Yugoslav Journal of Operations Research
26 (2016), Number 1, 17–32
DOI: 10.2298/YJOR140721004C

DIAMETER CONSTRAINED RELIABILITY OF LADDERS AND SPANISH FANS

Héctor CANCELA
Facultad de Ingeniería, Universidad de la República
cancela@fing.edu.uy

Franco ROBLEDO
Facultad de Ingeniería, Universidad de la República
frobledo@fing.edu.uy

Pablo ROMERO
Facultad de Ingeniería, Universidad de la República
promero@fing.edu.uy

Pablo SARTOR
Universidad de Montevideo
psartor@um.edu.uy

Received: July 2014 / Accepted: January 2015

Abstract: We are given a graph $G = (V, E)$, terminal set $K \subseteq V$ and diameter $d > 0$. Links fail stochastically and independently with known probabilities. The diameter-constrained reliability (DCR for short) is the probability that the K-diameter is not greater than d in the subgraph induced by non-failed links.

The contributions of this paper are two-fold. First, the computational complexity of DCR-subproblems is discussed in terms of the number of terminals $k = |K|$ and diameter d. Here, we prove that if $d > 2$, the problem is $\mathcal{NP}$-Hard when $K = V$, and second, we compute the DCR efficiently for ladders and Spanish fans.

Keywords: Diameter Constrained Reliability, Computational Complexity, Graph Theory.

MSC: 90B15, 90B18, 90B25.

1. INTRODUCTION

The diameter-constrained reliability (DCR for short) is a metric that subsumes the classical network reliability (CLR for short). Both metrics serve to model the reliability of communication networks. Vast literature exists for the CLR, see [5]. The DCR was introduced by [8], inspired in delay-sensitive applications over the Internet infrastructure. It can model situations where there are limits in the acceptable number of hops. Among these we have latency-sensitive contexts, point-to-point and voice-over-IP applications. Arnon Rosethal proved that the CLR is inside the class of $\mathcal{NP}$-Hard problems [10]. As a corollary, the general DCR is $\mathcal{NP}$-Hard as well, hence intractable unless $\mathcal{P} = \mathcal{NP}$.

Consider a network modeled by a simple graph $G = (V, E)$ with $|V| = n$ nodes and $|E| = m$ links, a distinguished set of nodes $K \subseteq V$ called *terminals*, $k = |K|$ and an integer d that represents the maximum acceptable number of hops. Links fail stochastically and independently, ruled by a vector $p = (p_1, \ldots, p_m)$.

The DCR is the probability that the terminals of the resulting subgraph remain connected by paths composed by d links, or less. This number is denoted by $R^d_{K,G}(p)$.

The focus of this paper is on the computational complexity of DCR subproblems, in terms of k and d. The article is organized in the following manner. Formal definitions of the CLR and DCR problems are provided as particular instances of stochastic binary systems in Section 2. The computational complexity of the DCR is discussed in Section 3. The main contribution of this paper is a new result on complexity theory, provided in Section 4. Specifically, we prove that the DCR is in the computational class of $\mathcal{NP}$-Hard problems in the all-terminal scenario ($k = n$) with a given diameter $d \geq 3$. In Section 5, the DCR is computed for certain elementary families. Polynomial time algorithms for the DCR evaluation are given for two specific families. Concluding remarks and open problems are summarized in Section 6.

2. TERMINOLOGY

We are given a system with m components. These components are either "up" or "down", and encoded by $x = (x_1, \ldots, x_m)$. Additionally we have a structure function $\phi : \{0,1\}^m \to \{0,1\}$ such that $\phi(x) = 1$ if and only if the system works under state x. When the components work independently and stochastically, with certain probabilities of operation $p = (p_1, \ldots, p_m)$, the pair (ϕ, p) defines a *stochastic binary system*, or SBS for short, following the terminology of Michael Ball [1]. An SBS is *coherent* whenever $x \leq y$ implies that $\phi(x) \leq \phi(y)$, where the partial order set $(\leq, \{0,1\}^m)$ is bit-wise (i.e., $x \leq y$ if and only if $x_i \leq y_i$ for all $i \in \{1, \ldots, m\}$). If $\{X_i\}_{i=1,\ldots,m}$ is a set of independent binary random variables with $P(X_i = 1) = p_i$ and $X = (X_1, \ldots, X_m)$, then $r = E(\phi(X)) = P(\phi(X) = 1)$ is the *reliability* of the SBS.

Now, consider a simple graph $G = (V, E)$, a subset $K \subseteq V$ and a positive integer d. Let us choose an arbitrary order of the link-set $E = \{e_1, \ldots, e_m\}$, $e_i \leq e_{i+1}$. For each subgraph $G_x = (V, E_x)$ with $E_x \subseteq E$, we identify a binary word $x \in \{0,1\}^m$, where $x_i = 1$ if and only if $e_i \in E_x$; this is clearly a bijection. A subgraph $G_x = (V, E_x)$ is *d-K-connected* if $d_x(u, v) \leq d, \forall \{u, v\} \subseteq K$, where $d_x(u, v)$ is the distance between nodes u and v in the graph G_x. Then, we define the structure $\phi : \{0,1\}^m \to \{0,1\}$ such that $\phi(x) = 1$ if and only if the graph G_x is *d-K*-connected. If we assume nodes are perfect but links fail stochastically and independently ruled by the vector $p = (p_1, \ldots, p_m)$, the pair (ϕ, p) is a coherent SBS. Its reliability, denoted by $R^d_{K,G}(p)$, is called *diameter constrained reliability*, or DCR for short. A particular case is $R^{n-1}_{K,G}(p)$, called *classical reliability*, or CLR for short.

In all coherent SBS, a *pathset* is a state x such that $\phi(x) = 1$. A *minpath* is a state x such that $\phi(x) = 1$ but $\phi(y) = 0$ for all $y < x$ (i.e., a minimal pathset). A *cutset* is a state x such that $\phi(x) = 0$, while a *mincut* is a state x such that $\phi(x) = 0$ but $\phi(y) = 1$ if $y > x$ (i.e., a minimal cutset).

3. COMPLEXITY

The class $\mathcal{NP}$ is the set of problems polynomially solvable by a non-deterministic Turing machine; see [6]. A problem is $\mathcal{NP}$-Hard if it is at least as hard as every problem in the set $\mathcal{NP}$ (formally, if every problem in $\mathcal{NP}$ has a polynomial reduction to the former). It is widely believed that $\mathcal{NP}$-Hard problems are intractable (i.e., there is no polynomial-time algorithm to solve them). An $\mathcal{NP}$-Hard problem is $\mathcal{NP}$-Complete if it is inside the class $\mathcal{NP}$. Valiant defines the class $\#\mathcal{P}$ of counting problems, such that testing whether an element has a property or not can be accomplished in polynomial time [12]. A problem is $\#\mathcal{P}$-Complete if it is in the set $\#$-$\mathcal{P}$ and it is at least as hard as any problem of that class.

