Diameter-Constrained Reliability: Complexity, Factorization and Exact computation in Weak Graphs

Eduardo Canale
Instituto de Matemática y Estadística (IMERL)
Julio Herrera y Reissig 565
Montevideo, Uruguay
canale@fing.edu.uy

Juan Piccini
Instituto de Matemática y Estadística (IMERL)
Julio Herrera y Reissig 565
Montevideo, Uruguay
piccini@fing.edu.uy

Franco Robledo
Instituto de Matemática y Estadística (IMERL)
Julio Herrera y Reissig 565
Montevideo, Uruguay
frobledo@fing.edu.uy

Pablo Romero
Instituto de Matemática y Estadística (IMERL)
Julio Herrera y Reissig 565
Montevideo, Uruguay
promero@fing.edu.uy

ABSTRACT

In this paper we address a problem from the field of network reliability, called diameter-constrained reliability. Specifically, we are given a simple graph $G = (V, E)$ with $|V| = n$ nodes and $|E| = m$ links, a subset $K \subseteq V$ of *terminals*, a vector $p = (p_1, \ldots, p_m) \in [0, 1]^m$ and a positive integer d, called *diameter*. We assume nodes are perfect but links fail stochastically and independently, with probabilities $q_i = 1 - p_i$. The *diameter-constrained reliability* (DCR for short), 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 general DCR computation is inside the class of $\mathcal{NP}$-Hard problems, since is subsumes the complexity that a random graph is connected. In this paper the computational complexity of DCR-subproblems is discussed in terms of the number of terminal nodes $k = |K|$ and diameter d. A factorization formula for exact DCR computation is provided, that runs in exponential time in the worst case. Finally, a revision of graph-classes that accept DCR computation in polynomial time is then included. In this class we have graphs with bounded co-rank, graphs with bounded genus, planar graphs, and, in particular, Monma graphs, which are relevant in robust network design. We extend this class adding arborescence graphs. A discussion of trends for future work is offered in the conclusions.

Categories and Subject Descriptors

C.4 [**Performance of Systems**]: Reliability, availability, and serviceability; D.2.8 [**Mathematical Software**]: Reliability and Robustness—*computational complexity, performance measures*

General Terms

Network Reliability

Keywords

Computational Complexity, Network Reliability, Diameter-Constrained Reliability

1. INTRODUCTION

The definition of DCR has been introduced in 2001 by Héctor Cancela and Louis Petingi, inspired in delay-sensitive applications over the Internet infrastructure [10]. Nevertheless, its applications over other fields of knowledge enriches the motivation of this problem in the research community [12].

We wish to communicate special nodes in a network, called *terminals*, by d hops or less, in a scenario where nodes are perfect but links fail stochastically and independently. The all-terminal case with $d = n - 1$ is precisely the probability that a random graph is connected, or *classical reliability problem* (CLR for short). Arnon Rosenthal proved that the CLR is inside the class of $\mathcal{NP}$-Hard problems [24]. As a corollary, the general DCR is $\mathcal{NP}$-Hard as well, hence intractable unless $\mathcal{P} = \mathcal{NP}$.

The focus of this paper is the computational complexity of DCR subproblems in terms of the number of terminals k and diameter d, and the efficient computation of the DCR for distinguished graph topologies.

In Section 4, a formal definition of DCR is provided as a particular instance of a coherent stochastic binary system. The computational complexity of the DCR is discussed in terms of the diameter and number of terminals in Section 3. The contributions of this paper are two-fold. First, we close the complexity analysis of the DCR problem in terms of k and d. Indeed, we prove in this section that the DCR is in the computational class of $\mathcal{NP}$-Hard problems in the all-terminal scenario ($k = n$) with a given diameter $d \geq 2$. The computational complexity for other possible pairs for k and d is already available from prior literature from this area.

Then, we provide an exact DCR computation by means of a factorization technique inspired in [20], in Section 4. Finally, we extend the class of known graphs that permit an efficient (i.e. polynomial time) computation for the DCR in Section 5. A particular but relevant family of these graphs are Monma graphs, which plays a key role in the design of robust network design [19, 23, 8]. 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 the binary state is captured by a binary word $x = (x_1, \ldots, x_m) \in \{0,1\}^m$. Additionally, we have a structure function $\phi : \{0,1\}^m \to \{0,1\}$ such that $\phi(x) = 1$ if the system works under state x, and $\phi(x) = 0$ otherwise. 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 from [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. A subgraph $G_x = (V, E_x) \subseteq G$ 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. Let us choose an arbitrary order of the edge-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. Therefore, we define the structure $\phi : \{0,1\}^m \to \{0,1\}$ such that $\phi(x) = 1$ if G_x is d-K-connected, and $\phi(x) = 0$ otherwise. 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 a 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). We will denote $\mathcal{O}^K_d(G)$ to the set of all d-K-connected subgraphs of a ground graph G.

