RELIABILITY OF COMMUNICATION NETWORKS WITH DELAY CONSTRAINTS: COMPUTATIONAL COMPLEXITY AND COMPLETE TOPOLOGIES

H. CANCELA and L. PETINGI

Received 22 June 2003

Let $G = (V, E)$ be a graph with a distinguished set of terminal vertices $K \subseteq V$. We define the K -diameter of G as the maximum distance between any pair of vertices of K . If the edges fail randomly and independently with known probabilities (vertices are always operational), the *diameter-constrained K -terminal reliability* of G , $R_K(G, D)$, is defined as the probability that surviving edges span a subgraph whose K -diameter does not exceed D . In general, the computational complexity of evaluating $R_K(G, D)$ is NP-hard, as this measure subsumes the classical K -terminal reliability $R_K(G)$, known to belong to this complexity class. In this note, we show that even though for two terminal vertices s and t and $D = 2$, $R_{\{s,t\}}(G, D)$ can be determined in polynomial time, the problem of calculating $R_{\{s,t\}}(G, D)$ for fixed values of D , $D \geq 3$, is NP-hard. We also generalize this result for any fixed number of terminal vertices. Although it is very unlikely that general efficient algorithms exist, we present a recursive formulation for the calculation of $R_{\{s,t\}}(G, D)$ that yields a polynomial time evaluation algorithm in the case of complete topologies where the edge set can be partitioned into at most four equi-reliable classes.

2000 Mathematics Subject Classification: 05C99, 90B25.

1. Introduction. The components of a communication network (e.g., nodes, communication links) may be subject to random failures. Failures may arise from natural catastrophes (e.g., hurricanes), component wearout, or action of intentional enemies.

A communication network can be modelled by a graph $G = (V, E)$, where V and E are the sets of vertices and edges, respectively, of G . Moreover, the probabilities of failure of the network components can be represented by assigning probabilities of failure to the vertices and/or edges of its underlying graph.

A widely used probabilistic model is the one where the edges fail randomly and independently with known probabilities, and where the vertices are always operational. Let $G = (V, E)$ be a graph with a distinguished set $K \subseteq V$. We define the K -diameter of G as the maximum distance between any pair of vertices of K . If the edges fail randomly and independently with known probabilities, in [2, 8], the *diameter-constrained K -terminal reliability* of G , $R_K(G, D)$, is defined as the probability that surviving edges span a subgraph whose K -diameter does not exceed D , or equivalently, as the probability that for each pair of vertices $\{u, v\} \subseteq K$, there exists an operating path between u and v of at most D edges.

One particular application of this measure is when transmissions between every two nodes of a terminal set K of a network (modelled by a graph $G = (V, E)$) are required

to experience a maximum delay DT , where T is the delay experienced at a single node or edge. The probability that after random failures of the communication links the surviving network meets the delay DT is $R_K(G, D)$.

As real networks are subject to failures, the diameter-constrained reliability can be useful in different contexts. For example, this measure gives an indicator of the suitability of an existing network topology to support good-quality voice over IP applications between a pair of terminals. In the case of a videoconference, we take K to be the set of the participating nodes, and the diameter-constrained reliability gives the probability that we can find short enough paths between all of them. Another potential case of interest is a number of protocols which, in order to avoid congestion by looping data, assign a timeout date or a maximum number of hops to each data packet, to control information; that is, the case for some peer-to-peer (P2P) networks, such as Freenet, Gnutella, or others [3, 7], in which a fixed maximum number of hops are allowed for communication between nodes. In these cases, the diameter-constrained unreliability (i.e., one minus the reliability) gives the probability that, due to failed links, there are some nodes of the network, which are not mutually reachable using these protocols.

In Section 2, we introduce some basic notation and definitions that will be used in the following sections, and we present $R_K(G, D)$ as a generalization of the classical reliability measure $R_K(G)$, allowing us to conclude that, in general, the complexity of evaluating $R_K(G, D)$ is NP-hard. In Section 3, we show that even though $R_K(G, D)$, for $|K| = 2$ and $D = 2$, can be calculated efficiently, this problem is NP-hard for fixed values of D , $D \geq 3$. In Section 4, we generalize this result for any fixed number of terminal vertices. In Section 5, we present a recursive formulation for the calculation of $R_{\{s,t\}}(G, D)$ that yields a polynomial-time evaluation algorithm for the case of complete topologies, whose edge set can be partitioned into at most four equi-reliable classes. Finally, in Section 6, we present some open problems and final remarks.