Computing the reliability of a coherent SBS is at least as hard as recognition and counting minimum cardinality mincuts/minpaths; see [1]. Arnon Rosenthal proved that the CLR is $\mathcal{NP}$-Hard, showing that the minimum cardinality mincut recognition is precisely Steiner-Tree problem, included in Richard Karp's list [10]. The CLR for both two-terminal and all-terminal cases are still $\mathcal{NP}$-Hard, as proved by reduction to counting minimum cardinality $s-t$ cuts in [9]. As a consequence, the general DCR is $\mathcal{NP}$-Hard as well. Later effort has been focused to particular cases of the DCR, in terms of the number of terminals $k = |K|$ and diameter d. When $d = 1$, every pair of terminals must be linked by a link, $R^1_{K,G} = \prod_{\{u,v\} \subseteq K} p_{u,v}$, where $p_{u,v}$ denotes the probability of operation of link $\{u,v\} \in E$, and $p_{u,v} = 0$ if $\{u,v\} \notin E$. The problem is still simple when $k = d = 2$. In fact, $R^2_{\{u,v\},G} = 1 - (1 - p_{u,v}) \prod_{w \in V - \{u,v\}} (1 - p_{u,w} p_{w,v})$. Héctor Cancela and Louis Petingi proved that the DCR is $\mathcal{NP}$-Hard when $d \geq 3$ and $k \geq 2$ is a fixed input parameter [4], in strong contrast with the case $d = k = 2$. The literature offers at least two proofs that the DCR has a linear-time algorithm when $d = 2$ and k is a fixed input parameter. Eduardo Canale et. al. present a recursive proof, while Pablo Sartor presents an explicit expression for $R^2_{K,G}$ that is computed in a linear time of elementary operations. Figure 1 summarizes the known results for the computational complexity of the DCR in terms of k and d.

4. MAIN THEOREM

The DCR belongs to the class of $\mathcal{NP}$-Hard problems in the all-terminal case with diameter $d \geq 3$. The main source of inspiration is [4], where the authors prove that the DCR is $\mathcal{NP}$-Hard when $d \geq 3$ and $k \geq 2$ is a fixed input parameter. First, they

d \ k	k (fixed): 2	k (fixed): 3 ...	$k = n$ or free
2	$O(n)$ [4]	$O(n)$ [3]	Unknown
3 ... $n-2$	$\mathcal{NP}$-Hard [4]		Unknown
$n-1$ ...	$\mathcal{NP}$-Hard [10]		$\mathcal{NP}$-Hard [9]

Figure 1: DCR Complexity in terms of the diameter d and number of terminals $k = |K|$

show that the result holds for $k = 2$, and they further generalize the result for fixed $k \geq 2$. For our purpose it will suffice to revisit the first part. Before, we state a technical result first proved in [2]. Recall that a *vertex cover* in a graph $G = (V, E)$ is a subset $V' \subseteq V$ such that V' meets all links in E. The graph G is bipartite if there exists a bipartition $V = V_1 \cup V_2$ such that $E \subseteq V_1 \times V_2$.

Lemma 4.1. *Counting the number of vertex covers of a bipartite graph is $\#\mathcal{P}$-Complete.*

Proposition 4.2. *The DCR is $\mathcal{NP}$-Hard when $k = 2$ and $d \geq 3$.*

Proof. Let $d' = d - 3 \geq 0$ and $P = (V(P), E(P))$ a simple path with node set $V(P) = \{s, s_1, \ldots, s_{d'}\}$ and link set $E(P) = \{\{s, s_1\}, \{s_1, s_2\}, \ldots, \{s_{d'-1}, s_{d'}\}\}$. For each bipartite graph $G = (V, E)$ with $V = A \cup B$ and $E \subseteq A \times B$, we build the following auxiliary network:

$$G' = \{(A \cup B \cup V(P) \cup \{t\}, E \cup E(P) \cup \{\{s_{d'}, a\}, a \in A\} \cup \{\{b, t\}, b \in B\}, \tag{1}$$

where all links of G are perfect but the ones from $I = \{\{s_{d'}, a\}, a \in A\} \cup \{\{b, t\}, b \in B\}$, which fail independently with identical probabilities $p = 1/2$. Consider the terminal set $K = \{s, t\}$. The auxiliary graph G' is illustrated in Figure 2. The reduction from the bipartite graph to the two-terminal instance is polynomial.

A set cover $A' \cup B' \subset A \cup B$ induces a cutset $I' = \{\{s_{d'}, a\}, a \in A'\} \cup \{\{b, t\}, b \in B'\}$ (i.e. if all links in I' fail, the nodes $\{s, t\}$ are not connected). Reciprocally, that cutset determines a set cover. Therefore, the number of cutsets $|C|$ is precisely the number of vertex covers of the bipartite graph $|\mathcal{B}|$. Moreover:

$$|\mathcal{B}| = 2^{|A|+|B|}(1 - R^d_{\{s,t\},G'}(1/2)).$$

Thus, the DCR for the two-terminal case is at least as hard as counting vertex covers of bipartite graphs. □

Now we prove our main result:

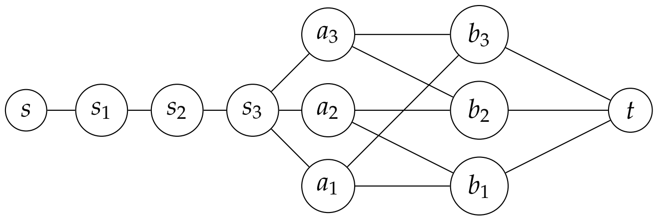

Figure 2: Example of auxiliary graph G'' with terminal set $\{s,t\}$ and $d = 6$, for the particular bipartite instance C_6.

Theorem 4.3. *The DCR is $\mathcal{NP}$-Hard when $k = n$ and $d \geq 3$.*

Proof. Extend the auxiliary graph $G' = (V', E')$ to $G'' = (V'', E'')$, where $V'' = V'$ and $E'' = E' \cup \{\{a, a'\}, a \neq a', a, a' \in A\} \cup \{\{b, b'\}, b \neq b', b, b' \in B\}$. In other words, just add links in order to connect all nodes from A, and all nodes from B. We keep the same probabilities of operation, where new links are perfect.

