

Polynomial-Time Computability of the Edge-Reliability of Graphs Using Gilbert's Formula

THOMAS J. MARLOWE^{*} and LAURA SCHOPPMANN[†]

*Department of Mathematics and Computer Science Seton Hall University,
South Orange NJ 07079.*

(Received 12 August 1996; Revised 15 October 1996)

Reliability is an important consideration in analyzing computer and other communication networks, but current techniques are extremely limited in the classes of graphs which can be analyzed efficiently. While Gilbert's formula establishes a theoretically elegant recursive relationship between the edge reliability of a graph and the reliability of its subgraphs, naive evaluation requires consideration of all sequences of deletions of individual vertices, and for many graphs has time complexity essentially $\Theta(N!)$. We discuss a general approach which significantly reduces complexity, encoding subgraph isomorphism in a finer partition by invariants, and recursing through the set of invariants.

We illustrate this approach using threshold graphs, and show that any computation of reliability using Gilbert's formula will be polynomial-time if and only if the number of invariants considered is polynomial; we then show families of graphs with polynomial-time, and non-polynomial reliability computation, and show that these encompass most previously known results.

We then codify our approach to indicate how it can be used for other classes of graphs, and suggest several classes to which the technique can be applied.

Keywords: Computational complexity; efficient recursive algorithms; networks; reliability; threshold graphs; Gilbert's formula; graphs; polynomial-time computations

Classification Categories: 90B12, 90B25

^{*} Marlowe is also an Affiliate Professor in the Department of Computer and Information Science, New Jersey Institute of Technology, Newark NJ 07102, and is a Research Associate of its Real-Time Computing Lab. Email: marlowe@homer.njit.edu.

[†] Email: schoppla@shu.edu

1 INTRODUCTION

Analysis of networks for use in communication, and both distributed and parallel computing, must examine not only cost and performance, but also vulnerability. While fault tolerance and recovery are also of interest, one important approach to diminishing vulnerability is failure avoidance, through provision of reliable networks. There are a number of graph models for reliability.

The edge reliability problem and its applications. One commonly-used performance measure of network reliability is the *all-terminal reliability* of an undirected graph $G = \langle V, E \rangle$. Given that the nodes of G are perfectly reliable, and the edges e fail independently and with known probability p_e , the *all-terminal edge reliability* $R[G]$ is the probability that the surviving edges constitute a spanning connected subgraph S of G . Reliability is used to measure vulnerability in communication and computer networks; this is then used either to compare alternative implementations, or to provide fault-tolerance, typically by adding edges, or via redundancy.

We restrict our attention to the case in which the failure probability is constant; that is, there exists a p so that for all e , $p_e = p$. In this case, we can parametrize the reliability by p , and obtain a *reliability polynomial*, $R[G](p)$. Henceforward, we use (edge) “reliability” to mean the all-terminal edge reliability polynomial.

The computation of all-terminal edge reliability for arbitrary graphs, even for constant failure probability, appears to be a computationally difficult problem; in (Provan [1]) it is shown that it is $\#P$ -complete. For this reason, significant efforts have been made to obtain upper and lower bounds on all-terminal reliability, and to determine subclasses of graphs for which reliability can be efficiently computed, and to find techniques for computing reliability for those classes.

Among the graphs known to have polynomial-time algorithms for reliability are acyclic graphs, complete graphs (Buzacott [2]), and a few classes of threshold graphs, as discussed below. Many of these results use variants of Gilbert’s formula (Colbourn [3]), which provides a recursive method for evaluation of the all-terminal edge reliability of simple graphs. We present a technique, related to the approach of (Buzacott [2]) for complete graphs, for efficient evaluation of reliability based on Gilbert’s formula. We apply this approach to threshold graphs, and in Section 8, to a number of other classes of graphs.

The contribution of this paper is, in our opinion, the technique at least as much as the particular results. At the heart of the technique is identification of an invariant encoding the embedding-isomorphic classes of subgraphs of a given graph. This encoding is a common approach to constructing recursive algorithms, and together with memoization of previously computed results, ensures that reliabilities are not repeatedly evaluated. Other ideas used repeatedly include memoization of previously computed results, use of sequences of differences, and, in some cases, construction of an auxiliary polynomial. We discuss limitations, and show a simple case in which the technique provably cannot provide a polynomial-time algorithm.

Conventions. We use the standard definitions and conventions of graph theory, combinatorics, and complexity theory (Cormen [4], Harary [5], Tucker [6]).

We use the terms “node”, “point”, “vertex” interchangeably. All graphs are simple, without parallel edges or self-loops. For simplicity in complexity computations, logarithms are assumed to be base 2; this has no effect on polynomial-time complexity classes.

For a vertex v in a graph G , the *neighborhood* of v , $N(v)$, is the set of all (other) vertices to which it is adjacent. A subgraph $H = \langle V_H, E_H \rangle$ of a graph $G = \langle V_G, E_G \rangle$ is *induced* if $E_H = E_G \cap (V_H \times V_H)$. A *spanning* subgraph is one for which $V_H = V_G$.

Given a graph G and a subgraph H , the cut set $\text{cut}(H, G)$ is the set of edges incident on both H and $G - H$. Two subgraphs H_1 and H_2 of G are *embedding-isomorphic* if they are isomorphic, and in addition, their cut-sets with respect to G have the same cardinality.

