

## Research Article

# Survivable Network Design and Optimization with Network Families

Brody Todd,<sup>1,2</sup> Abiose Ibigbami,<sup>1</sup> and John Doucette<sup>1,2</sup>

<sup>1</sup> Department of Mechanical Engineering, University of Alberta, Edmonton, AB, Canada T6G 2G8

<sup>2</sup> TRILabs, Edmonton, AB, Canada T5K 2M5

Correspondence should be addressed to John Doucette; [john.doucette@ualberta.ca](mailto:john.doucette@ualberta.ca)

Received 15 April 2014; Accepted 9 June 2014; Published 1 July 2014

Academic Editor: Liansheng Tan

Copyright © 2014 Brody Todd et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

In modeling communication networks for simulation of survivability schemes, one goal is often to implement these schemes across varying degrees of nodal connectivity to get unbiased performance results. Abstractions of real networks, simple random networks, and families of networks are the most common categories of these sample networks. This paper looks at how using the network family concept provides a solid unbiased foundation to compare different network protection models. The network family provides an advantage over random networks by requiring one solution per average nodal degree, as opposed to having to solve many, which could take a significant amount of time. Also, because the network family looks at a protection scheme across a variety of average nodal connectivities, a clearer picture of the scheme's performance is gained compared to just running the simulation on a single network or select few networks.

## 1. Introduction and Background

Communication technology is a fundamental component of modern societies, allowing rapid exchange of knowledge, collaboration, and much more. Underpinning this fundamental component of our society is a robust reliable network. As such, the need to drive reliability and availability in the core communication infrastructure is a constant concern. The focus of this work is on network survivability design where the most common failures occur in network spans. This is due to the highly uncontrolled environment in which these spans exist and the significant cost of redundancy and because fibre links comprising a span are generally quite lengthy (hundreds of kilometres in terrestrial long-haul networks and even longer in some undersea networks) and are often routed through remote locations that are difficult to fully secure. Even relatively small urban networks are subject to a high rate of cable cuts (one estimate claims 13 cuts per 1000 miles of fibre [1]), caused by construction work crew mistakes, vehicle accidents, boat anchors (for undersea cables), damaged water mains, and natural disasters, as well as deliberate actions such as vandalism and other criminal activities (terrorism is now also a concern).

In the past two decades, many network survivability techniques have been developed to provision spare capacity on a network in such a way that it can withstand failure of one or more of its spans [2–4]. An important aspect to consider when designing, comparing, or otherwise evaluating these network design models is how their performance compares over a variety of network configurations. This work will investigate how proper choice of network topologies will allow for improved comparative evaluations of network design methods over varying network connectivities. A preliminary version of this work previously appeared in ICC 2008 [5].

*1.1. Network Survivability Models.* The most common network survivability techniques include 1+1 *automatic protection switching* (1+1 APS) [6], *demand-wise shared protection* (DSP) [3], *shared backup path protection* (SBPP) [7, 8], span restoration [2], path restoration [9], *p*-cycles [10], and *self-healing rings* [11]. The survivability models we have used herein are 1+1 APS, span restoration, path restoration, SBPP, and *P*-cycles. For a brief review of these techniques, see [5].

Implementing survivability schemes can take a number of approaches. Common design approaches include heuristic

algorithms, including genetic algorithms, tabu search, and custom heuristics; however, to guarantee a certain level of optimality, *integer linear programming* (ILP) techniques are required [2, 12, 13]. One class of ILP based methods is what we can consider an *arc path* approach, where we preprocess a network's topology to enumerate a set of distinct eligible logical routes and the ILP model selects a subset of the paths, upon which to route restoration paths. This is the design method we will use herein. More specifically, we use the ILP design formulations described in [4] for each of the above survivability techniques. We ask the interested reader to refer to that document for detailed descriptions of those ILP formulations, as space restrictions prevent us from providing them all here.

*1.2. Network Families.* When designing or evaluating network survivability schemes, we must first select a set of test networks to evaluate those schemes. Common approaches for selecting these test networks include using topologies modeled after real networks and selecting topologies used elsewhere in literature or simulated random topologies. The difficulty in the selection of the test networks is with the ability to evaluate the survivability schemes' performance across a variety of topologies and network connectivities. As will be shown later, network connectivity is a key dimension in evaluating schemes for providing a more complete picture of each scheme's behaviour. Another approach that has been taken is to create what we call network families [4, 5, 14] which contain a common underlying topology, differing only as much as needed to create a sequence of networks that vary in connectivity.

A network family is a set of networks varying only in the number of spans that each contains, where all nodes and a set of common spans keep a consistent configuration. These families are commonly designed by starting with a highly connected master network (Figure 1(a)) with similar characteristics to real networks (such as connectivity locality). A sequence of networks is then created by successively removing spans (Figures 1(b) and 1(c)), creating a series of networks that keep consistent demand patterns and nodal configurations. Span removals were done in a pseudorandom manner ensuring that biconnectivity was maintained. The removal process iteratively selects a random span for possible removal, ensures that removing it does not break biconnectivity in the network, and then either marks the span as essential or removes it. This process is repeated until there are no more spans that can be removed from the network. Alternatively, the opposite approach could just easily create a family of networks, starting with a minimally connected base network; pseudorandom span additions are made until the network reaches some desired high level of connectivity. It is important to note that the selection of a network family is dependent on the initial topology, along with the demand matrix and nodal configuration that is representative of the types of networks being studied. The selection of a good master network is dependent on defining these characteristics in a representative manner of the class of network that the survivability scheme is designed for. The type of master network selected in this work is discussed in Section 2.

The purpose of the study behind this paper is to evaluate the use of network families in comparison to two other methods of selecting test networks, stand-alone representations of real networks, and sets of randomly generated topologies.

*1.3. Elements Affecting Network Redundancy.* A common method of evaluating an optimally (or near optimally) designed network is its redundancy, though various capacity measures can easily be used as a surrogate (total working plus spare capacity, total spare capacity, etc.). This redundancy metric is affected by a number of characteristics, five of which are of significance to the work herein. These characteristics include the demand matrix (traffic volume between each node pair in the network), the number of nodes in the network and their configuration, and the number of spans in the network and their configuration. It should be noted that, for full comparison across this range of factors influencing survivability scheme performance, each degree of freedom should be tested. It is the intention of this work to demonstrate a mechanism to isolate the number of spans (average nodal degree), as this can be a major differentiator in topology characteristics globally [15].

The demand matrix selected for a network topology obviously has a significant impact in the capacity and redundancy required. The impact of the demand matrix can come from a number of factors including overall volume and distribution, where distribution includes variability and locality. A demand matrix that is highly localized will require less capacity than the one that has a wide ranging demand pattern. Variability in demand patterns can also have a significant implication on overall capacity and redundancy as there would be a few large capacity demands that would dominate the requirements for redundant capacity and reduce the capability of the survivability scheme to share spare capacity. The effects and implications of demand patterns are difficult to predict, and, as such, keeping them consistent is an important consideration when comparing survivability schemes.

The nodal topology of a network also has a significant effect in the redundancy of a network, specifically in configuration and localized density. Consider the impact that modifying a single node in a network would have in path routing. This would impact path lengths and layouts and could introduce or remove trap topologies (Figure 1) and other characteristics that would have a significant impact on the redundancy of the network design. The addition or removal of nodes obviously would also affect the network design in a similar manner.