Consider now the all-terminal case $K = V''$ for G'', and given diameter $d \geq 3$. The key is to observe that the cutsets in the all-terminal scenario for G'' are precisely the $s-t$ cutsets in G', and they have the same probability. Indeed, each pair of terminals from the set A are directly connected by perfect links; the same holds in B. The distance between s and $s_{d'}$ is $d' = d - 3 < d$, so these nodes (and all the intermediate ones) respect the diameter constraint. Finally, if there were an $s - t$ path (i.e. a path from s to t), the diameter of the resulting subgraph of G'' would be exactly d. Therefore, $R^d_{\{s,t\},G'} = R^d_{V'',G''}$, and again:

$$|\mathcal{B}| = 2^{|A|+|B|}(1 - R^d_{\{s,t\},G'}(1/2)) = 2^{|A|+|B|}(1 - R^d_{V'',G''}(1/2)).$$

Thus, the DCR for the all-terminal case is at least as hard as counting vertex covers of bipartite graphs. □

5. POLYNOMIAL ALGOROTHMS FOR SPECIAL TOPOLOGIES

In this section, we provide expressions for computing the DCR in elementary graphs (this is, graphs with maximum degree 2), as well as two efficient algorithms to find the DCR in non-elementary families, to know, ladders and Spanish fans.

5.1. *Elementary Graphs*

An elementary graph G has maximum degree $\Delta_G = 2$. If G is not connected, its DCR is null. The other possible graphs are either paths or cycles.

An elementary path $P_n = \{x_1, \ldots, x_n\}$ with terminal nodes $K = \{x_1, x_n\}$ has a trivial expression for the DCR. Indeed, $R^d_{\{x_1,x_n\},P_n} = \prod_{i=1}^{n-1} p(x_i x_{i+1}) 1_{\{d \geq n-1\}}$, where $1_{\{x\}}$ is one if and only if x is true, and 0 otherwise.

In the cycle $C_n = (\{x_1, \ldots, x_n\}, \{x_n, x_0\} \cup \{x_i, x_{i+1}\}_{i=1,\ldots,n-1})$ with $n \geq 3$ and $K = \{x_0, x_i\}$ for some $i \leq \lfloor n/2 \rfloor$, a straightforward discussion leads to the conclusion that $R^d_{\{x_0,x_i\},C_n} = 0$ when $d < i$, $R^d_{\{x_0,x_i\},C_n} = \prod_{j=0}^{i-1} p(x_j x_{j+1})$ when $i \leq d < n - i$ and $R^d_{\{x_0,x_i\},C_n} = 1 - (1 - \prod_{j=0}^{i-1} p(x_j x_{j+1}))(1 - p(x_n, x_0) \prod_{j=i}^{n-1} p(x_j x_{j+1}))$ when $d \geq n - i$, being the latter expression identical to the classical two-terminal reliability for the cycle.

We remark that the the general DCR computation of bipartite graphs is, in general, $\mathcal{NP}$-Hard. The hardness is a corollary of Theorem 4.3. An open problem is to find the DCR in complete graphs.

5.2. DCR in ladders and Spanish fans

Now we illustrate the exact computation of the DCR for particular graph topologies. We start by a topology that is simple yet not trivial, namely the "ladder graphs". This allow us to see an example of how the existence of diameter constraints adds complexity to the problem when compared to the CLR. We will present an algorithm to find the two-terminal CLR for ladders with any length l. Then, we extend it to compute the two-terminal DCR. Finally, we develop a similar discussion for another family that we call "Spanish fan".

5.2.1. Ladders

Let us work with a family of networks L_l whose topologies, shown in Fig. 3, are defined by a natural parameter $l > 1$, called the number of "steps" of the ladder.

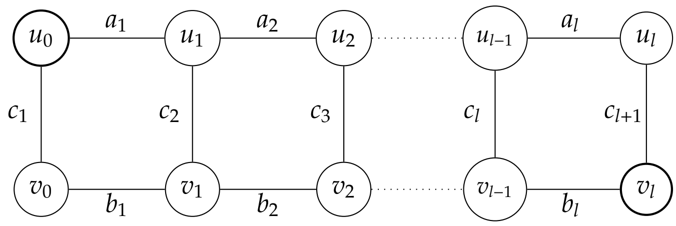

Figure 3: The ladder network

We set as terminal nodes $K = \{u_0, v_l\}$ and we want to find the connectedness probability for nodes u_0 and v_l (with and without length constraints). We assume nodes are perfect but links fail independently with known probabilities. We denote $\alpha_1, \ldots, \alpha_l$ the reliability of links $a_1, \ldots, a_l$; $\beta_1, \ldots, \beta_l$ the reliability of links $b_1, \ldots, b_l$; and $\gamma_1, \ldots, \gamma_{l+1}$ the ones from links $c_1, \ldots, c_{l+1}$. Finally, $\bar{\alpha}_i, \bar{\beta}_i, \bar{\gamma}_i$ denote respectively the unreliabilities $1 - \alpha_i, 1 - \beta_i, 1 - \gamma_i$.

5.2.2. CLR in ladders

Now we develop an algorithm that returns $R_{\{u_0,v_l\},L_l}$. First, we introduce the following definitions.

- We denote as $u \underset{i}{\rightsquigarrow} v$ the event where there is a path that connects nodes u and v using only links whose labels have subindices between 1 and i;

- $\mathcal{A}_i = (u_0 \underset{i}{\rightsquigarrow} u_i) \wedge \neg(u_0 \underset{i}{\rightsquigarrow} v_i)$;
- $\mathcal{B}_i = (u_0 \underset{i}{\rightsquigarrow} v_i) \wedge \neg(u_0 \underset{i}{\rightsquigarrow} u_i)$;
- $C_i = (u_0 \underset{i}{\rightsquigarrow} u_i) \wedge (u_0 \underset{i}{\rightsquigarrow} v_i)$.

Observe that for each $i = 1 \dots l$, the three events $\mathcal{A}_i$, $\mathcal{B}_i$ and C_i are pairwise disjoint. For example, the network configuration depicted in Fig. 4 belongs to the events C_1, $\mathcal{B}_2$ and C_3 but not to the events $\mathcal{A}_1$, $\mathcal{B}_1$, $\mathcal{A}_2$, C_2, $\mathcal{A}_3$ nor $\mathcal{B}_3$.