Threshold graphs. A threshold graph can be thought of as composed of a single large clique, together with a number of points adjacent only to clique vertices, called *cone points*. Most definitions of threshold graphs require them to be *nested*, so that the cone points can be ordered so that their neighborhoods are nested.

Threshold graphs form a class close enough to a complete graph that the problem of computing their reliability has been considered tractable. Threshold graphs also play an important role in providing lower bounds for the reliability of a graph. In (Satyanarayana [7]), it is shown that threshold graphs minimize reliability, that is, for any graph G , there is a threshold graph T with the same number of points and edges so that

$R[T](p) \leq R[G](p)$ for all $0 \leq p \leq 1$. In (Petingi [8]), it is shown that these lower bounds can be achieved by threshold graphs whose reliability is easily computed, namely, the balloon graphs.

Nonetheless, the question of when a class of threshold graphs, or of connected graphs in general, has a all-terminal reliability polynomial which is easily (polynomial-time) computable remains of significant interest in understanding graph reliability. At present, only limited classes of threshold graphs are known to have polynomial-time-computable reliability.

PROPOSITION 1 *The following classes of threshold graphs have polynomial-time-computable reliability: (1) balloon graphs (one cone point); (2) proper threshold graphs (all cone points of same degree); (3) a proper graph together with one other cone point.*

Proof (1) (Petingi [8]). (2) (Schoppmann [9]). (3) (Saccoman [10]).

2 RELATED APPROACHES TO EDGE RELIABILITY

2.1 Gilbert's Formula

Original formulation. Edge reliabilities for (connected) graphs can be computed recursively using Gilbert's formula (Colbourn [3], Gilbert [11]). Let G be a connected graph, and v be a vertex of G . Then the reliability polynomial $R(G)$ is given by:

$$R[G](p) = 1 - \sum_{v \in G' \subset G} (1-p)^{|cut(G', G)|} R(G') \quad (1)$$

where G' is a proper, connected, induced subgraph of G , $cut(H, K)$ is the cut set of H and $K - H$, and the edge reliability of $\{v\}$ is 1. It is fairly easy to see that the degree of $R[G]$ is $|E(G)|$. Vertex v is called the *pivot* of the expansion.

The complexity of computing the reliability polynomial through a sequence of pivots is essentially independent of the pivot chosen at each step, so we consider a recursive version of Gilbert's formula using the same pivot vertex for every expansion. Gilbert's formula can thus be considered as a formula for *rooted connected graphs*.

We henceforward assume that we are dealing with the class of rooted connected graphs, that all subgraphs (unless otherwise stated) contain the

root, and that graph isomorphism preserves the root. We typically suppress the root, and write G instead of $\langle G, v \rangle$. We also assume that all subgraphs are induced, unless otherwise stated.

Complexity for Threshold Graphs. Naive recursive application of Gilbert's formula will have unacceptably high time complexity, as well as disappointing space complexity.

THEOREM 1 *Let G be a threshold graph, v a cone vertex of highest degree and w a clique point of highest degree (that is, a neighbor of the cone points of lowest degree). Let G have n clique points and m cone points, and $e = O(n^2 + nm)$ edges. Then the reliability polynomial computation by Gilbert's formula on the rooted graph $\langle G, v \rangle$ takes time $\Omega((n+m-2)!e)$ and space $\Omega((n+m-2)e)$.*

Proof For the time complexity, consider the sequences of graphs resulting when one vertex at a time is deleted.

For the space complexity, consider one such sequence of deletions. The state of the computation must include the current index in each of the $n+m-2$ summations, and each of the partially computed polynomials. \square

By considering a path P_n or the star graph $K_{1,n}$, we can see that Gilbert's formula is not always computationally preferable, even with optimizations. At best, versions of Gilbert's formula will require $O(n^2)$ evaluations, but other techniques will use only $O(n)$ operations.

2.2 Buzacott's Approach for Complete Graphs

Although techniques are presented in (Buzacott [2]) for computing edge reliabilities of complete graphs even for non-constant edge probabilities, we here discuss only his refinement of Gilbert's formula for the constant-probability case.

Note that, for K_n , all subgraphs with k nodes are isomorphic, and have cut sets of the same size. Modifying his notation slightly, we have:

$$R[K_n] = 1 - \sum_{k=1}^n \binom{n-1}{k-1} (1-p)^{k(n-k)} R[K_k]$$

where the $n-1$ excludes the root. Inductive application results in computation of reliability of subgraphs lying in complete graphs, for all pairs of sizes (k, r) , $1 \leq k < r \leq n$. This requires $O(n^2)$ computations.

3 A GENERAL APPROACH TO COMPUTING RELIABILITY

We would like to generalize Buzacott's approach to more general graphs. First, we remember subgraphs and reliabilities as they are seen – note that this requires being able to identify when we are revisiting a graph. Second, if we could tell when two subgraphs are (embedding-)isomorphic, and how many subgraphs are isomorphic to a given subgraph of a given graph, then we could recurse on isomorphic subgraphs instead of all subgraphs. Unfortunately, detecting graph isomorphism is known to be hard in general, and will likely be messy even for “nice” families of graphs. What we can do, however, is try to estimate isomorphism and inclusion by an invariant I with an order \leq with the following properties. (Condition (3) is relaxed somewhat in (Marlowe [12]).)

