December 2003

Articulating User Preferences About Survivability in a Network Planning DSS

Steven Cosares
Hofstra University

Meral Binbasioglu
Hofstra University

ARTICULATING USER PREFERENCES ABOUT SURVIVABILITY IN A NETWORK PLANNING DSS

Steven Cosares
Hofstra University
acsstc@hofstra.edu

Meral Benbasioglu
Hofstra University
acsmxb@hofstra.edu

Abstract

Decision support systems for planning today's high capacity networks, like those in telecommunications, must include features that model "survivability". This property assures that traffic is protected from the potential damages associated with some disruption, like a break in a network link. When a network provides "100% survivability", the property is fairly well defined: all of the traffic is protected against any network failure. However, if a planner is willing to accept some smaller percentage of survivability, e.g., to lower the cost of the network, then some ambiguities arise. There are a variety of ways that partial survivability can be interpreted, measured, and provided. Currently available decision support systems for network planning are limited in their ability to model these different options. In this paper, we proposed some features that could be added as enhancements to a DSS, so that it can better satisfy a planner's needs regarding survivability. We describe features to model the alternative schemes available to measure partial survivability. These would give the planner the opportunity to articulate more meaningful objectives regarding network survivability. The enhanced DSS could then invoke the appropriate models and algorithms to help the planner identify what she would consider to be superior network solutions. In addition, such a system could help the planner evaluate the trade-offs between network survivability and network cost.

Introduction: The Network Planning Domain

An individual responsible for the planning of some commodity distribution network, like those arising in transportation, power generation, or telecommunications, must consider a variety of criteria in order to evaluate her designs. First, she has to be confident that the network she plans will be able to accommodate the demands for traffic between locations in a cost effective manner. Then additional qualities are considered. Even networks of modest size give rise to some very large, very complex problems that require the planner to employ a sophisticated computer-based decision support system, like the ones described in Cosares, et. al. (1995) and Mikhail, et. al. (1996). The planner would use the DSS to store data about the point-to-point demands, (i.e., the number of units of the commodity requested between each pair of locations), the structure of the underlying network, (represented as a set of links and nodes), and the costs of establishing connections and expanding network capacity. Then the DSS would help the planner determine, for some time horizon, which links and nodes in the network to activate, how much capacity to purchase for each, and how to route the traffic flow over the resulting network. To match the specific domain of the planner, a telecommunications system would suggest placements of transmission systems and purchases of cable and equipment, while a transportation system would suggest a schedule for the purchase of new vehicles and how to route them to make their deliveries.

The criteria that a network planner uses to identify a good plan from amongst the multitude of feasible possibilities often come in conflict with each other. This means that the planner must evaluate trade-offs between key objectives. For example, a telecommunications planner may wish to build a network with extra capacity to assure that unexpected fluctuations in demand can be accommodated. However, this would subvert his or her objective of building the network at minimal cost. A transportation planner may wish to take advantage of economies of scale by using vehicles with larger capacities, but risks severe losses should a vehicle be delayed or rendered inoperable because of an accident. A challenge to those who design network-planning systems is to develop features that help a planner evaluate network plans and identify superior planning solutions, based on the decision criteria she finds most important. Since different planners may express different objectives, the DSS would need a variety of models and algorithms; the ones that best address a particular planner's preferences would be invoked.

In this paper we focus on network "survivability". This is defined to be the extent to which uninterrupted service is provided, even in the event of an unexpected failure in some network component(s). Telecommunications networks that use fiber-optics based technologies, for example, employ high capacity systems capable of moving billions of bytes of data and voice traffic between locations every second. If any link or node in this network were rendered inoperable, the financial consequences could be severe. Large financial services firms claim to lose millions of dollars each time their service is disrupted for more than a few seconds. There are also incalculable societal costs if, say "911" service is suspended for any appreciable length of time. Hence, network planners often want the DSS to help them develop plans that provide some level of survivability. Unlike costs, which can be measured by counting the dollars spent in a plan, the amount of survivability provided is harder to quantify. There are a number of reasonable ways to interpret and measure survivability. The DSS would need features to help a planner identify precisely what she wants when she requests a certain amount of survivability. Then the DSS would invoke the appropriate models and algorithms to measure the survivability provided in a plan and possibly suggest reasonable alternative planning solutions.