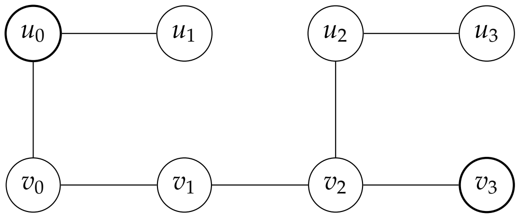

Figure 4: Example of a configuration of the ladder L_3

Using this notation we can write $R_{\{u_0,v_l\},L_l}$ as

$$R_{\{u_0,v_l\},L_l} = \Pr(\mathcal{A}_l \wedge c_{l+1}) + \Pr(\mathcal{B}_l) + \Pr(C_l) = \gamma_{l+1}\Pr(\mathcal{A}_l) + \Pr(\mathcal{B}_l) + \Pr(C_l) \quad (2)$$

Next, for ease of notation, we denote as x and $\bar{x}$ the events in which a link labeled x is operating or failing respectively, and as xz the intersection of two such events x and z. Observe that, for $i = 1, \dots, l$, it holds that

$$\begin{aligned} \mathcal{A}_i = \mathcal{A}_{i-1} &\wedge \left(a_i\bar{b}_i\bar{c}_i \vee a_ib_i\bar{c}_i \vee a_i\bar{b}_ic_i\right) \vee \\ \mathcal{B}_{i-1} &\wedge a_i\bar{b}_ic_i \vee \\ C_{i-1} &\wedge \left(a_i\bar{b}_i\bar{c}_i \vee a_i\bar{b}_ic_i\right) \end{aligned} \quad (3)$$

$$\begin{aligned} \mathcal{B}_i = \mathcal{A}_{i-1} &\wedge \bar{a}_ib_ic_i \vee \\ \mathcal{B}_{i-1} &\wedge (\bar{a}_ib_i\bar{c}_i \vee a_ib_i\bar{c}_i \vee \bar{a}_ib_ic_i) \vee \\ C_{i-1} &\wedge (\bar{a}_ib_i\bar{c}_i \vee \bar{a}_ib_ic_i) \end{aligned} \quad (4)$$

$$\begin{aligned} C_i = \mathcal{A}_{i-1} &\wedge a_ib_ic_i \vee \\ \mathcal{B}_{i-1} &\wedge a_ib_ic_i \vee \\ C_{i-1} &\wedge (a_ib_i\bar{c}_i \vee a_ib_ic_i)\,. \end{aligned} \quad (5)$$

The unions and intersections of events involving the links a_i, b_i and c_i were written in canonical form, to evidence each of the eight possible conjoint states,

Procedure CLR-Ladder$(l, \vec{\alpha}, \vec{\beta}, \vec{\gamma})$

$a \leftarrow 1; b \leftarrow 0; c \leftarrow 0$
for all $i \in 1, \dots, l$ **do**
  $a' \leftarrow a\,(\alpha_i\bar{\beta}_i\bar{\gamma}_i + \alpha_i\beta_i\bar{\gamma}_i + \alpha_i\bar{\beta}_i\gamma_i) + b\alpha_i\bar{\beta}_i\gamma_i + c\,(\alpha_i\bar{\beta}_i\bar{\gamma}_i + \alpha_i\bar{\beta}_i\gamma_i)$
  $b' \leftarrow a\,(\bar{\alpha}_i\beta_i\gamma_i) + b\,(\bar{\alpha}_i\beta_i\bar{\gamma}_i + \alpha_i\beta_i\bar{\gamma}_i + \bar{\alpha}_i\beta_i\gamma_i) + c\,(\bar{\alpha}_i\beta_i\bar{\gamma}_i + \bar{\alpha}_i\beta_i\gamma_i)$
  $c' \leftarrow a\,(\alpha_i\bar{\beta}_i\bar{\gamma}_i + \alpha_i\beta_i\bar{\gamma}_i + \alpha_i\bar{\beta}_i\gamma_i) + b\,(\alpha_i\bar{\beta}_i\gamma_i) + c\,(\alpha_i\bar{\beta}_i\bar{\gamma}_i + \alpha_i\bar{\beta}_i\gamma_i)$
  $a \leftarrow a'; b \leftarrow b'; c \leftarrow c'$
end for
return $b + c + a\gamma_{l+1}$

Figure 5: Algorithm for computing the s,t-CLR of a ladder graph

yielded by the individual states of the links, and how they contribute to $\mathcal{A}_{i+1}$, $\mathcal{B}_{i+1}$ and C_{i+1}. Moreover, all unions involve disjoint events, and all intersections involve independent events. Therefore, the following identities for the probabilities $\Pr(\mathcal{A}_i)$, $\Pr(\mathcal{B}_i)$ and $\Pr(C_i)$ hold:

$$\begin{aligned}\Pr(\mathcal{A}_i) = \Pr(\mathcal{A}_{i-1})\,&(\alpha_i\bar{\beta}_i\bar{\gamma}_i + \alpha_i\beta_i\bar{\gamma}_i + \alpha_i\bar{\beta}_i\gamma_i) + \\ &\Pr(\mathcal{B}_{i-1})\alpha_i\bar{\beta}_i\gamma_i + \\ &\Pr(C_{i-1})\,(\alpha_i\bar{\beta}_i\bar{\gamma}_i + \alpha_i\bar{\beta}_i\gamma_i)\end{aligned} \tag{6}$$

$$\begin{aligned}\Pr(\mathcal{B}_i) = \Pr(\mathcal{A}_{i-1})\,&(\bar{\alpha}_i\beta_i\gamma_i) + \\ &\Pr(\mathcal{B}_{i-1})\,(\bar{\alpha}_i\beta_i\bar{\gamma}_i + \alpha_i\beta_i\bar{\gamma}_i + \bar{\alpha}_i\beta_i\gamma_i) + \\ &\Pr(C_{i-1})\,(\bar{\alpha}_i\beta_i\bar{\gamma}_i + \bar{\alpha}_i\beta_i\gamma_i)\end{aligned} \tag{7}$$

$$\begin{aligned}\Pr(C_i) = \Pr(\mathcal{A}_{i-1})\,&(\alpha_i\beta_i\gamma_i) + \\ &\Pr(\mathcal{B}_{i-1})\,(\alpha_i\beta_i\gamma_i) + \\ &\Pr(C_{i-1})\,(\alpha_i\beta_i\bar{\gamma}_i + \alpha_i\beta_i\gamma_i)\,.\end{aligned} \tag{8}$$

Note that the probabilities with subindices i can be computed just employing probabilities that involve the subindex $i-1$. The base cases are given by $\Pr(\mathcal{A}_0) = 1$, $\Pr(\mathcal{B}_0) = 0$ and $\Pr(C_0) = 0$. So, we can write an algorithm that is linear in l (shown in Fig. 5), to compute $R_{u_0,v_l}(L_l)$. The expression in the `return` statement corresponds to Eq. (2).