The last two factors, network connectivity, and how that connectivity is arranged can also significantly affect redundancy in a network. The ability to diversely route traffic is directly related to the spatial diversity and number of spans in a network; we are able to more directly route survivable traffic and/or better share spare capacity in a richly connected network, reducing required redundancy. The configuration of the spans in the network also plays a role in network's capacity requirements. If demand in a network is broadly distributed, but the connectivity is centered around relatively few clusters, the effect on route diversity and spare

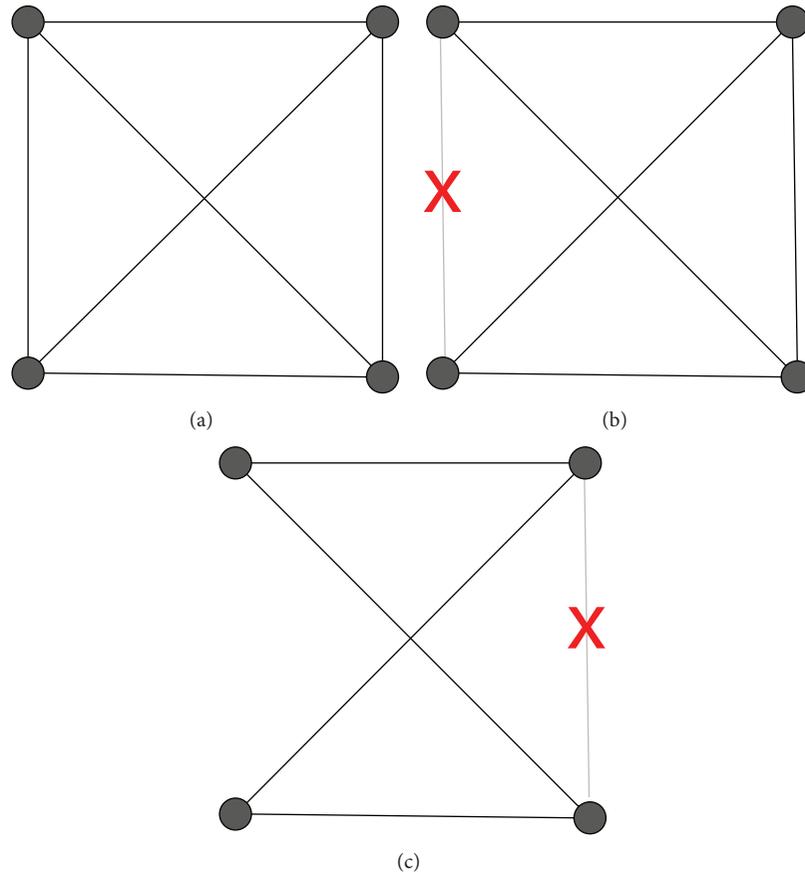


FIGURE 1: An example for node network family.

capacity sharing would increase redundancy and capacity requirements when compared to a network where average nodal degree was more consistent across nodes. This implies that although connectivity is a good descriptor of a network, the layout and configuration of the nodes and spans can have an impact on key metrics when evaluating a survivability scheme.

In summary, all of these characteristics are difficult to normalize, and each can have a significant effect on the redundancy or capacity cost of a survivable network design. The selection of test networks should account for these variables in order to evaluate the performance of the selected schemes across topologies and demands.

#### 1.4. Using Network Families to Reduce Degrees of Freedom.

The above discussion is significant in that it frames the motivation to use network families. Network families are able to restrict the above dimensions to allow comparison of survivability schemes along one of the most important dimensions in a network, average connectivity.

Network families directly limit variability in demand matrices, nodal topology, and nodal density. Variability in span configuration and network connectivity are also tightly controlled. With each successive network varying only by the addition or removal of a single span, the effects of span configuration remain as consistent as possible, and the

control variable becomes network connectivity. This control provides a greater level of confidence that changing levels of redundancy are due to characteristics of the survivability scheme and not an artifact of the differences between stand-alone test networks.

If random or unrelated networks were used in a comparative evaluation of a survivability scheme [16–19], there would be little confidence that the results or observations were based on the scheme or if they were due to characteristics of specific networks. The resulting redundancies could be affected by changing demand matrices, node configuration, and/or span configuration in addition to the changes in network connectivity. Since the effects of the characteristics mentioned above are difficult to characterize and normalize, the ability to compare the survivability schemes across topologies and connectivities is limited.

**1.5. Goals and Motivation.** The purpose of this study is to investigate the use of network families in comparative evaluations of network survivability schemes or models across a range of network connectivities. Two common approaches also used in comparative analysis include using selected real world networks and using large sets of random networks. This study looks to quantitatively compare these approaches to test network selection in order to confirm that network families represent a valid selection method.

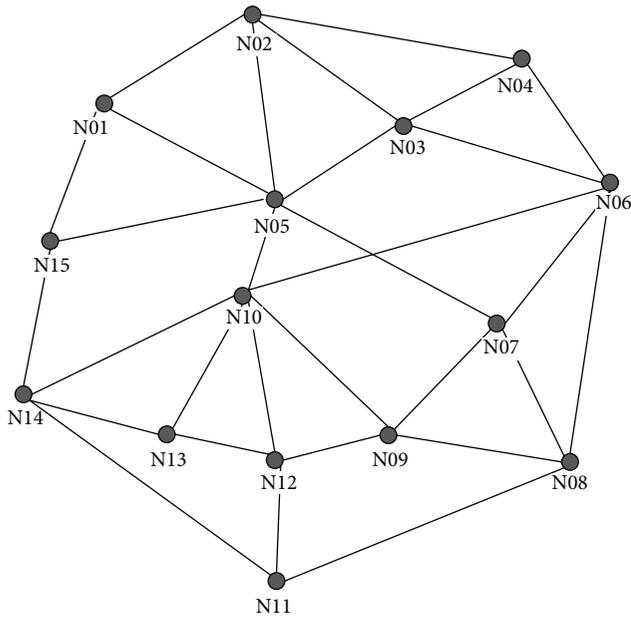


FIGURE 2: 15-node master network in our test case family.

## 2. Experimental Setup

Our tests used three sets of networks; one set comprises a 15-node network family, a second set consists of 15-node stand-alone random networks, and a third set consists of a variety of actual networks. Our network family consisted of 15 related but distinct 15-node networks ranging in average nodal degree from  $\bar{d} = 2.13$  (16 spans) up to  $\bar{d} = 4.0$  (30 spans). The master network used to create this family can be seen in Figure 2. Subsequent networks were created using the process described in Section 1.2. The network was designed to have an evenly dispersed nodal configuration, typical of what is seen in large scale backbone networks [15]. The demand matrix between nodes was designed to be randomly dispersed in a symmetric fashion.

The set of random networks consists of 300 pseudo-randomly generated networks, with 20 completely unrelated networks for each nodal degree in the network family. These networks were considered to be pseudorandom as they were derived to have characteristics that could be conceivably implemented in reality (with respect to node density and degree). Effort was made to keep the scale of the networks consistent (i.e., the point-to-point distances between the most distant nodes); however, there were variations that had to be accounted for in the analysis (Section 3.1.1).

The actual networks used in this study were composed of data that is representative of networks that are currently implemented. We collected network maps of some service providers, regional networks, and international carriers from various online sources. Table 1 provides a summary of some parameters of these networks. The average nodal degree of the real world test case networks ranged between  $\bar{d} = 2.76$  and  $\bar{d} = 5.25$ . No actual demand information was publicly available for these networks; hence, the demand matrix we

used for each of the networks is a full mesh, with the demand values ranging from 1 to 10 units generated randomly with a uniform distribution (as with all other networks in our study).

Three of the real networks (BLCR, NOR, and NYC) were later selected and used as master networks for their respective network families with the average nodal degree ranging from  $\bar{d} = 2.0$  to  $\bar{d} = 4.0$ . This enabled us to observe the survivability mechanisms across different network families derived from real networks.

