High-Reliability Architectures for Networks under Stress

Guy E. Weichenberg Vincent W. S. Chan Muriel Médard
Laboratory for Information and Decision Systems
Massachusetts Institute of Technology

***Abstract*— In this paper, we consider the task of designing a physical network topology that meets a high level of reliability using unreliable network elements. We are motivated by the use of networks, and in particular optical networks, for high-reliability applications which involve unusual and catastrophic stresses. Our network model is one in which nodes are invulnerable and links are subject to failure, and we consider both the case of statistically independent and statistically dependent link failures. Our reliability metrics are the common all-terminal connectedness measure and and the less commonly considered two-terminal connectedness measure. We compare in the low and high stress regimes, via analytical approximations and simulations, common commercial architectures designed for all-terminal reliability when links are very reliable with alternative architectures which consider *both* of our reliability metrics. Furthermore, we show that for independent link failures network design should be optimized with respect to high stress reliability, as low stress reliability is less sensitive to graph structure; and that under high stress, very high node degrees are required to achieve moderate reliability performance.**

***Index Terms*—system design, graph theory.**

I. INTRODUCTION AND MOTIVATION

Network reliability has become an especially important issue, as optical networks are currently being considered for high-reliability applications. For example, when used for the transport of control signals of jet engines and control surfaces, networks need to provide virtually uninterrupted communication.

When network components fail in a benign fashion with small probability, sparsely connected networks, such as those used in most commercial networks today, can provide adequate levels of reliability. This is because in such scenarios, only single failures typically need to be dealt with at any given time. In the event of a catastrophic stress, however, where a large portion of a network has failed, a high degree of connectedness in a network is required to maintain communication, since many links are needed to backup primary communication paths.

We consider networks which are highly connected. The cost of rich connectedness is a secondary issue in local-area networks (LANs) in contrast to wide-area networks (WANs), where connectedness is hampered by the high cost of fiber runs.

The network reliability synthesis problem considered here is the design of a network which achieves a prescribed level of reliability (in a sense to be defined later) under stress situations, while minimizing the number of components used. Most reliability studies to date have focused on the analysis and design of networks, with emphasis on all-terminal reliability, when links are very reliable. This is appropriate when modelling benign component failures due to low stress, such as normal wear of components. However, the design of networks when links are unreliable, owing to high stress, which is addressed in this paper, is interesting for several reasons. In situations where the probability that a network is connected is quite small, some degree of connectedness in the network could still allow for important functions to be carried out, such as relaying emergency signals in times of distress. For example, in an aircraft application, even a small probability of connectedness could allow for adequate time for the aircraft to fail gracefully should it come under catastrophic stress. This work is thus a step towards bridging the gap between theory and practice by providing design insights which are of immediate value in the planning of high-reliability networks.

Furthermore, the independence assumption in the majority of previous work is inappropriate for situations where, for example, LANs found in automobiles and aircrafts are subjected to environmental stresses that cause localized, correlated failures, or when network components share a common piece of equipment. This paper thus explores reliability models which permit statistical dependency among component failures. While the results obtained for such models are preliminary, they do develop

The research in this paper was supported by: Defense Advanced Research Projects Agency, "Robust Architectures for Multi-Service, Multi-Level Reliability, Multi-Level Service and Multi-Priority WDM Local Area Networks", MDA972-02-1-0021; and National Science Foundation ITR/SY, "High Speed Wavelength-Agile Optical Networks", 008963-001.

intuition for the critical factors in reliable network design, and represent a first step towards the formulation of a general design methodology for networks.

The model we will be using in this paper, where nodes are invulnerable and links are vulnerable is relevant to optical networks, and in particular, all-optical networks. In such networks, the highly-reliable passive optics in network nodes are modelled as graph nodes, and fiber links and transmitter/receivers, which are significantly more prone to failures, are modelled as graph edges. In optical networks, lightpath diversity [1] can be used in place of alternate routing to guarantee critical message delivery deadlines. In this scheme, a power-limited optical transmitter splits its transmitted data along multiple disjoint optical paths. The signals from multiple paths are then recombined at the receiver and decoded. An additional benefit of highly connected optical network topologies is that these networks substantially reduce hop counts and thus save on expensive optical switching equipment, such as OXC's.

In this work, we consider both the case of low and high stress. In low stress situations, we assume that link failures occur with probability 0.2 or below and can be statistically dependent. In high stress situations, link failures occur with probability 0.8 or above and can again be statistically dependent. It should be noted that in this latter assumption of high link failure probability, we are not assuming that networks normally operate in this mode. Rather, high link failure probabilities are assumed given that a catastrophic stress has occurred.

While network reliability metrics such as throughput or delay may be relevant to some network applications [2], connectedness measures are useful in situations where network performance is considered satisfactory as long as the network remains connected, or when the network's ability to provide a minimal level of service is of interest. In addition, connectedness is the relevant metric in many high-reliability applications, where capacities of network components are over-designed, such that connectedness of nodes ensures acceptable network performance. We will thus be principally concerned with connectedness measures of a network.

The main contributions of this work are:

- Establishing general graph-theoretic principles for the design of networks under stress for high all- and two-terminal reliability connectedness criteria.
- Comparing analytically the graph-theoretic design criteria for low stress and high stress regimes.
- Studying and comparing numerically common commercial network designs with designs optimized for high stress and low stress regimes.

Some necessary background is provided in Section II. Section III discusses network design when statistically independent link failures are assumed. In this section, we propose and justify a design methodology, and carry out a series of simulations to gain design insights. In Section IV, we consider network design with statistically dependent link failures. We introduce a simple Markov model, and then carry out approximate reliability analyses of special network topologies.

II. Reliability metrics and reliable network topologies

In this work, networks will be modelled as undirected graphs. Two distinct nodes in such a graph are *connected* if there exists a path between the nodes. An undirected graph is *connected* if there exists a path between every pair of distinct nodes. A (minimal) set of edges in a graph whose removal disconnects the graph is a *(prime) edge cutset*. A (minimal) set of nodes which has the same property is a *(prime) node cutset*. The minimum cardinality of an edge cutset is the *edge connectivity* or *cohesion* $\lambda(G)$. The minimum cardinality of a node cutset is the *node connectivity* or *connectivity* $\chi(G)$. Analogous two-terminal metrics are the edge-connectivity $\lambda_{sd}(G)$ and node-connectivity $\chi_{sd}(G)$ with respect to a pair of nodes s and d. The two-terminal edge (respectively, node) connectivity of a graph is the minimum number of edges (respectively, nodes) whose removal disconnects the node pair.

A. Deterministic metrics

Two rudimentary, deterministic, all-terminal reliability criteria are the cohesion and connectivity of the graph underlying a network. An n-node, e-edge graph having maximum cohesion is a *max*-λ graph. Similarly, an n-node, e-edge graph having maximum connectivity is a *max*-χ graph. The following bounds relate connectivity and cohesion to the basic parameters of a graph [3]:

$$\chi \leq \lambda \leq \delta \leq \frac{1}{n}\sum_{i=1}^{n} d_i = 2e/n. \tag{1}$$

Harary has shown [4] that the bounds in (1) can be achieved, through the construction of *Harary graphs*. More refined deterministic criteria for network reliability can also be defined, such as the number of edge or node cutsets of order λ or χ in a *max*-λ or *max*-χ graph, respectively. A graph is *super*-λ if it is max-λ and every edge disconnecting set of order λ isolates a point of degree λ.

An alternative measure of a graph's ability to remain connected is the number of spanning trees it possesses. The characterization of graphs with a maximum number of trees has been solved for sparse graphs when the number of edges is at most $n+3$, and for dense graphs when the number of edges is at most $n/2$ less than that of the complete graph K_n [5–7], the n node graph which has all of its nodes adjacent.