Isomorphism – refining : If $I(G_1) = I(G_2)$ then $G_1 \simeq G_2$. (1)

Order – preserving : If $G_1 \subset G_2$ then $I(G_1) \leq I(G_2)$. (2)

Embedding Invariant : If $G_1, G_2 \subset G_3$, and $I(G_1) = I(G_2)$,
then $|cut(G_1, G_3)| = |cut(G_2, G_3)|$ (3)

Ideally, the invariant should “minimally” represent the set of embedding-isomorphism classes of subgraphs, in the sense that almost all $I' < I(G)$ represent subgraphs, and that unequal invariants are “likely” to represent non-embedding-isomorphic subgraphs. We say that an invariant is *tight* if (A) the number of invariants of subgraphs of G is polynomial if and only if the number of embedding-isomorphic subgraphs is polynomial, and (B) there is a constant $c > 0$ so that if there are N invariants $I' \leq I$, then at least $c \cdot N$ of them correspond to (rooted, connected, induced) subgraphs of G .

If we can find a tight order-preserving invariant, which, for a family of subgraphs \mathcal{G} , takes on only a polynomial number of values in the size of the graph, and if the reliability polynomial for each subgraph corresponding to an invariant can be computed in polynomial time, then the reliability computation for graphs in \mathcal{G} will be computable in polynomial time, *and conversely*.

We illustrate this principle in the following sections, using threshold graphs.

4 AN INVARIANT FOR THRESHOLD GRAPHS

Definition of the invariant. Let G be a threshold graph consisting of a clique K_n and m cone vertices v_1, v_2, \dots, v_m where

$$N(v_i) \subseteq N(v_{i+1}) \subseteq K_n \text{ for } 1 \leq i \leq m - 1.$$

Distinguish $v_r = r$ as the root of the rooted connected graph $\langle G, r \rangle$.

Note that permutation of cone points with the same degree (and therefore the same neighborhoods), other than r , produces a graph isomorphism, as does exchange of two clique points with the same (possibly no) cone point neighbors. We therefore partition the set of cone points as

$$\begin{aligned} W_i &= \{w_{i1}, w_{i2}, \dots, w_{i,p_i}\} \\ &\text{for } 1 \leq i \leq t \\ W^* &= \{r\} \end{aligned}$$

where $p_i > 0$ for all $1 \leq i \leq t$, $N(w_{ij}) = N(w_{ik})$ for all i, j, k and $N(w_{i1}) \subset N(w_{i+1,1})$ for all i . Note that $\sum p_i + 1 = m$.

This gives a corresponding partition of clique $K_n = H_1 \cup H_2 \cup \dots \cup H_t \cup H^* \cup H_0$ where $H_1 = N(w_{11})$, $H_{i+1} = N(w_{i+1,1}) - N(w_{i1})$ for $1 \leq i \leq t - 1$, $H^* = N(r) - N(w_{t1})$, and $H_0 = K_n - \cup H_i - H^*$.

If we let $|H_i| = h_i$, then $h_1 = \text{deg}(w_{11})$, $h_{i+1} = \text{deg}(w_{i+1,1}) - \text{deg}(w_{i1})$, $h^* = \text{deg}(r) - \text{deg}(w_{t1})$, and $h_0 = n - \text{deg}(r)$. Note that all the h_i , except perhaps h^* and h_0 , are positive, and that their sum is (trivially) n .

LEMMA 1 $|E_G| = \binom{n}{2} + \sum_{i=1}^t p_i (\sum_{j=1}^i h_j) + (\sum_{j=1}^i h_j + h^*)$.

Proof The formula counts, respectively, the clique edges, the edges adjacent to each W_i , and the edges adjacent to r . □

We now define an invariant I via $I(G) = \langle p_1, \dots, p_t, h_1, \dots, h_t, h^*, h_0 \rangle$, and for any subgraph G' of G , $I(G') = \langle \alpha_1, \dots, \alpha_t, \beta_1, \dots, \beta_t, \beta^*, \beta_0 \rangle$, where $\alpha_i = |G' \cap W_i|$, $\beta_i = |G' \cap H_i|$, and likewise for β^* and β_0 . $I(G')$ will be the invariant for G' as a graph unless some of the α_i or β_i (other than β^* or β_0) are 0. We define the order \leq on the set of invariants component-wise.

Even when there are zeroes, there is a straightforward conversion between the invariant for G' as a subgraph and its invariant as a graph; namely, while there are still zeroes, if $\alpha_i = 0$, delete it, and set $\beta'_i + 1 = \beta_i + \beta_{i+1}$ and delete β_i ; if $\beta_i = 0$, do the same, exchanging α and β .

EXAMPLE Consider the graph illustrated in Figure 1(a), which can be seen to have invariant $I = \langle 2, 1, 2, 2, 1, 1 \rangle$.

One of the subinvariants of I is $I' = \langle 0, 1, 2, 0, 0, 1 \rangle$. We construct a subgraph H for this invariant as follows: $2t + 2 = 6$, so $t = 2$. Thus there is one non-root cone point, and three points in the clique, two of which are neighbors of the cone point (and the root). Thus I' is an invariant for H in Figure 1(b). Figure 1(c) shows H embedded in G , allowing the reader to check subsequent claims about cut sets, etc.