**2.1. Survivable Design Methods.** As already discussed above, design models corresponding to five common network survivability schemes were used; these schemes were 1+1 APS, span restoration, path restoration, SBPP, and  $p$ -cycles. The specifics of each model are described in [4], where models corresponding to *spare capacity allocation* (SCA) and *joint capacity allocation* (JCA) are provided. As in [4], we use both approaches (JCA and SCA) herein, enforcing integrality in all capacity-related decisions variables but without considerations for modularity in link capacity allocation [19].

The path-restorable and SBPP design models each were provided with the five shortest eligible working routes and the ten shortest eligible restorations routes for each lightpath demand (i.e., node pair). The 1+1 APS designs routed all demands on the single shortest pair of routes between their end nodes (no formal optimization needed). The span-restorable design model was provided with the five shortest eligible working routes per lightpath demand and the ten shortest eligible restorations routes for each span failure scenario. The  $p$ -cycle design model was provided with the five shortest eligible working routes per lightpath demand and the 1000 shortest eligible cycles that can be drawn in the network (where they exist). All eligible routes and cycles were enumerated via a depth-first-search method.

In terms of the demand matrix, a full mesh was used (with each node pair assigned some demand), with each node pair exchanging a randomly generated number of demands between 1 and 10 (using a uniform distribution). The network families and random networks used the exact same demand matrix, ensuring consistent demand volumes across the networks. The real networks used distinct, randomly generated demand matrices using the same parameters mentioned above.

All ILP formulations were implemented using AMP, and were solved using CPLEX 10.1 on a Sun 1.6 GHz quad CPU machine with 16 GB of RAM. The results are based on full CPLEX terminations with *mipgap* settings of 0.001, meaning solutions are guaranteed to be within 0.1% of optimal.

## 3. Experimental Results

**3.1. Comparison of Families and Random Networks.** Comparing network families with random networks was done using a normalized mean and maximum and minimum of the random networks with the normalized results of the network family, as shown in Figure 3 through Figure 7. In each of those charts, each data point in the network family curves

TABLE I: List of implemented networks used in this study.

Network	# Nodes	# Spans	$\bar{d}$	Notes
BNDCN	21	29	2.76	Bandcon International Network [20]
NG	17	24	2.82	Nobel-Germany National Network [21]
IN	11	16	2.91	ITC Deltacom Regional Network [22]
DELTA	59	89	3.01	ITC Deltacom Regional Network [23]
C&W	24	38	3.16	Cable & Wireless USA IP Backbone [24]
COX	21	38	3.61	Cox National IP Backbone Network [25]
BLCR	15	28	3.73	Bellcore Network [26]
NOR	27	52	3.85	Norway National Network [21]
BTNA	35	72	4.11	British Telecom North American Network [27]
NYC	16	42	5.25	New York Regional Network [21]

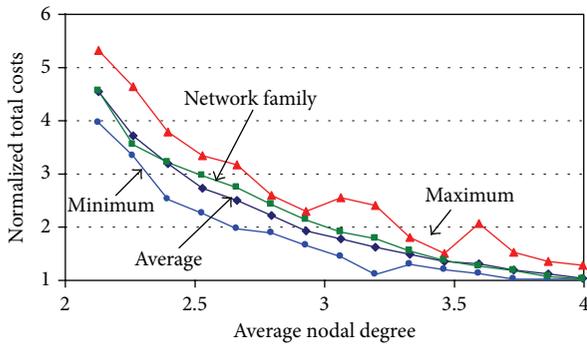
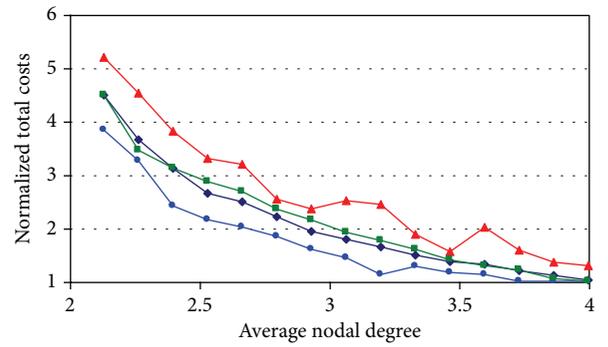
FIGURE 3: Normalized  $p$ -cycle JCA total capacity costs.

FIGURE 4: Normalized span restoration JCA total capacity costs.

(green curve with square markers) represents the normalized total working plus spare capacity costs for the member of the family with the indicated average nodal degree. The three other curves correspond to the aggregated results from the 300 pseudorandom networks. Each data point on the dark blue curve with diamond markers represents the mean normalized total capacity averaged over all random networks with the indicated average nodal degree. Each data point on the light blue curve with circular markers represents the minimum normalized total cost across the 20 random networks at the indicated average nodal degree, and each data point on the red curve with triangle markers represents the maximum normalized total cost across the random networks at the indicated average nodal degree. The results are presented according to the five survivability mechanisms used,  $p$ -cycle (Figure 3), span restoration (Figure 4), SBPP (Figure 5), path restoration (Figure 6), and 1+1 APS (Figure 7).

**3.1.1. Normalization of Results.** In order to compare the capacity costs of designs on different topologies, the results had to be normalized. This normalization removed the effects of slight variations in network scale. The network family capacity costs were normalized to the cost of the network corresponding to a nodal degree of 4.0 for each survivability scheme. The pseudorandom networks had to be normalized in a slightly different manner. The result for each random network was divided by the total unit cost (the cost of one unit of capacity allocated to each span) for all of the spans in

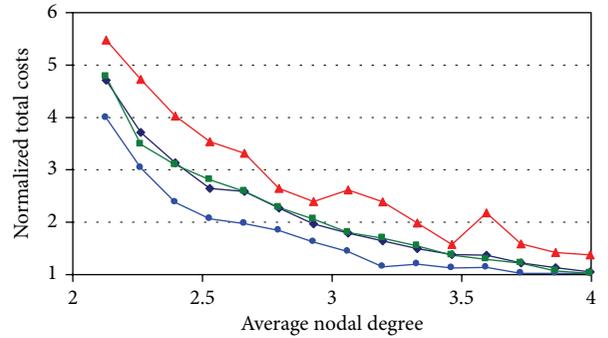


FIGURE 5: Normalized SBPP JCA total capacity costs.

the network and then normalized to the lowest cost network in the set of networks with  $\bar{d} = 4.0$ . Although the random networks were designed with similar scales, variations were still present. Normalizing each to their unit cost first provided a way to account for these variations.

All curves shown follow a similar general behaviour, and the network family results were quite similar to the average of the pseudorandom networks. It is also obvious that there is a large gap between the maximum and minimum results for each average nodal degree. The mean difference for  $p$ -cycles was 32%, while for span restoration it was 34%, and for the SBPP results it was 37%. The largest difference between the maximum and minimum was 54% in the  $p$ -cycle curve

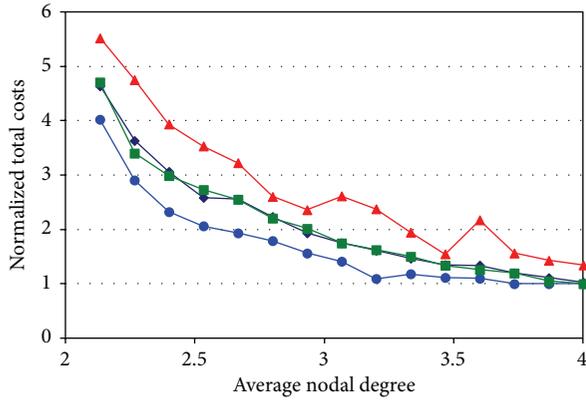


FIGURE 6: Normalized path restoration JCA total capacity costs.