The notation in this paper follows that of Harary [5], unless otherwise noted.

2. Preliminaries. The following notation and definitions will be used in the sequel.

(i) Let $G = (V, E, \mathcal{P}(E))$ be a probabilistic graph with a distinguished set $K \subseteq V$, $|K| \geq 2$, and $D \in \mathbb{Z}^+$, with $1 \leq D \leq n - 1$, where $n = |V|$ and where $\mathcal{P} : E \rightarrow [0, 1]$ are the operational probabilities of the set of edges E . For ease of notation, we represent the operational probability of an edge $e \in E$ as $p(e) = 1 - q(e)$ ($q(e)$ is the probability of failure).

(ii) Let the sample space Ω represent the set of all possible subsets of E corresponding to sets of operational links (i.e., $\Omega = 2^E$).

(iii) Under the assumption of independent edge failures, each $H \in \Omega$ has occurrence probability

$$P(H) = \prod_{e \in H} p(e) \prod_{e \notin H} q(e). \quad (2.1)$$

(iv) $H \in \Omega$ is a *pathset* or *operating state* if H spans a subgraph whose K -diameter is at most D .

(v) Let $\mathcal{C}_K^D(E) = \{H \in \Omega : H \text{ is a pathset}\}$.

(vi) An operating state H of $\mathcal{O}_K^D(E)$ is called a *minpath* if $H - \{e_i\} \notin \mathcal{O}_K^D(E)$, for all $e_i \in H$.

(vii) $H \in \Omega$ is a *failure state* if H spans a subgraph whose K -diameter is greater than D (if H spans a subgraph where two vertices of K belong to different connected components, then its K -diameter is infinite).

(viii) Let $\bar{\mathcal{O}}_K^D(E) = \{H \in \Omega : H \text{ is a failure state}\}$.

From the definition of $R_K(G, D)$ and definition (v), one gets

$$R_K(G, D) = \Pr(\mathcal{O}_K^D(E)) = \sum_{H \in \mathcal{O}_K^D(E)} \prod_{e \in H} p(e) \prod_{e \notin H} q(e). \quad (2.2)$$

Similarly, (2.2) can be rewritten in terms of the failure states:

$$R_K(G, D) = 1 - Q_K(G, D) = 1 - \sum_{H \in \bar{\mathcal{O}}_K^D(E)} \prod_{e \in H} p(e) \prod_{e \notin H} q(e). \quad (2.3)$$

A widely used probabilistic measure (see [1, 4, 11, 12]) is the classical *K-terminal reliability* of a graph G , $R_K(G)$, defined on the same probabilistic model (i.e., edges fail randomly and independently with known probabilities and the vertices are always operational). The measure $R_K(G)$ is the probability that for every pair of vertices $u, v \in K$, there exists an operating path between u and v . In this case, there are no length restrictions of the paths joining the vertices of K , and by noting that the maximum length of a path joining a pair of vertices is of at most $n - 1$ edges, where n is the number of vertices of G , then

$$R_K(G) = R_K(G, n - 1). \quad (2.4)$$

This generalization of the classical reliability parameter allows us to reflect more stringent performance objectives by restricting the maximum length of a path in a network.

Let $G = (V, E)$ and let K be a set of terminal vertices of G . For the classical reliability measure, computations of the K -terminal reliability (see [10]) and the specific cases when $|K| = 2$ (see [13]) and $K = V$ (see [6, 9]) were shown to be NP-hard. From these results and the fact that $R_K(G) = R_K(G, n - 1)$, $R_{\{s, t\}}(G) = R_{\{s, t\}}(G, n - 1)$, and $R_V(G) = R_V(G, n - 1)$, where n is the number of vertices of the graph G , and by restricting $D = n - 1$, we have the following theorem.