B. Probabilistic metrics

Deterministic reliability metrics sometimes do not provide adequate measure of the susceptibility of networks to disconnection because these metrics do not account for the reliability of network components. Probabilistic reliability criteria, on the other hand, require knowledge of deterministic network properties, in addition to the reliability of network components, and thus yield a more meaningful measure of network reliability. For this reason, this work will primarily be concerned with probabilistic reliability criteria.

Probabilistic reliability metrics require the concept of a probabilistic graph. A *probabilistic graph* is an undirected graph where each node has an associated probability of being in an operational state and likewise for each edge. In probabilistic reliability analyses, networks under stress are modelled as probabilistic graphs.

Almost all approaches to probabilistic reliability analysis have focused on the probability that a subset of nodes in a network are connected when links are very reliable. Thus, the all-terminal reliability of a probabilistic graph can be defined as the probability that any two nodes in the graph have an operating path connecting them. If links fail in a statistically independent fashion with probability p, then the all-terminal reliability $P_c(G,p)$ is given by:

$$
\begin{aligned}
P_c(G,p) &= \sum_{i=n-1}^{e} A_i (1-p)^i p^{e-i} && (2)\\
&= 1 - \sum_{i=\lambda}^{e} C_i p^i (1-p)^{e-i} && (3)
\end{aligned}
$$

where A_i denotes the number of connected subgraphs with i edges, and C_i denotes the number of edge cutsets of cardinality i. For values of p sufficiently close to zero, $P_c(G,p)$ can be accurately approximated by $1 - C_\lambda p^\lambda (1-p)^{e-\lambda}$. In this case, an optimally reliable graph — one that achieves the maximum $P_c(G,p)$ over all graphs with the same number of nodes and edges — has a minimum number of cutsets of size $\lambda = \lfloor 2e/n \rfloor$. Therefore, in this regime of p, optimally reliable graphs are super-λ graphs. For values of p sufficiently close to unity, $P_c(G,p)$ can be accurately approximated by the first term in (2), $A_{n-1}(1-p)^{n-1}p^{e-n+1}$, where $A_{n-1} = t(G)$. Therefore, for values of p sufficiently close to unity, an optimally reliable graph has a maximum number of spanning trees.

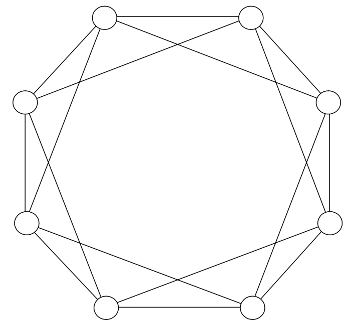

Fig. 1. The $H(8,4)$ Harary graph.

The two-terminal reliability of a probabilistic graph is the probability that a given pair of nodes, s and d, have an operating path connecting them:

$$
\begin{aligned}
P_c^{sd}(G,p) &= \sum_{i=w_{sd}}^{e} A_i^{sd} (1-p)^i p^{e-i} && (4)\\
&= 1 - \sum_{i=\lambda_{sd}}^{e} C_i^{sd} p^i (1-p)^{e-i} && (5)
\end{aligned}
$$

where w_{sd} is the shortest path length between nodes s and d, A_i^{sd} is the number of subgraphs with i edges that connect nodes s and d, λ_{sd} is the minimum number of edge failures required to disconnect nodes s and d, and C_i^{sd} is the number of cutsets with respect to nodes s and d of cardinality i. If we wish to maximize $\min_{s,d}\left[P_c^{sd}(G,p)\right]$ when $p \approx 0$, then it is apparent from (5) that the property of super-λ is a necessary condition. This is because $\lambda = \min_{s,d}[\lambda_{sd}]$, and for super-$\lambda$ graphs, $C_{\lambda_{sd}}^{sd}$ attains the minimum bound of two.

C. Harary graphs and circulants

As previously mentioned, Harary graphs, first presented in [4], achieve the bounds in (1). This result implies that Harary graphs also achieve the maximum value of $\min_{s,d}[\lambda_{sd}]$ and $\min_{s,d}[\chi_{sd}]$ over all graphs with n nodes and e edges. In a $H(n,\Delta)$ Harary graph where Δ is even, each node i, $0 \le i \le n-1$, is adjacent to nodes $i \pm 1, i \pm 2, \ldots, i \pm \lfloor \Delta/2 \rfloor (\bmod n)$; and if Δ is odd, then each node $i = 1, \ldots, \lfloor (n-1)/2 \rfloor$ is also adjacent to node $i + \lfloor n/2 \rfloor$. See Figure 1 for an example of a Harary graph.

Harary graphs belong to a more general family of graphs known as *circulants*. The circulant graph $C_n\langle a_1, a_2, \ldots, a_h\rangle$, or more compactly, $C_n\langle a_i \rangle$, where $0 < a_1 < a_2 < \ldots < a_h < (n+1)/2$, has $i \pm a_1, i \pm a_2, \ldots, i \pm a_h (\bmod n)$ adjacent to each node i. Owing to a theorem by Mader [8], which proves that every connected node-symmetric[1] graph has $\lambda = \Delta$, all connected circulants are max-λ. Furthermore, the only circulants which are not super-λ are the cycles and the graphs $C_{2m}\langle 2, 4, \ldots, m-1, m\rangle$ with $m \geq 3$, and m an odd integer [9].

In [10], Wang and Yang derive a useful result for the number of spanning trees in circulant graphs. In [9], Boesch and Wang examine the diameter properties of circulants and derive lower diameter bounds for the family of graphs. In [11], the same authors determined that even degree Harary graphs possess the fewest number of edge cutsets of cutset cardinality i, when $\lambda \leq i \leq 2\Delta - 3$. Each cutset in the above range of cardinalities was shown to isolate a single node in the Harary graph.

D. Cages and Moore graphs

We now discuss regular graphs which, for a given number of nodes and edges, achieve maximum girth. The problem of finding such graphs is equivalent to the well-studied *cage* problem — finding regular graphs of degree Δ and girth g with the minimum number of nodes $n(\Delta, g)$. The search for cages with degrees exceeding three and girths exceeding five has proven to be very difficult with few results obtained.

Any graph which achieves the *Moore bound*, a lower bound for $n(\Delta, g)$, is known as a *Moore graph*. Moore graphs are, by definition, cages. A well-known property of Moore graphs is that they have minimum diameter k, which grows as the logarithm of n, over all regular graphs of the same degree having the same number of nodes. See Figure 2 for a diagram of the Moore graph with $g = 5$ and $\Delta = 3$, also known as the Petersen graph.

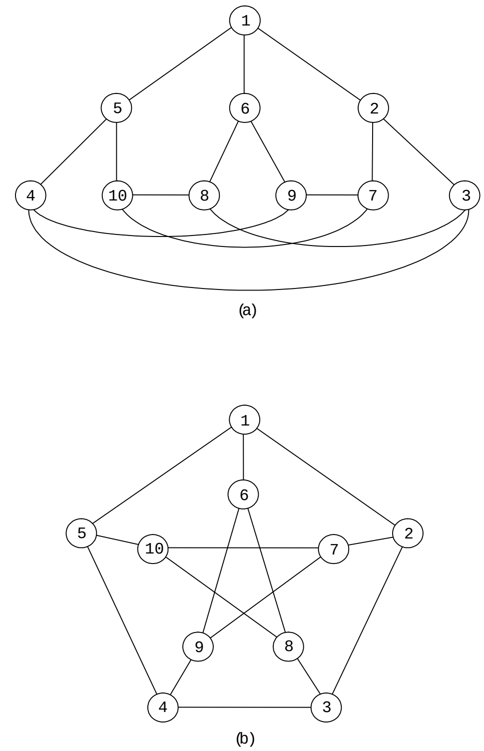

Fig. 2. Two representations of the $g = 5, \Delta = 3$ Moore graph, also known as the Petersen graph. The upper diagram (a) is the full tree representation using node 1 as the root node. For any Moore graph, a full-tree representation using any node as the root is possible.