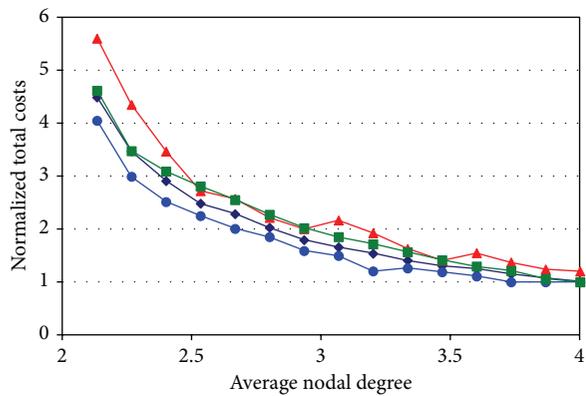


FIGURE 7: Normalized 1+1 APS JCA total capacity costs.

at  $\bar{d} = 3.2$ . The implication of this variability is significant. A pseudorandom network could require a capacity cost anywhere between the maximum and minimum values (or beyond, as these are just the limits of the sample set we used and not true maximums or minimums possible). There could be, in theory, a set of networks ranging from  $\bar{d} = 2.67$  to  $\bar{d} = 3.6$  that all had the same capacity costs for their  $p$ -cycle design (see Figure 3 where a normalized cost of 2.0 falls between the maximum and minimum for that range of connectivities). If the set of simulated networks fell into a scenario such as this, there could be erroneous conclusions about the performance of the survivability model with regard to network connectivity, when the full set of data says something different.

As demonstrated in Figure 3 to Figure 7, there is variability in the capacity costs of the pseudorandom networks. To get an idea of how this variability is distributed, we look more closely at the SBPP results in Figure 8, which displays data points corresponding to SBPP design costs for all of our random networks. As can be observed in the figures, there is a wide distribution at each network connectivity level. This implies then that a sequence of pseudorandom networks could correspond to any number of the sequential data points shown. They would show the observed trend; however, it would be quite nonuniform. In contrast, the network family

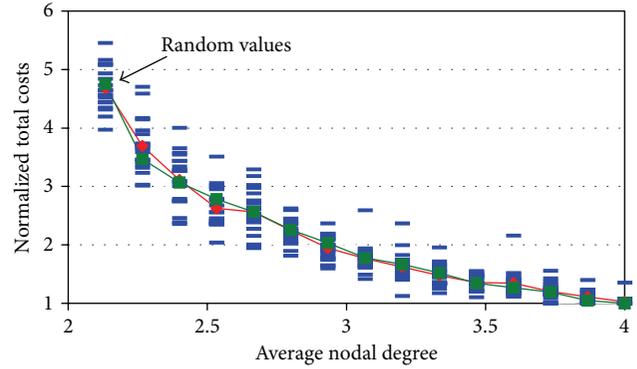
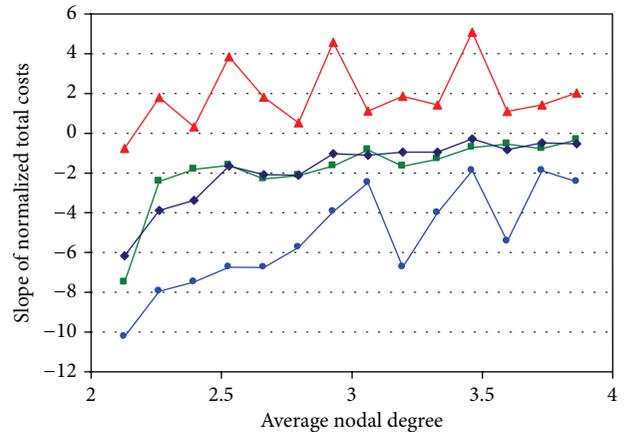


FIGURE 8: Normalized SBPP total costs of random networks with average and network family.

FIGURE 9: Slope of normalized  $p$ -cycle JCA total capacity costs.

curves were quite smooth and run closely to the average of the pseudorandom networks, suggesting that they offer a better representation of how a particular restoration might behave with respect to network connectivity.

To account for the high level of variability in pseudorandom networks, a large number of test networks are required for each connectivity level (we used 20 and were able to produce a reasonably smooth average curve). Of significance, however, is that one network family of 15 networks produced similar results to the average of 300 networks. To more closely examine how closely a network family could reproduce the behaviour of the average of the pseudorandom networks, the first derivative (or slope) of the normalized total costs was compared. Figures 9, 10, and 11 show the slopes of the network family cost curves, along with the maximum, minimum, and average slopes for the pseudorandom networks for the  $p$ -cycle, span restoration, and SBPP data. For space considerations, we omit the 1+1 and path restoration results; they are consistent with the data shown in Figure 9 to Figure 11. These additional results are available on request.

The comparison of the slopes of the pseudorandom network averages and the network family highlight how well one network family represents the behaviour of many random networks when compared across a range of network

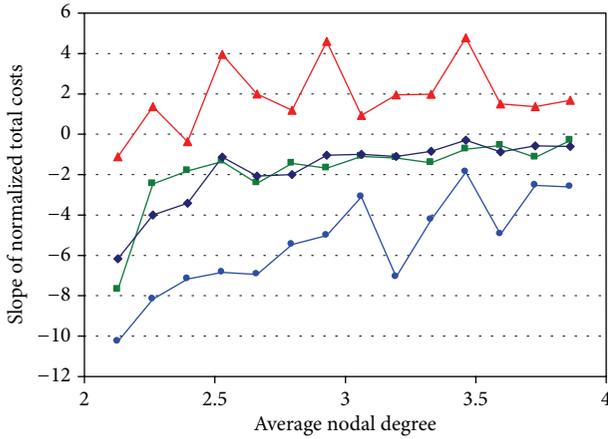


FIGURE 10: Slope of normalized span restoration JCA total capacity costs.

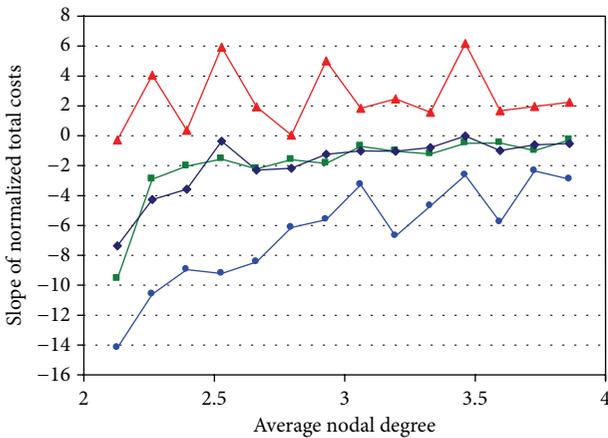


FIGURE 11: Slope of normalized SBPP JCA total capacity costs.

connectivities. This implies that using a network family to observe the behaviour of a survivability scheme and/or network design model over varying connectivities produces a result that is similar to the average of a large number of networks. When dealing with models that may take days, weeks, or months to find a solution, this provides significant time savings; one could use a single network family rather than many stand-alone (i.e., random) networks.

Next, we compare  $p$ -cycle capacity costs of the 15 networks in our network family with that of 15 networks randomly selected from the set of pseudorandom networks (one at each average nodal degree). Figure 12 plots the normalized total costs of the network family (green curve with square markers) alongside the costs of the randomly selected networks (blue curve with diamond markers). Correspondingly, Figure 13 plots the slopes of the curves in Figure 12. For brevity, we displayed only the results from the  $p$ -cycle designs; however, these results are representative of the other models as well.