THEOREM 2.1. *For the diameter-constrained reliability, the computational complexity of computing $R_K(G, D)$, $R_{\{s, t\}}(G, D)$, and $R_V(G, D)$ is NP-hard.*

Even though it is very unlikely that $R_K(G, D)$ can be evaluated efficiently, we cannot preclude that it is the case when fixed values of the diameter parameter D are under consideration. We address this question in Sections 3 and 4.

3. Evaluating $R_K(G, D)$ when $|K| = 2$ and when D is a constant value. In this section, we establish the computational complexity of computing $R_K(G, D)$ when K is composed of two terminal vertices s and t , and for fixed values of diameter parameter D .

In [8], an efficient formulation was given for the evaluation of diameter-constrained two-terminal reliability of a network when terminals s and t should be connected by operating paths of at most two edges (i.e., $D = 2$).

THEOREM 3.1. *Let $G = (V, E)$ be a simple graph where each edge $(u, v) \in E$ operates independently with probability $p_{(u,v)}$, and let $\mathcal{N}_{\{s,t\}} = \{u_1, u_2, \dots, u_l\}$ be the common neighborhood of terminal nodes s and t ; then*

$$R_{\{s,t\}}(G, 2) = \begin{cases} 1 - R' & : (s, t) \notin E, \\ 1 - (1 - p_{(s,t)})R' & : (s, t) \in E, \end{cases} \quad (3.1)$$

where

$$R' = \prod_{i=1}^l (1 - p_{(s,u_i)} p_{(u_i,t)}). \quad (3.2)$$

Even though $R_{\{s,t\}}(G, 2)$ can be computed in time linearly on the number of nodes of G , we next show that the complexity of evaluating $R_{\{s,t\}}(G, D)$, for fixed values of D , is NP-hard.

For ease of notation, instead of representing a state (operational or failure) as a set of edges of a graph G , we represent it as a subgraph of G spanned by this set.

THEOREM 3.2. *Evaluating $R_{\{s,t\}}(G, D)$, for fixed $D \geq 3$, is NP-hard.*

PROOF. We show that counting the number of vertex covers of a bipartite graph is polynomial-time Turing reducible to counting the number of failure states of an undirected graph with terminal vertices s and t and fixed diameter parameter $D \geq 3$.

An instance of the bipartite vertex cover consists of a bipartite graph $G = (V, E)$; let X and Y be the classes in the bipartition of V . A vertex cover is a set of vertices $C = C_X \cup C_Y$, $C_X \subseteq X$, and $C_Y \subseteq Y$ such that every edge of E has at least one endpoint in C . The problem of counting the number of vertex covers of a bipartite graph was shown to be #P-complete by Provan and Ball [9].

Let $D = 3 + d$ and let $P = (V(P), E(P))$ be a path on $d + 1$ vertices, where $V(P) = \{s, s_1, \dots, s_d\}$ and $E(P) = \{(s, s_1), (s_1, s_2), \dots, (s_{d-1}, s_d)\}$ (it is understood that if $d = 0$, P is composed of an isolated vertex $s_d = s$). We construct a probabilistic graph $G' = (V', E', \mathcal{P}(E'))$ from G . The vertex set V' consists of V , $V(P)$, and t . The edge set E' consists of E , $E(P)$, the edge sets $E_X = \{(s_d, x) : x \in X\}$ and $E_Y = \{(t, y) : y \in Y\}$ (see Figure 3.1). The edges' operational probabilities are $p(e) = 1$ if $e \in E$, $p(e) = 1$ if $e \in E(P)$, and $p(e) = 1/2$ otherwise. The terminal set is $K = \{s, t\}$.

It is clear that this transformation is polynomial on the size of the input of the bipartite graph, since D is a constant value.

From this construction, a state H of G' consists of the bipartite graph G , the path P , the vertex s_d , possibly joined to some vertices of X , and the vertex t , possibly joined to some vertices of Y of the bipartite graph G .

We then construct a one-to-one correspondence $Z : W \mapsto \overline{\mathcal{O}}_{\{s,t\}}^D(E')$ from the set of vertex covers W of the bipartite graph to the failure states of G' as follows: if $w = (S \cup T) \in W$, where $S \subseteq X$ and $T \subseteq Y$, is a vertex cover of the bipartite graph, then

 FIGURE 3.1. Graph G' constructed from bipartite G and constant d .

