

A NEW TECHNIQUE IN THE GLOBAL RELIABILITY
OF CYCLIC COMMUNICATIONS NETWORK

Jon A. Sjogren (USAARTA/AVSCOM)
NASA Langley Research Center
Hampton, VA 23665-5225

Article Proposed for the
Microelectronics and Reliability
(An International Journal & World Abstracting Service)

Abstract.--The global reliability of a communications network is the probability that given any pair of nodes, there exists a viable path between them. A characterization of connectivity, for a given class of networks, can enable one to find this reliability. Such a characterization is described for a useful class of undirected networks called "daisy-chained" or "braided" networks. This leads to a new method of quickly computing the global reliability of these networks. Asymptotic behavior in terms of component reliability is related to geometric properties of the given graph. Generalization of the technique is discussed.

(NASA-TM-103479) A NEW TECHNIQUE IN THE
GLOBAL RELIABILITY OF CYCLIC COMMUNICATIONS
NETWORK (NASA) 32 p CSCL 17B

N90-26214

Unclass

65/32 0292255

1. INTRODUCTION

In a distributed computing system, correct performance may well require *connectivity* of the network. That is, each processor (node) must be able to exchange information with any other processor. A simple line network then has this property, since a message can always find a path, as long as all the internode *links* (or edges) are properly functioning (*up* or *viable*). If we assume that there is a non-zero probability that a given link is failed (non-functioning) at a given time, more links may be added to the network. This increases the probability that enough paths will exist for connectivity.

The *daisy-chain ring*, or *braided ring*, communications structure is mentioned in Pradhan's [2] survey article, and discussed by Grnarov et al. [3], and Hafner et al. [4]. An implementation of this architecture has been built at the NASA Langley Research Center as a candidate system for inter-process communication on the space station. We consider both the *daisy-chain ring* and *line* configurations, and give a characterization of connectivity for each. This provides understanding of the idea of connectedness in these networks, and leads to a new method for computing their *global reliability* (probability of connectedness). This type of reliability is also known as *all-to-all* reliability (Provan [6], Ball and Provan [1]) and *all-terminal* reliability (Politof and Satyanarayana [7]). We consider that all nodes are perfectly reliable and that the links have equal unreliability.

Other ways of finding this reliability exist; classically one may attempt to use path-set or cut-set methods, (see [5]) but this fails to exploit the considerable symmetry of the given situation and will be much more expensive in terms of memory and number of arithmetic operations. Even using sophisticated tools such as boolean-expression analyzers, a 10-node system may take several days to run (on a VAX 11/780), with considerable input pre-processing overhead. Results from such a vast number of computations must be considered suspect because of cumulative error. An approach such as series-parallel reduction [7], may be used, but the technique of *blocking sets* introduced here is well suited to a clear, explicit derivation, and is suggestive of further generalization.

Using this new characterization of connectivity, a polynomial-time computation for the global reliability of these networks (the braided line and ring) is determined. This has been implemented as a program that recursively calculates ring and line connectedness probabilities for increasing values of n , the number of nodes. An explicit polynomial in U , the link unreliability, can easily be written down by using this algorithm. The technique (disjoint blocking sets) also permits a reliability analysis of related communications architectures, such as a daisy-chain ring where several links are *known* to be failed (or missing) and where several others are known to be functioning. Our results agree with numbers produced by an exhaustive spanning-tree approach for 5, 6, 7, and 8 nodes. Answers for less than 30 nodes are produced almost instantly on a VAX 11/750 machine.

After defining the notation, in section 2 we define the networks we are concerned with, give a didactic analysis of the simple line and ring problems, and define the q -blocking set. In sections 3 and 4 the main theorems are stated relating non-connectivity to the existence of certain q -blocking sets. In section 5, a natural decomposition of the sample space is given with an explanation of why this is pivotal to our approach. In section 6, recursive formulas

for the braided-line and -ring connectivity are derived using conditional probabilities and a recursive analysis. This argument is simple for the case of the braided-line, but more subtle for the ring case. In section 7, numerical results are given and some asymptotic properties are noted. Some interesting relations between asymptotic values of the reliability and the number of trees in the graph are noted. Also, generalizations for applying the method to sub-ring architectures are examined, along with limitations on trying to generalize further.

Notation

n	number of nodes
l_n	simple line with n nodes
r_n	simple ring with n nodes
L_n	braided n -line
R_n	braided n -ring
$Q(n)$	reliability of L_{n-1}
$Q_R(n)$	reliability of R_n
U	link failure probability
α, β	links (edges)
$l_{\beta,n}$	$l_n - \{\beta\}$
$(i, i+1)$	primary edge
$(i, i+2)$	moon
$[i, j]$	interval in n -line
$[i, j] \pmod{n}$	interval in n -ring
E	set of edges
G, H	n -graph, network with failures (edges removed)
F_{L_n}	probability space for L_n
σ	outcome in probability space
$B_{i,j}, C_{i,j}$	blocking sets (blockers)
B_q	blocker of length q
\bar{B}	extension of blocker to R_n
π	path in graph
$P(B)$	probability of B occurring
$(3,4)^-$	the edge (3,4) is down
T, W	events; subset of F_{L_n} or F_{R_n}
G_l	global blocker
Ξ	global 'interval'
γ_i, ζ_j	blocking probabilities
$\phi(n)$	number of flops
M	number of spanning trees
Λ	exotic network
P_α, P_β	failure probabilities of R_n with α or β removed
$x \equiv y$	approximate equality

Other, standard notation is given in "Information for Readers & Authors" at rear of each issue.

2. PROBLEM STATEMENT

Introductory Example

Consider the *simple line* network with n nodes l_n . [Figure 1]. Its *links*, or edges, may be given by

$$e_i = (i) \rightarrow (i+1)$$

for node number $1 \leq i \leq n-1$. One may also write e_i as $(i, i+1)$. The simple line may be turned into the *simple ring* r_n by adding a link $e_n = (n, 1)$. Let U be the unreliability of a link: the probability of its being failed (inoperative) at a given time. Then the simple line is *disconnected* (has more than one path-connected component) if and only if one or more of the links are failed (*down*). The probability of this *event* may be computed in several different ways. Firstly, it is seen that the event of being disconnected is complementary to the event (actually a single outcome) of all links being "up" (functioning). The probability of this last event is of course

$$1 - P_D(n) = (1-U)^{n-1}.$$

Here, $P_D(n)$ is the probability of the line being disconnected. Another method for computing $P_D(n)$, which seems a bit artificial but serves to introduce concepts and notation that will be of use later, is as follows. The n -line l_n can be written as

$$l_n = [1, \dots, n].$$

Select α , a link in l_n , say $\alpha = e_k = (k, k+1)$. Then when α is *removed*, two lines remain, namely

$$\begin{aligned} l_{first} &= [1, \dots, k], \\ l_{last} &= [k+1, \dots, n]. \end{aligned}$$