Using a straightforward statistical analysis of the  $p$ -cycle results showed that the network families were never more than one standard deviation away from the average of the

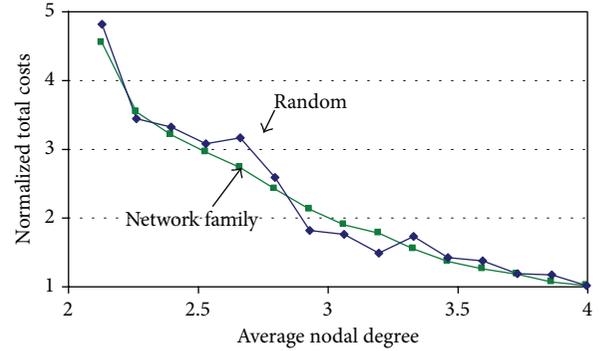


FIGURE 12: Normalized  $p$ -cycle JCA total capacity costs for a random subset of networks compared to the network family.

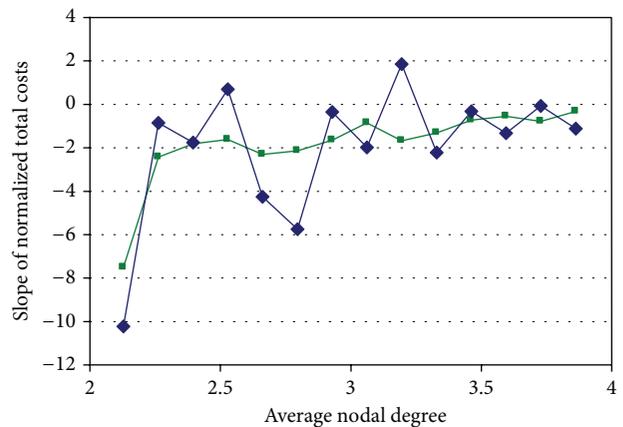


FIGURE 13: Slope of the normalized  $p$ -cycle JCA total capacity costs for a random subset of networks compared to the network family.

pseudorandom networks. This was the same across all the survivability models used (75 data points in total), except for two data points. In comparison, the set of 15 randomly selected networks exceeded one standard deviation from the average 6 out of 14 times for the  $p$ -cycle data alone and 28 out of 70 times across all the survivability models. For further emphasis on the ability of the network family to resemble the behaviour of the average of the pseudorandom networks, the mean difference between the two sets and the average were calculated. When compared to the average normalized total cost, the network family had a mean difference of 0.092, while the set of randomly selected networks was 0.16. The difference was even more pronounced when the difference in slope was compared. The network family had a mean difference of 0.041 relative to the pseudorandom network average, and the random set of networks had a difference of 1.34. Emphasized here is the ability for network families to model the behaviour of the average of a large set of pseudorandom networks when used as test case networks for a network design model across varying network connectivities.

3.2. Comparison of Network Families with Real Networks. The results obtained for SCA and JCA for path restoration mechanism are shown in Figures 14 and 15, respectively.

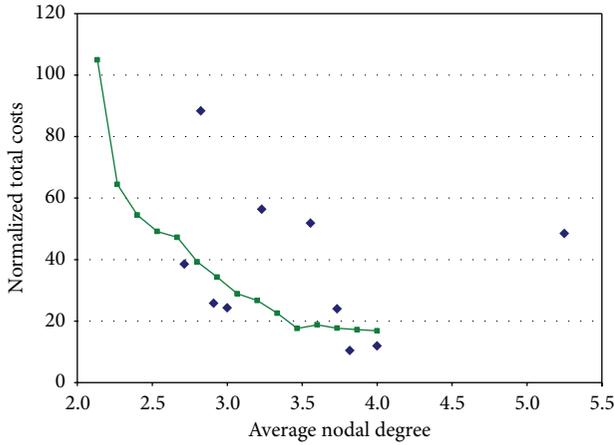


FIGURE 14: Normalized path restoration SCA spare capacity costs.

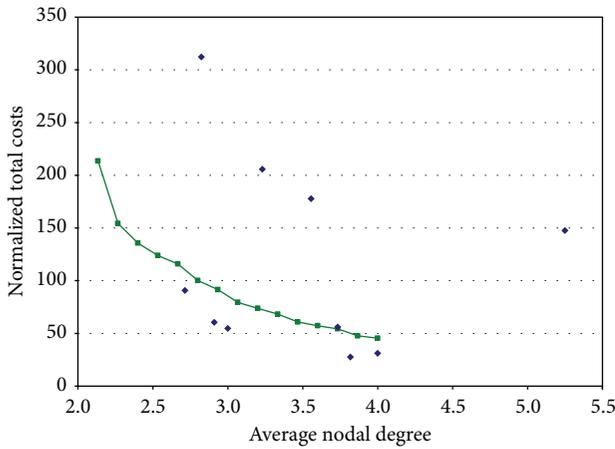


FIGURE 15: Normalized path restoration JCA total capacity costs.

The green curve with square data points corresponds to the average normalized spare capacity cost of the 15-node network families. This provides a benchmark for our comparison. Each data point plotted on this curve represents the average cost of designing the members of the network families with the indicated average nodal degree. The data plotted as individual points using the blue diamond markers scattered about are the normalized spare capacity cost we obtained from the real world networks. From Figure 14, we observe that, at average nodal degrees of 2.7, 2.9, 3.0, and 4.0, the normalized spare costs for the real networks fall relatively close to those from the network families with deviations of 15%, 26%, 19%, and 29%, respectively. We also observe that the largest gap between results from network families and real networks occurred at  $\bar{d} = 2.8$  for path restoration SCA with about 124% deviation.

Similarly, for path restoration JCA, in Figure 15, we observe that, at average nodal degrees of 2.7, 2.9, 3.0, and 4.0, the normalized total costs for the real networks also fall relatively close to those from the network families, with deviations of 15%, 33%, 31%, and 32%, respectively. We observe that the largest gap between results from network

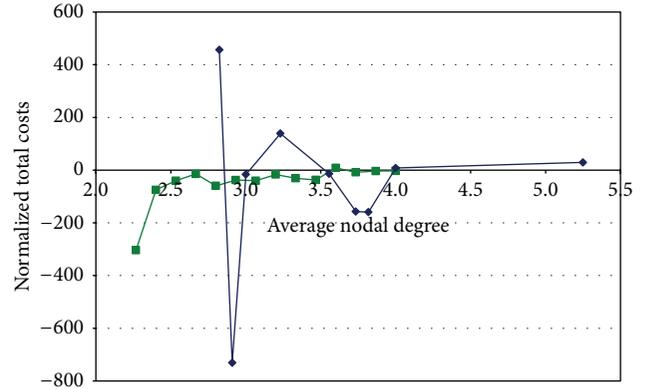


FIGURE 16: Slope of path restoration SCA normalized spare capacity cost.

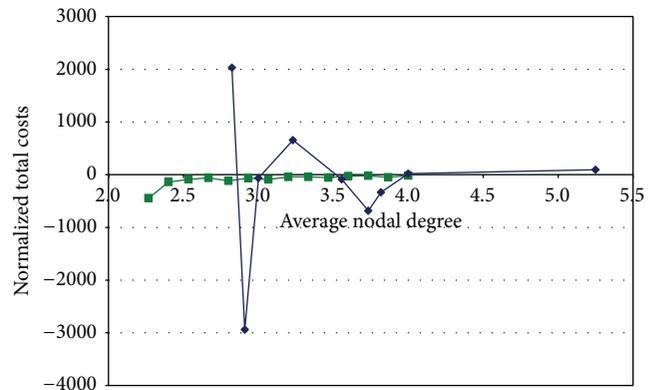


FIGURE 17: Slope of path restoration JCA normalized total capacity cost.

families and real network occurred at  $\bar{d} = 2.8$  with a deviation of approximately 210%. It is therefore clear that there is a much greater variability in the design costs across real networks that possess diverse characteristics than in network families.