We recall a bit of terminology coming from graph theory, which will be used throughout this treatment. A graph $G = (V, E)$ is *bipartite* if there exists a bipartition $V = V_1 \cup V_2$ such that $E \subseteq \{\{x,y\} : x \in V_1, y \in V_2\}$. A *vertex cover* in a graph $G = (V, E)$ is a subset $V' \subseteq V$ such that V' meets all links in E.

Given two graphs G_1 and G_2 with the same vertex set V, $f : V \to V$ is a *K-isomorphism* from G_1 to G_2 if it is an isomorphism that fixes the set K. In that case G_1 and G_2 are *K-isomorphic*. Given a simple graph $G = (V, E)$ and $e = \{x,y\} \in E$, an *elementary division* of e is a couple of edges $e_1 = \{x,z\}$ and $e_2 = \{z,y\}$ that replace e in G, where $z \notin V$. Two graphs G_1 and G_2 are *homeomorphic* if there exists a graph G such that G_1 and G_2 can be obtained from G by means of a sequence of elementary divisions. If $P = \{V_1, \ldots, V_c\}$ is a partition of V, the *quotient graph* is $G' = (P, E')$, where $\{V_i, V_j\} \in E'$ if and only if $i \neq j$ and there exists an edge from a vertex of V_i to a vertex of V_j in E. We say v_j is *reachable* from v_i either when $v_i = v_j$ or there is a path from v_i to v_j. In a simple graph G, reachability is an equivalence relation, and c, the number of classes in the quotient graph, is the number of *connected components*. Given a simple graph $G = (V, E)$ with $n = |V|$ vertices and $m = |E|$ edges, its *rank* is $r(G) = n - c$, while its *co-rank* is $c(G) = m - n + c$. A *connected graph* verifies $c = 1$; then $r(G) = n - 1$ and $c(G) = m - n + 1$. In topological graph theory, the *genus of a graph* G is the least natural g such that G can be drawn without crossing itself in a surface with genus g. A *planar graph* verifies $g = 0$.

3. COMPUTATIONAL COMPLEXITY

The class $\mathcal{NP}$ is the set of problems polynomially solvable by a non-deterministic Turing machine [15]. 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 $\mathcal{NP}$. Stephen Cook proved that the joint satisfiability of an input set of clauses in disjunctive form is an $\mathcal{NP}$-Complete decision problem; in fact, the first known problem of this class [13]. In this way, he provided a systematic procedure to prove that a certain problem is $\mathcal{NP}$-Complete. Specifically, it suffices to prove that the problem is inside the class $\mathcal{NP}$, and that it is at least as hard as an $\mathcal{NP}$-Complete problem. Richard Karp followed this hint, and presented the first 21 combinatorial problems inside this class [16]. Leslie Valiant defines the class #$\mathcal{P}$ of counting problems, such that testing whether an element should be counted or not can be accomplished in polynomial time [26]. 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.

Recognition and counting minimum cardinality mincuts/minpaths are at least as hard as computing the reliability of a coherent SBS [1]. Arnon Rosenthal proved the CLR is $\mathcal{NP}$-Hard [24], showing that the minimum cardinality mincut recognition is precisely Steiner-Tree problem, included in Richard Karp's list. The CLR for both two-terminal and all-terminal cases are still $\mathcal{NP}$-Hard, as Michael Ball and J. Scott Provan proved by reduction to counting minimum cardinality $s - t$ cuts [22]. 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 $d = 1$ all terminals must have a direct link,

$$R^1_{K,G} = \prod_{\{u,v\} \subseteq K} p(uv),$$

where $p(uv)$ denotes the probability of operation of link $\{u,v\} \in E$, and $p(uv) = 0$ if $\{u,v\} \notin E$. The problem is still simple when $k = d = 2$. In fact,

$$R^2_{\{u,v\},G} = 1 - (1 - p(uv)) \prod_{w \in V - \{u,v\}} (1 - p(uw)p(wv)).$$

Héctor Cancela and Louis Petingi rigorously proved that the DCR is $\mathcal{NP}$-Hard when $d \geq 3$ and $k \geq 2$ is a fixed input parameter [11], in strong contrast with the case $d = k = 2$.