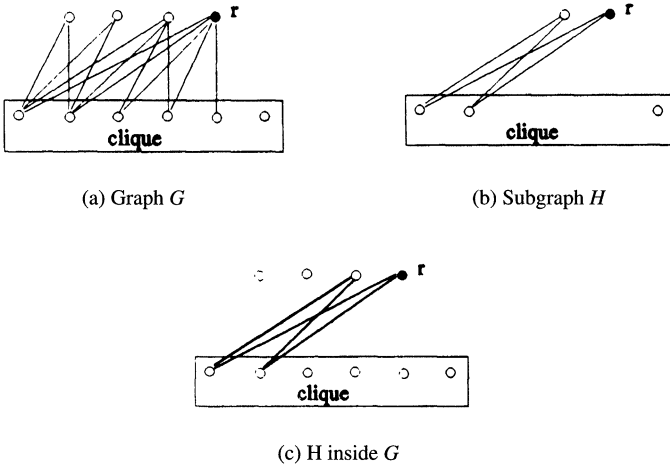


FIGURE 1 A rooted threshold graph and a subgraph

I is an invariant. Note that, trivially, two subgraphs with the same invariant are isomorphic, and a subgraph of the original graph, or of one of its subgraphs, has a strictly smaller invariant. We can also see that I is embedding-invariant, and has one more useful property, namely, that it is easy to recognize the invariants corresponding to connected subgraphs.

LEMMA 2 *Let I' be an invariant. Suppose that I' meets the following conditions for all i : (a) If $\alpha_i \neq 0$ then $\sum_{j=1}^i \beta_j > 0$, and (b) if $\sum_{j=1}^i \beta_j + \beta^* = 0$, then $\beta_0 = 0$. Then there is a (connected) subgraph of G with subgraph-invariant I' .*

Conversely, if these conditions fail, then I' does not represent a connected subgraph.

The following allows us to specify reliability in terms of invariants:

LEMMA 3 For $G_1 \subset G_2$, $|cut(G_1, G_2)|$ depends only on $I(G_1)$ and $I(G_2)$. We can therefore refer unambiguously to $|cut(I_1, I_2)|$. Further, $|cut(I_1, I_2)|$ can be determined in time polynomial in t .

Proof Consider the two invariants and look at the edges between G_1 and $G_2 - G_1$. The numbers of clique-clique, clique-cone, and cone-clique edges then clearly depend only on combinatorial information encoded in the invariants. Using the δ_i , each computation involves at worst $O(t)$ operations. (For details, see (Marlowe [12]).) □

I is tight. Proving tightness is somewhat more difficult.

LEMMA 4 Given an invariant $I = \langle p_1, \dots, p_p, h_1, \dots, h_p, h^*, h_0 \rangle$, there are precisely

$$\Pi = \prod_{i=1}^t (p_i + 1) \cdot \prod_{i=1}^t (h_i + 1) \cdot (h^* + 1) \cdot (h_0 + 1)$$

sub-invariant sequences of I .

COROLLARY 1 There are $\Gamma =$

$$\Pi \left[1 - \frac{1}{p_1 + 1} \left[1 - \frac{1}{h_1 + 1} \left[1 - \frac{1}{p_2 + 1} \left[\dots \left[1 - \frac{1}{h_t + 1} \left[1 - \frac{1}{h^* + 1} \left[1 - \frac{1}{h_0 + 1} \dots \right] \dots \right] \right] \right] \right] \right] \geq \frac{1}{2} \Pi$$

invariants corresponding to (rooted connected) subgraphs of G .

LEMMA 5 Let $t \geq 4$.

Let $N_1 = \{I = (\alpha_1, \dots, \alpha_p, \beta_1, \dots, \beta_p, \beta^*, \beta_0)\}$ such that for $1 \leq i \leq t$, $\alpha_i > 0$, $\beta_i > 0$, and $\beta_0 \neq 1$, and $I \neq I(G)$.

Likewise, let $N_2 = \{I = (\alpha_1, \dots, \alpha_p, 1, \dots, 1, 0, 0)\}$, where $\alpha_i \in \{0, 1\}$, and $I \neq I(G)$.

Let $G(N_j)$, $j = 1, 2$ be a set of graphs consisting of one representative of each invariant in N_j . Then $G(N_j)$ is a set of non-isomorphic graphs.

Proof Each I in N_j represents a connected graph, by lemma 2. Since $t \geq 4$, each graph in $G(N_1) \cup G(N_2)$ contains at least four clique points, and so has a unique maximal clique consisting entirely of clique points.

In $G(N_1)$, each graph has $t + 1$ classes of cone points, where the classes, and their nesting order, are distinguishable. Any pair of invariants must differ in at least one position; but considering the α_i , β_i , β^* , and β_0 in that order shows that in each case the underlying graphs also differ.

In $G(N_2)$, $\alpha_i = 1$ if and only if there is a non-distinguished cone point with precisely i neighbors, and different invariants must differ in some α_i position. □

THEOREM 2 *Let $S = S(G, r)$ be the set of nonisomorphic subgraphs of $\langle G, r \rangle$. Then $|S| \in \Omega(\Pi^{1/3})$. Thus $|S|$ is polynomial in $n + m$ if and only if Π is.*

Proof If $t < 4$, then $|S| \leq \Pi \leq (m + 1)^3 (n + 1)^3 \cdot n \cdot n$, which is polynomial in $n + m$. So assume $t \geq 4$.

Consider N_1 and N_2 from the previous lemma, where clearly $N_1 \cup N_2 \subset S$, and thus $|S| \geq \max(|N_1|, |N_2|)$. Also note that $|N_1 \cap N_2| = 1$.