III. Network design with statistically independent link failures

In this section, we model networks as probabilistic graphs with the following properties:

- Nodes are invulnerable;
- Edges fail in a statistically independent fashion with probability p;
- Edge capacities are assumed to be sufficiently large to carry any possible network flow;
- Once an edge fails it cannot be repaired.

A. Design of reliable networks

Ideally, a network design methodology should appeal to a single, simple family of graphs for all possible network configurations. The family of circulant graphs is an ideal candidate for such a reliability methodology for a number of reasons. The circulant family of graphs is rich — a circulant graph can be defined for most combinations of number of nodes and degree. In addition, circulants inherently possess good reliability properties. For example, in our discussion of circulants in Section II-C, we indicated that nearly all circulants are super-λ. In addition, in a recent work by Sawionek, Wojciechowski and Arabas [12], the family of circulant graphs were shown to *most probably* contain a uniformly optimally reliable graph when such a graph exists, except for when $e \leq n + 3$. Figure 3 summarizes our design results.

1) Designing for all-terminal reliability when p is low: When p is low and we would like to design a network for a prescribed level of all-terminal reliability, then we know that the class of optimal graphs is restricted to those that are super-λ. Intuitively, this is because super-λ graphs minimize the number of most likely disconnection scenarios. In [13], Bauer et al. derive an explicit bound on p for which super-λ graphs are optimal. In [13], Bauer et al. also derive somewhat complicated conditions which ensure that $P_c(G,p) > (1+\epsilon)P_c(\widehat{G},p)$.

Within the class of super-λ graphs, even degree Harary graphs were shown to be especially good when p is low, since they achieve the fewest number of cutsets of cardi-

[1]Two nodes u and v in a graph are *similar* if there is an automorphism which maps u onto v. A graph in which all nodes are similar is *node-symmetric*.

	All-terminal reliability	Two-terminal reliability
low p	Super-λ *Harary graphs, other super-λ graphs*	Super-λ *Harary graphs, Moore graphs, other super-λ graphs*
high p	Max. number of trees *Max. tree circulants*	Min. diameter *Moore graphs, min. diameter circulants*

Fig. 3. Summary of design results. The top line in each quadrant is a necessary condition for optimality with respect to the corresponding vulnerability region and reliability metric. The lines below are the types of graphs suggested by our methodology.

nality i, when $\lambda \leq i \leq 2\Delta - 3$. Thus, if we are principally concerned with all-terminal reliability in the low p regime, then we should design networks as Harary graphs. In [14, 15], we derive several new results for the family of Harary graphs which allow us to develop closed form bounds for all-terminal reliability which are tight when p is low.

2) Designing for two-terminal reliability when p is low: We now consider the task of designing a network with n nodes which meets an objective value of $\min_{s,d}\left[P_c^{sd}(G,p)\right]$ when p is low. As in the all-terminal case, a necessary condition for optimality with respect to two-terminal reliability when $p \approx 0$ is the super-λ property. Furthermore, for even degree Harary graphs any (not necessarily prime) cutset of cardinality i for for $\lambda \leq i \leq 2\Delta - 3$ isolates either s or d alone. Hence, Harary graphs are a good design choice when two-terminal reliability is of principal interest. The derivation of tight, closed form bounds for the two-terminal reliability of Harary graphs when p is low is also derived in [14, 15].

3) Designing for all-terminal reliability when p is high: As discussed in Section II-B, when we are interested in optimizing the design of an n-node network with respect to all-terminal reliability in the p high regime, we seek an architecture which maximizes the number of spanning trees. Intuitively, this is because graphs with a maximum number of trees maximize the number of most-likely graph connection scenarios. In [14], we derive an upper bound on the number of spanning trees for an n-node, e-edge graph, which can be used to obtain an estimate of the required degree Δ. After determining an estimate for Δ from this bound, we search the finite space of n node circulants with degree Δ for the configuration with the largest number of spanning trees. The number of spanning trees of a circulant is easily computed using Wang and Yang's result in [10].

We note that there seems to exist a relationship between a graph's diameter and its number of spanning trees, although the precise relationship is unclear. In most instances, regular graphs with small diameters have a large number of spanning trees. However, in general, a smaller diameter does not imply a larger number of spanning trees, or vice versa. The intuition behind this trend is that for the same number of nodes and edges, the nodes of a symmetric graph with a larger diameter are generally more distant from one another. The result is that there are fewer combinations of edges of the graph that could form spanning trees since there are more constraints on the edges in order that more distant nodes be connected. Hence, the number of spanning trees generally decreases with diameter when the number of nodes and edges is held constant. Therefore, if we wish to design a network with a large number of spanning trees, it is reasonable to alternatively design a network with a small diameter (which is the figure of merit when designing for two-terminal reliability when p is high). Thus, if a configuration for a minimum diameter circulant is readily available, an exhaustive search over all candidate circulant graphs could be avoided.

4) Designing for two-terminal reliability when p is high: We now consider the task of designing a network with a constraint on the two-terminal reliability metric $\min_{s,d}\left[P_c^{sd}(G,p)\right]$ when p is high. A simple lower bound for $\min_{s,d}\left[P_c^{sd}(G,p)\right]$ is:

$$(1-p)^{k(G)} \leq \min_{s,d}\left[P_c^{sd}(G,p)\right] \qquad (6)$$

which is just the probability that the shortest path between the most distant node pair is available.

Using this inequality, we first determine a value for the diameter k. The value chosen for Δ should be as small as possible, while still sufficiently large to ensure that a circulant with the specified values of n, k and Δ can be constructed. The relationship among n, k and Δ for circulant graphs was investigated in [9] by Boesch and Wang. In [14], we show that in the best case, the diameters of even degree Δ circulants grow as the $\left(\frac{\Delta}{2}\right)^{\text{th}}$ root of the number of nodes n; and in the best case, the diameters of odd degree Δ circulants grow as the $\left(\frac{\Delta-1}{2}\right)^{\text{th}}$ root of the number of nodes n. On the other hand, we recall from our discussion in Section II-D that the diameters of Moore graphs grow with the logarithm of the number of nodes n. However, for networks of 50 nodes or less the difference in the minimum degree required when the diameter is held constant, is usually zero or one and occasionally two. Furthermore, recall that with the exception of a few configurations, Moore graphs are not realizable. We therefore

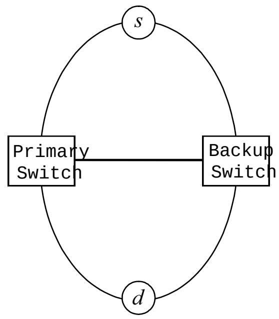

Fig. 4. Dual-homed switch topology (Ethernet).

conclude that circulant graphs which achieve the diameter bounds in [14] are optimal or nearly optimal with respect to two-terminal reliability when the number of nodes is on the order of tens, which is the case for most networks of interest.

B. Simulation results

1) Commercial networks versus our candidate topologies: We now conduct a comparison among Harary graphs — one of our candidate topologies — and some topologies employed in commercial networks — dual-homed switch graphs, rings, and multi-rings.