Now the probabilistic "event" ' l_n is down' corresponds to $E_a \cup E_b$, where E_a is the event 'link α is down', and E_b is the event 'link α is functioning and either l_{first} or l_{last} is down'. Then E_a and E_b are disjoint, $E_a \cap E_b = \emptyset$ since α cannot be simultaneously up and down. One can further write E_b as a disjoint union

$$E^{1,2} \cup E^{1,\bar{2}} \cup E^{\bar{1},\bar{2}}$$

where, for example, $E_b^{1,\bar{2}}$ is the event where α is down, l_{first} is up and l_{last} is down. We then have

$$\begin{aligned} (2.1) \quad P_D(n) &= U + (1-U)\{P_D(k)(1-P_D(n-k-1)) + (1-P_D(k))P_D(n-k-1)\} \\ &= U + (1-U)(P_D(k) + P_D(n-k-1) - P_D(k)P_D(n-k-1)). \end{aligned}$$

Here we assume that $P_D(k)$ for values of k less than n are already computable. This recursion is complete when

$$\begin{aligned} P_D(0) &= 0, \\ P_D(1) &= U. \end{aligned}$$

For the case of the simple ring, disconnection requires *two* links to be down. (When only one link is down, there is either a clockwise or a counterclockwise path between two

given nodes.) There are $\binom{n}{c}$ ways of choosing a set of cardinality c from $E =$ set of all n edges. Then the probability of disconnection is

$$(2.2) \quad \sum_{i=2}^n \binom{n}{c} U^c (1-U)^{n-c}.$$

Alternatively, note that for the ring r_n to be up, it is sufficient that

- a) all links be up, OR
- b) at most one link be down.

The probability of the event in a) is $(1-U)^n$, but it can also be generated as follows. Picking a particular link β , that which remains when it is removed is an n -line l_n whose up-probability is already known (by the above discussion). Therefore for the event a), we have $(1-U)(1-P_D(n))$, which checks out to $(1-U)^n$. The probability in b) can be given by considering a link β . If β is down and the "complementary n -line" $l_{\beta,n}$ is up, the event is satisfied. But the probability of this occurring is

$$U \cdot (1-P_D(n)).$$

There are n choices for β here, and the choices lead to pairwise disjoint events. Finally adding the probabilities for a) and b), we obtain

$$(2.3) \quad 1-P_D(r_n) = (1-U)(1-P_D(n)) + nU \cdot (1-P_D(n)),$$

which equals $nU \cdot (1-U)^{n-1} + (1-U)^n$ which is consistent with formula (2.2).

Braided Lines and Rings

We now define the *braided line* network with n nodes L_n . The network consists of two types of link:

- 1) primary edges $e_i: (i) \rightarrow (i+1)$, $i = 1, \dots, n-1$, which can be written $(i, i+1)$, and
- 2) secondary edges, or "moons" (due to a fancied resemblance to a rising moon). These edges connect alternating nodes, and are in turn of two types: upper moons $u_{2j-1}: (2j-1) \rightarrow (2j+1)$ and lower moons $u_{2j}: (2j) \rightarrow (2j+2)$. These links are also given by $(2j-1, 2j+1)$ and so forth. See Figure 2 for L_7 .

The braided-ring (daisy-chain) network R_n is similar with an additional primary edge $(n, 0)$ and secondary edges $(n, 1)$, $(n-1, 0)$. See Figure 3 for R_{12} . To analyze the respective connectedness probabilities for these two types of network, a precise definition of disconnection, and a necessary and sufficient condition for disconnection will next be given.

3. CONNECTIVITY AND BLOCKING

Given two nodes a and b in a graph G , a *path* π from a to b is a subset E_1 of E , the edges of G , with the following properties. Each edge has two vertices (boundary nodes). For our purposes, there is no loss in assuming that these are distinct. Then there must exist functions

$$\phi_i: E_1 \rightarrow \{\text{Nodes}\},$$

$$\phi_t: E_1 \rightarrow \{\text{Nodes}\},$$

such that $\phi_i(e)$ is a vertex of e , called the initial vertex, and $\phi_t(e)$ is a different vertex (if there are two distinct ones), called the terminal vertex. There must also be an ordering of $E_1 = (e_1, \dots, e_k)$ so that $\phi_i(e_1) = a$, $\phi_t(e_k) = b$, and for $j = 1, \dots, k-1$, $\phi_t(e_j) = \phi_i(e_{j+1})$.

Thus a path from a to b is an ordered set of oriented edges where terminal and initial nodes of successive edges are equal, beginning at a and ending at b . What we will be considering are subgraphs of L_n with the same nodes, which we call an n -graph G . We will look at G in three slightly different ways. Firstly, G can be viewed as L_n , with only certain edges removed. Secondly, G can be considered as a mapping from the edges of L to the set {up, down}. Thirdly, G is an *outcome*, an element of the event F_{L_n} . Thus a path in G is simply a path in L which has no down edges (for the particular up-down assignment). In the diagrams, down edges will be represented by a short segment through them, as in Figure 2.

There are certain subsets (of the set of up and down edges) of L_n which are closely tied to the connectivity of a graph G . Such a *blocking* (or barrier) set is uniquely specified by an initial node i and a final node j . If $q = j-i$, we may call it a q -blocking set, written B_q , or we may write $B_{i,j}$. The defining conditions (rules) of a blocking set are:

- 1) All the primary edges $(i, i+1)^-, \dots, (j-1, j)^-$ are *down*, AND
- 2) Any upper moon u whose node-ends are both in (i, \dots, j) is *up*. (That is, if $u = (k, k+2)$, then $i \leq k$ and $k+2 \leq j$.) Similarly, any lower moon v whose ends are in (i, \dots, j) must be *up*, AND
- 3) If $i > 1$, the moon $(i-1, i+1)$ is *down*. If $i = 1$, no corresponding condition applies (this moon does not exist). If $j < n$, then the moon $(j-1, j+1)$ is *down*. If $j = n$, no condition applies.

Formally, a blocking set B is a certain type of subset of $E \cup E^-$, the union of all edges as considered "up" with all edges as considered "down". A blocking set gives rise to an event $F_B \subset F_E$. For certain edges, their up-down status is specified by the three rules above; all other edges are not specified, and each generates two separate outcomes. We will often not distinguish between the blocking set B_q and its interpretation as an event. Thus we can write $B_q \subset L_n$, $P(B_q)$, and so on. In fact, the probability of $B_{i,j}$ is fairly easily found:

$$P(B_{i,j}) = \begin{cases} U^{q+2}(1-U)^{q-1} & \text{if } i \neq 1 \text{ and } j \neq n, \\ U^{q+1}(1-U)^{q-1} & \text{if } i = 1 \text{ or } j = n \text{ but not both,} \\ U^q(1-U)^{q-1} & \text{if } i = 1 \text{ and } j = n \end{cases}$$

where U again is the probability of link failure. The differences arise due to the presence of moons on either end of the blocking set. In $B_{i,j}$, the moon $(i-1, i+1)$ is called a *left-moon* if it exists (when $i \neq 1$, thus contributing to $P(B)$), otherwise it is a *virtual left-moon* and contributes nothing. Similarly, there are right-moons and virtual right-moons. All *right* and *left*-moons must be considered *down*. The other moons of $B_{i,j}$ are called *inner-moons* and should be considered *up*.