FIGURE 3.2. Failure state constructed from a vertex cover of bipartite graph (black vertices represent a vertex cover).

$Z(w) = (V', E \cup E(P) \cup \{(s_d, x) : x \in X - S\} \cup \{(t, y) : y \in Y - T\})$ is a failure state of G' (see Figure 3.2).

To show that this is a one-to-one correspondence, let $w = (S \cup T)$ be a vertex cover of the bipartite graph G , and suppose $x \in X - S$ and $y \in Y - T$. Clearly, (x, y) is not an edge in $Z(w)$, otherwise if (x, y) is an edge, this edge is not covered by w . Thus, there are no paths of length D joining s and t in $Z(w)$. Therefore, $Z(w)$ is a failure state of G' .

Conversely, consider a failure state H which, as remarked previously, must include all the edges of path P , all the edges of set E , and possibly some edges joining vertex s_d to vertices in X and joining vertex t to some vertices in Y . As H is a failure state, the vertices of X and Y adjacent to s_d and t must form an independence set (if that was not the case, there would exist a path of length D joining s and t). Therefore, the remaining vertices of the bipartite graph form a vertex cover which is $Z^{-1}(H)$.

From (2.3), we obtain

$$\begin{aligned} R_{\{s,t\}}(G', D) &= 1 - \sum_{i=0}^{|X|+|Y|} F_i(G', \{s, t\}, D) \left(\frac{1}{2}\right)^i \left(\frac{1}{2}\right)^{|X|+|Y|-i} (1)^{|E|+|E(P)|} \\ &= 1 - \frac{\sum_{i=0}^{|X|+|Y|} F_i(G', \{s, t\}, D)}{2^{|X|+|Y|}}, \end{aligned} \quad (3.3)$$

where $F_i(G', \{s, t\}, D)$ is the number of failure states of G' with $|E| + |E(P)| + i$ edges. But $\sum F_i(G', \{s, t\}, D)$ is equal to the number of vertex covers of the bipartite graph G , thus the result follows. $\square$

In the next section, we generalize this result for any fixed number of terminal vertices.

4. Evaluating $R_K(G, D)$ for any fixed number of terminal vertices K and when D is a constant value, $D \geq 3$. In this section, we show that complexity of calculating $R_K(G, D)$ is NP-hard when any fixed number of terminal vertices are under consideration and when D is a constant value, $D \geq 3$.

THEOREM 4.1. *Evaluating $R_K(G, D)$ for a fixed number of terminal vertices and for fixed D , $D \geq 3$, is NP-hard.*

PROOF. Let $k = |K|$ with $k > 2$, and $D = 3 + d$. We construct a probabilistic $G'' = (V'', E'', \mathcal{P}(E''))$ obtained from the probabilistic graph $G' = (V', E', \mathcal{P}(E'))$ mentioned in the proof of [Theorem 3.2](#) as follows:

- (a) for each j , $1 \leq j \leq k-2$, let $P_j = (V_j = \{u_1^j, u_2^j, \dots, u_D^j\}, E_j = \{(u_1^j, u_2^j), \dots, (u_{D-1}^j, u_D^j)\})$ be a path on D vertices;
- (b) let the set of edges $E_s = \{(s, u_1^j) : 1 \leq j \leq k-2\}$ (i.e., we make the terminal vertex s adjacent to each vertex u_1^j of P_j , $1 \leq j \leq k-2$);
- (c) let the set of edges $E_t = \{(t, u_D^j) : 1 \leq j \leq k-2\}$ (i.e., we make the terminal vertex t adjacent to each vertex u_D^j of P_j , $1 \leq j \leq k-2$);
- (d) let $G'' = (V' \cup (\bigcup_{j=1}^{k-2} V_j), E' \cup E_s \cup E_t \cup (\bigcup_{j=1}^{k-2} E_j), \mathcal{P}(E''))$ (see [Figure 4.1](#)).

For the operational probabilities of G'' , if e is an edge in which one of the endpoints is s_d or t and the other endpoint is a vertex of the bipartite graph G , then $p(e) = 1/2$, otherwise let $p(e) = 1$. Let also the terminal set $K = \{s, t\} \cup (\bigcup_{j=1}^{k-2} \{u_1^j\})$.