We can also look at the slopes of the SCA and JCA capacity cost curves, as a metric for the reduction in total cost on a network as the average nodal degree increases; we plot those slopes in Figures 16 and 17, respectively. Compared with the network families, we see that the real networks did not provide a consistent trend since we would expect that the capacity cost of the networks should decrease with average nodal degree.

We can use the same representation for data plotted in Figures 18, 19, 20, 21, 22, and 23, which shows SCA and JCA results for  $p$ -cycle, SBPP, and span restoration. Similar to earlier charts, data points for the real networks are scattered all over the charts with only a few falling close to those from the network families, again supporting our claims that network families permit better and unbiased evaluation of a network design model. For space considerations, we omit the 1+1 APS data as well as the slope analysis for all schemes, as they are all consistent with the data shown.

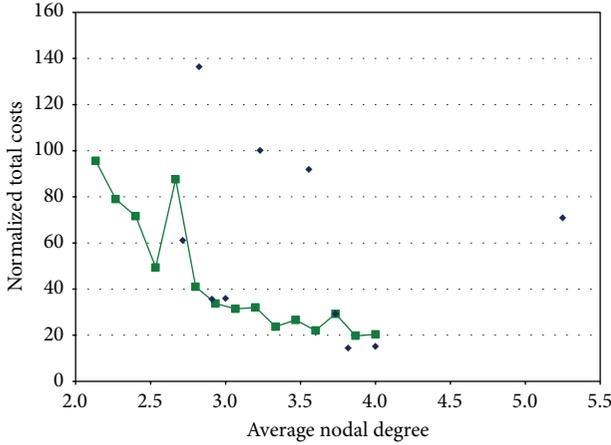
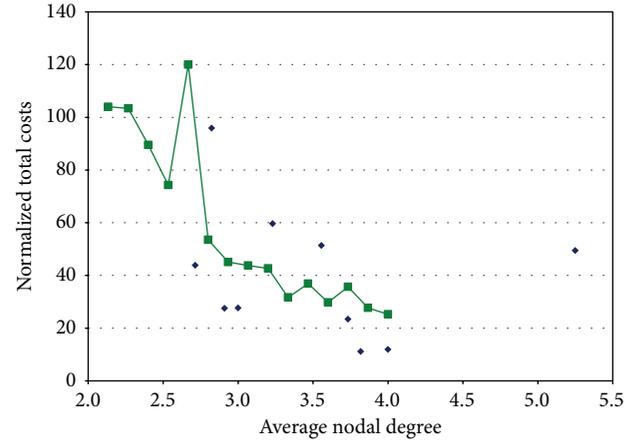
FIGURE 18: Normalized  $p$ -cycle restorable SCA spare capacity costs.

FIGURE 20: Normalized SBPP SCA spare capacity costs.

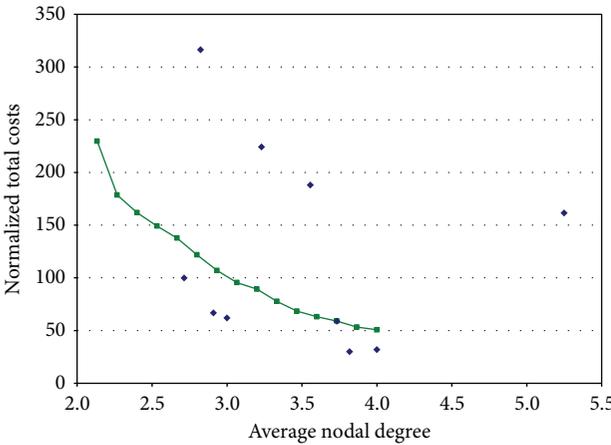
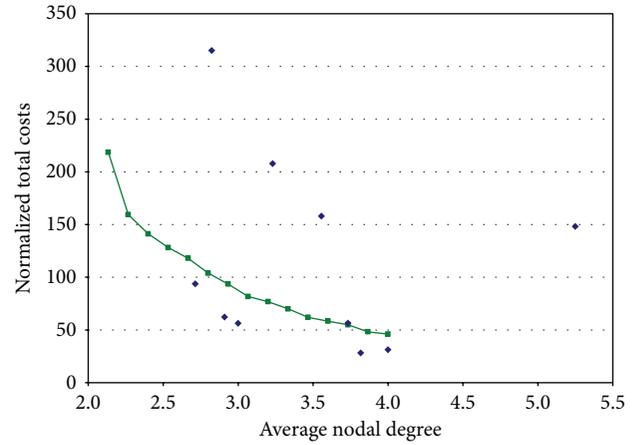
FIGURE 19: Normalized  $p$ -cycle restorable JCA total capacity costs.

FIGURE 21: Normalized SBPP JCA total capacity costs.

### 3.3. Using Real Networks as a Basis for Network Families.

In the next part of our analysis, we generated network families from the 16-node NYC and the 27-node NOR real networks. For each network, we successively added and/or removed individual spans one at a time to produce a series of related networks of average nodal degrees ranging from slightly above 2.0 (subject to maintaining biconnectivity) to 4.0. We plot the results of this analysis in Figures 24, 25, 26, 27, 28, and 29, where we show data SCA spare capacity costs and JCA total capacity costs, respectively, for network design models corresponding to each of the  $p$ -cycle, SBPP, and span restoration survivability schemes, respectively. Each data point in those figures represents the objective function value (i.e., total spare capacity costs for the SCA designs and total working plus spare capacity costs for the JCA designs) of the member of the indicated network family at the specified average nodal degree for the  $p$ -cycle, SBPP, and span restoration survivability schemes, as denoted in the figure captions.

As we can observe from these figures, the behaviour of the various curves provides a closer semblance to that of the 15-node network families than we had observed earlier

with the original stand-alone real networks. Except for only a few abnormalities, particularly in the NOR-based network family, these new families permit a much better view of the behaviour of the various survivability schemes and design models. Upon closer inspection of the abnormalities, they appear to be due to peculiarities in the eligible routes and  $p$ -cycles that are enumerated using the method described in Section 2; they are artifacts that arise from the limitations of our experimental approach. If we had enumerated a more complete set of eligible routes and/or  $p$ -cycles, these abnormalities diminish in severity or even disappear altogether. Another factor contributing to these abnormalities is the shortest path routing approach for working routing, at least in the SCA designs. Because working routes are the shortest path routed, the working capacities in a particular test case are occasionally less than those in the next sparser member of the same network family

In any case, the trend observed in these curves is as we expected; if network families are generated from real networks rather than just selecting stand-alone real networks randomly for testing, we are able to overcome the irregularities that would otherwise have arisen and potentially

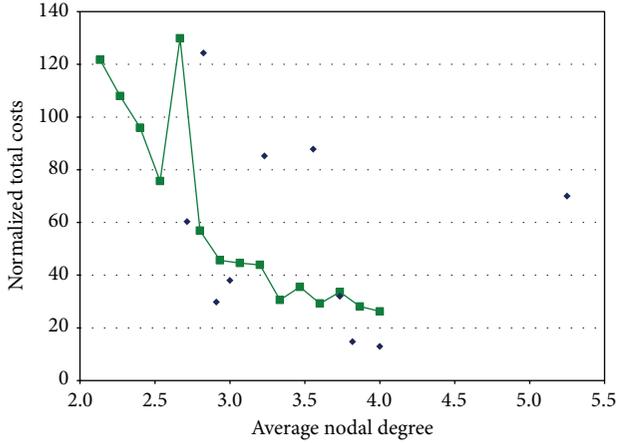


FIGURE 22: Normalized span restoration SCA spare capacity costs.