Any blocking set $B_{i,j}$ corresponds to a *contiguous set* of q primary edges, namely $(i,i+1), (i+1,i+2), \dots, (j-1,j)$. Any such collection of edges gives rise to a unique blocking set. An edge from the collection then *belongs to* that blocking set, and the blocking set *covers* that primary edge.

Definition. In an n -graph G (subgraph of L_n), node a is disconnected from node b if and only if there is no path of edges in G (up edges) with initial node a and terminal node b .

This is consistent with our terminology on connectedness and n -graphs.

Definition. The n -graph G contains the blocking set $B_{i,j}$ if the outcome $G \in F_{L_n}$ satisfies $G \in F_{B_{i,j}}$ as well. In other words, the edges whose "up" and "down" status is specified by rules 1) - 3) that define B retain their status in G . Figure 4a shows a 5-graph G containing $B_{2,4}$.

The main result we seek is that an n -graph G is disconnected if and only if it contains a blocking set. Now any 1-graph is connected, so we take $n \geq 2$ unless otherwise indicated.

Lemma 1. Let G be an n -graph ($n \geq 2$). Then node 1 is disconnected from node 2 if and only if the primary edge $(1,2)^-$ is *down* and belongs to a blocking set B of G (G contains B).

Proof. First assume that $(1,2)$ belongs to a blocking set B . We wish to show that node 1 is disconnected from node 2 in G . In the case $n = 2$, the only blocking set is $B_{1,2}$, and the only edge is $(1,2)^-$ which is down by rule 1). Thus 1 and 2 are disconnected. If $n > 2$, then the blocking set extends to $[1, \dots, j]$. (The corresponding primary edges are covered.) There are no paths in G originating at 2 and extending to the left. Any edge leaving 2 is a moon extending to the right. But if the *right* end-node of this moon lies outside of B , then

- 1) this moon is *down* by rule 3), OR
- 2) the moon is virtual (when $n = 3$).

In either case, node 2 is isolated and we are finished. Thus the right vertex 4 of this moon is in B . The primary edge $(3,4)^-$ is down by rule 1). So, node 4 has no edges leading to the left except for $(2,4)$. If this moon is part of a path to node 1, then the above argument (about edges from 2) shows that there is a shorter path $4 \rightarrow 1$. Similarly, nodes 6, 8, \dots are in B , or there is *no* path $\pi: 1 \rightarrow 2$. But B is finite and thus the conclusion holds.

Next we must show that for 1 to be disconnected from 2 implies that some $B_{1,k}$ is contained in G . Clearly $(1,2)^-$ must be down. If $n = 2$, this down edge already satisfies the rules for a blocking set so we are finished. If $n > 2$, if $(2,3)$ is up, then $(1,3)^-$ must be down, or there is a connection. But then $B_{1,2}$ is the required blocking set with a virtual left-moon. Take $(2,3)^-$ as down and $(1,3)$ as up. If n happens to be 3, then $B_{1,3}$ already is a blocking set having $(1,2)^-$ as a primary edge and a virtual right-moon. Now if $n > 3$, the moon $(2,4)$ may be taken as *up*, since otherwise $B_{1,3}$ would still fit the definition of a blocking set with the required property. But then $(3,4)^-$ must be down, else we have a path $1 \rightarrow 3 \rightarrow 4 \rightarrow 2$.

We continue *rightwards* in the following manner. On reaching a new node j ,

- i) If $j = n$ we have constructed *down* primary edges and *up* moons, so we have a blocking set $B_{1,n}$ satisfying the lemma. OR,
- ii) Consider the edge $(j,j+1)$. If this edge is up then the moon $(j-1,j+1)$ must be down or there will be a connection (path) $1 \rightarrow 2$. (Consider the cases $j = \text{even, odd}$.)

If the moon $(j-1, j+1)^-$ is down, then $B_{i,j}$ satisfies the lemma. On the other hand, if the edge $(j, j+1)$ is down go to step i) or ii) with j replaced by $j+1$, and so on. By finiteness, the lemma is proved.

Observation. If $a, b \in G$, with $a < b$, are disconnected, then there exists a' with $a \leq a' < b$, such that a' and $a'+1$ are disconnected.

Proof of Observation. If $b = a+1$, we are finished. If not, pick a node c distinct from a and b lying between them. Then c is disconnected from either a or b . Rename this new disconnected pair a and b and proceed as before. This eventually results in a' with the desired property.

The following two lemmas are generalizations of lemma 1 and imply the main result on line networks almost immediately.

Lemma 2. Suppose a and $a+1$ are disconnected in G . Then there is a blocking set $B_{s(a), t(a)}$ which covers $(a, a+1)$; that is, $(a, a+1)$ is a primary edge of B .

Proof. Clearly $(a, a+1)^-$ is down. If $a = 1$, we are finished according to lemma 1. Else consider Figure 4b. If the moon $(a-1, a+1)$ is up, then the $(a-1, a)^-$ primary edge is down. Now a can be neither the right nor the left endpoint of a blocking set; that would violate rule 2. It follows immediately that $(a, a+1)$ is a primary edge in a blocking set if and only if $(a-1, a)$ is a primary edge in the *same* blocking set. Also, $a-1$ is disconnected from a or there would be a connection from a to $a+1$ through $a-1$. But we know that if $a-1$ were equal to 1, there would be (by lemma 1) a blocking set covering $(a-1, a, a+1)$. If not, by induction on a , there is a B that covers $(a-1, 1)$ and must cover $(a, a+1)$ as well, since we have just seen that these two edges are both primary edges in the same blocking set.

Thus we may assume that the moon $(a-1, a+1)^-$ is down.

At this point, the simplest way to argue is to consider the $n-a+1$ -graph H , which is the subgraph of G obtained by throwing away all nodes $1, \dots, a-1$ and any edges incident to them. Clearly $a (= 1_H)$ and $a+1 (= 2_H)$ are disconnected in H , and $(a, a+1)$ is covered by a blocking set $C_{1,k}$ of H . But since $(a-1, a+1)^-$ of G is down, this shows that $B_{a, k+a-1}$ in G is a blocking set covering $(a, a+1)$. This completes the proof of lemma 2.

Corollary. If $a, b \in G$, with $a < b$, are disconnected, then there exist a', b' with $a \leq a' < b' \leq b$ such that a' and b' are disconnected and there is a blocking set $B_{s,t}$ in G where $s \leq a' < b' \leq t$. In particular B covers part of $[a, \dots, b]$.

Proof. The corollary follows immediately from lemma 2 and the *Observation*.