We describe a set of enhancements that can be made to a network planning DSS that allow it to present comprehensive information to the planner regarding the survivability provided in a plan. Based on some articulation of the planner's priorities, (developed, e.g., using the techniques described in Sun and Liu (2001)), the appropriate models could be selected for the DSS to help the planner evaluate and rank any proposed network solutions and/or evaluate the trade-off between survivability and other decision criteria, like cost. Though we focus on telecommunications, the approaches we develop can be applied to other types of networks that support the secure delivery of commodities.

Domain Component: Survivability

After the data regarding the network structure, the demands, and the costs is loaded, a DSS for survivable network planning helps the user develop a plan that is assumed to consist of (at least) the following components:

- A set of equipment and cable purchases, which establish (expand) the capacity of the links and nodes in the network. The system's models usually enforce a number of technical constraints that govern how demand is feasibly assigned to this capacity, (e.g., high bandwidth demands can only be assigned to high capacity transmission systems, appropriate terminating and cross-connection equipment needs to be purchased). These decisions directly affect the total cost of the plan.

- A routing of the demands that is consistent with the capacity decisions made above. A route consists of a sequence of links (transmission systems) from one location to another, inter-connected by node (switching) equipment. A demand is satisfied when it is assigned to one or more feasible routes between its origin and destination having sufficient total capacity. Clearly, the routing decisions impact the capacity purchase decisions and vice versa.

- A contingency plan for directing traffic, in the case of every potential network failure scenario. If some link (or node) is incapacitated, new feasible paths have to be found for the affected demands. Some transmission systems are built with dedicated capacity to protect its traffic. In other cases, flexible switching equipment helps redirect traffic to avoid failed locations, (see Wu (1992) and Deemster, et. al. (1999)). Like the initial "working" routes determined above, the selection of these "protection" routes must be consistent with the capacity available in the network. If there is insufficient capacity or some network elements' physical constraints limit their ability to provide protection, then some demands may be (temporarily) unserved during a network disruption.

We say that the plan provides "100% survivability" if no demand that needs protection is left unserved in any of the network failure scenarios considered. (We point out that if a node fails, any demand originating or terminating at the node cannot be recovered, so those demands are not counted in the calculation of the survivability.) The network plan components are illustrated with the following very small example (see Figure 1).

The values next to the links in Figures 1a and 1b represent the capacities. Since additional capacity can be purchased wherever needed, it should be evident that there is a trade-off between maximizing survivability and minimizing costs. Figure 1a above shows that demand #1, for service between nodes A and G, is for 5 units. To reduce its vulnerability, the demand is split so that 3 of the units are assigned to the route from A to B to C to G, while 2 of the units are assigned to the route from A to D to E to F to G. Demand #2, between nodes D and F, is for 3 units that are all assigned to the route from D to E to F. Notice that there is sufficient capacity in each of the links to accommodate the working routes for the demands.