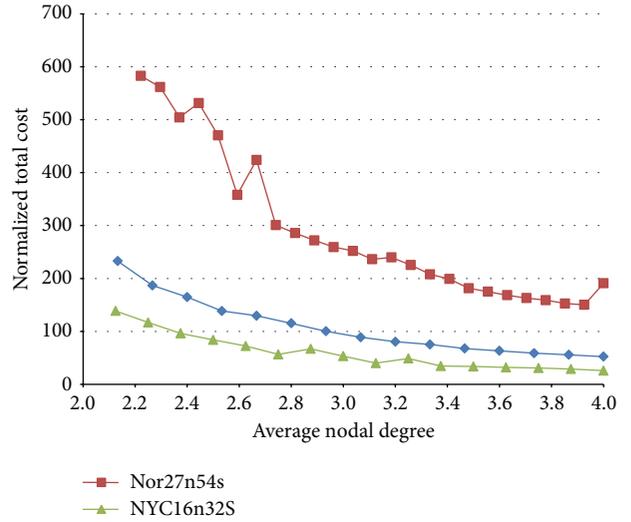


FIGURE 25: Normalized  $p$ -cycle restoration JCA costs for network families based on real network.

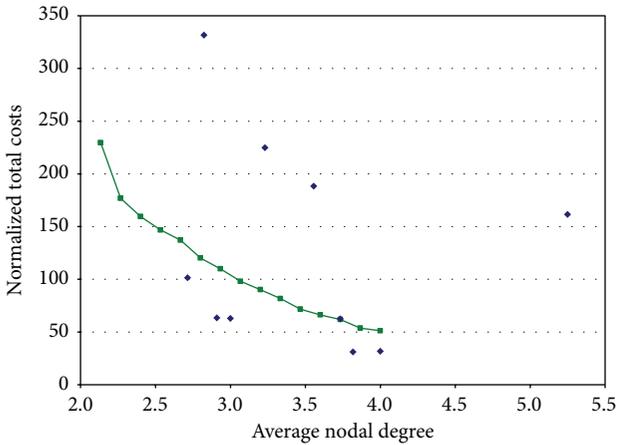


FIGURE 23: Normalized span restoration JCA total capacity costs.

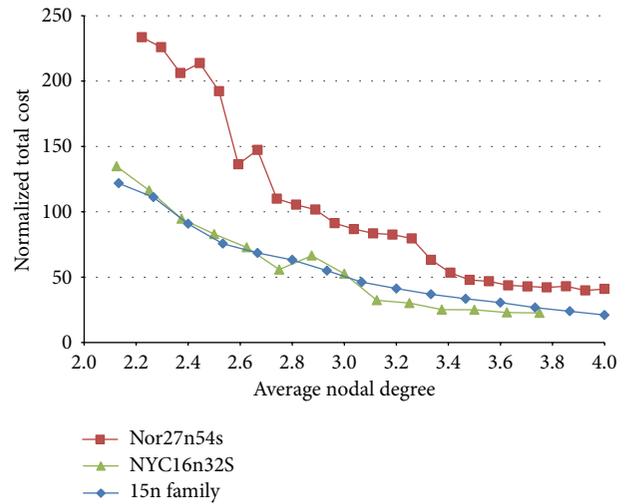


FIGURE 26: Normalized SBPP restoration SCA costs for network families based on real network.

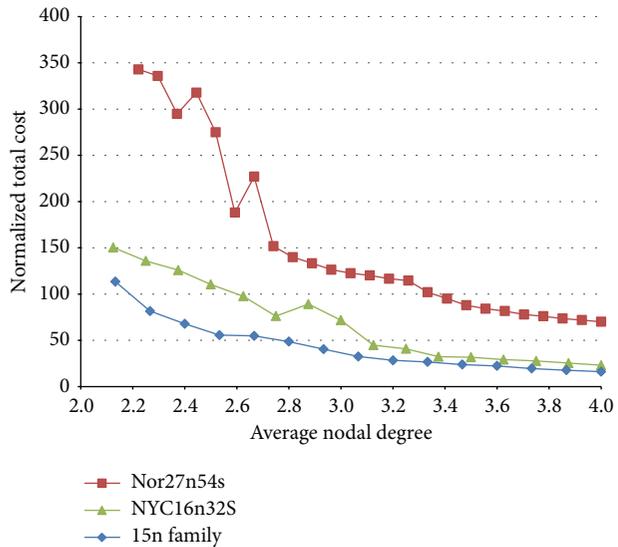


FIGURE 24: Normalized  $p$ -cycle restoration SCA costs for network families based on real network.

obscured the behaviour(s) we sought to characterize. We can therefore assert that a comparative evaluation of design models over varying connectivity (as achieved via network families) allows us to better understand how a model behaves irrespective of the actual scale of cost. As a result, we do not need to run ILP design models with potentially hundreds of stand-alone test case networks; rather, it can be run just a few times over one or two network families. Hence, we are able to reduce overall experimental runtime while obtaining results that are more meaningful.

#### 4. Conclusion

Evaluating survivability schemes and network optimization and design models across a variety of network connectivities

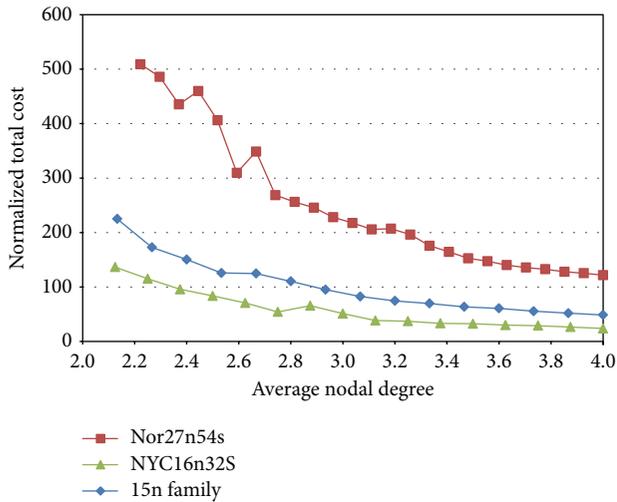


FIGURE 27: Normalized SBPP restoration JCA costs for network families based on real network.

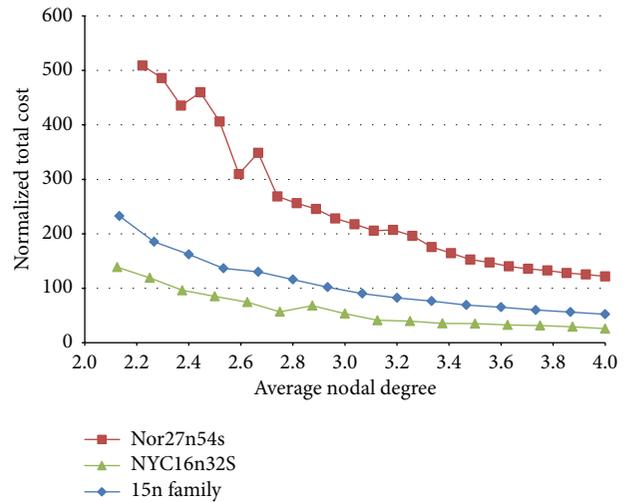


FIGURE 29: Normalized span restoration JCA costs for network families based on real network.