Lemma 3. If the blocker $B_{i,j}$ is contained in $G \subset L_n$, an n -graph, then node i is disconnected from node $i+1$.

Proof. If $i = 1$, this follows from lemma 1. If $i > 1$, suppose there is a path $\pi: i \rightarrow i+1$. If this path only visited nodes $k \geq i$, we might as well be in the graph $H = G \cap (i, \dots, n)$, with irrelevant edges deleted. That would give a connection between $1_H = i_G$ and $2_H = (i+1)_G$ which violates lemma 1. Otherwise, some edge of π reaches some node $k < i$ for the first time. This must be along edge $(k = i-1, i+1)$ (the directed edge $i+1 \rightarrow i-1$), which is impossible since this is *down* by rule 3), or edge $(k = i-1, i)$ (directed edge $i \rightarrow i-1$). In the latter case, the path will subsequently return to node i . This gives a loop in π , touching only nodes $k \leq i$ which may be deleted, giving a shorter path which eventually gives a contradiction.

Theorem 1. A braided-line network $G \subset L_n$ is disconnected if and only if it contains a blocking set $B_{i,j}$, $1 \leq i < j \leq n$.

Proof. Suppose that G contains a blocking set $B_{i,j}$. Then i is disconnected from $i+1$ by lemma 3, so the network is disconnected. Conversely, if G is assumed disconnected, let a and b be nodes with $a < b$ which have no connection-path. Then the corollary to lemma 2 implies the existence of a blocking set B that covers part of the interval $[a,b]$. This blocking set satisfies the conclusion of the theorem.

4. THE RING GRAPHS

In this section we discuss the similarities and differences between conditions on the braided line and braided ring for (dis-)connectivity. The definition of H , an n -ringed-graph is similar to that of an n -graph. It may be considered to be a subset of the set of edges of R_n , or as an assignment of the value "up" or "down" to each edge of R_n .

The blocking sets of R_n are defined as follows. Consider a consecutive (contiguous) set of edges in R_n . If this set is considered to have a beginning node i and an end-node j , the blocking set $B_{i,j}$ satisfies rules 1), 2) and 3) of section 3. Thus we have edges $(i,i+1), (i+1,i+2), \dots, (j-1,j) \pmod{n}$ as primary edges belonging to $B_{i,j}$. Note that in case $i = j+1 \pmod{n}$, the moons $(j-1,i)^-$ and $(j,i+1)^-$ are down. The case $i = j$ is also possible. In this case all primary edges of the ring R_n are down and belong to the blocking set $B_{i,i+n}$, and the moon $(j-1,i+1)^-$ is down. The remaining case is when the blocking set B has no beginning or end. This implies that every moon of the ring R_n has end-nodes interior to B and must be considered up. For each n , there is only one such blocking set; if H contains such a B , then in fact $H = B$, hence H corresponds to an *outcome*, an event with a single element. All of the primary edges and moons of R_n are specified: the primary edges are down and the moons are up. This may be called the global n -blocking set $G1$. It is easy to see that if n is odd, ≥ 5 , the global blocker gives rise to a connected n -ring graph, whereas for n even, $n \geq 4$, the global blocker is disconnected with two components.

A q -line $L_{i,j}$, $j-i = q \pmod{n}$, can be embedded in R_n for $q \leq n$. Given a blocking set $B_{i,k} \subset H \cap L_{i,j}$ considered as a q -line blocking set, we say that B *extends to a ring blocking set* C , if C is indeed a ring blocking set as just defined, $C \subset H$, and $C \cap L_{i,j} = B_{i,k}$. Similarly if we have $B_{l,j} \subset H \cap L_{i,j}$ or even $B_{i,j}$, then C may be a blocking set with end-nodes, or the global blocker. For the main result on connectedness of a braided ring network, we need a few more lemmas. In Figure 5 we have a blocking set $B_{1,3}$ in $L_{1,4}$ which does not extend to a blocking set in R_6 .

Lemma 4. Consider B_q , a blocker of length q , where q is odd ($= 1 \pmod{2}$), in an n -graph G . Then if $B_q = B_{i,j}$, node s of G disconnects from node t whenever $s \leq i < j \leq t$.

Proof. A path π from s to t consists only of up primary edges and moons, so it must meet node i or node $i+1$. In particular, the directed path beginning at s (choosing the correct orientation) meets the interval $[i, \dots, j]$ for the first time in one of those points, i or $i+1$. This must be i since the moon $(i-1,i+1)^-$ is down. Similarly the path with opposite orientation from t meets $[i, \dots, j]$ for the first time in j . But by rule 2) since q is odd, j is connected to $i+1$. Thus π gives a connection between i and $i+1$, which contradicts lemma 3.

Lemma 5. A blocking set $B_q \subset G$, an n -graph, satisfies the following when q is even ($= 0 \pmod{2}$)

mod 2), and $q \geq 2$. If $B_q = B_{i,j}$, $i < j$, then given $s \leq i$ and $t \geq j$, s disconnects from $i+1$ and t disconnects from $i+1$.

Proof. As in the proof of lemma 4, a path from s to $i+1$ must meet node i which is impossible by lemma 3. By symmetry there is no path from t to $j-1$. But $i+1$ is connected to $j-1$ since q is even, hence there is no path from t to $i+1$ either.

Definition. We say that two ring-blocking sets B^1 and B^2 are *disjoint* if they have no primary edges in common.

Theorem 2. Let $H \subset R_n$ be an n -ringed graph. Then H is disconnected if and only if there are two disjoint blocking sets C^1, C^2 in H , or one blocking set C_q where the length q is even ($q \equiv 0 \pmod{2}$).

Proof. Assume that there exists a blocking set $C^1 = C_{a,b} \subset H$. Let $q_1 = \text{length of } C^1 = b-a$, and consider C^1 as a q_1 -graph in L_{q_1} . Since a and $a+1$ are disconnected in C^1 , if there is a path $\pi: a \rightarrow a+1$, it must enter a node of $H - C^1$ for the first time. The node it reaches is always $a-1$ (resp. $b+1$) (mod n), by rule 3. If the path subsequently enters C^1 at a (resp. b), the path may be shortened by removing this loop. If the path instead enters C^1 at b (resp. a), we obtain a connection between a and b as part of π . But if q_1 is even, that part of π can be replaced by a path entirely within C^1 . Eventually we get a path $a \rightarrow b$ entirely within $C^1 \subset L_{q_1}$, which contradicts lemma 5. In this case H is disconnected. On the other hand, if q_1 were odd, consider $C^2 = C_{f,g}$. If $q_2 = g-f$ is even, H is disconnected by what was just shown. But if q_2 is odd, we just saw a connection $a \rightarrow b$ entirely contained in $H - C^1$. Considering $C^2 \subset H - C^1 \subset L_{n-q_1+1}$, this gives a connection between $1 \in L_{n-q_1+1}$, $a \in R_n$, and $n-q_1+1, b \in R_n$, which is impossible by lemma 4, due to the existence of the odd blocking set C^2 . We conclude that two odd blocking sets or one *even* blocking set disconnect a ring network.