The literature offers two proofs that the DCR has a polynomial-time algorithm when $d = 2$ and k is a fixed input parameter [25, 6]. Pablo Sartor et. al. present a recursive proof [25], while Eduardo Canale et. al. present an explicit expression for $R^2_{K,G}$ that is computed in a polynomial time of elementary operations [7].

Here, we will prove that the DCR is inside the class of $\mathcal{NP}$-Hard problems in the all-terminal case with diameter $d \geq 2$. The main source of inspiration for the first result is the article authored by [11], where they proved that the DCR is $\mathcal{NP}$-Hard when $d \geq 3$ and $k \geq 2$ is a fixed input parameter. There, the authors prove first 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:

Lemma 1. *Counting the number of vertex covers of a bipartite graph is #$\mathcal{P}$-Complete [2].*

Proposition 1. *The DCR is $\mathcal{NP}$-Hard when $k = 2$ and $d \geq 3$ [11].*

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 edge 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:

$$G' = (A \cup B \cup V(P) \cup \{t\}, E \cup E(P) \cup I\}, \qquad (1)$$

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

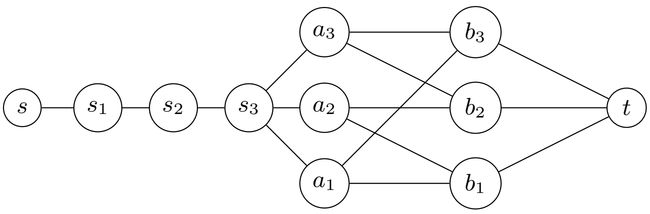

Figure 1: Example of auxiliary graph G'' with terminal set $\{s, t\}$ and $d = 6$, for the bipartite instance C_6.

A vertex cover $A' \cup B' \subseteq A \cup B$ 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 vertex cover. Therefore, the number of cutsets $|\mathcal{C}|$ is precisely the number of vertex covers of the bipartite graph $|\mathcal{B}|$. When $p = 1/2$, all cutsets are equally likely, and the source-terminal reliability evaluation at $p = 1/2$ is:

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

Finally, using the fact that $|\mathcal{B}| = |\mathcal{C}|$ and by substitution:

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

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

The result for $d \geq 3$ is a corollary of Proposition 1.

Theorem 1. *The DCR is $\mathcal{NP}$-Hard when $k = n$ and $d \geq 3$.*

PROOF. Consider the auxiliary graph $G' = (V', E')$ from Fig. 1. Extend G' furthermore, and consider $G'' = (V'', E'')$, 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 words, just add links in order to connect all nodes from A, and all nodes from B. We keep the same probabilities of operation that in G', and the 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 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. ♠

Theorem 2. *The DCR is $\mathcal{NP}$-Hard when $k = n$ and $d = 2$.*

Given a graph $G = (V, E)$, consider $G' = (V \cup \{a, b\}, E \cup \{\{x, a\}, \{x, b\}, \forall x \in V\})$. By its definition, G' has diameter $d = 2$. All links are perfect, except the ones incident to a, with $p(ax) = 1/2$. Consider the DCR for G'. We will show that the number of minimum cardinality pathsets in G' is precisely the number of vertex covers in G'. Since counting minimum cardinality pathsets is at least as hard as computing the reliability of a coherent SBS [1], the result will follow.

A minimum cardinality pathset in G' contains all perfect links, and $\{a, x_1\}, \ldots, \{a, x_r\}$ for certain nodes $x_i \in V$. Since H is a minimum cardinality pathset, the graph $G_H = (V, H)$ has diameter 2, but the diameter is increased under any link deletion. Let $N_a = \{x : \{a, x\} \in H\}$ the set of neighbor vertices for the terminal node a. The key is to observe that *vertex a reaches every node in two steps if and only if N_a is a vertex cover*.

Indeed, suppose a reaches every node in two steps. Then, for any $x \in V \setminus N_a$ there exists a path xya, so $y \in N_a$ and thus N_a is a vertex cover. Conversely, if N_a covers V, let $x \in V$. Then, either $x \in N_a$ and x is adjacent with a, or $x \in V \setminus N_a$ and there exists $y \in N_a \cap N_x$, so xya is a path of two hops between x and a.

The minimality of N_a as a cover follows from the minimality of H as a pathset.

The whole picture of DCR complexity is provided in Fig. 2, which closes the complexity analysis for different independent pairs (k, d).