5.2.3. DCR in ladders

Now we show how the previous method can be extended to find $R^d_{\{u_0,v_l\},L_l}$. A technical lemma will be useful.

Lemma 5.1. *Suppose that* $u_0 \underset{i}{\rightsquigarrow} u_i$ *and* $u_0 \underset{i}{\rightsquigarrow} v_i$. *The shortest paths that connect* u_0 *to* u_i *and* u_0 *to* v_i $(i = 1, \dots, l)$, *using only links whose labels have subindices not above* i, *have lengths that differ exactly in one.*

Proof. Let p_1 and p_2 be examples of the shortest paths that connect u_0 to u_i and v_i, respectively (using only links with subindex not above i). Both of them will have the general structure of Fig. 6 where, walking from u_0 to u_i and v_i, they meet a certain number of times, before they reach u_i and v_i. Let w be the last common node. The part p_1' of p_1 going from u_0 to w has the same length as the part p_2' of p_2 going from u_0 to w; otherwise, one of p_1 or p_2 could be made shorter by replacing p_1' by p_2' or vice versa. Suppose that $w = u_j$ with $j \in \{0, i\}$. Then the parts of p_1 and p_2 going from w to u_i and v_i must be as shown in Fig. 6, where clearly the difference in length is one; otherwise w would not be the last common node. The same applies in case that $w = v_j$ with $j \in \{0, i\}$. □

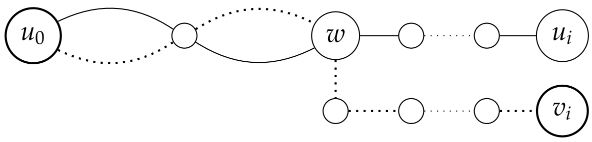

Figure 6: Lemma 5.1

In the case of the DCR, for each $i = 1, \dots, l$, we have to deal with more events than in the case of the CLR ($\mathcal{A}_i$, $\mathcal{B}_i$ and $\mathcal{C}_i$). Now the events are defined by the pairs of distances from u_0 to u_i and v_i. Let us denote by $\mathcal{E}_i^{m,n}$ the event where $dist(u_0, u_i) = m$ and $dist(u_0, v_i) = n$ when considering only the operational links with subindices not above i. In light of Lemma 5.1 and the specific topology of the ladder, no event other than $\mathcal{E}_i^{\infty,\infty}$, $\mathcal{E}_i^{m,\infty}$, $\mathcal{E}_i^{m+1,m}$, $\mathcal{E}_i^{m,m+1}$, $\mathcal{E}_i^{\infty,m}$ can be feasible, being $i = 0, \dots, l$ and $m = i, \dots, 2i+1$. As we did for the CLR, we can express the events for i by using only events with subindices $i-1$:

$$\mathcal{E}_i^{m,\infty} = \mathcal{E}_{i-1}^{m-1,\infty} a_i \overline{b_i c_i} \vee \mathcal{E}_{i-1}^{m-1,m} a_i \bar{b}_i \vee \mathcal{E}_{i-1}^{\infty,m-2} a_i \bar{b}_i c_i \vee \mathcal{E}_{i-1}^{m-1,m-2} a_i \bar{b}_i \tag{9}$$

$$\mathcal{E}_i^{m+1,m} = \mathcal{E}_{i-1}^{\infty,m-1} a_i b_i c_i \vee \mathcal{E}_{i-1}^{m,m-1} a_i b_i \tag{10}$$

$$\mathcal{E}_i^{m,m+1} = \mathcal{E}_{i-1}^{m-1,\infty} a_i b_i c_i \vee \mathcal{E}_{i-1}^{m-1,m} a_i b_i \tag{11}$$

$$\mathcal{E}_i^{\infty,m} = \mathcal{E}_{i-1}^{\infty,m-1} b_i \overline{a_i c_i} \vee \mathcal{E}_{i-1}^{m,m-1} \bar{a}_i b_i \vee \mathcal{E}_{i-1}^{m-2,\infty} \bar{a}_i b_i c_i \vee \mathcal{E}_{i-1}^{m-2,m-1} \bar{a}_i b_i. \tag{12}$$

Similar argument on disjointness of unions and independence of intersections apply to the events involved in these equations. The only base case ($i = 0$) with non-zero probability is $\Pr(\mathcal{E}_0^{0,\infty}) = 1$. Now we can build an algorithm (shown in Fig. 7) to compute $R^d_{\{u_0,v_l\},L_l}$. It starts by considering the trivial cases where the threshold distance d between u_0 and v_l is too low to be reachable (the lowest possible one is $l+1$). Similarly, if d is above $2l-1$, then the operating configurations are exactly the same as for the CLR, since given the topology of the network, no configuration exists where the distance between u_0 and v_l is above $2l-1$. Next, the algorithm defines a square array, initialized with zeros, large enough to store the probability of all the events $\mathcal{E}$. For ease of notation, we assume that any reference involving out-of-range subindices returns zero. The constant t is set to $2l+3$ to represent the array index for ∞. The probability of the base case $\mathcal{E}_0^{0,\infty}$ is set to one. Then the algorithm proceeds by sequentially processing the links, one "ladder step" at a time. For each step, only those values of m that can accumulate probability are considered; they range from i to $2i$. Once the array is computed, the algorithm builds a vector $\vec{p}$ that stores the probability that the distance between u_0 and v_l is an integer m, for every index m ranging from $l+1$ to $2l+1$. Finally, the algorithm returns the cumulative probability that the distance is any value between $l+1$ and the parameter d.

The execution time is dominated by the nested iteration, thus the algorithm has a complexity in time

$$O(\sum_{i=1}^{l}\sum_{m=i}^{2i} 1) = O(\sum_{i=1} l(i+1) = O((l+1)l/2+l) = O(l^2)$$

which is quadratic with respect to l. We implemented the algorithm, detailed in Figure 7, in C++ and tested on an Intel Core2 Duo T5450 CPU machine with 2 GB RAM. Table 1 shows the values of $R^d_{\{u_0,v_l\},L_l}$ obtained for $1 \leq l \leq 20$ and $2 \leq d \leq 41$, with identical link reliabilities of $p = 3/10$. Table 2 shows the values for $R^d_{\{u_0,v_l\},L_l}$ when $1 \leq l \leq 40$ and $2 \leq d \leq 81$, with identical reliabilities of $p = 9/10$. Figure 9 shows (left) the elapsed time for computing L_1 to L_{150} and (right) the square root of these times, making clear their quadratic evolution. Finally, four tests were run for the network L_{20}, and their results are shown in Figure 10. The four charts show the probabilities that the random variable "distance between u_0 and v_{20}" has the values that range from 21 to 41. (These are the possible finite values it can have for L_{20}). The difference between 1 and the cumulative probability of each chart is the probability that both nodes are disconnected. This kind of probability distribution functions are used by D. Migov to compute the two-terminal DCR of networks with junction points by means of the convolution operator [7].