Conversely, suppose that H is disconnected. Then take, by an argument similar to the *Observation* above, nodes a and $a+1$ which are disconnected from one another. Pick any line graph $L_s \subset R_n$ which covers $(a, a+1)$. Since a and $a+1$ are disconnected in L_s , there is a blocking set B covering $(a, a+1)$ in L_s . Now B extends to a blocking set in R_n . We see this, for example, in the case $B = B_{1,k}$ in L_s ; a connection between nodes 1 and 2 of B leads to a connection between a and $a+1$. If the moon $(n, 2)^-$ is down, B satisfies rule 3) on the left. Otherwise the primary edge $(n, 1)$ must be down or a connection $1 \rightarrow 2$ exists. Then continue on the left until

- i) a moon is down, OR
- ii) node k is visited.

If necessary, we extend B to the right as well, obtaining a ring-blocking set $\hat{B} \subset H$. If $\text{length}(\hat{B}) = q$ is even we are done. If $q = n$ is odd, all nodes are connected (whether \hat{B} is a global blocking set or not), and this contradicts the hypothesis. Thus we may assume that $\hat{B} = \hat{B}_{i,j}$ where $(j, j+1, \dots, i) \pmod{n}$ consists of at least two nodes. Clearly j and i are disconnected in H , so the interval $[j, i]$ contains a blocking set $C \subset L_{i-j}$ which may be extended on both sides to a blocking set \hat{C} of R_n .

5. FAMILIES OF DISJOINT EVENTS

In this section, we give a natural decomposition of the probability space (set of outcomes) associated with the connectivity problem for the braided-line and -ring networks. In the following section, this decomposition will be used to compute the all-to-all reliabilities (connectedness probabilities) that we seek. The method involves finding a set of events $\{T\}$ such that

- i) every outcome $\sigma \in T$ corresponds to a connected network,
- ii) every such 'connected' outcome belongs to some such T (exhaustion),
- iii) $i \neq j \Rightarrow T_i \cap T_j = \emptyset$ (disjointness).

Then it follows that if X is the event "the network is connected", we have

$$(5.1) \quad P(X) = \sum_i P(T_i).$$

In section 4, the concept of disjoint blocking set was introduced. In fact, we now prove that disjointness generally holds.

Proposition 1. Let B^1 and B^2 be any blocking sets in a given network $H \subset L_n$ or $\subset R_n$. Then $B^1 = B^2$ or B^1, B^2 are disjoint (have no primary edges in common).

Proof. If B^1 and B^2 are both global blocking sets (in R_n), then they are the same. Thus if $B^1 \neq B^2$, we might as well assume that $B^1 = B_{i,j}$ is not global. Then if we are in L_n , we take $B^2 = B_{f,g}$. If we are in R_n , and B^2 is global = G1, then $(i-1, i+1)^-$ is down because of B^1 , but $(i-1, i+1)$ is up due to G1. No outcome H (specified network configuration) can satisfy these conditions simultaneously, so $B^2 \neq G1$. If $i = f, j = g$, then $B^1 = B^2$ by the identical specification given in rules 1) - 3). Suppose then that $i < f < j$ (the case $i < g < j$ is handled similarly). But then $(f-1, f+1)$ is down due to $B_{f,g}$ and up due to $B_{i,j}$ since $(f-1, f+1)$ is an inner moon unless $f-1 = j-1$ which is ruled out. All other cases are handled by symmetry and lead to the disjointness property which was to be proved.

Recall from section 2 how we characterized the *disconnection* event for the simple line and simple ring. One takes a family of events which is exhaustive and whose members are pairwise mutually exclusive and selects the members corresponding to the 'event' in question, in this case line- or ring- disconnectedness. In fact one may take as this family $F = \{S\}$, where S = some subset of the set of edges of this line (or ring) network. A subset S is chosen if

- i) it contains at least one edge (simple line case), OR
- ii) it contains at least two edges (simple ring case).

The probabilities for the various events were given in section 2. Now we construct analogously the families of events involved in the *braided*-line and -ring 'events'. Let T be a set of disjoint intervals. Each interval is a set of consecutive primary edges and may be written $[g, g+1, \dots, h]$ or $[g, h]$, modulo n in the case of the n -ring. Additionally in the case of the n ring, an 'interval' Ξ representing the entire ring comes about (no beginning or end).

We wish to make correspond to such a set T an event $W \subset F(L_n)$, each element of W being a set of up-down configurations of the n -line, resp. n -ring. This correspondence is given by $\phi: T \mapsto W$, where W is the event consistent with the following two criteria.

(5.2) Correspondence Criteria

- i) if $[g, h] \in T$, the event $F(B_{g,h})$ holds,
- ii) if $[g, h], [k, m] \in T$, and no $[a, b] \in T$ exists, where $h \leq a < b \leq k$, the event: "the sub-graph $L_{h,k}$ of L_n (or R_n) is connected" holds.

In other words, all of the events determined by criteria i) and ii) may be collected and their intersection formed to give $\phi(T) = W$. Also, $\phi(\Xi) = G1$, the event of the global blocking set, which of course consists of a single outcome. The following observation is stated as a theorem.

Theorem 3. The set of events $\{W\}$ is exhaustive and pairwise disjoint, that is, $T_1 \neq T_2 \Rightarrow W_1 \cap W_2 = \emptyset$, and $\bigcup_T \phi(T) = F(L_n)$.

Demonstration. The set is exhaustive since for any outcome $\sigma \in F(L_n)$ or $F(R_n)$, which is essentially an n -graph or n -ringed graph H , we can examine interval to see whether they satisfy the blocking rules 1) - 3). When we find such an interval, remove it and repeat the process on the complementary interval. Finally no further such blocking set will be found. We are left with a number of blocking intervals and a number of 'complementary' intervals. These complementary intervals satisfy criterion ii) of (5.2) above, since any interval which contains no blocking set is connected by theorem 1. Next we must show that σ belongs to exactly one such event. Suppose $\sigma \in \phi(T_1) \cap \phi(T_2)$. Without belaboring the obvious, it is clear that if the sets $\{B^1\}$ and $\{B^2\}$ of blocking sets are different, then either

- 1) there exist B_a^1 from W_1 and B_b^2 from W_2 whose primary edges overlap but $B_a^1 \neq B_b^2$. This is impossible by proposition 1. OR,
- 2) some B' from W_1 has all of its primary edges within an 'up' (connected) interval $[i, j]$ according to W_2 .

In the second case, the n -graph H corresponding to the outcome σ , has a blocking set $B \subset [i, j]$. But since $\hat{B} = B \cap [i, j]$ is a blocking set in the n -line $L_{i,j}$, $H \cap [i, j]$ cannot be connected according to lemma 3. We have therefore constructed $\{W\}$, an exhaustive and pairwise mutually exclusive family of events on $F(L_n)$ or $F(R_n)$, as required by the theorem.