However, by elementary counting techniques,

$$|N_1| \geq \prod_{i=1}^t p_i \prod_{i=1}^t h_i (h^* + 1) h_0 - 1 \text{ and } |N_2| \geq 2^t - 1.$$

But since $|S| \leq \Pi = \prod_{i=1}^t (p_i + 1) \prod_{i=1}^t (h_i + 1) (h^* + 1) (h_0 + 1)$, we have

$$\begin{aligned} \Pi &\leq \prod_{i=1}^t (2p_i) \prod_{i=1}^t (2h_i) (h^* + 1) (2h_0) \\ &= \prod_{i=1}^t p_i \prod_{i=1}^t h_i (h^* + 1) h_0 2^{2t+1} \quad \square \\ &\leq 2(|N_1| + 1)(N_2 + 1)^2 \leq |S|^3 \end{aligned}$$

5 COMPUTING THE RELIABILITY POLYNOMIAL FOR THRESHOLD GRAPHS

Gilbert’s Formula for a Good Invariant. We can now restate Gilbert’s formula in terms of invariants.

$$R[I] = 1 - \sum_{I' < I} \chi(I', I) N(I', I) (1 - p)^{|cut(I', I)|} R[I'] \tag{2}$$

where $\chi(I', I)$ is 1 if I' represents connected subgraphs of I , and 0 otherwise, and $N(I', I)$ is the number of such subgraphs if $\chi = 1$ (and can be anything if $\chi = 0$).

This formula bears a close resemblance to Buzacott's; the difficulty lies not in expressing the formula, but in determining the invariants, and obtaining expressions for the necessary coefficients.

Assume we have computed reliability for all smaller invariants, and that we have precomputed and stored all binomial coefficients, and can access them in constant time. The following results follow from counting and previous lemmas.

LEMMA 6 *Let \mathcal{G} be the class of threshold graphs, and $G \in \mathcal{G}$. Given two invariants $I_1 \leq I_2$ representing subgraphs of G , where $I_j = (\alpha_1^j, \dots, \alpha_i^j, \beta_1^j, \dots, \beta_i^j, \beta^{*,j}, \beta_0^j)$, for each subgraph G_2 corresponding to I_2 , there are precisely*

$$\prod_{i=1}^t \binom{\alpha_i^2}{\alpha_i^1} \prod_{i=1}^t \binom{\beta_i^2}{\beta_i^1} \binom{\beta^{*,2}}{\beta^{*,1}} \binom{\beta_0^2}{\beta_0^1}$$

subgraphs of G_2 with invariant G_1 .

LEMMA 7 *For the invariant of Section 4 on threshold graphs, the functions $\chi(I', I)$, $N(I', I)$, and $|\text{cut}(I', I)|$ are known and computable in time $O(t)$.*

LEMMA 8 *Assume I is a subgraph invariant of $G \in \mathcal{G}$, and that the reliability polynomials $R[I']$ for all $I' \leq I$ are available, and can be accessed in time proportional to the degree of the polynomial. Let $\Pi(I)$ be the set of invariants less than I . Then the reliability polynomial $R[I]$ can be computed in time $O(t + (\text{deg}(R[I]))^2 |\Pi(I)|)$.*

Proof The values $\chi(I', I)$, $N(I', I)$, and $|\text{cut}(I', I)|$ are each computable in time $O(t)$. Each polynomial $(1 - p)^{|\text{cut}(I', I)|}$ can be expanded in time proportional to its degree, and each multiplication takes at worst the product of the factor degrees — at most $(\text{deg}(R[I]))^2$. There are $|\Pi(I)|$ factors; terms can be added as computed, at constant cost; the final subtraction from 1 is insignificant. Thus total cost is $O(t + (\text{deg}(R[I]))^2 |\Pi(I)|)$. \square

THEOREM 3 *The total time complexity of computing the reliability of a threshold graph G is $O((n + m)^4 \Pi^2)$, and the space complexity is $O((n + m)^4 + (n + m)^2 \Pi)$.*

Proof Precomputation of binomial coefficients $\binom{p}{q}$ for $0 \leq q \leq p \leq e$ = $|E_G|$ requires time and space each $O(e^2) = O((n+m)^4)$.

The time complexity result then follows from Lemma 8, since $\deg(R(I)) \leq e \leq (n+m)^2$ for each I , and $|\Pi(I)| \leq \Pi$, and a total of Π different polynomials must be computed.

For space complexity, note that all the binomial coefficients, and all of the Π polynomials, of degree up to $e \leq (n+m)^2$, have to be stored, but everything else can be computed on the fly in time $O(t) = O(m)$. □

COROLLARY 2 *The time and space complexity of computing the reliability polynomial for a family of threshold graphs \mathcal{H} with the invariant form of Gilbert's formula is polynomial if and only if the number of invariants Π is.*

6 A CHARACTERIZATION OF POLYNOMIAL-TIME RELIABILITY FOR THRESHOLD GRAPHS

From the previous section, our computation of reliability polynomials for a family of threshold graphs \mathcal{H} is polynomial if and only if the number of invariants is polynomial. This provides examples of polynomial and non-polynomial families, and almost always quickly resolves complexity for any other family of threshold graphs.