The dual-homed switch architecture is illustrated in Figure 4. In this topology, each node is connected to a primary and a secondary switch through a dedicated link. In addition, the two switches are bridged. Communication between a node pair, although normally first attempted through the primary switch, can be carried out via any available path. Switched Ethernet is a very common example of the dual-homed switch architecture [16], and we will therefore refer to the dual-homed switch architecture simply as Ethernet. In an m multi-ring graph, there are m undirected edges between nodes that would otherwise have one undirected edge in a regular ring graph.

In our comparison, each graph supports 14 nodes and the degree of the multi-ring and the Harary graph is four. We further assume that nodes, including the two switches in the Ethernet topology, are invulnerable, and that the Ethernet bridge reliability is identical to that of the other links in the network.

Figure 5 depicts the performance of the topologies when $p \leq 1/2$. Between Ethernet and the ring, which are the degree two topologies, Ethernet exhibits better all- and two-terminal reliability. Ethernet's superior performance can be attributed to the fact that it scales weakly with the number of nodes in the graph. For example, for all-terminal reliability, the number of cutsets of order two is $n = 14$ in Ethernet, whereas it is $\binom{n}{2} = 91$ in the ring. Similarly, for two-terminal reliability, the number of cutsets of order two is two in Ethernet, whereas it is

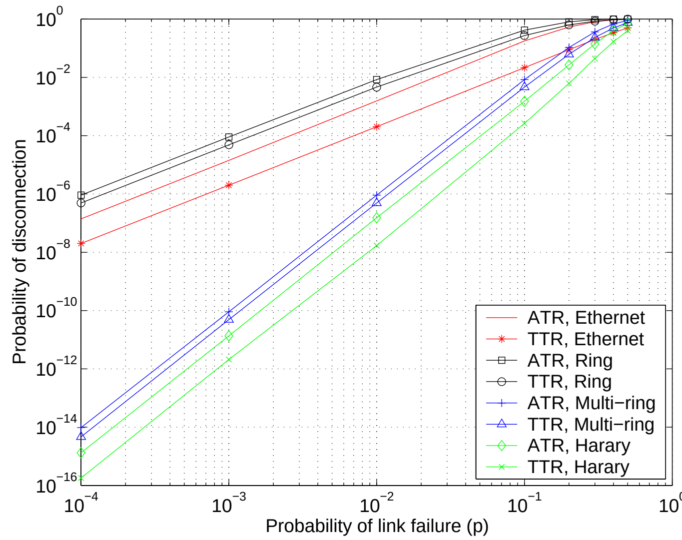

Fig. 5. Probability of disconnection versus p for the 14 node Ethernet, ring, double-ring and $H(14,4)$ graphs when $p \leq 1/2$.

$n^2/4 = 49$ in the ring. The same scalability explanation also applies when accounting for the superior performance of $H(14,4)$ relative to the double ring, which is also a degree four graph. With respect to all-terminal reliability, $H(14,4)$, since it is super-λ, possesses $n = 14$ cutsets of order four, whereas the double ring possesses $\binom{n}{2} = 91$ cutsets of order four. For two-terminal reliability, the number of cutsets of order two is two in $H(14,4)$, whereas it is $n^2/4 = 49$ for the double ring.

In Figure 6, the performance of the topologies is plotted when $p \geq 1/2$. With respect to all-terminal reliability, it is easy to see that Ethernet has far more spanning trees than the ring, which only has $n = 14$, thus accounting for its superior reliability performance. Similarly, $H(14,4)$ has 1.9898×10^6 spanning trees, whereas the double ring has $n2^{n-1} = 1.1469 \times 10^5$ spanning trees. Hence, we expect $H(14,4)$ to perform better than the double ring, which is indeed the case. With respect to two-terminal reliability, the performance difference between Ethernet and the ring is great. This is because Ethernet has a diameter of two, whereas the ring has a diameter of $\lfloor n/2 \rfloor = 7$. The two-terminal reliability difference between $H(14,4)$ and the double ring is also significant, owing to the fact that $H(14,4)$ has a diameter of four, whereas the double ring has a diameter of $\lfloor n/2 \rfloor = 7$.

We conclude that the reliability of rings is consistently poorer than that of the Ethernet topology. Of course, the price paid for this superior reliability is the cost of the switches. We also conclude that multi-rings have poor reliability performance relative to super-λ circulants of the same degree, such as Harary graphs. This indicates that there is a significant reliability advantage in strategic positioning of link capacity rather than adding redundant

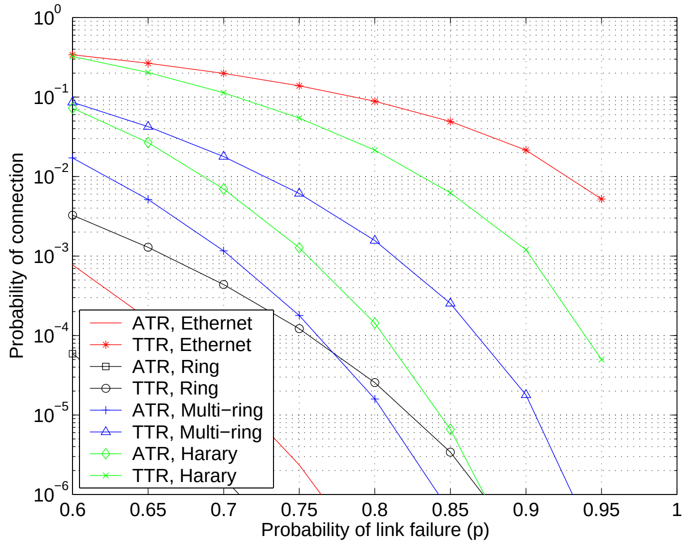

Fig. 6. Probability of connection versus p for the 14 node Ethernet, ring, double-ring and $H(14,4)$ graphs when $p \geq 1/2$.

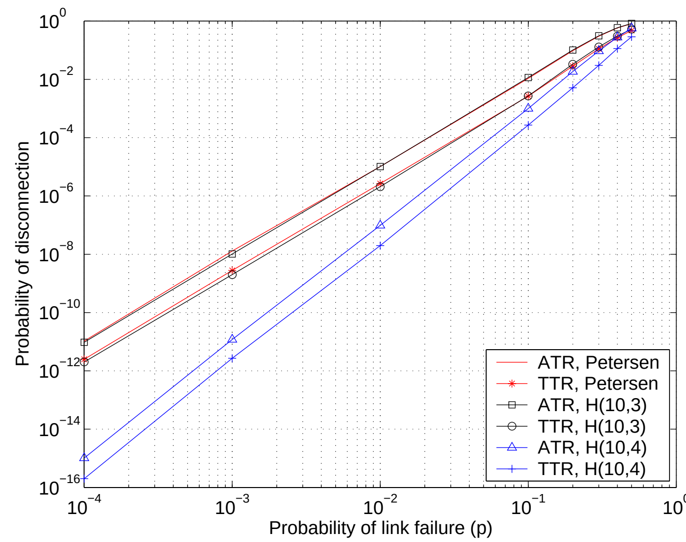

Fig. 7. Probability of disconnection versus p for the Petersen graph, $H(10,3)$ and $H(10,4)$ when $p \leq 1/2$.

backup links. This justifies our pursuit of alternative network topologies with high degrees of connectedness in order to achieve high levels of reliability. We next conduct simulation comparisons among a variety of such topologies.

2) Comparison among candidate topologies: In this section, we present simulation results for several network designs. These results verify our previous insights, and also shed light on the relative performance of different network configurations.