6. THE RELIABILITY COMPUTATION

We introduce some notation: $Q(n)$ is the probability that a braided $(n+1)$ -line, with $n+1$ nodes and n links and link failure probability fixed at U , is connected (functioning). Thus L_{n+1} satisfies global reliability with probability $Q(n)$. It is clear that

(6.1)

$$\begin{aligned} Q(0) &= 1 \\ Q(1) &= 1 - U. \end{aligned}$$

We next derive a recursive formula for $Q(n)$. As in the "artificial" derivation of the simple line reliability in section 2, which resulted in formula (2.1), we focus upon a particular primary edge α . In fact we take $\alpha = (n, n+1)$ and perform a case-by-case analysis of the up-and-down status of links associated with α .

If $\alpha = (n, n+1)^-$ is down and the moon $(n-1, n+1)^-$ is down, then the node n is isolated, so this possibility contributes nothing to $Q(n)$. See Figure 6. If α is down while $(n-1, n+1)$ is up, then it is easy to see that L_{n+1} will be connected if and only if L_n is connected. The following contribution is obtained:

$$(6.2) \quad U(1-U) \cdot Q(n-1).$$

The other cases occur when $\alpha = (n, n+1)$ is up. Then clearly if $L_n = [1, \dots, n]$ is all-to-all connected (for a particular up-down choice on all its edges), then so is L_{n+1} (for that same choice on edges common with L_n). The abuse of language explained by the parenthetical remarks will be resorted to without further comment. However, under certain conditions L_n could be disconnected and L_{n+1} still be functioning. According to theorem 1, this can happen only when B , a blocking set of L_n , is no longer a blocking set as a subset of L_{n+1} . But this happens precisely when $B = B_{i,n}$, $1 \leq i \leq n-1$, is a blocker which abuts the end of L_n . Such a blocker *will not* be a blocker in L_{n+1} as long as α is up and $(n-1, n+1)$ is also up, on account of blocking rule 3). Of course if there are other blockers in L_n besides this B , L_{n+1} will also be disconnected, and no contribution to $Q(n)$ will be made. We have obtained an expression

$$(6.3) \quad (1-U)\{Q(n-1) + (1-U) \sum_{i=1}^{n-1} \gamma_i Q(n-i-1)\},$$

$$\text{where } \gamma_i = P(B_{n-i-1, n-1}) = \begin{cases} U^{i+1}(1-U)^{i-1} & i \neq n-1 \\ U^{n-1}(1-U)^{n-2} & i = n-1 \end{cases}$$

Then $Q(n)$ is given by adding (6.2) and (6.3) and using (6.1). It is of course simple to implement these formulas by a computer program.

Next we consider the case of the braided n -ring R_n , having n nodes and n primary links, whose global reliability is written as $Q_R(n)$. It should not cause confusion that $Q(n)$ refers to the braided line with $n+1$ nodes. According to theorem 2, R_n is functioning if and only if

- i) there are no blocking sets, OR
- ii) there is at most one blocking set, and it covers an odd number of primary intervals.

Now it follows from theorem 3 that the following primary intervals lead to disjoint events by the correspondence of section 5, and their union contains all outcomes for which R_n is connected:

- I) the empty interval (leading to *no* blocking sets in R_n),
- II) an interval $[i, \dots, j] \pmod{n}$ where $j-i$ modulo n is odd,
- III) the entire ring R_n when n is odd, leading to the global blocker G1.

We consider the contribution of each of these events to $Q_R(n)$. The contribution from case III) is simplest and yields

$$(6.4) \quad U^n(1-U)^n.$$

Next consider the contribution from II). It is $\sum_{i,j} P(B_{i,j}) Q(n-j+i)$ where $1 \leq i, j \leq n$, and $n-j+i$ is taken modulo n . These number-theoretic obscurities are quickly cleared up if we convert the expression into

$$(6.5) \quad n \cdot \sum_{q=1}^n \gamma_q Q(n-q),$$

$$\text{where } \gamma_q = \begin{cases} U^{q+2}(1-U)^{q-1} & q \neq n \\ U^{q+1}(1-U)^{q-1} & q = n \end{cases}$$

Formula (6.2) follows since there are precisely n distinct intervals $[i, \dots, j]$ with $j-i = q \pmod{n}$ ($1 \leq q \leq n$).

Finally we consider case I), which is handled similarly to the braided-line recursion, with the use of conditional probabilities. Here, as in section 2, we look at a particular primary edge α , and the up-down status of itself and of its neighboring moons. Let α be the primary edge $(n, 1)$. The associated moons, called "wings" are $(n, 2)$ and $(n-1, 1)$. In the first instance, suppose that α is up and both the wings are down. Then every blocker in the line $L_{n-1} = [1, \dots, n]$ remains a blocker in R_n . Since no new blockers are formed, the contribution to $Q_R(n)$ is

$$(6.6) \quad (1-U)U^2 \cdot Q(n-1).$$

If on the other hand, α is up and exactly one wing is up (which can occur in 2 ways), blockers in L_n which abut the "down" wing remain blockers. Blocking sets from L_n which abut the "up" wing do not become blocking sets in R_n . By similar reasoning to the line case above, we obtain a contribution of

$$(6.7) \quad 2 \cdot U^2(1-U)[Q(n-1) + \sum_{q=1}^{n-1} \gamma_q Q(n-1-q)],$$

$$\text{where } \gamma_q = \begin{cases} U^{q+1}(1-U)^{q-1} & q \neq n-1 \\ U^q(1-U)^{q-1} & q = n-1 \end{cases}$$

The trickiest case arises when α and both the wings $(n-1, 1)$ and $(n, 2)$ are up. In this case blockers from L_n abutting either end are no longer blocking sets in R_n . In fact, two blocking sets of this kind can exist in L_n without R_n containing any blockers (by placing one at either end). The contribution is

$$(6.8) \quad (1-U)^3 \cdot [Q(n-1) + \sum_{k=0}^{n-1} \zeta_k Q(n-1-k-j)],$$

$$\text{where } \gamma_k \text{ is defined as in (6.6), but where } \zeta_j = \begin{cases} U^{j+1}(1-U)^{j-1} & j \neq n-k-1 \\ U^j(1-U)^{j-1} & j = n-k-1 \end{cases}$$

The case $n = 12, k = 5, j = 6$ is illustrated in Figure 7. We see that since $B_{12,6}, B_{6,11}$ abut at node 6, the moon $(5, 7)^-$ should not be counted as a "down" wing for both blocking sets. This is the meaning of the exceptional case in the definition of ζ_j .