4. FACTORIZATION IN DCR

Let us consider a network $G = (V, E)$ with perfect nodes and identical probabilities of operation $p_e = p \, \forall e \in E$. Denote $n = |V|$ and $m = |E|$ the respective number of nodes and links in the network. Totaling exhaustive and mutually disjoint events, Michael

	k (fixed)		$k=n$ or free
d \ k	2	$3\ldots$	
2	$O(n)$	$O(n)$	$\mathcal{NP}$-Hard
$3, \ldots, n-2$	$\mathcal{NP}$-Hard		$\mathcal{NP}$-Hard
$n-1, \ldots$	$\mathcal{NP}$-Hard		$\mathcal{NP}$-Hard

Figure 2: DCR Complexity in terms of the diameter d and number of terminals $k = |K|$

Ball and Scott Provan observed that [2]

$$R_{V,G}(p) = \sum_{i=0}^{n} F_i p^{m-i}(1-p)^i, \tag{2}$$

being F_i the number of connected subgraphs $H = (V, E')$ for G such that $|E'| = m - i$. Therefore, the problem can be reduced to counting subgraphs. In particular, if c denotes the minimum cardinality cutset (mincut) then F_{m-c} is the number of those cutsets.

The classical reliability problem, CLR, is $\mathcal{NP}$-Hard (see Section 3 for a discussion of computational complexity). Since DCR is an extension of CLR, it is $\mathcal{NP}$-Hard as well. Once these classical problems are known to be computationally hard, the research community delved into the development of exact exponential algorithms, close approximations and polynomially solvable subclasses of the CLR.

The literature is vast, and we are forced to choose inspirational and most cited works. Remarkably, Moskowitz [21] proposed series parallel reductions and deletion of irrelevant edges, as well as the deletion-contraction principle (or Factoring Theorem): let $e = \{x, y\} \in E$ be an arbitrary edge, $G - e = (V, E - e)$ represents edge-deletion, $G * e$ is an edge contraction (the nodes $\{x, y\}$ are both identified with x, and the graph $G * e$ has possibly multiple edges), and K' is the new terminal-set after the identification of nodes x and y. Then:

$$R_{K,G} = (1 - p_e)R_{K,G-e} + p_e R_{K',G*e}. \tag{3}$$

A notorious computational method for rough estimations of $R_{K,G}$ is Crude Monte Carlo (CMC) and its enhancements [14]. The key idea is to pick N independent random graphs $G_1, \ldots, G_N$ that respect the correct probability law for the links, and set a binary random variable X_i to 1 if the desired condition is met or $X_i = 0$ otherwise. By Kolmogorov's strong law, the average random variable $\overline{X_N}$ converges almost surely to $P(X = 1)$, precisely the target probability (K-connectedness in the CLR, for instance). This estimation is unbiased, and its error can be reduced linearly with the sample size N. Unbiased estimations for $R_{K,G}$ are usually compared with respect to efficiency, which considers both expected square error (i.e., variance) and computational effort. Héctor Cancela and Mohammed El Khadiri developed a Recursive Variance Reduction (RVR) estimation for $R_{K,G}$, with clearly winning efficiency with respect to CMC [9]. Other valuable approximation methods are cross-entropy [18], antithetic variables and uniform bounds [5]. Here we just touched on the surface of CLR. The curious reader can find a recent survey in [4].

The DCR additionally requires the terminals to be connected by path composed not more that d hops. The new parameter d is called the *diameter*, and the target probability is denoted by $R^d_{K,G}$, following the terminology of Héctor Cancela and Louis Petingi [10]. An analogous expression of (2) for the case of homogeneous links holds:

$$R^d_{K,G}(p) = \sum_{i=0}^{n} F_i^{(K,d)} p^{m-i}(1-p)^i, \tag{4}$$

where now $F_i^{(K,d)}$ is the number of d-K-connected subgraphs $H = (V, E')$ for G such that $|E'| = m - i$, and the terminals are linked by paths with d hops or less. Since DCR is a generalization of CLR (the CLR occurs when $d \geq |V| - 1$), the general DCR problem is $\mathcal{NP}$-Hard as well. Special care is needed to adapt Expression 3 for the diameter constrained measure, since the node-contraction operation does not preserve distances. The reader can find an adaptation of factorization the diameter-constrained measure in [10]. There, the authors need to identify all paths that include the selected link.