In our first set of simulations, we consider the Petersen graph (the Moore graph with $g = 5$ and $\Delta = 3$), and the Harary graphs, $H(10,3)$ and $H(10,4)$. When p is low, Figure 7 indicates the expected result that $H(10,4)$ possesses a lower probability of disconnection by a factor of approximately p relative to the Petersen graph and $H(10,3)$. Perhaps an unexpected finding is the closeness of the performance of the Petersen graph and $H(10,3)$ when p is low. In fact, all- and two-terminal reliability can be well-approximated by np^{Δ} and $2p^{\Delta}$, respectively, when p is low. Thus, with respect to all- and two-terminal reliability when p is low, the sparse family of Moore graphs offers little or no benefit over the richer family of super-λ graphs. It is only when two-terminal reliability in the p high regime is of interest that Moore graphs present a significant advantage over other competing topologies as they possess smaller graph diameters. Thus, when designing a network topology we should focus on optimizing the network structure with respect to high stress reliability, as low stress reliability is virtually unchanged provided that the underlying graph is super-λ, which is the case for nearly all circulants.

In our next set of simulations, we investigate the effect of node degree on reliability in the $p \geq 1/2$ regime. Specifically, we are interested in determining the node degrees required to achieve all- and two-terminal reliabilities in the useful range of 0.1 to 1. For our simulations, we consider a variety of 14 node circulants. Figure 8 depicts the all-terminal performance of these graphs in the $p \geq 1/2$ regime. As expected, the all-terminal reliability increases with node degree. Another observation is that the performance difference among graphs with the same node degree is more pronounced at lower node degrees than at higher node degrees. Intuitively, this is because structural changes in sparser graphs can more dramatically affect the relative reliability properties of graphs than in denser graphs. Unfortunately, these simulations also indicate that to achieve all-terminal reliabilities in the range of 0.1 to 1 when $p \geq 1/2$, very high node degrees are required. In fact, when p exceeds roughly 0.87, even the complete graph K_{14} cannot achieve a reliability above 0.1. Furthermore, in line with our previous observation, once we realize that a high node degree is required to achieve a reliability in the range of 0.1 to 1, the graph's actual structure is not very important.

Our simulation results for the two-terminal reliability of the same seven graphs in the $p \geq 1/2$ regime are illustrated in Figure 9. The trends observed in this figure are similar to those discussed above. In fact, for two terminal-reliability, these trends are even more apparent. For example, the performance difference of graphs with the same node degree is quite significant at lower degrees, while the performance of the graphs is virtually indistinguishable at higher degrees. Intuitively, this is because topological idiosyncrasies (i.e. diameter) of graphs can be magnified in a graph's two-terminal reliability figure since the connect-

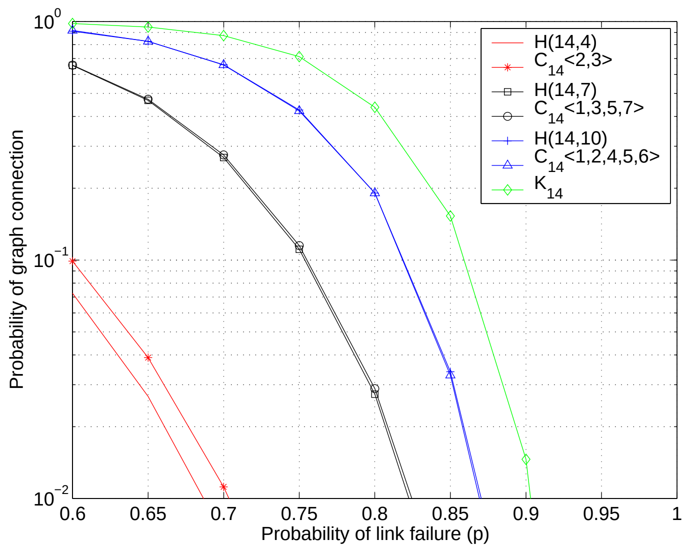

Fig. 8. Probability of graph connection versus p for $H(14,4)$, $C_{14}\langle 2,3\rangle$, $H(14,7)$, $C_{14}\langle 1,3,5,7\rangle$, $H(14,10)$, $C_{14}\langle 1,2,4,5,6\rangle$ and the complete graph K_{14}, when $p \geq 1/2$.

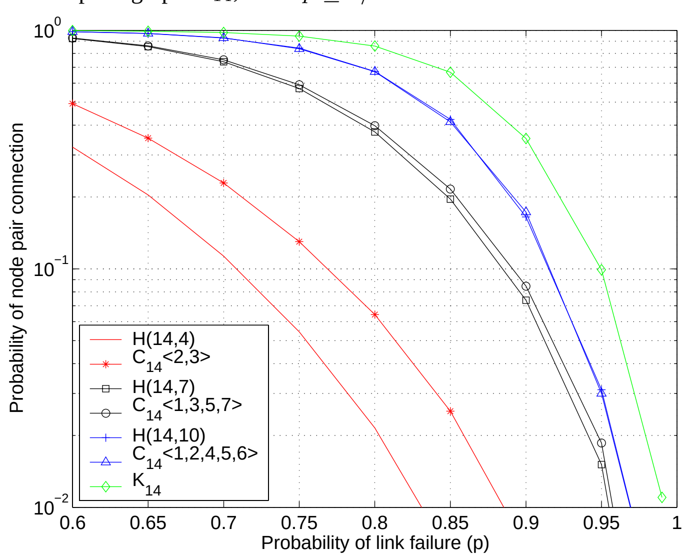

Fig. 9. Worst-case probability of node pair connection versus p for $H(14,4)$, $C_{14}\langle 2,3\rangle$, $H(14,7)$, $C_{14}\langle 1,3,5,7\rangle$, $H(14,10)$, $C_{14}\langle 1,2,4,5,6\rangle$ and the complete graph K_{14}, when $p \geq 1/2$.

edness of the worst node pair is only considered; whereas all-terminal reliability is a global connectedness measure. The simulation results also indicate that two-terminal reliabilities above 0.1 can be achieved at low to moderate node degrees. For example, the minimum diameter circulant $C_{14}\langle 2,3\rangle$ of degree four, achieves two-terminal reliabilities above 0.1 when p is approximately less than 0.75.

IV. Network Design with Statistically Dependent Link Failures

As discussed in the introductory section, many situations arise for which the modelling assumption that network links fail in a statistically independent fashion is inappropriate. As we shall see, modelling link failures in a statistically independent fashion can lead to dangerously optimistic conclusions regarding the reliability of a network.

In this section, we carry out simple, approximate reliability analyses of special network topologies based on existing dependent link failure models as well as a new Markov model introduced here. Unfortunately, the different assumptions used in each of these models preclude a detailed comparison among these topologies, except when small correlation among link failures is present. These models, however, may be applied in comparisons among graphs belonging to the same family.

A. *Markov model*

In order to illustrate our Markov model, consider m links which are of interest, $l_1, l_2, \ldots, l_m$. Now, let us assume that there is a Markovian failure dependency among these m links; that is, conditioned on the state of link $j-1$, link j is independent of the states of links $1, 2, \ldots, j-2$. Let l_j denote the event that link j is operational, and let $\bar{l}_j$ denote that link j is not operational. Let us further assume that the marginal probability distributions of the states of each of the links is identical (i.e. $\Pr\left(\bar{l}_j\right) = p$), and that $\Pr\left(\bar{l}_j|\bar{l}_{j-1}\right)$ is also identical for all j. Thus, the probability that all m links have failed is given by:

$$\Pr\left(\bar{l}_1\bar{l}_2\ldots\bar{l}_m\right) = p\left[\Pr\left(\bar{l}_j|\bar{l}_{j-1}\right)\right]^{m-1}.$$

Alternatively, if only the correlation coefficient ρ of the states of adjacent links is available, then the probability that all m links fail can be shown to be:

$$\Pr\left(\bar{l}_1\bar{l}_2\ldots\bar{l}_m\right) = p\left[\rho(1-p)+p\right]^{m-1}. \tag{7}$$

Note that if p is low and $\rho \gg p$, then the probability that all m links have failed is approximately equal to $p\rho^{m-1}$.