PROPOSITION 2 *For each class $\mathcal{G}_{a,b,c,d,e,f}$ of threshold graphs (for fixed integer constants a, b, c, d, e, f) for which*

1. *The number of classes of non-equivalent cone points is at most $t \leq a + b \log(n+r)$;*
2. *There are at most c invariants α_i for which $\alpha_i > d$;*
3. *There are at most e invariants $\beta_i, 1 \leq i \leq t$, for which $\beta_i > f$; $\mathcal{G}_{a,b,c,d,e,f}$ has a polynomial-time algorithm for edge reliability.*

Proof For simplicity, we count the $c + e$ invariants twice.

$$\begin{aligned} \Pi &\leq d^{a+b \log n} \cdot n^c \cdot f^{a+b \log n} \cdot n^e \cdot n \cdot n \\ &= (df)^a n^{2+c+e+b(\log d + \log f)} \\ &= n^{O(1)} \quad \square \end{aligned}$$

As a corollary, we get the previously known examples of polynomial-time computable edge-reliability. Note that these classes trivially cover the examples of Proposition 1.

COROLLARY 3 *The following families of graphs have polynomial-time-computable edge reliabilities (where all parameters are fixed integers).*

1. *The class \mathcal{G}_c of threshold graphs with cone points of at most c different degrees.*
2. *The class \mathcal{G}'_c of threshold graphs with at most $O(\log n)$ cone points, of degree at most c .*

However, not all polynomial-time-computable-reliability families fit into this class.

PROPOSITION 3 *Let c be a fixed integer parameter.*

1. *The class $\mathcal{G}_c^\#$ of threshold graphs with at most $c \frac{\log n}{\log \log n}$ cone points, so that each invariant (except possibly β^* and β_0) is at most $(\log n) - 1$, has polynomial-time-computable reliability.*
2. *The class \mathcal{G}_c^* of threshold graphs with at most $\log \log n$ cone points, such that each invariant (except possibly β^* and β_0) is at most $n^{\frac{c}{\log \log n}} - 1$, has polynomial-time-computable reliability.*

Proof

$$\begin{aligned} (1) \quad \Pi &\leq n^2 \cdot (\log n)^{2c \frac{\log n}{\log \log n}} \\ &= n^2 2^{2c \log n} = n^{2+2c} = n^{O(1)} \end{aligned}$$

$$\begin{aligned} (2) \quad \Pi &\leq n^2 \cdot (n^{\frac{c}{\log \log n}})^{2 \log \log n} \\ &= n^{2+2c} = n^{O(1)} \quad \square \end{aligned}$$

Nonetheless, the line between polynomial and non-polynomial can be seen to be fairly sharp. For example:

PROPOSITION 4 *Let $N = n + m$.*

1. *Suppose a family of threshold graphs \mathcal{H} contains graphs of arbitrarily large size for which the number of distinct degrees of cone points is $\Omega(f(N) \log N)$, where $\lim_{N \rightarrow \infty} f(N) = \infty$. Then \mathcal{H} does not have polynomial-time-computable reliability.*
2. *Suppose a family of threshold graphs \mathcal{H} contains graphs of arbitrarily large size for which there are at least $\Omega(f(N) \frac{\log N}{\log \log N})$ non-equivalent cone points for which the corresponding β_i is at least $\log N$, where $\lim_{N \rightarrow \infty} f(N) = \infty$. Then \mathcal{H} does not have polynomial-time-computable reliability.*

Proof

$$(1) \quad \prod \geq 2^{f(n) \log n} = n^{f(n)}.$$

$$2) \quad \prod > (\log n)^{f(n) \frac{\log n}{\log \log n}} = 2^{f(n) \log n} = n^{f(n)}. \quad \square$$

7 APPLICABILITY OF THE TECHNIQUE

In (Marlowe [12]), we present results on additional examples, all closely related either to threshold graphs or complete graphs, of the use of our technique. Two extend threshold graphs to allow a number of cone-point cross edges, and imperfect nesting, respectively. The others show that our technique applies to families of graphs considered in reliability theory and for fault-tolerance: a complete graph minus a matching, k -partite and related graphs, and sunflowers. We present a summary of those results in Section 8. In contrast, we show below that wheel graphs do not have polynomial-time-computable reliability using Gilbert's formula.

Wheels are a problem for Gilbert's formula. We now show that no version of Gilbert's formula can yield a polynomial-time algorithm for the reliability of a wheel W_{n+1} (a cycle C_n with a central vertex (the *hub*) adjacent to all of the cycle vertices (the *rim*)). Since we believe there is a polynomial-time algorithm using expansion by minors, this shows that a Gilbert's formula-based method may be strongly suboptimal even for graphs which are not essentially acyclic, or, like the cycle, acyclic after the deletion of a few vertices or edges.

Let \mathcal{W} be the family of wheels. Without loss of generality, we use the hub as the root vertex.

PROPOSITION 5 *Let h be the hub of a wheel W with $2n$ rim vertices. Then there are at least $\text{Parts}(n)$ non-isomorphic subgraphs of $\langle W, h \rangle$ of size n (where $\text{Parts}(n)$ is the number of partitions of the integer n (Tucker [6])).*

Proof Except for W itself, each (rooted induced) subgraph consists of a set of disjoint paths, P_1, P_2, \dots , with each path vertex connected to the hub (see Figure 2). As with the sunflower, distinct sequences of non-decreasing path lengths correspond to non-isomorphic graphs.

Let $\pi = (a_1, a_2, \dots, a_\pi)$ be a partition of n . Suppose, without loss of generality, the rim vertices are labeled v_1, v_2, \dots, v_{2n} , and adjacency is equivalent to adjacency of indices mod $2n$. We construct a subgraph of W as follows.