Observe that if a link $e \in E$ fails, the DCR event corresponds to network $G - e$, where all link reliabilities are kept the same (but $p_e = 0$). On the other hand, if e operates, we should find the DCR of network G^e, that is precisely G but $p_e = 1$. A similar factorization formula for the DCR is the following:

$$R^d_{K,G} = p_e R^d_{K,G^e} + (1 - p_e) R^d_{K,G-e} \tag{5}$$

It is worth to notice that the recursion provided by Equation (5) iteratively *deletes* or *consolidate* links. As a consequence, the iterative procedure finishes in non-connected instances, or, on the other hand, in "strong" networks, where all links are perfect. In the latter, the network is either d-K-connected (and the DCR equals 1) or not (where the DCR equals 0). Future work is required to test exact DCR computations in polynomial time using this novel factorization technique.

5. DCR IN SPECIAL GRAPHS

So far, an efficient (polynomial-time) computation of the DCR is available only for special graphs, to know, paths, cycles, ladders, generalized ladders and spanish fans [25]. The reader can appreciate from Figure 2 that an efficient computation is also feasible for diameter $d = 2$ and a fixed number of input terminals k [6]. An explicit expression for $R^d_{K,G}(p)$ is provided by [7].

In this article, we will extend the previous list, adding Weak graphs, Monma graphs, Tree graphs and Arborescence graphs.

Definition 1. *Let $G = (V, E)$ a simple graph, $K \subseteq V$ and d a positive integer. The graph G is d-K-r weak if $G - U$ is d-K disconnected, for every set $U \subseteq E$ with $|U| \geq r$.*

In words, "r-weakness" states the network fails (i.e. is not d-K connected) whenever we remove an arbitrary set of r links (or more). We consider an analogous notion of strong graphs.

Definition 2. *Let $G = (V, E)$ a simple graph, $K \subseteq V$ and d a positive integer. The graph G is d-K-s strong if $G - U$ is d-K connected, for every set $U \subseteq E$ with $|U| \leq s$.*

Theorem 3. *Let $G = (V, E)$ a d-K-r weak graph, for some r independent of n. Then, the DCR can be found in polynomial time in n.*

PROOF. Given an arbitrary configuration $G' = (V, H) \subseteq G$, we can decide in polynomial time whether G' is d-K-connected or not. Let us denote $\mathcal{O}^r$ to the set of all configurations (V, H), with $|H| \geq m - r$, where $m = |E|$. Since G' is d-K-r weak, summing the probability of disjoint events with positive probability we get that

$$R^d_{K,G} = \sum_{G' \in \mathcal{O}^r} 1_{\{G' \in \mathcal{O}^K_D(G)\}} \prod_{e \in E(G')} p(e) \prod_{e \notin E(G')} (1 - p(e)), \tag{6}$$

where $1_{\{x\}} = 1$ if x is true, and $1_{\{x\}} = 0$ otherwise. It suffices to show that the number of terms in the sum is polynomial with respect to n. In fact, by Sum-rule, the cardinality $|\mathcal{O}^r|$ is precisely:

$$|\mathcal{O}^r| = \sum_{i=0}^{r-1} \binom{m}{m-i} \sim m^{r-1}, \tag{7}$$

where the symbol $\sim$ means that both real sequences are equivalent when m tends to infinity. Observe that $m < n^2$ holds for all connected graphs. Therefore, $|\mathcal{O}^r| \sim m^{r-1} \leq n^{2r-2}$, and the number of terms from Expression (7) is bounded by a polynomial in n. Thus, $R^d_{K,G}$ can be found in a polynomial number of elementary operations in n. ♠

An analogous argument holds for strong graphs.

Corollary 1. *Let $G = (V, E)$ a d-K-s strong graph, for some s independent of n. Then, the DCR can be found in polynomial time in n.*

PROOF. There is a polynomial number of potential non-d-K-connected subgraphs in G. As a consequence, the complement of the DCR, $1 - R^d_{K,G}$, can be found in polynomial time in n. ♠

Corollary 2. *The DCR in connected graphs $G = (V, E)$ with bounded co-rank $c(G) = m - n + 1$ can be found in polynomial time in n.*

PROOF. Consider a simple graph $G = (V, E)$, with bounded co-rank $c(G)$, a terminal set $K \subseteq V$ and diameter d. If we delete an arbitrary link set $U \subseteq E$ of cardinality $c(G) + 1$, the resulting subgraph has less links than a tree. Then, $G - U$ us disconnected, and G is d-K-$(c(G)+1)$ weak. Since $c(G)+1$ is a constant bound, Theorem 3 applies, and the DCR can be found in polynomial time in n. ♠

Corollary 3. *If the number of faces of a connected graph G of genus g has a constant bound, the diameter-constrained reliability can be computed in polynomial time.*

PROOF. Follows from the fact that the number of faces f of a graph of genus g is $f = m - n + 2 - 2g = c(G) - (2g - 1)$. Then, the co-rank $c(G) = f + 2g - 1$ has a constant bound. ♠

Corollary 4. *Consider a graph G with genus g and a constant bound on its faces. Then, if we consider an arbitrary arborescence for G, its diameter-constrained reliability can be computed in polynomial time.*

PROOF. Trees do not add faces, and Corollary 3 holds for any arborescence of G. ♠

The relevance of the following corollary comes from the fact that most telecommunication networks are planar.