5.2.4. Spanish fans

Let us see another example based on a topology that we call "Spanish fan". This family of networks F_l is shown is Fig. 8, parameterized by $l \in \mathcal{Z} : l > 1$. It can be seen as a ladder in which one of the sides was collapsed into a single node (v). Again, we set as terminal two nodes, $K = \{u_0, u_l\}$ and so we want to compute the probability that the nodes u_0 and u_l are connected by a path with length not above a certain integer d.

It is easy to see that, in any configuration, when only considering links with subindices up to i, the following statements hold for all $i = 1, \ldots, l$:

Procedure DCR-Ladder$(l, d, \vec{\alpha}, \vec{\beta}, \vec{\gamma})$

if $d \leq l$ **then**
    return 0
else if $d \geq 2l - 1$ **then**
    return CLR-Ladder$(l, \vec{\alpha}, \vec{\beta}, \vec{\gamma})$
end if
$t \leftarrow 2l + 3$
$e(\cdot, \cdot) \leftarrow$ array $(t + 1, t + 1)$ initialized with 0's
$e(0, t) \leftarrow 1$
$\vec{x_1}, \vec{x_2}, \vec{x_3}, \vec{x_4} \leftarrow$ vectors with subindices $1, \dots, 2l + 1$
for all $i \in 1, \dots, l$ **do**
    for all $m \in i, \dots, 2i + 1$ **do**
        $x_1(m) \leftarrow e(m-1, t)\alpha_i\overline{\beta_i\gamma_i} + e(m-1, m)\alpha_i\bar{\beta}_i + e(t, m-2)\alpha_i\bar{\beta}_i\gamma_i + e(m-1, m-2)\alpha_i\bar{\beta}_i$
        $x_2(m) \leftarrow e(t, m-1)\beta_i\overline{\alpha_i\gamma_i} + e(m, m-1)\bar{\alpha}_i\beta_i + e(m-2, t)\bar{\alpha}_i\beta_i\gamma_i + e(m-2, m-1)\bar{\alpha}_i\beta_i$
        $x_3(m) \leftarrow e(t, m - 1)\alpha_i\beta_i\gamma_i + e(m, m - 1)\alpha_i\beta_i$
        $x_4(m) \leftarrow e(m - 1, t)\alpha_i\beta_i\gamma_i + e(m - 1, m)\alpha_i\beta_i$
    end for
    for all $m \in i, \dots, 2i$ **do**
        $e(m, t) \leftarrow x_1(m)$
        $e(t, m) \leftarrow x_2(m)$
        $e(m + 1, m) \leftarrow x_3(m)$
        $e(m, m + 1) \leftarrow x_4(m)$
    end for
    $e(i - 1, t) \leftarrow 0; e(t, i - 1) \leftarrow 0; e(i, i - 1) \leftarrow 0; e(i - 1, i) \leftarrow 0$
end for
$\vec{p} \leftarrow$ vector with subindices $l + 1, \dots, 2l + 1$
for all $m \in l + 1, \dots, 2l + 1$ **do**
    $p(m) \leftarrow \left(\sum_{j \in \{m-1, m+1, t\}} e(j, m)\right)$
    $p(m) \leftarrow p(m) + (e(m - 1, m - 2) + e(m - 1, t))\,\gamma_{l+1}$
end for
return $\sum_{m=l+1,\dots,d} p(m)$

Figure 7: Algorithm for computing the s, t-DCR of a ladder graph

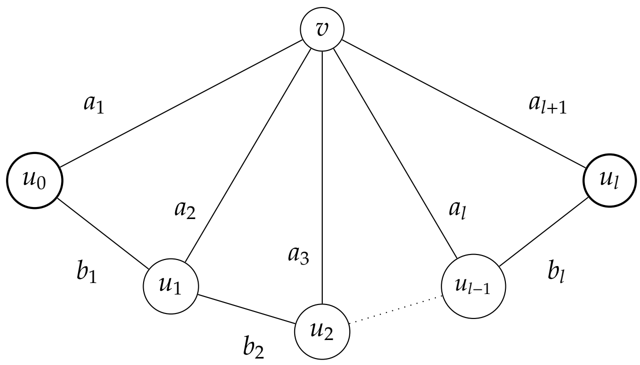

Figure 8: Spanish fan network

- the distance between u_0 and u_i can be any value in the range of $\min(3, i), \ldots, i+1$ or infinite;
- the distance between u_0 and v can be any value in the range of $1, \ldots, i$ or infinite.

Now, analogously to Section 5.2.2, we can define $\mathcal{E}_i^{m,n}$ as the event where $dist(u_0, v) = m$ and $dist(u_0, u_i) = n$ when considering only the operational links with subindices not above i. For a given value of $i \in 1, \ldots, l$, the only elementary events that can occur with probabilities other than zero, and at the end contribute to $R_{u_0,u_l}(F_l, d)$, are:

- $\mathcal{E}_i^{\infty,i}$
- $\mathcal{E}_i^{m,\infty} : m \in 1, \ldots, i$
- $\mathcal{E}_i^{m,n} : n \in 1, \ldots, i \wedge m \in 1, \ldots, n$
- $\mathcal{E}_i^{m,i+1} : m \in 1, \ldots, i.$

Therefore, we can write $R_{u_0,u_l}(F_l, d)$ as a sum of probabilities due to event disjointness,

$$R_{u_0,u_l}(F_l, d) = \sum_{i=1}^{d} \Pr(\mathcal{E}_l^{\infty,i}) + \alpha_{l+1} \sum_{i=1}^{d-1} \Pr(\mathcal{E}_l^{i,\infty}) + \sum_{j=1}^{d} \sum_{i=1}^{j} \Pr(\mathcal{E}_l^{i,j}) \tag{13}$$

where α_{l+1} is the probability that edge a_{l+1} operates.