B. *Reliability of the Ethernet graph*

We now compute the reliability of the Ethernet graph when dependence among link failures is present. We assume that nodes, including switches, are invulnerable, and that link failures are correlated only if the corresponding links are incident at the same non-switch node. Using this dependent failure model, the probability of connection of the graph is:

$$\begin{aligned} P_c(G,p,p_b) = {} & \left[2(1-p)-(1-p)\left(\Pr\left(l_j|l_{j-1}\right)\right)\right]^n(1-p_b) \\ & + \left(2(1-p)^n - \left[(1-p)\left(\Pr\left(l_j|l_{j-1}\right)\right)\right]^n\right)p_b, \end{aligned}$$

and the two-terminal probability of connection is:

$$P_c^{sd}(G,p,p_b) = 2(1-p)^2 + 2(1-p)^2(1-p_b) - 4(1-p)^2(1-p_b)\left[\Pr\left(l_j|l_{j-1}\right)\right] - (1-p)^2\left[\Pr\left(l_j|l_{j-1}\right)\right]^2 + 2(1-p)^2(1-p_b)\left[\Pr\left(l_j|l_{j-1}\right)\right]^2.$$

where p_b denotes the probability of failure of the bridge joining the switches.

C. *Reliability of ring and multi-ring graphs*

In this subsection, we we develop closed-form expressions for the all- and two-terminal reliability of rings and multi-rings, assuming that nodes are invulnerable and that links fail in a statistically dependent fashion in accordance with our Markov model.

1) Ring graph: Assuming that nodes are invulnerable, the probability that a ring remains connected is the probability that zero links or exactly one link fails in the ring. These two probabilities can be computed using a chain rule expansion along consecutive links around the ring. Note that the state of the final link in the expansion is influenced by its neighboring links on either side, rather than by just one link. This last probability term must therefore be specified in order to complete the model. For an n node ring, the probability of graph connection can thus be expressed as:

$$P_c(G,p) = (1-p)\left[\Pr\left(l_j|l_{j-1}\right)\right]^{n-2} \left[1 + (n-1)\Pr\left(\bar{l}_j|l_{j-1}, l_{j+1}\right)\right].$$

To compute the two-terminal reliability, we note that the probability that a node pair remains is connected is equal to the probability that all of the links on at least one of the two disjoint paths between the nodes remain operational. Hence, for a diameterically-spaced pair of nodes on an n ring graph, the two-terminal probability of connection is:

$$P_c(G,p) = ((1-p)\left[\Pr\left(l_j|l_{j-1}\right)\right]^{\lfloor n/2\rfloor - 1} + (1-p)\left[\Pr\left(l_j|l_{j-1}\right)\right]^{\lceil n/2\rceil - 1} - (1-p)\left[\Pr\left(l_j|l_{j-1}\right)\right]^{n-2}\Pr\left(l_j|l_{j-1}, l_{j+1}\right).$$

2) Multi-ring graph: We now generalize the above analysis to multi-rings. As in the independent failure model, we only need to replace the parameter p in the above equations with a parameter which reflects the probability of the m parallel links failing in an m multi-ring. We may incorporate statistical dependence into this parameter by using our Markov model to replace p by $p\left[\Pr\left(\overline{l'}_j|\overline{l'}_{j-1}\right)\right]^{m-1}$. Note that the conditional probability in this expression for parallel links is different from the previous conditional probability for consecutive links in the ring.

D. *Reliability of Harary graphs when p is low*

In this subsection, we state approximate expressions for the all- and two-terminal reliability of Harary graphs when link failure dependencies are present. We use the basic idea of the $\epsilon-$model presented in [17] in conjunction with the Harary graph analysis developed in [15]. In [15], we note that every graph disconnection scenario can be viewed as a partitioning of the graph into two subsets of j and $n-j$ nodes S_j and S_{n-j}, respectively, which are disconnected; and that a partition of j consecutive nodes minimizes the number of edges joining S_j to S_{n-j}. Since the edges joining S_j to S_{n-j} are in "closest" proximity when the nodes in S_j are consecutive, we reason that a conservative estimate for the reliability of Harary graphs can be obtained by treating each possible S_j as a consecutive partition of nodes. We cannot rigorously state that such an estimate would be a lower bound for the probability of graph connection because in order to do so, we would need a complete probability distribution for the states of all links in the graph.

When $p_{\hat{\epsilon}} = p\hat{\epsilon} \approx 0$, where a link is $\hat{\epsilon}$ more likely to fail when one or more of the links joining S_j and S_{n-j} have failed, the probability of graph and node pair disconnection for Harary graphs is approximately $\frac{n}{\hat{\epsilon}}p_{\hat{\epsilon}}^{\Delta}$ and $\frac{2}{\hat{\epsilon}}p_{\hat{\epsilon}}^{\Delta}$, respectively. The analogous expressions for the independent failure model are np^{Δ} and $2p^{\Delta}$, respectively. In order to get a feeling for the difference in these two sets of expressions, let us consider a 20 node, degree four Harary graph with probability of link failure 10^{-2} and $\hat{\epsilon} = 5$. Our dependency model yields the values 2.5×10^{-5} and 2.5×10^{-6} for the all- and two-terminal probabilities of disconnection, respectively. On the other hand, the independence model yields the values 2×10^{-7} and 2.5×2^{-8}, respectively.

For large values of $\hat{\epsilon}$, the above asymptotic expression are no longer good estimates of the all- and two-terminal reliabilities. In the limit of $\hat{\epsilon} = p^{-1}$ (equivalently, $\rho = 1$), these estimates do not approach p, which is the expected probability of disconnection. We attribute the diminishing accuracy of the asymptotic expressions to the fact that these estimates are union bounds on prime failure events. As $\hat{\epsilon}$ increases, the probability that multiple prime failure events simultaneously increases, thereby making the union bound loose.

E. *All-terminal reliability when p is high*

In this subsection, we state a simple, approximate expression for the all-terminal reliability of a graph when p is high. Recall that when p is high, the probability that a graph remains connected is approximately equal to the probability that the operational links in the graph form a spanning tree. In the independent failure model, each of the $t(G)$ spanning trees of a graph has a probability of $(1-p)^{n-1}p^{e-n+1} \approx (1-p)^{n-1}$ of occurring. However, when statistical dependence among link failures is present, the probabilities of occurrence of the different spanning trees are in general not the same, as they depend upon the exact structure of the spanning trees. Nevertheless, we can approximate the all-terminal reliability by assuming that links remain operational in a similar manner as assumed for Harary graphs in Section IV-D. That is, the probability of a spanning tree occurring is given by $(1-p)\left[\Pr\left(l_j|l_{j-1}\right)\right]^{n-2}$. Hence, the all-terminal reliability of a graph can be approximated by:

$$P_c(G,p) \approx t(G)(1-p)\left[\Pr\left(l_j|l_{j-1}\right)\right]^{n-2} \tag{8}$$

where we recall that $\Pr\left(l_j|l_{j-1}\right)$ denotes the probability that link j is operational given that an adjacent link $j-1$ is operational.

As ϵ increases to values near $(1-p)^{-1}$, the all-terminal reliability estimate exceeds unity, whereas it should approach $(1-p)$. Thus, (8) is a reasonable all-terminal reliability estimate for only small values of ρ. The diminishing accuracy of the estimate as ρ increases is expected, since (8) is a union bound on the spanning tree events. As ρ increases, we are increasing the probability of occurrence of each spanning tree event, and the union bound becomes looser because the probability of occurrence of multiple spanning tree events is no longer insignificant.