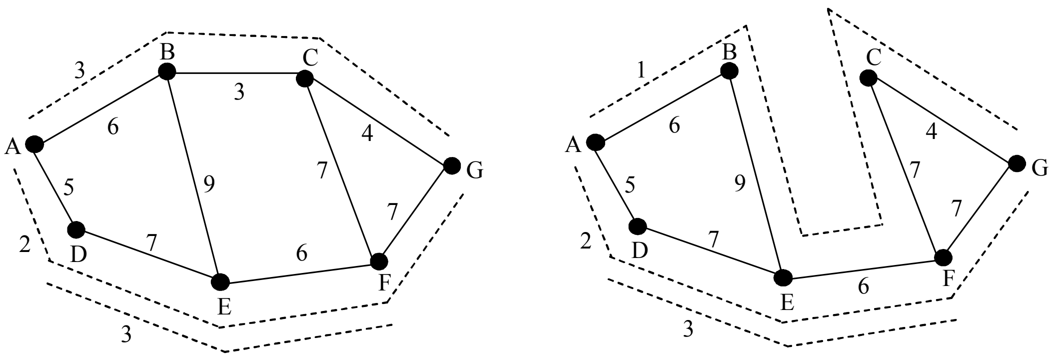

1a) Working Routes **1b) Protection Routes**

Figure 1

Now suppose that the link between nodes B and C were disabled. Then 3 out of the 5 units of demand #1 are affected by this failure. If the appropriate switching equipment has been put into the network, the contingency plan may state that demands affected by this link failure be re-routed as shown in Figure 1b. Because there are only 6 units of capacity on the link between nodes E and F, the amount of demand #1 that can be restored is limited to only 1 unit out of the 3 units affected by the failure. Hence this plan *does not* provide 100% survivability. Additional capacity would have to be purchased in order to accommodate a planner preference for greater survivability.

The central questions are "If 100% survivability is not provided in this plan, then what percentage adequately describes how much survivability *is* provided?" and "If a planner states she will accept 90% survivability how do we know when this constraint has been met?"

While 100% survivability is a well-defined concept – protect all of the demands from any failure scenario – partial survivability is more ambiguous. If a planner is willing to accept 90% survivability, she may be placing a constraint on every demand or a constraint that involves some single aggregate measure representing the entire network plan. Furthermore, the constraint could be associated with the worst-case failure rather than some average case failure event. Alternatively, the intent of this percentage may be to constrain the amount of overall damage caused by any single network failure event.

While there is much ambiguity, the example in Figure 1 suggests that a DSS can make the following reasonable calculations. For each scenario *s*, representing the failure of some single link or node in the network, let the value β_{ks} be the proportion of the total demand #*k* that is affected. For example, in the failure scenario described in Figure 1b, the values of β are 60% for demand #1 and 0% for demand #2. Let the value α_{ks} be the proportion of the total demand #*k* that are not restored in the plan. In the scenario in Figure 1b, the values of the α are 40% for demand #1 and 0% for demand #2.

A DSS with enhancements for partial survivability would calculate these values for only those demands that the planner specifies as requiring protection. In a typical network-planning situation, certain demand types lack the importance to justify the extra expenditures that protection requires. In addition, the DSS would allow the planner to identify the set of failure scenarios to consider. This may include scenarios where more than one network component fails at the same time. A planner preference for 100% survivability would require that the values for α be 0 for *all* of the demands considered, under *all* of the failure scenarios considered.

The values for the β measure the "preventative" survivability in the network plan. They indicate the extent to which the (working) routes for the network demands are exposed to the dangers associated with a failure scenario. If the working routes for a demand are diversified, i.e., the demand is spread out over multiple (disjoint) paths wherever possible, then only a fraction of the units will be present at any link or node in the network. However, the demand will be present at more network locations, so a greater number of failure scenarios will impact the demand, giving a greater number of positive β. However, the severity of any single failure incident will be reduced, so the magnitudes of the β will be smaller. Conversely, if a demand uses a single route, it decreases the probability, but increases β (the severity), of a link or node failure along its route. Notice that by definition that the α values, which represent the affected demands that are not be saved, must be less than or equal to the β values. If the α values are closer to the β values than to 0, then that would indicate that very little survivability is provided in the plan.

DSS Feature: Presenting Data to the Network Planner

Various levels of support are provided by existing network planning systems to help the user develop a survivable plan. The DSS described in Mikhail, et. al. (1996), for example, is used to evaluate a network plan that is provided by the user. In addition to calculating the costs and other characteristics of the plan, it is designed to indicate when the goal of 100% survivability has not been met. This type of system supports a planner who performs "what-if" analysis; it is most appropriate in cases where she has a high level of knowledge about network planning. The system can be used to "fine tune" a solution, or to determine where capacity should be added or removed, or to provide justifications for planning decisions that may have already been made, (e.g., for reporting to regulating agencies). Certainly, such a system would provide additional benefits were it to include diagnostics regarding *partial* survivability as well.