Starting with v_1 , let P_1 contain a_1 vertices, and delete the next vertex. Then let P_2 contain the next a_2 vertices, and delete the next, and so on. Finally, after constructing P_π and deleting its next neighbor, delete any remaining rim vertices. (This must be possible since $\pi \leq n$ and $\sum a_i = n$.) The graph just constructed has invariant π . \square

Figure 2 shows an induced subgraph of a wheel W_n , where $n \geq 13$. The subgraph will have invariant $(1, 2, 2, 3)$.




FIGURE 2 A subgraph of a wheel

COROLLARY 4 *The number of non-isomorphic subgraphs of W_n is super-polynomial. Therefore, no algorithm for the reliability of a wheel based strictly on Gilbert's formula can compute the reliability in polynomial time.*

Proof By (Hardy [13]), $|\text{Parts}(n)| \sim e^{\sqrt{n}}$. \square

8 EXTENSIONS AND GRAPHS OF OTHER TYPES

In (Marlowe [12]), we present results on additional examples, all closely related either to threshold graphs or complete graphs, of the use of our technique. Two extend threshold graphs to allow a number of cone-point cross edges, and imperfect nesting, respectively. The others show that our technique applies to families of graphs considered in reliability theory and for fault-tolerance: a complete graph minus a matching, k -partite and related graphs, and sunflowers. We present a summary of those results in Table I.

TABLE I Summary of Results

<i>Example</i>	<i>Parameters</i>	<i>Polynomial-Time?</i> <i>(using Gilbert's Formula)</i>
Acyclic Graph		Yes (previously known) (expansion by minors)
Star	vertices n	Yes (previously known)
Complete Graph	vertices n	Yes (previously known)
K^n - matching	vertices n	Yes
k -layered graph	fixed k	Yes
Threshold graph	p_i, h_i, h^* (see Section 5)	Linear-time test
Sunflower	core size n_o clique sizes n_i	Polynomial-time test
k -layered graph	k , set sizes $ S_i $	
	C_k, P_k	like threshold
	K_k	like sunflower
Ladder graph		Some cases
Threshold w/cross edges		Some cases
Wheel	size n	No

Each of these examples also illustrates subtleties in the technique. The cross-edge example shows that the \leq relation is not always straightforward. The ladder shows that formulating a good invariant can be tricky. The complete graph minus matching example illustrates the combinatorial complexity which can emerge in counting the number of subgraphs with a given invariant; k -layered graphs illustrate issues of connectivity; and the sunflower shows the problem when large numbers of invariants corre-

spond to isomorphic subgraphs. The cross-edge example, and the complete graph minus a matching also require extensions to handle varying sizes of cut sets.

In formulating general invariants, two new subtleties should be emphasized. First, the order on the invariants need not be component-wise, trading complication for expressiveness. Second, the number and set of values of an invariant less than or equal to a given invariant value has to be computable in time polynomial in the size of that set; this is unlikely to be a problem in practice.

Definitions. In a *ladder threshold graph*, cone points need not quite have nested neighborhoods; we can think of the resulting graph as a (not necessarily disjoint) union of the non-root cone points of b threshold graphs with the same clique and root. For $b = 2$, we can define an invariant which essentially considers a two-dimensional analogue of the threshold invariant defined by pairwise intersections of differences of clique sets. A family of ladder threshold graphs will have polynomial-time complexity if the two component threshold graphs do, and either all of the clique sets are small, or for each clique set differences in one component, there are at most a constant number of non-empty intersections with the other component. This approach trivially generalizes to larger b .

k-partite graphs are good candidates for this approach, since they are regular, and in many ways like complete graphs. There is a natural k -tuple invariant (after accounting for the root). For k fixed or slowly increasing, the entire class of k -partite graphs has polynomial-time-computable reliability. k -layered graphs are a generalization of k -partite graphs. In a k -partite graph, there are k sets of independent vertices, all of which are connected to vertices in every other set. We can obtain other graphs by starting with k sets of independent vertices and a graph G on k vertices, and connecting the vertices in sets i and j precisely if there is an edge (v_i, v_j) in G . Our results extend to cases in which G is P_k or C_k .

A *sunflower* S is a union of cliques, K_{k_i} , $1 \leq i \leq s$, whose pairwise intersection is a constant clique K_k for $k > 0$, where, of course, $k < k_i$ for all i . K_k is called the *core* of the sunflower, and the remaining vertices of each clique form its *petals*. By construction, there are no edges between vertices in different petals. We distinguish one of the core vertices as the root. Clearly, any induced (connected, rooted) subgraph of S is again a sunflower, although possibly with fewer petals (perhaps zero petals, if vertices

are either in the core or in a single petal). While we can form an invariant I by defining $I(S) = \langle k_1 - k, k_2 - k, \dots, k_s - k, k - 1 \rangle$ and extending to subgraphs by intersection, it is far from a tight invariant, since any permutation of the first s components for which each component is at most $k_i - k$ is a legal invariant for an isomorphic graph. In fact, the reader can verify that if $k_i = k + 1$ for all i , the number of invariants will be $k 2^s$, but the number of non-isomorphic subgraphs will be $k(s - 1)$. This problem can be overcome to some extent by using ordered invariant sequences (that is, with elements in increasing order), but this complicates counting significantly, requiring inclusion-exclusion. Using this new invariant, we show that, for a given family of sunflowers, if the number of distinct petal sizes is small in all invariants (so that in particular inclusion-exclusion is fast), then reliability can be computed in polynomial time.