Now we turn to the cases having α^- down. If one wing is up and the other down, it is remarkable but true that all blockers of L_n remain blockers in R_n . Thus we get a contribution

$$(6.9) \quad 2 \cdot U^2(1-U)Q(n-1).$$

If both wings are up, no "end blocker" (one that includes the node 1 or n) of L_n remains a blocking set with one exception. This is the blocker corresponding to the entire line $[1, \dots, n]$, which turns into the global blocking set GI! However, one must notice that if there are blockers abutting both sides of α , they merge into one large blocker in R_n . Thus we allow blockers abutting either end in L_n but not

- a) pairs of such blockers (one at each end), NOR
- b) the blocker $B_{1,n}$.

This gives a contribution of

$$(6.10) \quad U(1-U)^2 \cdot \{Q(n-1) + 2 \cdot \sum_{q=1}^{n-2} \gamma_q Q(n-1-q)\}.$$

Note the upper limit of the sum; hence the exceptional case in the definition of γ_q never occurs.

Finally, if α is down and both wings are down, this constitutes a blocking set $B_{n,1}$ *per se* so no contribution to case I) is made. Now adding (6.4) through (6.10) gives a recursive expression for $Q_R(n)$ which may easily be implemented on a computer.

Such an implementation was done in the 'C' language, and was run in many cases. See Figure 8 for a comparison of different architectures. The results were compared in several cases to answers given by an exhaustive method. [A.L. White & K. Dotson, personal communication]. A "spanning tree" in a graph is a tree whose nodes are all the nodes of the graph. If any spanning tree is "up", the graph is connected. All spanning trees were enumerated and given as input to a boolean solver, which computed the probability of the union of the events by exhaustive consideration of possible outcomes.

Results found are given in table 1 (Link Unreliability = 0.7)

Table 1.
System Unreliability

	Exhaustive Boolean Method	Blocking Set Method
R_5	0.743739	0.7437395224
R_6	0.822945	0.8229446323

There is agreement to within 5×10^{-7} or about $7 \times 10^{-5}\%$.

In the case of the braided six ring, the exhaustive method consumed several hours of run-time, whereas the current method gave its result without perceptible passage of time.

Solution of a general network reliability problem ought to involve computations which increase in number exponentially with the number of nodes. For a given topology, it is

possible to find a closed polynomial solution once and for all; hence it is meaningful to give the number of operations needed to find this polynomial, which is of degree $2n$ for the braided ring graph.

Thus, suppose that to find $Q(n-1)$ takes $\phi(n-1)$ floating point operations (flops). That is, after $\phi(n-1)$ operations, the value of $Q(n-1)$ has been found and stored for later use, along with that of $Q(i)$, $1 \leq i < n-1$. Then formulas (6.2) and (6.3) give

$$\phi(n) = 2 + \phi(n-1) + 2 + \sum_{i=1}^{n-1} (2i+1)$$

$= 2 \cdot n(n-1) + (n-1) + 4 + \phi(n-1)$ roughly. So in the long run, $\phi(n) \cong \phi(n-1) + 2n^2$. Thus we may expect

$$(6.11) \quad \phi(n) \cong 2 \cdot \sum_{i=1}^n i \cong \frac{2 \cdot n^3}{3},$$

for moderate to large n .

Figure 9 shows how the actual number of flops varies with the size of the braided line. This was done using a MATLAB implementation. Although "start-up costs" make this number larger than suggested by (6.11) for small n , when n becomes larger the actual growth seems to be greater than $O(n^2)$ but less than $O(n^3)$. Similar analysis of the ring computations (6.4) - (6.10) indicate that computational complexity growth (given about n storage locations) is on the order of n^3 or less.

7. APPROXIMATIONS AND GENERALIZATIONS

Asymptotic Analysis

When the link failure probability U is very small, it is possible to get quite accurate estimates for system reliability by performing only a trivial amount of computation. Consider the braided n -line L_n . Let U be small and $1-U$ "close to" 1. Consider an event characterized by some blocking sets and their complementary intervals. The "up" probability for the complementary intervals is roughly equal to 1. Several factors of U and $(1-U)$ enter into the blocking set probability calculation; terms with the fewest factors of U will be dominant. In fact, the blocking sets $B_{1,2}$ and $B_{n-1,n}$ give rise to failure probabilities $\cong U^2$ each. This gives a first-order approximation

$$(7.1) \quad A(n-1) = 1 - 2 \cdot U^2,$$

which is independent of n (!) In reality, making n larger will increase the unreliability. This means that U may no longer be discrepancy is recovered by a second-order approximation. The singleton blockers $B_{i,i+1}$, for $i = 2, \dots, n-2$, contribute about U^3 apiece to unreliability. In addition, $B_{1,3}$ and $B_{n-2,n}$ contribute a like amount leading to

$$(7.2) \quad P_D(L_n) = 2 \cdot U^2 + (n-1) \cdot U^3$$

in our older notation.

For the ring case, in R_n an even 2-blocker $B_{i,i+2}$, $i=1, \dots, n \pmod{n}$ contributes $n \cdot (1-U)U^4$. The pair of abutting 1-blockers $\{B_{i,i+1}, B_{i+1,i+2}\}$, $i=1, \dots, n$ contributes nU^5 , and adding gives nU^4 . This expression could also have been arrived at by noting that four down edges can isolate a single point in this way only (all edges to a certain node are down). The next power which contributes anything is U^6 and therefore we expect $P_D(R_n) \cong nU^4$ to be a good approximation for small U .

A different situation arises when U is large ($1-U$ is small). Then we have an approximation

$$(7.3) \quad 1 - P_D(R_n) \cong M \cdot (1-U)^{n-1},$$

for *reliability*, where M is the number of spanning trees of R_n . Viability of a spanning tree is "rare" on the order of $(1-U)^{n-1}$, but the coincidence of two distinct trees is "negligibly" rarer, therefore we may add the probabilities of these events as if they were disjoint.

As an "application" of this observation we took $n=5$, $U=.99$, computed $P_D(R_n)$ by our algorithm and found reliability to be 1.2×10^{-6} and $M \cong 119.82$. For $U=.999$, $1-P_D(R_n) = 1.245 \times 10^{-10}$ and $M \cong 124.47$. The correct answer is $M=125$ which may be found by the Binet-Cauchy formula [8], p. 145. For $n=11$, $U=.99$, we get $M=86043$ by our approximation when the answer should be 87131. Larger values of U do not avail in this algorithm as "reliability" is approaching machine precision. The program could be written to perform the divisions by $(1-U)$ at suitable junctures, instead of all at once, to allow further test of (7.3) should anyone wish to do so.

Other Networks

The solution technique employed in this article can be adapted to more general classes of networks. Two cases of this are, for the braided n -ring R_n , what is the overall reliability given that a certain link (primary edge or moon) is known to be down or when some particular link is known to be *functioning*. Let us apply these ideas to the calculation of the difference between the reliabilities of the resulting networks when a primary link, or a "moon", respectively, fails. Such knowledge could become a consideration in reconfiguration strategy. For example, the IAPSA architecture [9] embodies fault-tolerant computers, such as the Advanced Information Processing System Fault-Tolerant Computer (AIPS-FTP) developed by Charles Stark Draper Laboratories, as nodes and an Input-Output Mesh network for internode communication. A reconfiguration strategy for the network in the face of two link failures might have to decide which of two links (say one a primary link, one a moon) to repair or replace.