Corollary 5. *If the number of faces of a planar graph G has a constant bound, the diameter-constrained reliability can be computed in polynomial time.*

PROOF. A planar graph has genus 0. ♠

The property is unaffected by elementary subdivisions of a graph:

Corollary 6. *If a family of graphs $\mathcal{F}$ has all its elements homeomorphic to a fixed graph, its diameter-constrained reliability can be computed in polynomial time with respect to its order.*

PROOF. Homeomorphic graphs have the same co-rank. ♠

Now we focus on a distinguished family of graphs coming from robust network design. Specifically, Clyde Monma et. al. studied the minimum cost two-connected network design problem, for the metric case spanning all nodes in the set V [19]. There, the authors prove that there exists a solution $G' = (V, H) \subseteq G$ such that every vertex in G' has degree 2 or 3, and the deletion of one or two links from G' leaves one bridge in one of the resultant connected components. Moreover, those graphs are either a Hamiltonian cycle in G or contain a *Monma graph* as an induced subgraph. The term Monma graph was introduced in later works with this family of graphs [8]. Figure 3 sketches a general Monma graph.

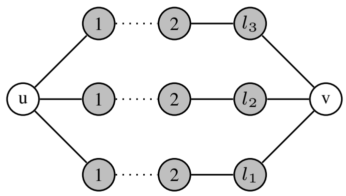

Figure 3: Monma's graph structure.

The following corollary provides the dimension of reliability in the study of Monma graphs:

Corollary 7. *The diameter-constrained reliability of Monma graphs can be computed in polynomial time respect to its order.*

PROOF. Monma graphs are those homeomorphic to the general graph consisting in two vertices and three edges joining them. ♠

Observe that Monma graphs are 3-K-weak for every selection of the terminal set K. Therefore, Theorem 3 also proves Corollary 7. All sub-trees in a Monma graph are obtained removing two links from different u-v-paths (see Figure 3). If we delete more than three links, the resulting subgraph is disconnected. In Appendix A we count the number of spanning trees in an arbitrary Monma graph. Also, trees are 1-K-weak for every subset of terminals K, and Theorem 3 states that the DCR computation in trees is feasible in polynomial time. Indeed, we show in Appendix B that it is linear in the order of the tree.

6. CONCLUSIONS

In this paper we address the diameter-constrained reliability. This measure joints is the probability that all distinguished terminals $K \subseteq V$ in a network $G = (V, E)$ remain connected by d hops or less, where links $e \in E$ may fail with certain probabilities $q_e = 1 - p_e$.

The DCR is $\mathcal{NP}$-Hard, since it subsume the probability that a random graph is connected. We summarize the computational complexity of DCR sub-problems in terms of the number of terminals $k = |K|$ and diameter d. It remains $\mathcal{NP}$-Hard in all cases but $d \leq 2$ and k finite.

Deletion-contraction formulas are available for the classical reliability problem (CLR). However, contractions modify the diameter. Therefore, we adapted this recursive technique with the introduction of a different factorization methodology.

An efficient (polynomial time) DCR computation is possible in special graphs. Indeed, from prior literature we know that the DCR in paths, cycles, ladders, generalized ladders and spanish fans can be found efficiently [25].

In this paper we extended the previous list, including weak and strong graphs, some graphs with bounded genus, arborescences, graphs with bounded co-rank and special classes, to know, Monma graphs and trees.

The best design (minimum cost) 2-node-connected metric network must be either Hamiltonian or it has a Monma graph a an induced subgraph. Then, this work connects reliability aspects of network design in a probabilistic context with robust network design.

As a future work, we wish to find the DCR in Halin graphs, which play a key role in robust network design (specially in 3-connected minimum cost network design). Furthermore, we will analyze local properties of the DCR (node contraction, link deletion and other local movements) that will enrich our understanding in this measure that connects quality (in hop-constrained applications) with reliability. A hint for this study is DCR factorization.