F. Two-terminal reliability when p is high

When p is high and we are interested in the two-terminal reliability of a graph, we use a variation of the simple bound stated in Section III-A.4:

$$\min_{s,d}\left[P_c^{sd}(G,p)\right] \geq (1-p)^{k(G)}$$

Using our Markov model along the shortest path between the worst-case node pair, the above expression becomes:

$$\min_{s,d}\left[P_c^{sd}(G,p)\right] \approx (1-p)\left[\Pr\left(l_j|l_{j-1}\right)\right]^{k(G)-1}$$

where, again, $\Pr\left(l_j|l_{j-1}\right)$ denotes the probability that link j is operational given that link $j-1$ is operational.

As an example, let us consider a 20 node, degree four Harary graph with probability of link operation 10^{-2} and conditional probability $\Pr\left(l_j|l_{j-1}\right) = 5 \times 10^{-2}$. Note that for Harary graphs, the diameter grows linearly with the number of nodes in the graph. In this example, the network diameter is five. Our dependency model yields a lower bound of 6.25×10^{-8}, whereas the independence model yields a lower bound of 1×10^{-10}.

G. Comparison of models and topologies

We conclude this section with a comparison of the models and topologies studied in the previous subsections. As we shall see, it is difficult to make a fair reliability comparison among the Ethernet, ring, multi-ring and Harary graphs studied in this chapter because the underlying dependent failure model is different in some of these cases.

1) All- and two-terminal reliability when p is low: In Figure 10, we plot the all- and two-terminal reliability performance of the ten node Ethernet, ring, double-ring and $H(10,3)$ graphs as a function of the correlation coefficient ρ, when $p = 10^{-4}$. When the correlation coefficient ρ is small – that is, when link failures are almost independent – the relative performance of the topologies is what we would expect from the independent failure model.

As ρ increases, the different assumptions in the different dependent failure models manifest themselves. For example, $H(10,3)$, which possesses the best reliability performance among all graphs when $\rho \approx 0$, exhibits increasingly poor performance relative to the other graphs as ρ increases to one. We attribute this to the conservative model developed for Harary graphs in Section IV-D when p is low. In this model, we first made the pessimistic assumption that every graph disconnection scenario is a partitioning of the graph into two consecutive subsets. We then made the additional pessimistic assumption that the links joining these two partitions are equally correlated. Thus, as ρ increases we expect the accuracy of our model to diminish. In fact, in the extreme scenario where $\rho = 1$, we require the all- and two-terminal reliabilities to reduce to p. However, as illustrated in Figure 10, our model yields probabilities of disconnection greater than p. Similarly, our dependent failure model for the Ethernet graph in Section IV-B becomes increasingly inaccurate as ρ increases. Again, in the extreme scenario where $\rho = 1$, we require the all- and two-terminal reliabilities to reduce to p. However, owing to our assumption that correlation only exists among the two links incident at each non-switch node, we obtain probabilities of disconnection greater than p in this extreme case. On the other hand, our model for the ring and multi-ring graphs in Section IV-C yields correct asymptotic reliabilities when $\rho \approx 1$, as illustrated in Figure 10.

2) All- and two-terminal reliability when p is high: Figure 11 illustrates the all-terminal reliability as a function of the correlation coefficient ρ for the ten node Ethernet, ring, double-ring, $H(10,3)$ and Petersen graphs when $p = 0.9$. The analysis underlying the performance of the $H(10,3)$ and Petersen graphs is that of Section IV-E, and for the ring and multi-ring topologies we follow Section IV-C. Lastly, for the Ethernet graph, the model used is

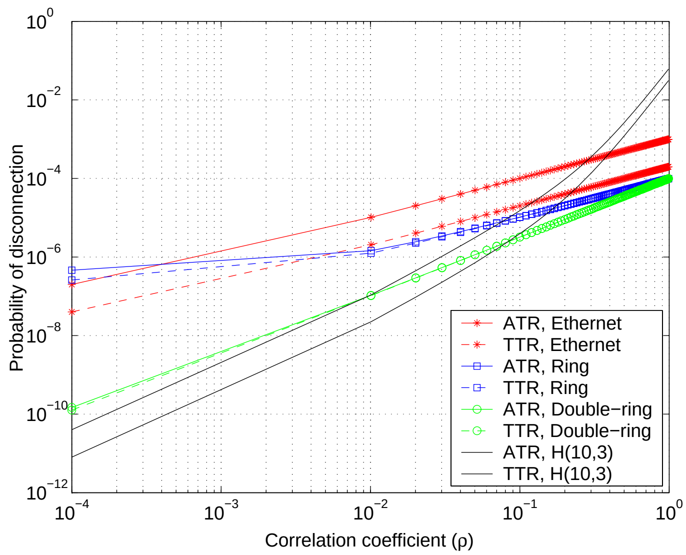

Fig. 10. Probability of disconnection versus correlation coefficient ρ for the ten node Ethernet, ring, double-ring and $H(10,3)$ graphs when $p = 10^{-4}$.

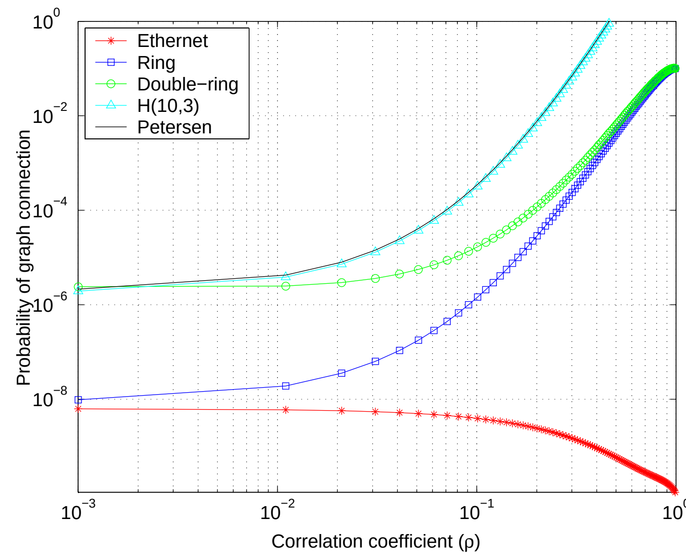

Fig. 11. Probability of graph connection versus correlation coefficient ρ for the ten node Ethernet, ring, double-ring, $H(10,3)$ and Petersen graphs when $p = 0.9$.

that of Section IV-B. When $\rho \approx 0$, the trends depicted in Figure 11 are what we expect from the independent failure model. As ρ increases, however, the modelling assumptions underlying the different topologies take effect. The ring and multi-ring graphs' all-terminal reliabilities converge to the correct value of $(1-p)$ as ρ approaches unity. The all-terminal reliability of the Harary and Petersen graphs exceeds unity as ρ approaches unity for the reasons discussed in Section IV-E.

The all-terminal reliability of Ethernet exhibits a peculiar downward trend as ρ increases. When $\rho \approx 0$, the all-terminal reliability of Ethernet is approximately $(1-p)^n\,[2+2^n(1-p)]$. If $2 \ll 2^n(1-p)$, then the all-terminal reliability is dominated by the probability of graph connection given that the bridge is operational. Conversely, if $2 \gg 2^n(1-p)$, then the all-terminal reliability is dominated by the probability of graph connection given that the bridge has failed. When $\rho \approx 1$, the all-terminal reliability of Ethernet is dominated by the probability of graph connection given that the bridge has failed, and is approximately $(1-p)^n$, which is at least a factor of two worse than the all-terminal reliability when $\rho \approx 0$. On the other hand, if all link failures in the Ethernet topology were correlated, then the all-terminal reliability would converge to $(1-p)$ as ρ approaches one. However, since our Ethernet link failure model assumes independence among different sets of link failures, the all-terminal probability converges to the probability that the two links incident at each non-switch node are operational, which is $(1-p)^n$.

