frac{k-\lambda}{2} \rfloor$.

Theorem 12 For $k \geq 7$ and for Simplified Design Problem with no edges E_2 :

$$N'(p, \lambda, k) \geq n + \frac{7n}{12} - c''_8,$$

with $c''_8 = \frac{k-\lambda}{2} + \frac{7(k-\lambda)}{24}$.

Notice that the difference between the number of switches of a (p, λ, k) – 8network and the lower bound is at most 8.

4 Constructions for any k and any λ . General (p, λ, k) –Networks

We present here General (p, λ, k) –Networks (see Definition 7). Their sizes are close to minimal for small k (see Remark 1). They are built with ν –boxes introduced in Definition 6.

4.1 Preliminaries: ν –boxes and ν –permutation networks




Fig. 5. A minimal 4–box.

The decisive property of ν –boxes is expressed in Lemma 2.

Definition 5 (ν –graph) A ν –graph ($\nu \geq 2$) is a pair (G, l) , where $G = (V, E)$ is a simple non oriented graph and l a positive integer function defined on V such that $l(V) = 2\nu$ and that, for every vertex x , one has $l(x) + \delta(x) = 4$.

Definition 6 (ν –box) A ν –box is a ν –graph such as for all integer function i defined on V with $0 \leq i \leq l$ and $i(V) = \nu$, there exist ν edge-disjoint paths such that every vertex x is the beginning of $i(x)$ paths and the end of $l(x) - i(x)$ paths.

Examples: For $\nu = 1$ a minimal ν –box is reduced to a vertex. For $\nu = 2$ a minimal ν –box consists of a triangle. A minimal 4–box can be seen in Figure 5. These examples are obtained from permutation networks. A ν –permutation network is a network that can route its ν inputs to any permutation of its ν outputs.

Proposition 4 A ν -permutation network is a ν -box.

Proofs of properties of ν -boxes are omitted here. There exist linear asymptotic constructions for ν -boxes. Nevertheless, for small ν , no constructions of ν -box smaller than the corresponding permutation network have been found. For $\nu \leq 6$, it has been proved that minimal permutation networks are optimal ν -boxes. For $\nu \leq 6$, these minimal networks are known as AS-Waksman Permutations networks see [7]. For these reason we choose AS-Waksman permutation networks for our ν -boxes for small k .

An useful tool to prove validity is:

Lemma 2 In a ν -box, for every subset $X \subseteq V$ we have

$$|\Gamma(X)| \geq \min(l(X), l(\bar{X})),$$

where $\Gamma(W)$ is the set of vertices of \bar{X} adjacent to a vertex of W .

The proof reduces to a flow problem.

4.2 General (p, λ, k) -Networks




Fig. 6. Scheme of General (p, λ, k) -Networks.

Definition 7 (General (p, λ, k) -Networks) A General (p, λ, k) -Networks (see Figure 6) is built with $\left\lceil \frac{p+k}{k} \right\rceil$ $\left\lceil \frac{k}{2} \right\rceil$ -boxes connected in circle. These boxes are connected with

- a maximal number, $\left\lfloor \frac{p+\lambda}{2} \right\rfloor$, edges of type E_2 ,

- 1 edge of kind E_1 if $p + \lambda$ is odd and 0 otherwise,
- $\lfloor \frac{k-\lambda}{2} \rfloor$ edges with one node with 2 outputs on it (edges of kind E'_2),
- 1 edge with one node with one output on it (edge of kind E'_1), if $k - \lambda$ is odd and 0 otherwise,
- the remaining of type E_0 $\left(\lfloor \frac{k}{2} \rfloor \lfloor \frac{p+k}{k} \rfloor - e_2 - e_1 - e'_2 - e'_1 \right)$.

Lemma 3 The number of switches of a general- (p, λ, k) -network is

$$n + \frac{B_{min}(\lfloor \frac{k}{2} \rfloor)}{2^{\lfloor \frac{k}{2} \rfloor}} n + c'_g - c,$$

with $B_{min}(\nu)$ the number of nodes of a minimal ν -box, with $c'_g = B_{min} \lfloor \frac{n \bmod 2^{\lfloor \frac{k}{2} \rfloor}}{2^{\lfloor \frac{k}{2} \rfloor}} \rfloor$ and $c = \lfloor \frac{k-\lambda}{2} \rfloor$.

k	3, 4	5, 6	7, 8	9, 10	13, 14
Size	$n + \frac{n}{4}$	$n + \frac{n}{2}$	$n + \frac{5n}{8}$	$n + \frac{8n}{10}$	$2n$

Fig. 7. Size of General (p, λ, k) -networks for small k .

Remark 1 The sizes of General (p, λ, k) -networks for small k using AS-Waksman permutation networks as $\lfloor \frac{k}{2} \rfloor$ -boxes can be seen in Figure 7. General (p, λ, k) -networks for small k are close to minimal networks (To compare with the networks of Section 3).

Proposition 5 A General (p, λ, k) -Network is a valid (p, λ, k) -Network.

5 Conclusion

There remain a lot of open problems in this general issue. For example, the demands of Alcatel include networks with few (six) to around thirty unused inputs that have to tolerate few (ten) to around twenty output failures. We present some networks close to minimal for k from one to a dozen. But it remains to find tighter bounds for them and to explore larger values of k .

References

1. Blum, M., Karp, R., Papadimitriou, C., Vornberger, O., Yannakakis, M.: The complexity of testing whether a graph is a superconcentrator. Inf. Proc. Letters **13** (1981) 164–167

2. Du, D.Z., Hwang, F.K., eds.: Notes on the Complexity of Switching Networks. Volume 42. Kluwer Academic Publishers (2000)
3. Pippenger, N.: Superconcentrators. *SIAM Journal on Computing* **6** (1977) 298–304
4. Bermond, J.C., Perennes, S., Tóth, C.D.: Fault tolerant on-board networks with priorities. Manuscrit (To be published)
5. Bermond, J.C., Delmas, O., Havet, F., Montassier, M., Perennes, S.: Réseaux de télécommunications minimaux embarqués tolérants aux pannes. In: *Algotel*. (2003) 27–32
6. Bermond, J.C., Havet, F., Tóth, D.: Fault tolerant on board networks with priorities. *Networks* **47**(1) (2006) 9–25
7. Beauquier, B., Darrot, E.: Arbitrary size Waksman networks and their vulnerability. *Parallel Processing Letters* **3-4** (2002) 287–296
8. Havet, F.: Repartitors, selectors and superselectors. Submitted to *Journal of Interconnection Networks* (2006)