7. REFERENCES

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

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

[3] N. Biggs. *Algebraic Graph Theory*. Cambridge Mathematical Library. Cambridge University Press, 1993.

[4] F. T. Boesch, A. Satyanarayana, and C. L. Suffel. A survey of some network reliability analysis and synthesis results. *Networks*, 54(2):99–107, sep 2009.

[5] J. Brown and C. Colbourn. Non-stanley bounds for network reliability. *Journal of Algebraic Combinatorics*, 5(1):13–36, 1996.

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

[7] E. Canale, H. Cancela, F. Robledo, and P. Sartor. The complexity of computing the 2-k-reliability in networks. *Information Processing Letters*, 114(9):457 – 461, 2014.

[8] E. Canale, P. Monzón, and F. Robledo. Global synchronization properties for different classes of underlying interconnection graphs for kuramoto coupled oscillators. In *Future Generation Information Technology*, volume 5899 of *Lecture Notes in Computer Science*, pages 104–111. Springer Berlin Heidelberg, 2009.

[9] H. Cancela and M. E. Khadiri. The recursive variance-reduction simulation algorithm for network reliability evaluation. *IEEE Transactions on Reliability*, 52(2):207–212, 2003.

[10] H. Cancela and L. Petingi. Diameter constrained network reliability: exact evaluation by factorization and bounds. In *ICIL'2001 (International Conference on Industrial Logistics)*, pages 359–366, 2001.

[11] H. Cancela and L. Petingi. Reliability of communication networks with delay constraints: computational complexity and complete topologies. *International Journal of Mathematics and Mathematical Sciences*, 2004:1551–1562, 2004.

[12] C. J. Colbourn. Reliability issues in telecommunications network planning. In *Telecommunications network planning, chapter 9*, pages 135–146. Kluwer Academic Publishers, 1999.

[13] S. A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the third annual ACM symposium on Theory of computing*, STOC '71, pages 151–158, New York, NY, USA, 1971. ACM.

[14] G. S. Fishman. A Monte Carlo sampling plan for estimating network reliability. *Operations Research*, 34(4):581–594, july-august 1986.

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

[16] R. M. Karp. Reducibility among combinatorial problems. In R. E. Miller and J. W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.

[17] G. Kirchoff. Über die auflösung der gleichungen, auf welche man bei der untersuchung der linearen verteilung galvanischer ströme geführt wird. *Ann. Phys. Chem.*, 72:497–508, 1847.

[18] D. P. Kroese, K.-P. Hui, and S. Nariai. Network reliability optimization via the cross-entropy method. *IEEE Transactions on Reliability*, 56(2):275–287, 2007.

[19] C. Monma, B. Munson, and W. Pulleyblank. Minimum-weight two-connected spanning networks. *Mathematical Programming*, 46(1-3):153–171, 1990.

[20] F. Moskowitz. The analysis of redundancy networks. *AIEE Transactions on Communication Electronics*, 39:627–632, 1958.

[21] F. Moskowitz. The analysis of redundancy networks. *American Institute of Electrical Engineers, Part I: Communication and Electronics, Transactions of the*, 77(5):627–632, 1958.

[22] S. J. Provan and M. O. Ball. The Complexity of Counting Cuts and of Computing the Probability that a Graph is Connected. *SIAM Journal on Computing*, 12(4):777–788, 1983.

[23] F. Robledo. *GRASP heuristics for Wide Area Network design*. PhD thesis, INRIA/IRISA, Université de Rennes I, Rennes, France, 2005.

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

[25] P. Sartor. *Propriétés et méthodes de calcul de la fiabilité diamètre-bornée des réseaux*. PhD thesis, INRIA/IRISA, Université de Rennes I, Rennes, France, december 2013.

[26] L. Valiant. The complexity of enumeration and reliability problems. *SIAM Journal on Computing*, 8(3):410–421, 1979.

APPENDIX

A

The complexity of a simple graph G is its number of spanning trees, denoted by $\kappa(G)$. Gustav Kirchhoff provided an efficient way to count $\kappa(G)$, by means of a determinant [17]. Indeed, $\kappa(G)$ is an arbitrary minor of the Laplacian matrix $L = A_G - \Delta$, being A_G de adjacency matrix of G (i.e. $A = (a_{i,j}) : a_{i,j} = 1$ if $\{i,j\} \in E$; or 0 otherwise) and Δ_G a diagonal matrix with the degrees of the vertices. The result is known as "Matrix-Tree Theorem", and it is a seminar result in the field called Algebraic Graph Theory [3].