Graphs consisting of a complete graph minus a matching (a set M of edges with each vertex adjacent to at most one edge of M) are of significant interest in reliability theory, where they provide upper bounds for edge reliability (Satyanarayana [7]). The family of all such graphs has polynomial-time-computable reliability. The invariant is of the form (a, b, c) , where a is the number of non-root clique vertices, b is the number of matched vertices (other than the neighbor r' of the root if matched), and c is a $(0,1)$ value indicating the absence or presence of r' . The difficulty in this case is that the invariant is not embedding-isomorphic. However, subgraphs with a given invariant are restricted in how they lie in the graph, and it is easy to define an auxiliary polynomial which divides them into embedding-isomorphic classes.

For any family of threshold graphs \mathcal{G} , consider the family \mathcal{G}' defined by adding edges between cone points such that each cone point, except perhaps the root r , has at most one such incident edge. Call this extension a *cone-point near matching*. The invariant approach now combines features used in threshold graphs, and in K_n minus a matching. We can show that (1) adding any number of edges incident on the root does not destroy polynomial-time-computability of reliability; and (2) neither does adding $O(\log n)$ other cross edges. Further, for a class of cone-point near-matchings whose underlying family of threshold graphs has at most c non-equivalent classes of cone points, for a fixed integer c , the entire class has polynomial-time-computable reliability.

Summary of results. We give the results summarized above in Table I.

9 CONCLUSIONS AND FUTURE DIRECTIONS

We have given a general approach which characterizes precisely the polynomial-time computability of the all-terminal-edge-reliability by means of Gilbert's formula for many classes of graphs, and have sketched a number of examples, including several classes (threshold graphs, K_n minus a matching, k -partite graphs) of significant interest in reliability theory. Unfortunately, we also show that in most cases the computation also inherently uses polynomial space; any non-PSpace computation would apparently have to use a different method, which would have to avoid the computation of reliability for most isomorphism classes of subgraphs.

There are at least three directions in which we would like to continue this work. We would like to more fully characterize families of graphs to which our approach applies; we would also like to have a better handle on isomorphism-rich cases like the sunflower; and we would like to determine when other techniques may be preferable to our approach. We also expect to characterize cases in which changing the root at some point in the computation may lead to a more natural description by invariants.

We would also like to extend our technique, or find other, compatible techniques, to apply to families defined by topological rather than combinatorial properties, such as planar or outer-planar graphs, and to multigraphs. We will consider a hybrid technique which handles bridge edges, and some other special cases, using expansion by minors or related techniques (Brown [14], Petingi [8], Satyanarayana [7], Schopmann [9]).

A principal open problem is how to extend these methods to multigraphs, even of a restricted form (e.g., all multiedges incident on the root). If we can extend our approach to handle multigraphs, then we should be able to apply hybrid techniques much more generally.

Acknowledgements

We have benefitted in the development of this paper by conversations with colleagues in our department, particularly Dan Gross and John T. Saccoman, and from the encouragement of Charles Suffel of Stevens Institute and Charles Colbourn of the University of Waterloo. Schopmann was partly supported by a Seton Hall University Research Council Award and by a Clare Boothe Luce Fellowship.

References

- [1] J. S. Provan, M. O. Ball, The complexity of counting of counting cuts and computing the probability that a graph is connected, *SIAM J. Computing* **12** (1983), 777–788.
- [2] J. A. Buzacott, A recursive algorithm for finding reliability measures related to the connection of nodes in a graph, *Networks* **10** (1980), 311–327.
- [3] C. J. Colbourn, *The Combinatorics of Network Reliability*, Oxford University Press, New York, 1987.
- [4] T. H. Cormen, C. E. Leiserson, R. L. Rivest, *Introduction to Algorithms*, McGraw-Hill, New York, 1990.
- [5] F. Harary, *Graph Theory*, Addison-Wesley, Reading MA, 1972.
- [6] A. Tucker, *Applied Combinatorics*, John Wiley and Sons, New York, 1980.
- [7] A. Satyanarayana, L. Schoppmann, C. L. Suffel, A reliability-improving graph transformation with applications to network reliability, *Networks* **22** (1992), 209–216.
- [8] L. Petingi, J. T. Saccoman, L. Schoppmann, Uniformly least reliable graphs, *Networks* **27** (1996), 125–131.
- [9] L. Schoppmann, Computation of Edge Reliability Using Graph Transformations, Ph.D. Thesis, Stevens Institute of Technology, May 1990.
- [10] J. T. Saccoman, Some Synthesis Results in Network Reliability Theory, Ph.D. Thesis, Stevens Institute of Technology, May 1995.
- [11] E. N. Gilbert, Random graphs, *Ann. Math. Stat.* **30** (1959), 1141–1144.
- [12] T. J. Marlowe, L. Schoppmann, On the polynomial-time computability of the edge-reliability of graphs using Gilbert's formula, New Jersey Institute of Technology Technical Report CIS-95-24, June 1995.
- [13] G. W. Hardy, E. M. Wright, *An Introduction to the Theory of Numbers*, Oxford, England: Clarendon Press, 1938.
- [14] J. I. Brown, C. J. Colbourn, J. S. Devitt, Network transformations and bounding network reliability, *Networks* **23** (1993), 1–13.



Hindawi

Submit your manuscripts at
<http://www.hindawi.com>