As i grows, the number of such feasible elementary events has order $O(i^2)$. In virtue of Lemma 5.1, these events outnumber those of the ladder network. Again, we can express the events for a given value of i just using events defined for $i-1$, as follows:

$$\mathcal{E}_i^{\infty,i} = \mathcal{E}_{i-1}^{\infty,i-1}\bar{a}_i b_i$$

$$\mathcal{E}_i^{m,\infty} = \left(\cup_{j=m}^{\infty}\mathcal{E}_{i-1}^{m,j}\bar{b}_i\right) \vee \mathcal{E}_{i-1}^{m,\infty}\bar{a}_i b_i \vee \mathcal{E}_{i-1}^{\infty,m-1}a_i\bar{b}_i \quad (m \in 1,\dots,i)$$

$$\mathcal{E}_i^{m,n} = \mathcal{E}_{i-1}^{m,n-1}b_i \vee \mathcal{E}_{i-1}^{\infty,n-1}(m = n)a_i b_i \quad (n \in 1,\dots,i; m \in 1,\dots,n)$$

$$\mathcal{E}_i^{m,i+1} = \mathcal{E}_{i-1}^{m,i}b_i \quad (m \in 1,\dots,i)$$

Again, the probabilities with subindices i can be computed just employing probabilities that involve the subindex $i-1$. This time the base case is given by $\mathcal{E}_0^{\infty,0} = 1$. Similarly as we did for the ladder, we can build an algorithm that computes the probability of the $O(i^2)$ relevant events with subindices i, using the corresponding probabilities for $i-1$. The result is an algorithm with a number of operations that is $O(i^3)$, therefore with time complexity $O(i^3)$; once again, polynomial in i.

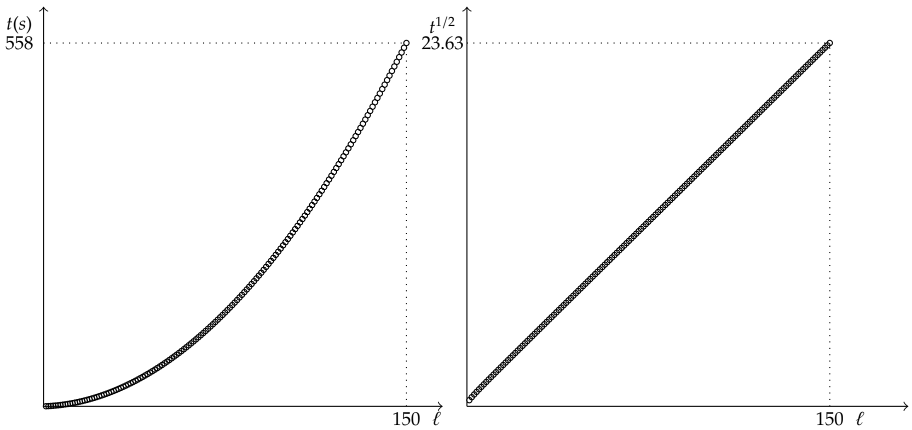

Figure 9: Elapsed time, and its square root, for computing L_1 to L_{150} using `DCR-Ladder`

6. CONCLUDING REMARKS

The reliability evaluation of a particular stochastic binary system has been discussed, namely the diameter-constrained network reliability (DCR). When the number of terminals k or the diameter parameter d are free inputs, the DCR is $\mathcal{NP}$-Hard, since it subsumes the classical network reliability problem. The case where

	l							
d	1	2	3	5	8	11	15	20
2	1.72E-01							
3	1.72E-01	7.56E-02						
4	1.72E-01	7.56E-02	2.99E-02					
5	1.72E-01	7.68E-02	2.99E-02					
6	1.72E-01	7.68E-02	3.14E-02	3.98E-03				
7	1.72E-01	7.68E-02	3.14E-02	3.98E-03				
8	1.72E-01	7.68E-02	3.14E-02	4.74E-03				
9	1.72E-01	7.68E-02	3.14E-02	4.74E-03	1.60E-04			
10	1.72E-01	7.68E-02	3.14E-02	4.75E-03	1.60E-04			
11	1.72E-01	7.68E-02	3.14E-02	4.75E-03	2.51E-04			
12	1.72E-01	7.68E-02	3.14E-02	4.75E-03	2.51E-04	5.73E-06		
13	1.72E-01	7.68E-02	3.14E-02	4.75E-03	2.58E-04	5.73E-06		
14	1.72E-01	7.68E-02	3.14E-02	4.75E-03	2.58E-04	1.24E-05		
15	1.72E-01	7.68E-02	3.14E-02	4.75E-03	2.58E-04	1.24E-05		
16	1.72E-01	7.68E-02	3.14E-02	4.75E-03	2.58E-04	1.37E-05	6.16E-08	
17	1.72E-01	7.68E-02	3.14E-02	4.75E-03	2.58E-04	1.37E-05	6.16E-08	
18	1.72E-01	7.68E-02	3.14E-02	4.75E-03	2.58E-04	1.37E-05	2.02E-07	
19	1.72E-01	7.68E-02	3.14E-02	4.75E-03	2.58E-04	1.37E-05	2.02E-07	
20	1.72E-01	7.68E-02	3.14E-02	4.75E-03	2.58E-04	1.37E-05	2.65E-07	
25	1.72E-01	7.68E-02	3.14E-02	4.75E-03	2.58E-04	1.37E-05	2.74E-07	1.78E-09
30	1.72E-01	7.68E-02	3.14E-02	4.75E-03	2.58E-04	1.37E-05	2.74E-07	2.05E-09
41	1.72E-01	7.68E-02	3.14E-02	4.75E-03	2.58E-04	1.37E-05	2.74E-07	2.05E-09

Table 1: $R^d_{\{u_0,v_l\},L_l}$ for $1 \leq l \leq 20$, $2 \leq d \leq 41$ and identical reliabilities $p = 3/10$.