Is the ring network more sensitive to a primary edge failure than to a moon failure?

We indicate the method in case the primary link $\alpha = (n,1)^-$ is down. As in section 6., we consider the three types of interval indicated in I), II), III). Case III), which concerns the global blocker, is similar to the analogous case handled in section 6. Case I), where there are no blocking sets, involves fewer terms, only (6.9) and (6.10), where the primary link α is down. In case II), the case of one blocker, we have to distinguish whether α is *in* the blocking set or in the complementary set. The first instance is routine, and in fact reduces to a line problem. The second case involves finding the probability that a braided line $[1, \dots, k]$ is up, given that $(i,i+1)^-$ is down. Using analysis similar to that needed for (6.3) and (6.7), this is readily found.

In the other case of interest, in view of symmetry we may as well take $\beta = (n, 2)^-$ as down. Without going into details of the implementation, suffice it to say that this problem presents a few peculiarities which confirm the flexibility of the solution method we employ. For instance, no global blockers are possible in this case, so case III) can be ignored. Since no blocker may extend over β , case II) becomes nearly like finding $Q(n-1)$, and for case I), a search is done of formulas (6.6) to (6.10) to see which of them contribute. (They are 6.6, 6.7, and 6.9.)

We let P_R, P_α, P_β denote, respectively, the failure probabilities of the (ordinary) n -ring, the n -ring with a primary edge failed, and the n -ring with a moon failed. When $n = 5, U = .9$, results were compared with the brute-force approach. They were $P_\alpha = .994780(278)$. The extra decimals were provided by our present method. This number was the computed value of P_β as well. This is as it should be since R_5 is a complete graph and $R - \{\alpha\}$ is isomorphic to $R - \{\beta\}$. This holds true for R_6 as well, but not as obviously. The first case where P_α may not equal P_β is for $n = 7$. Results are shown in Figure 10. It is somewhat surprising that the failure of a moon *degrades the reliability more* than the failure of a primary edge.

We observe that $P_\beta - P_R$ is greater than $P_\alpha - P_R$ by about 0.2% for $U = .05$, increasing to 5.5% for $U = .95$. As $U \rightarrow 0$ we expect P_β/P_α to approach 1, since both are dominated by terms $n \cdot U^4$. When $U \rightarrow 1$, we expect $\frac{(1-P_\beta)}{(1-P_\alpha)}$ to go to M_β/M_α , where M is the number of spanning trees in each respective graph. Taking $U = .99$, we get very small reliabilities, but $\frac{(1-P_\beta)}{(1-P_\alpha)} = .962737$. By the Binet-Cauchy formula, $M_\beta = 663$ and $M_\alpha = 689$; thus $\frac{M_\beta}{M_\alpha} = .962264$, which is reasonably close. In general R_α is a more reliable network than R_β because it has more spanning trees (see Table 2).

Finally, consider one approach to analyzing a more general type of network. If in the new network Λ , the edges are like the primary edges $(i, i+1)$ and include in addition (secondary) edges of the form $(i, i+3)$, we might try to define a blocking set as follows. A (blocking) interval $[i, \dots, j]$ has all of its primary edges down, all *internal* secondary edges up, and all secondary edges that *leave* the interval, down. This is a reasonable generalization of our previous construction. But consider Figure 11; both $[2, 4]$ and $[3, 5]$ satisfy this definition of blocking set. But their intersection is a non-empty interval. Thus proposition 1 does not hold. It is harder to apply an analogy of theorem 2, and the combinatorial analysis would seem difficult.

Table 2.
Number of Spanning Trees

Vertices	Trees	Moon Down	Primary Down
5	125	75	75
6	384	224	224
7	1183	663	689
8	3528	1932	2037
9	10404	5576	5984
10	30250	15950	17325
11	87131	45301	49751
12	248832	127872	141696
13	705757	359053	400993
14	1989806	1003574	1128361
15	5581500	2793800	3159800
16	15586704	7749924	8810949
17	43356953	21430143	24477219
18	120187008	59090912	67773152
19	332134459	162519651	187095569
20	915304500	445948800	515120925

REFERENCES

- [1] M.O. Ball, J.S. Provan, "Calculating Bounds on Reachability and Connectedness in Stochastic Networks", *Networks*, vol 13, 1983, pp 253-278.
- [2] D.K. Pradhan, "Fault-tolerant Multiprocessor and VLSI-based System Communication Architectures" in *Fault-Tolerant Computing: Theory and Techniques*, Prentice-Hall, 1986, pp 467-576.
- [3] L. Grnarov, L. Kleinrock, M. Gerla, "A Highly Reliable Distributed Loop Network Architecture", *Dig. 10th Annu. Int. Symp. Fault-Tolerant Comput.*, 1980 Oct.
- [4] E.R. Hafner, Z. Nenadal, "Enhancing the Availability of a Loop System by Meshing", *Proc. Int. Zurich Semin. Digit. Commun.*, 1976 Mar.
- [5] P.A. Jensen, M. Bellmore, "An Algorithm to Determine the Reliability of a Complex System", *IEEE Trans. on Reliability*, vol R18, 1969 Nov, pp 169-174.
- [6] J.S. Provan, "Polyhedral Combinatorics and Network Reliability", University of North Carolina at Chapel Hill Technical Report No. UNC/ORSA/TR-8219, 1981 Dec.
- [7] T. Politof, A. Satyanarayana, "Efficient Algorithms for Reliability Analysis of Planar Networks - A Survey", *IEEE Trans. Reliability*, vol R-35, 1986 Aug, pp 252-259.
- [8] M.N.S. Swamy, K. Thulasiraman, *Graphs, Networks, and Algorithms*, Wiley, New York, 1981.

- [9] G.C. Cohen, C.W. Lee, L.D. Brock, J.G. Allen, 'Design/Validation Concept for an Integrated Airframe/Propulsion Control System Architecture (IAPSA II)', NASA Contractor Report 178084, June 1986.

SIMPLE LINE NETWORK WITH N NODES



Figure 1.

SEVEN-NODE BRAIDED LINE NETWORK




Figure 2.

12-NODE BRAIDED RING NETWORK




Figure 3.

B2,4 Contained in L5




Figure 4a.

Construction of a Blocking Set Covering an Interval




Figure 4b.

B1,3 Contained in L1,4
But is not Blocking in R6




Figure 5.

Node $(n+1)$ is Isolated.




Figure 6.

RING BLOCKING SETS




Figure 7.

Computational Cost




Figure 9. Number of Nodes

Various Ring Networks





Figure 8. Link Failure Probability

Ring Networks with Missing Link



An $n \rightarrow n+3$ Braided Network
with Overlapping Intervals




Figure 11.