This construction is also polynomial on the size of the input of the bipartite G , since both $|K|$ and D are constant values.

If s and t ought to be connected by a path of at most D edges, then this path must comprise the path P of G' (see the proof of [Theorem 3.2](#)). In a failure state H' of G'' , every pair of terminal vertices is connected by paths of at most D edges, with the exception of s and t ; thus

$$\begin{aligned} H &= (V(H), E(H)) \text{ is a failure state of } G' \\ \Leftrightarrow H' &= \left(V(H) \cup \left(\bigcup_{j=1}^{k-2} V_j \right), E(H) \cup E_s \cup E_t \cup \left(\bigcup_{j=1}^{k-2} E_j \right) \right) \text{ is a failure state of } G''. \end{aligned} \quad (4.1)$$

Also, it follows from [\(4.1\)](#) that the occurrence probability (see [Section 2](#), definition (iii)) of the edge set $E(H')$ of a failure state H' of G'' is equal to the occurrence probability of the edge set $E(H)$ of its corresponding failure state H of G' as

$$P(E(H')) = P(E(H)) \cdot 1^{|E_s|} \cdot 1^{|E_t|} \cdot 1^{|\bigcup_{j=1}^{k-2} E_j|} = P(E(H)). \quad (4.2)$$

FIGURE 4.1. Graph G'' constructed from bipartite graph G .

Therefore, it follows from (2.3), (3.3), (4.1), and (4.2) that

$$R_K(G'', D) = R_{\{s,t\}}(G', D) = 1 - \frac{\sum_{i=0}^{|X|+|Y|} F_i(G', \{s, t\}, D)}{2^{|X|+|Y|}}. \quad (4.3)$$

But, as was shown in the proof of [Theorem 3.2](#), $\sum F_i(G', \{s, t\}, D)$ is equal to the number of vertex covers of the bipartite graph G , thus the result follows. $\square$

5. Two-terminal reliability of complete topologies with at most four edge reliability values. In this section, we define a class of complete graphs, whose edges can be partitioned into four elementary reliability values. We will denote these graphs by $G_A(n, q_{st}, q_s, q_t, q)$, where n is the number of intermediate nodes between the terminals s and t , $q_{st} = 1 - r_{st}$ is the unreliability of the edge connecting s and t , $q_s = 1 - r_s$ is the unreliability of the edges connecting s to the other n intermediate nodes, $q_t = 1 - r_t$ is the unreliability of the edges connecting t to the n intermediate nodes, and $q = 1 - r$ is the unreliability of the edges whose endpoints are intermediate nodes (see [Figure 5.1\(a\)](#)). The reliability for this network, $R_{\{s,t\}}(G_A(n, q_{st}, q_s, q_t, q), D)$, is a multinomial in q_{st}, q_s, q_t, q .

We will express the reliability of network G_A by taking into account the independence between the edge (s, t) and the other paths, and by conditioning on the number of operational edges between s and the intermediate nodes (exploiting the symmetry between those nodes). If $D = 1$, the only feasible path is the edge (s, t) ; then trivially $R_{\{s,t\}}(G_A(n, q_{st}, q_s, q_t, q), 1) = 1 - q_{st}$. We can now look at the case where $D > 1$. As a first step, we use the fact that (s, t) is a feasible path and that it is independent from all other simple paths from s to t . Then we can write

$$R_{\{s,t\}}(G_A(n, q_{st}, q_s, q_t, q), D) = (1 - q_{st}) + q_{st} R_{\{s,t\}}(G'_A(n, q_s, q_t, q), D), \quad (5.1)$$

FIGURE 5.1. $G_A(n, q_{st}, q_s, q_t, q)$ and $G'_A(n, q_s, q_t, q)$ networks.

where $G'_A(n, q_s, q_t, q)$ is the probabilistic network obtained from $G_A(n, q_{st}, q_s, q_t, q)$, by deleting the edge (s, t) (see Figure 5.1(b)).