	l							
d	1	3	6	10	15	20	30	40
2	9.64E-01							
3	9.64E-01							
4	9.64E-01	9.30E-01						
5	9.64E-01	9.30E-01						
6	9.64E-01	9.48E-01						
7	9.64E-01	9.48E-01	8.36E-01					
8	9.64E-01	9.48E-01	8.36E-01					
9	9.64E-01	9.48E-01	9.14E-01					
10	9.64E-01	9.48E-01	9.14E-01					
11	9.64E-01	9.48E-01	9.15E-01	6.86E-01				
12	9.64E-01	9.48E-01	9.15E-01	6.86E-01				
13	9.64E-01	9.48E-01	9.15E-01	8.66E-01				
14	9.64E-01	9.48E-01	9.15E-01	8.66E-01				
15	9.64E-01	9.48E-01	9.15E-01	8.73E-01				
16	9.64E-01	9.48E-01	9.15E-01	8.73E-01	5.07E-01			
17	9.64E-01	9.48E-01	9.15E-01	8.73E-01	5.07E-01			
18	9.64E-01	9.48E-01	9.15E-01	8.73E-01	7.93E-01			
19	9.64E-01	9.48E-01	9.15E-01	8.73E-01	7.93E-01			
20	9.64E-01	9.48E-01	9.15E-01	8.73E-01	8.22E-01			
22	9.64E-01	9.48E-01	9.15E-01	8.73E-01	8.23E-01	3.60E-01		
24	9.64E-01	9.48E-01	9.15E-01	8.73E-01	8.23E-01	7.05E-01		
26	9.64E-01	9.48E-01	9.15E-01	8.73E-01	8.23E-01	7.71E-01		
28	9.64E-01	9.48E-01	9.15E-01	8.73E-01	8.23E-01	7.75E-01		
30	9.64E-01	9.48E-01	9.15E-01	8.73E-01	8.23E-01	7.75E-01		
32	9.64E-01	9.48E-01	9.15E-01	8.73E-01	8.23E-01	7.75E-01	1.67E-01	
34	9.64E-01	9.48E-01	9.15E-01	8.73E-01	8.23E-01	7.75E-01	5.04E-01	
36	9.64E-01	9.48E-01	9.15E-01	8.73E-01	8.23E-01	7.75E-01	6.59E-01	
40	9.64E-01	9.48E-01	9.15E-01	8.73E-01	8.23E-01	7.75E-01	6.89E-01	
50	9.64E-01	9.48E-01	9.15E-01	8.73E-01	8.23E-01	7.75E-01	6.89E-01	6.11E-01
81	9.64E-01	9.48E-01	9.15E-01	8.73E-01	8.23E-01	7.75E-01	6.89E-01	6.12E-01

Table 2: $R^d_{\{u_0,v_l\},L_l}$ for $1 \leq l \leq 40$, $2 \leq d \leq 81$ and identical reliabilities $p = 9/10$.

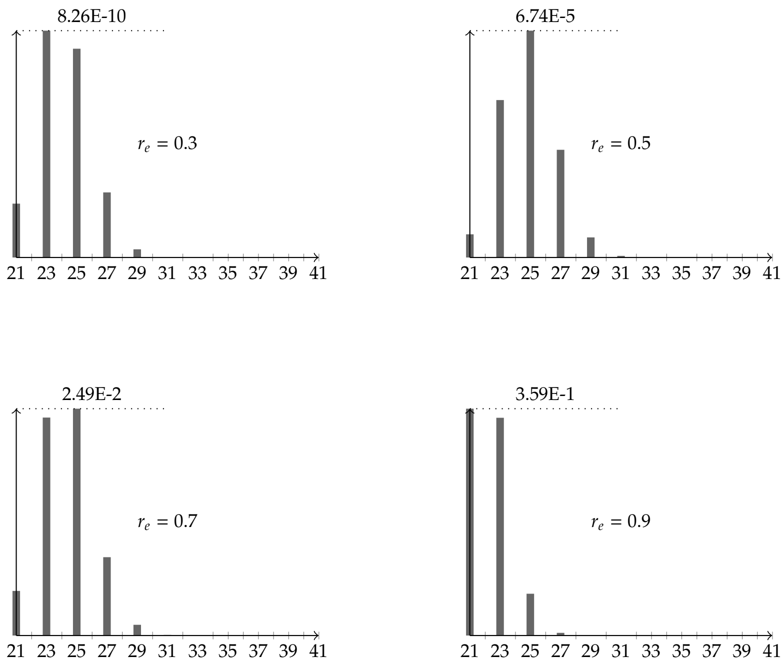

Figure 10: Probability distributions for the distance between u_0 and v_{20} in L_{20}

$d = 1$ and the case where $d = 2$ and k is fixed belong to the set $\mathcal{P}$ of problems solvable in time polynomial in the number of nodes n. Several subproblems, given by particular network topologies or constraints in n or k are polynomially computable. We computed some elementary cases and gave linear algorithms for two specific topologies. The DCR turns to be $\mathcal{NP}$-Hard when $k \geq 2$ is fixed and $d \geq 3$. The complexity of the case where $k = n$ and $d \geq 3$ was not determined in prior literature. In this article, we proved that it belongs to the $\mathcal{NP}$-Hard complexity class. As a corollary, the result also holds when $d \geq 3$ and k is a free parameter for the complexity analysis.

The computational complexity of the DCR remains unknown when $d = 2$ and $k = n$. It is worth to notice that when all links fail independently with identical probability $p = 1/2$, all graphs occur with the same probability. Therefore, the DCR when $d = 2$ and $k = n$ is at least as hard as counting all subgraphs with diameter up to 2 for a given graph of order n.

Acknowledgements: We wish to express our most sincere gratitude to the reviewers for their valuable comments and corrections, as well as to Dr. Eduardo Canale for his useful suggestions.

REFERENCES

[1] Ball, M. O., Computational complexity of network reliability analysis: An overview, *IEEE Transactions on Reliability*, 35(3) (1986) 230-239.

[2] Ball, M.O., and Provan, J.S., "The complexity of counting cuts and of computing the probability that a graph is connected", *SIAM J. Computing*, 12 (1983) 777-788.

[3] Canale, E., and Cancela, H., and Robledo, F., and Rubino, G. and Sartor, P., "On computing the 2-diameter-constrained K-reliability of networks", *International Transactions in Operational Research* 20(1) (2013) 49-58.

[4] Cancela, H., and Petingi, L., "Reliability of communication networks with delay constraints: computational complexity and complete topologies", *International Journal of Mathematics and Mathematical Sciences*, 29 (2004) 1551-1562.

[5] Colbourn, C.J., "Reliability issues in telecommunications network planning". Chapter 9 in Sansó, B. and Soriano, P. (eds.) *Telecommunications network planning*, Kluwer Academic Publishers, London/Dordrecht/Boston (1999) 135-146.

[6] Garey, M.R., and Johnson, D.S., *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman and Company, New York, NY, USA (1979).

[7] Migov, D., "Computing diameter constrained reliability of a network with junction points", *Automation and Remote Control*, 72 (2011) 1415-1419.

[8] Petingi, L., and Rodriguez, J., "Reliability of networks with delay constraints", *Congressus Numerantium*, 152 (2001) 117-123.

[9] Provan, S.J., and Ball, M.O., "The Complexity of Counting Cuts and of Computing the Probability that a Graph is Connected", *SIAM Journal on Computing* 12(4) (2001), 777-788.

[10] Rosenthal, A., "Computing the reliability of complex networks", *SIAM Journal on Applied Mathematics*, 32(2) (1977) 384-393.

[11] Sartor, P., "Propriétés et mèthodes de calcul de la fiabilité diamètre-bornée des réseaux", *PhD thesis, INRIA/IRISA, Université de Rennes I, Rennes, France*, (2013).

[12] Valiant, L., "The complexity of enumeration and reliability problems", *SIAM Journal on Computing* 8(3) (2009) 410-421.