In this appendix, we will study the complexity of Monma graphs, $\kappa(M_{l_1,l_2,l_3})$. To the best of our knowledge, even its simplicity this is the first place where a closed expression for $\kappa(M_{l_1,l_2,l_3})$ is available.

We invite the reader to see Figure 3. All Monma graphs have co-rank 2. As a consequence, in order to find spanning trees it is required to delete precisely two links. If both links are removed from the same independent path between nodes u and v, the resulting subgraph is disconnected. On the other hand, if two links from different paths are removed, a tree is obtained. Then, the complexity of M_{l_1,l_2,l_3} is:

$$\kappa(M_{l_1,l_2,l_3}) = l_1 l_2 + l_1 l_3 + l_2 l_3 \tag{8}$$

Expression (8) has the following combinatorial interpretation: it is the number of ways to remove two balls from different bins, where we have exactly l_i balls in bin i, where $i \in \{1,2,3\}$.

It is clear that $l_1 + l_2 + l_3 = m$, the number of links from Monma graph, and that $\kappa(M_{l_1,l_2,l_3}) \leq \binom{m}{2}$, since some deletion of pairs of links result in a tree. Now, we will find a tighter bound for the complexity of Monma graphs. For that purpose, we will study the structure of Monma graphs with m links that maximize the complexity. Consider the following combinatorial optimization problem:

$$\max_{l_1,l_2,l_3} f(l_1,l_2,l_3) = l_1 l_2 + l_1 l_3 + l_2 l_3$$
$$s.t.$$
$$l_1 + l_2 + l_3 = m$$
$$l_1, l_2, l_3 \in \mathbb{N}$$

By the symmetry of function f, we will assume that $l_1 \geq l_2 \geq l_3$ without loss of generality. We will prove that the maximum is attained when $l_1 = l_2 = l_3$ if $m = 3k$ for some $k \geq 1$, $l_1 = l_2 = l_3 - 1$ if $m = 3k-2$, for some $k \geq 1$, or $l_3 = l_2 = l_1 + 1$ otherwise. In the combinatorial interpretation is the following: "the number of ways to remove two balls from different bins is maximized when the number of balls in each bin is balanced".

Indeed, if $l_1 \geq l_2 - 2$ then $f(l_1 - 1, l_2 + 1, l_3) = f(l_1, l_2, l_3) + l_1 - l_2$. Therefore, we subtract a unit from l_1 and it to l_2, and the objective is increased, respecting the constraint $l_1 + l_2 + l_3 = m$. Therefore, the integers l_1, l_2 and l_3 that achieve the maximum cannot differ in more than one unit. The reader can appreciate that if we choose (w.l.o.g.) $l_1 \geq l_2 \geq l_3$ then $l_1 = l_2 = l_3$ if m is a multiple of 3, or they differ in one unit, as mentioned before.

A graph reading is the following: "the complexity of Monma graphs with a fixed size is maximized when the three independent paths have roughly the same length". This maximum is roughly $3(\frac{m}{3})^2 = \frac{m^2}{3} < \binom{m}{2}$. In this case, roughly two-thirds of all pair deletion of links are trees.

B

We already know that trees are 1-weak (i.e. an arbitrary link deletion disconnects them). In this paragraph, we will reinforce this result: The DCR of a tree can be computed in linear time with its order.

PROOF. Let $T = (V,E)$ be a tree and $K \subseteq E$ the terminal set. Since T is a tree, given two terminals $u, v \in K$ there is precisely one path P_{uv} that connects them. All those links must be operational, and the length of P_{uv} must be smaller than d. The links not included the set $P = \cup_{u,v \in K} P_{uv}$ are irrelevant. Let d' be the diameter of P. Therefore:

$$R^d_{K,T} = 1_{\{d' \leq d\}} \prod_{e \in P} p_e, \tag{9}$$

being p_e the probability of operation of link e, $1_{\{x\}}$ equals one if x is true and 0 otherwise. The set P can be found linearly in $|V|$ using breadth first search (BFS) with an arbitrary terminal $u \in K$ as the root node (the process finishes when all terminals are reached). Let $x \in K$ be the terminal that is farthest away from u during the BFS process. If we apply BFS again starting from x as the root node and y is farthest away we get $d' = d(x,y)$. So, the diameter d' can be obtained in linear time with $|V|$. Since the number of products in Expression (9) is $|P| \leq |E| \leq |V|$, the whole computation of $R^d_{K,T}$ be obtained in order $|V|$ elementary operations. ♠