As this network is completely symmetric with respect to the edges between s and the intermediate nodes (all nodes except s and t), we will define a partition of the probability state space for the network based on the number of those edges that fail or are operational. We define A_k to be the event where k edges from s to intermediate nodes fail and the remaining $n - k$ are operating; its probability is

$$\Pr\{A_k\} = \binom{n}{k} (q_s)^k (1 - q_s)^{n-k}. \quad (5.2)$$

The set $\{A_k : 0 \leq k \leq n\}$ is a partition of the probability space, as the events are pairwise-disjoint and their union has probability one. Applying the total probability theorem, we then have

$$\begin{aligned} R_{\{s,t\}}(G'_A(n, q_s, q_t, q), D) \\ &= \Pr\{\exists \text{ operational path of length } \leq D \text{ from } s \text{ to } t \text{ in } G'_A\} \\ &= \sum_{i=0}^n \Pr\{\exists \text{ operational path of length } \leq D \text{ from } s \text{ to } t \text{ in } G'_A \mid A_i\} \Pr\{A_i\}. \end{aligned} \quad (5.3)$$

We must now find an expression for the general term with $k \leq n$. The leftmost network shown in Figure 5.2 corresponds to this event, where k edges between s and intermediate nodes fail (i.e., can be removed from the network) and the remaining $n - k$ are operational, which are presented with a bolder trace. Finding an operational path of length less than or equal to D in this network corresponds to finding a path of length less than or equal to $D - 1$ in the network shown at the center of Figure 5.2, which is obtained by identifying s and the intermediate nodes to which it is unconditionally

FIGURE 5.2. $G'_A(n, q_s, q_t, q)$ when $n - k$ edges between s and intermediate nodes work and the rest fail.

connected into one node. In addition, as edges fail independently, the operational probability of a bank of parallel edges is the complement (to one) of the products of their failure probabilities. Thus, the set of parallel edges between s and t shown at the center of Figure 5.2 can be replaced by a single edge with reliability $1 - q_t^{n-k}$. Similarly, the set of parallel edges between s and an intermediate node of K_k can be replaced by a single edge with reliability $1 - q^{n-k}$.

The network resulting from this last operation is shown at the right of Figure 5.2, and corresponds to a topology, where the operational paths must have length at most $D - 1$. From (5.1), (5.2), and (5.3) and the latest fact, we can express the reliability of the original network G_A by the following formula:

$$\begin{aligned} R_{\{s,t\}}(G_A(n, q_{st}, q_s, q_t, q), D) \\ = (1 - q_{st}) + q_{st} \sum_{k=0}^n \binom{n}{k} (1 - q_s)^{n-k} q_s^k R_{\{s,t\}}(G_A(k, q_t^{n-k}, q^{n-k}, q_t, q), D - 1). \end{aligned} \quad (5.4)$$

The recursive application of this formula gives a multinomial on q_{st} , q_s , q_t , q .

It can be observed that the direct application of the recursive formula leads to evaluating the multinomial many times for values of $k \leq n$ and $d < D$, but with different values to be substituted for the parameters q_{st} , q_s , q_t , and q . This observation leads us to the following alternative formulation. We define a subclass of the class of networks G_A . These networks will be denoted by $G_B(m, h, q_t, q)$, with $G_B(m, h, q_t, q) = G_A(m, q_t^h, q^h, q_t, q)$. This means that these networks have the same complete graph topology, but that the reliabilities of the four-edge classes are defined by only two integer and two real parameters. It is trivial to observe that

$$R_{\{s,t\}}(G_B(m, h, q_t, q), D) = R_{\{s,t\}}(G_A(m, q_t^h, q^h, q_t, q), D). \quad (5.5)$$

Substituting in (5.4), we have that

$$\begin{aligned} R_{\{s,t\}}(G_A(n, q_{st}, q_s, q_t, q), D) \\ = (1 - q_{st}) + q_{st} \sum_{k=0}^n \binom{n}{k} (1 - q_s)^{n-k} q_s^k R_{\{s,t\}}(G_B(k, n - k, q_t, q), D - 1) \end{aligned} \quad (5.6)$$

and that if $D > 1$,

$$R_{\{s,t\}}(G_B(m, h, q_t, q), D) = (1 - q_t^h) + q_t^h \sum_{k=0}^m \binom{m}{k} (1 - q^h)^{m-k} q^{hk} R_{\{s,t\}}(G_B(k, m-k, q_t, q), D-1). \quad (5.7)$$

If $D = 1$, then

$$R_{st}(G_B(m, h, q_t, q), D) = 1 - q_t^h. \quad (5.8)$$

Moreover, it is also noted that the last term (i.e., $k = m$) of (5.7) is null since $G_B(m, 0, q_t, q)$ represents a network with $m + 2$ nodes in which the edges from s to the intermediate nodes fail, and where s is adjacent to the terminal node t by an edge whose reliability is 0 (i.e., $1 - q_t^0$).