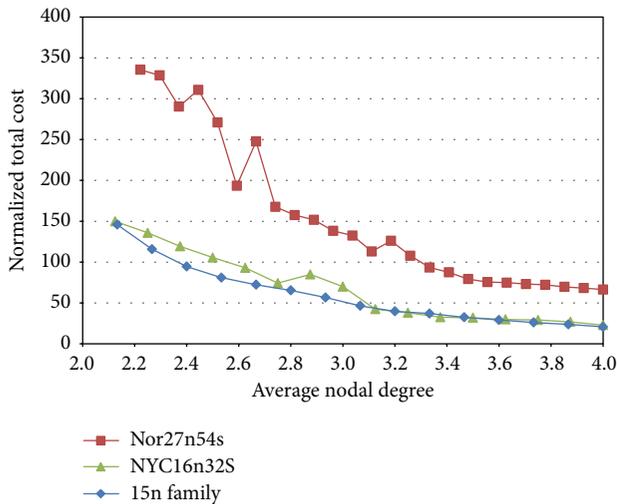


FIGURE 28: Normalized span restoration SCA costs for network families based on real network.

provides a method to generalize the behaviour and performance of the model, and proper selection of a set of test case networks is important. Network families provide a way to evaluate generalized behaviour across varying network connectivities without having to incur the potentially significant computational time required if a large set of random networks were used. It was shown that there was little bias between the average of many pseudorandom networks and a sample network family.

With regard to the use of network families over real network representations, rather than pseudorandom artificial network topologies, the results provide much more confidence when generalized. We observed that use of stand-alone real networks can include peculiarities that could show wildly varying capacity design costs as we sweep through their various connectivity levels. Use of network families prevents this by evaluating results over a consistent set of networks.

## Conflict of Interests

The authors declare that there is no conflict of interests regarding the publication of this paper.

## References

- [1] A. J. Vernon and J. D. Portier, "Protection of optical channels in all-optical networks," in *Proceedings of the National Fiber Optic Engineers Conference (NFOEC '02)*, pp. 1695–1706, Dallas, Tex, USA, September 2002.
- [2] M. Herzberg, S. J. Bye, and A. Utano, "Hop-limit approach for spare-capacity assignment in survivable networks," *IEEE/ACM Transactions on Networking*, vol. 3, no. 6, pp. 775–784, 1995.
- [3] A. M. C. A. Koster, A. Zymolka, M. Jäger, R. Hülsermann, and C. Gerlach, "Demand-wise shared protection for meshed optical networks," in *Proceedings of the Workshop on the Design of Reliable Communication Networks (DRCN '03)*, pp. 85–92, Banff, Canada, October 2003.
- [4] J. Doucette, *Advances on design and analysis of mesh-restorable networks [Ph.D. dissertation]*, University of Alberta, Edmonton, Canada, 2004.
- [5] B. Todd and J. Doucette, "Use of network families in survivable network design and optimization," in *Proceedings of the IEEE International Conference on Communications (ICC '08)*, pp. 151–157, Beijing, China, May 2008.
- [6] R. Bhandari, *Survivable Networks: Algorithms for Diverse Routing*, Kluwer Academic Publishers, 1998.
- [7] J. Doucette, M. Clouqueur, and W. D. Grover, "On the availability and capacity requirements of shared backup path-protected mesh networks," *Optical Networks Magazine*, vol. 4, no. 6, pp. 29–44, 2003.
- [8] S. Sengupta and R. Ramamurthy, "Capacity efficient distributed routing of mesh-restored lightpaths in optical networks," in *Proceedings of the IEEE Global Telecommunications Conference (GLOBECOM '01)*, pp. 2129–2133, San Antonio, Tex, USA, November 2001.
- [9] Y. Xiong and L. G. Mason, "Restoration strategies and spare capacity requirements in self-healing ATM networks,"

- IEEE/ACM Transactions on Networking*, vol. 7, no. 1, pp. 98–110, 1999.
- [10] D. Stamatelakis and W. D. Grover, “Theoretical underpinnings for the efficiency of restorable networks using preconfigured cycles (“p-cycles”),” *IEEE Transactions on Communications*, vol. 48, no. 8, pp. 1262–1265, 2000.
- [11] M. W. Maeda, “Management and control of transparent optical networks,” *IEEE Journal on Selected Areas in Communications*, vol. 16, no. 7, pp. 1008–1023, 1998.
- [12] B. Gavish, P. Trudeau, M. Dror, M. Gendreau, and L. Mason, “Fiberoptic circuit network design under reliability constraints,” *IEEE Journal on Selected Areas in Communications*, vol. 7, no. 8, pp. 1181–1187, 1989.
- [13] W. D. Grover and J. Doucette, “Topological design of survivable mesh-based transport networks,” *Annals of Operations Research*, vol. 106, pp. 79–125, 2001.
- [14] J. Doucette and W. D. Grover, “Comparison of mesh protection and restoration schemes and the dependency on graph connectivity,” in *Proceedings of the Workshop on Design of Reliable Communication Networks (DRCN '01)*, pp. 121–128, Budapest, Hungary, October 2001.
- [15] W. D. Grover, *Mesh-Based Survivable Networks: Options and Strategies for Optical, Mpls, Sonet, and Atm Networking*, Prentice Hall Professional, 2004.
- [16] B. G. Józsa, D. Orincsay, and A. Kern, “Surviving multiple network failures using shared backup path protection,” in *Proceedings of the IEEE International Symposium on Computers and Communication (ISCC '03)*, Antalya, Turkey, July 2003.
- [17] M. Menth, A. Reifert, and J. Milbrandt, “Self-protecting multipaths—a simple and resource-efficient protection switching mechanism for MPLS networks,” in *Proceedings of the 3rd International IFIP-TC6 Networking Conference (Networking '04)*, pp. 526–537, Athens, Greece, May 2004.
- [18] G. Banerjee and D. Sidhu, “Label switched path restoration under two random failures,” in *Proceedings of the IEEE Global Telecommunications Conference (GLOBECOM '01)*, pp. 30–34, San Antonio, Tex, USA, November 2001.
- [19] D. Schupke and R. Prinz, “Performance of path protection and rerouting for WDM networks subject to dual failures,” in *Proceedings of the Optical Fiber Communications Conference (OFC '03)*, vol. 1, pp. 209–210, Atlanta, Ga, USA, March 2003.
- [20] Bandcon International, “International Network,” September 2010, <http://www.bandcon.com/the-system-suite/network/>.
- [21] S. Orłowski, M. Piro, A. Tomaszewski, and R. Wessly, [sndlib](http://sndlib.zib.de), <http://sndlib.zib.de>.
- [22] Internode, “Domestic IP Network,” September 2010, <http://www.internode.on.net/pdf/network/internode-domestic-ip-network.pdf>.
- [23] Deltacom, “Network Map,” 2010, [http://deltacom.com/doc/NetworkReach\\_011110.pdf](http://deltacom.com/doc/NetworkReach_011110.pdf).
- [24] Bytepile, “Network Maps,” 2010, <http://bytepile.com/network.php>.
- [25] COX Business Services, “COX National IP Backbone,” <http://www.coxbusiness.com/backbone/index.html>.
- [26] Optical-Networ, “Network Topologies,” <http://www.optical-network.com/topology.php>.
- [27] British Telecommunication Global Services, “BT North-America,” September 2010, [http://www.bt.net/info/north\\_america.shtml](http://www.bt.net/info/north_america.shtml).



**Hindawi**

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