Figure 12 depicts the two-terminal reliability as a function of the correlation coefficient ρ for the ten node Ethernet, ring, double-ring, $H(10,3)$ and Petersen graphs, when $p = 0.9$. The analysis underlying the performance of the $H(10,3)$ and Petersen graphs is that of Section IV-F, in which we conservatively only account for the probability that the shortest path between the node pair exists. The model underlying the ring and multi-ring topologies is that of Section IV-C. Lastly, for the Ethernet graph, the model used is that of Section IV-B, which implies that link failures along the shortest path between the node pair are statistically independent.

When $\rho \approx 0$, the trends depicted in Figure 12 are what we expect from the independent failure model. Specifically, the relative performance of the topologies is largely governed by their respective diameters. As ρ increases, however, the effect of these different network diameters diminishes. As can be seen from Figure 12, the reliability performances of the $H(10,3)$, Petersen, ring and multi-ring graphs converge to the expected value of $(1-p)$. Ethernet, however, owing to the assumptions of its model, exhibits a similar downward trend as in the case of all-terminal reliability. When $\rho \approx 0$, the two-terminal reliability is approximately $2(1-p)^2$, which is approximately equal to the probability of one of the two-hop paths between the source and destination being operational. On the other hand, when $\rho \approx 1$, the two links from each non-switch node act as one link and there is effectively only one two-hop path between the source and destination. In this case, the two-terminal reliability is approximately $(1-p)^2$.

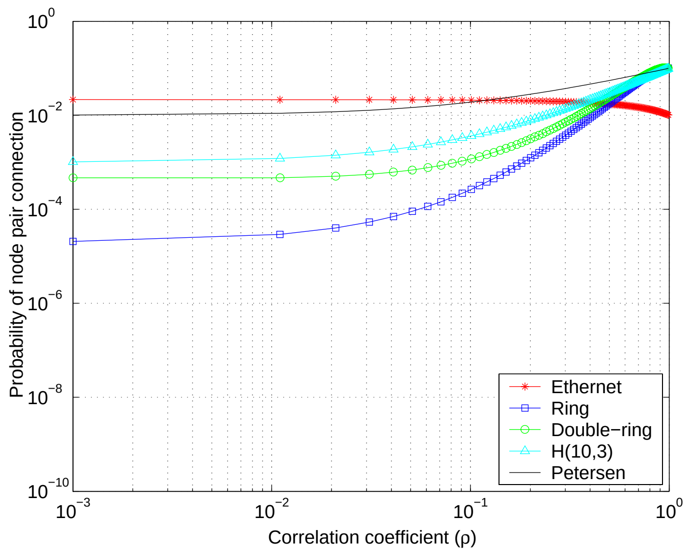

Fig. 12. Worst-case probability of node pair connection versus correlation coefficient ρ for the ten node Ethernet, ring, double-ring, $H(10,3)$ and Petersen graphs when $p = 0.9$.

V. Conclusion

In this work, we first considered the design of networks with statistically independent link failures and invulnerable nodes. We outlined and justified a design methodology in which circulant graphs were the principal candidate topologies. We found that: (i) When designing a reliable network topology, we should focus on optimizing the network structure with respect to high stress reliability, as low stress reliability is virtually unchanged provided that the underlying graph is super-λ (i.e. it achieves the minimum number of edge cutsets of maximum cardinality), which is the case for nearly all the circulant graphs proposed by our design methodology. (ii) To obtain all- and two-terminal reliabilities in the 0.1 to 1 range when links are unreliable, very large node degrees are required and that for such high node degree graphs, the actual graph structure is not very important.

We then broadened the scope of this work by allowing for the possibility of statistical dependency among link failures. We conducted approximate dependent failure analyses of several special topologies – Ethernet, ring, multi-ring and Harary graphs — using existing models and our simple Markov model, and have shown the danger in relying on an independent link failure model. In fact, using our Markov model, the probability of failure of m links with correlation coefficient ρ was shown to be approximately $p\rho^{m-1}$. Unfortunately, the models developed for the topologies we compared rested upon different assumptions, thereby making detailed comparisons among families of graphs difficult. On the other hand, these models may find use in comparisons among graphs belonging to the same family.

More work needs to be done with respect to dependent component failure models. Consistent, approximate models which strike a good balance between simplicity and applicability to a variety of topologies need to be developed. In addition, these models should possess intuitive inputs which are readily available to the network designer. Subsequently, optimality conditions for different regions of component vulnerability and dependency, akin to those developed for the independent failure model, need to be pursued.

References

[1] Y. Wen and V. W. S. Chan, "Ultra-reliable communication over unreliable optical networks via lightpath diversity: characterization and optimization," *To appear in IEEE Globecom 2003*.

[2] C. J. Colbourn, "Reliability issues in telecommunications network planning," in *Telecommunications network planning*, P. Soriano and B. Sanso, Eds. Kluwer Academic Publishers, 1999, ch. 9, pp. 135–146.

[3] F. Harary, *Graph Theory*. Addison-Wesley, 1969.

[4] ——, "The maximum connectivity of a graph," *Proceedings of the National Academy of Sciences USA*, vol. 48, pp. 1142–1146, 1962.

[5] F. T. Boesch, X. Li, and C. Suffel, "On the existence of uniformly optimally reliable networks," *Networks*, vol. 21, pp. 181–194, 1991.

[6] G. Wang, "A proof of Boesch's conjecture," *Networks*, vol. 24, pp. 277–284, 1994.

[7] A. Satyanarayana, L. Schoppmann, and C. Suffel, "A reliability-improving graph transformation wiht applications to network reliability," *Networks*, vol. 22, pp. 209–216, 1992.

[8] W. Mader, "Minimale n-fach kantenzusammenhängende graphen," *Math. Ann.*, vol. 191, pp. 21–28, 1971.

[9] F. T. Boesch and J. Wang, "Reliable circulant networks with minimum transmission delay," *IEEE Transactions on Circuits and Systems*, vol. 32, no. 12, pp. 1286–1291, Dec. 1985.

[10] J. F. Wang and C. S. Yang, "On the number of spanning trees in circulant graphs," *International Journal of Computer Mathematics*, vol. 16, pp. 229–241, 1984.

[11] F. T. Boesch and J. Wang, "Super line-connectivity properties of circulant graphs," *SIAM J. Algebraic and Discrete Methods*, vol. 7, pp. 89–98, 1986.

[12] B. Sawionek, J. Wojciechowski, and J. Arabas, "Synthesis of reliable networks in the presence of line failures," ser. Proceedings of IEEE International Symposium on Circuits and Systems, Geneva, May 2000, pp. IV–649–IV–652.

[13] D. Bauer, F. T. Boesch, C. Suffel, and R. Van Slyke, "On the validity of a reduction of reliable network design to a graph extremal problem," *IEEE Transactions on Circuits and Systems*, vol. 34, no. 12, pp. 1579–1581, Dec. 1987.

[14] G. E. Weichenberg, "High-reliability architectures for networks under stress," S. M. thesis, Massachusetts Institute of Technology, 2003.

[15] G. E. Weichenberg, V. W. S. Chan, and M. Médard, "Reliable architectures for networks under stress," *To appear in DRCN 2003*.

[16] R. Breyer and S. Riley, *Switched, Fast, and Gigabit Ethernet*. New Riders Publishing, 1998.

[17] S. N. Pan and J. Spragins, "Dependent failure reliability models for tactical communications networks," ser. Proceedings of International Conference on Communications, 1983, pp. 765–771.