The next theorem states the space and time complexity of evaluating $R_{\{s,t\}}(G_A(n, q_{st}, q_s, q_t, q), D)$.

THEOREM 5.1. *Evaluating $R_{\{s,t\}}(G_A(n, q_{st}, q_s, q_t, q), D)$ takes $O(n^3)$ and $O(Dn^3)$ space and time complexity, respectively.*

PROOF. We first start with the standard assumption that the edges' reliabilities are rational numbers p/q , where both p and q are of length $O(n)$, and that arithmetic operations on such rational numbers take constant time (i.e., $O(1)$) (see [4, page 4]).

We next note, from the right-hand side of (5.7), that evaluation of $R_{\{s,t\}}(G_B(m, h, q_t, q), d)$ for specific values of m and h depends on powers of the edges' unreliabilities q_t and q . Since the possible values for m and h are in $[0, n]$, we first proceed to evaluate and store the values q_t^i , for $0 \leq i \leq n$, and q^i , for $0 \leq i \leq n$. This procedure requires $O(n^2)$ space complexity (n values to be stored, each having length $O(n)$) and $O(n^2)$ time complexity. We then calculate and store all possible values q^{ij} , for $0 \leq i, j \leq n$. This can be accomplished by utilizing an $n \times n$ two-dimensional array and values already stored for powers of q . This procedure takes $O(n^3)$ space and time complexity. Values for $(1 - q^i)^j$ are stored in a similar fashion.

Finally, we proceed to evaluate $R_{\{s,t\}}(G_B(m, h, q_t, q), d)$ for specific values of m , h , and d . In order to execute this last step, we note, from (5.7), that an evaluation of $R_{\{s,t\}}(G_B(m, h, q_t, q), d)$ requires at most n values (i.e., n terms in the summation) of $R_{\{s,t\}}(G_B(m, h, q_t, q), d-1)$. An efficient evaluation is, for example, to first evaluate and store $R_{\{s,t\}}(G_B(m, h, q_t, q), d)$, for $d = 1$, and for all values of m (with $0 \leq m \leq n$) and h (with $0 \leq h \leq n$) (i.e., $O(n^3)$ space and $O(n^2)$ time complexity). Next we determine and store the values of $R_{\{s,t\}}(G_B(m, h, q_t, q), d)$ for $d = 2$ and for each pair (m, h) , by utilizing already stored values; by the previous assumptions, this step will take $O(n)$ execution time. We sequentially repeat this last procedure for all possible values of d , $1 \leq d < D$, yielding an $O(Dn^3)$ space and time complexity. We can also improve the storage to $O(n^3)$ by noting that to evaluate $R_{\{s,t\}}(G_B(m, h, q_t, q), d)$, it is only necessary to store values for $d-1$. $\square$

TABLE 5.1. Execution times for evaluating complete graph reliability.

n	D	Execution time (CPUs)	
		Iterative method	Recursive method
10	2	< 0.010	< 0.010
10	5	< 0.010	< 0.010
10	10	< 0.010	0.150
30	2	< 0.010	< 0.010
30	5	0.033	0.067
30	10	0.067	329.120
30	20	0.150	> 36000
30	30	0.233	> 36000
60	2	0.017	0.017
60	5	0.200	1.100
60	10	0.517	> 36000
60	20	1.150	> 36000
60	30	1.750	> 36000

We implemented this iterative procedure as well as the recursive method which is obtained by directly applying (5.4).

Table 5.1 shows the execution times (obtained on an Intel Celeron PC) for both the iterative and the recursive methods, for different values of n and D . Some entries are marked > 36000, corresponding to runs which exceeded 10 CPU hours (36000 seconds) and were aborted. It is easy to see that the iterative method has a much better behavior than the recursive formulation, leading to smaller execution times (especially for high values of D).

6. Conclusions and future work. In this paper, we investigated the computational complexity of the diameter-constrained K -terminal reliability.

In particular, this measure subsumes the classical K -terminal reliability measure when there is no restriction on the length of the paths connecting the terminal vertices. This equivalence between these two reliability measures allows us to conclude that, in general, the computational complexity of the diameter-constrained K -terminal reliability is NP-hard. Nevertheless, the problem of determining the computational complexity of $R_K(G, D)$, for specific cases of K and D , was open. We showed that even though for two terminal vertices s and t and $D = 2$, $R_{\{s,t\}}(G, D)$ can be determined in polynomial time, the problem of calculating $R_{\{s,t\}}(G, D)$ for fixed values of D , $D \geq 3$, is NP-hard. Moreover, we generalized this result for any arbitrary number of terminal vertices. As a consequence, this result justifies the implementation of approximation methods for the evaluation, and the determination of bounds for $R_K(G, D)$. Another relevant open problem is to determine subclasses of graphs for which polynomial-time evaluation algorithms exist. In this paper, we showed that for the case of complete topologies where the edge set can be partitioned into at most four equi-reliable classes, $R_{\{s,t\}}(G, D)$ can be computed efficiently.

REFERENCES

- [1] F. T. Boesch, X. Li, and C. Suffel, *On the existence of uniformly optimally reliable networks*, Networks **21** (1991), no. 2, 181–194.
- [2] H. Cancela and L. Petingi, *Diameter constrained network reliability: exact evaluation by factorization and bounds*, Proceedings of the International Conference on Industrial Logistics (Japan), 2001, pp. 359–366.
- [3] I. Clarke, O. Sandberg, B. Wiley, and T. W. Hong, *Freenet: a distributed anonymous information storage and retrieval system*, International Workshop on Design Issues in Anonymity and Unobservability, The International Computer Science Institute, California, 2000, pp. 311–320.
- [4] C. J. Colbourn, *The Combinatorics of Network Reliability*, International Series of Monographs on Computer Science, The Clarendon Press, Oxford University Press, New York, 1987.
- [5] F. Harary, *Graph Theory*, Addison-Wesley, Massachusetts, 1969.
- [6] M. Jerrum, *On the complexity of evaluating multivariate polynomials*, Ph.D. thesis, Department of Computer Science, University of Edinburgh, Edinburgh, 1981.
- [7] G. Pandurangan, P. Raghavan, and E. Upfal, *Building low-diameter P2P networks*, 42nd IEEE Symposium on Foundations of Computer Science (Las Vegas, Nev, 2001), IEEE Computer Society, California, 2001, pp. 492–499.
- [8] L. Petingi and J. Rodriguez, *Reliability of networks with delay constraints*, Congr. Numer. **152** (2001), 117–123.
- [9] J. S. Provan and M. O. Ball, *The complexity of counting cuts and of computing the probability that a graph is connected*, SIAM J. Comput. **12** (1983), no. 4, 777–788.
- [10] A. Rosenthal, *A computer scientist looks at reliability computations*, Reliability and Fault Tree Analysis (Conf., Univ. California, Berkeley, Calif, 1974), Society for Industrial and Applied Mathematics, Pennsylvania, 1975, pp. 133–152.
- [11] A. Satyanarayana and J. N. Hagstrom, *A new algorithm for the reliability analysis of multi-terminal networks*, IEEE Trans. Reliab. **30** (1981), 325–334.
- [12] D. R. Shier, *Network Reliability and Algebraic Structures*, Oxford Science Publications, The Clarendon Press, Oxford University Press, New York, 1991.
- [13] L. G. Valiant, *The complexity of enumeration and reliability problems*, SIAM J. Comput. **8** (1979), no. 3, 410–421.

H. Cancela: Instituto de Computación, Facultad de Ingeniería, Universidad de la República, J. Herrera y Reissig 565, 11300 Montevideo, Uruguay
E-mail address: cancela@fing.edu.uy

L. Petingi: Computer Science Department, College of Staten Island, City University of New York, 2800 Victory Boulevard, Staten Island, NY 10314, USA
E-mail address: petingi@postbox.csi.cuny.edu