Other types of systems, like those described in Cosares, et. al. (1995) and in a recent press release by RSoft (2002), have functionally that "completes" a partial plan provided by the user. Generally speaking, these systems use heuristics to find some low cost feasible solution that satisfies constraints set by the user. With rare exceptions, however, these systems allow few (if any) constraints regarding partial survivability. For example, the technologies considered in Cosares, et. al. (1995) are assumed to provide complete protection, so partial survivability levels are not modeled at all. Myung, et. al. (1999) use a measure based on the traffic lost in the worst-case failure. This allows the planner to evaluate the trade-off between cost and some single (limited) measure of survivability, but falls short of the flexibility we believe is required. In this and in other cases, the selected definition for partial survivability is not representative of a planner's preferences, but rather matches some fixed model structure or algorithm that is embedded in the DSS. We believe that this approach puts the cart before the horse. First planners need to be better informed about the variety of options available for measuring survivability. Network planning systems need to be enhanced with features that help the user disambiguate the domain aspects to provide her with relevant survivability data and to guide her in setting appropriate survivability constraints. In addition, the systems need be enhanced with functionality to give them the ability to flexibly invoke appropriate algorithms to generate solutions based on the particular survivability measures selected by the planner.

In this section, we describe how the first of these requirements can be met. The values for α_{ks}, which represent the fraction of each demand lost in every network failure scenario, provide the basis for any survivability measurement used by the planner. Individually, these values have little relevance. The DSS enhancements would include features that present a variety of reasonable descriptive statistics that organize and summarize these data in a way that gives the planner a clear understanding about how much survivability a particular network plan has. (Similar descriptions could be provided that are based on the β values as well). These measures would also provide a *language* through which a planner can set meaningful survivability parameters for DSS algorithms that generate proposed network planning solutions.

Demand-Based Measures

Many systems allow the user to select which individual point-to-point demands to provide with 100% protection. A natural extension would be for the DSS to allow the user to instead supply *any* desired percentage for each demand. In order for such constraints to be meaningful, the planner would need to be comfortable with the following diagram which represents some demand #k. It is generated by ordering the scenarios *s* by increasing value of α_{ks}.

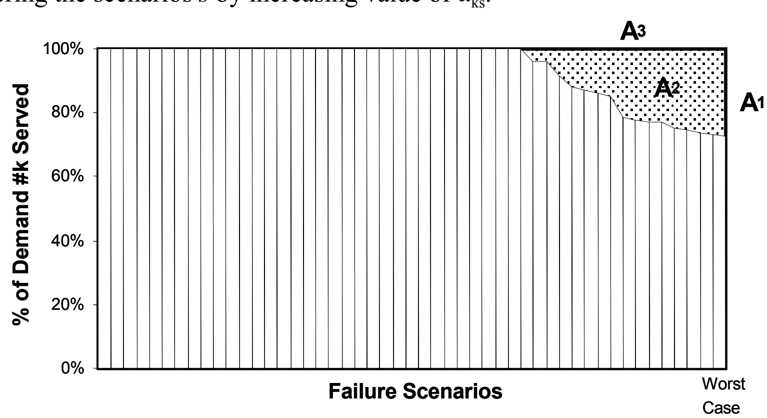

Figure 2. The Protection of Demand #k

If the demand were provided with 100% protection, the shaded region on the upper right corner would not exist. The region could be made smaller if additional capacity were effectively placed in the network. With the DSS providing this visual representation, the planner can see that there are three alternative ways to measure the percentage of protection that is provided to demand #k. The value A_1 represents the percentage of demand units that would be lost in the worst-case failure scenario. (Other demands may have a different failure scenario as its worst case). The value A_2 represents the average percentage of units of demand #k lost, over the set of scenarios considered. The value A_3 represents the percentage of scenarios in which some amount of demand #k is lost. It can also be interpreted as the probability that the demand would be impacted, should some failure event occur. Any of these measures is appropriate in some network-planning context, so the DSS should make all three values available.

The picture above represents the case where each failure scenario included is assumed to be equally likely. If the planner has some probability information about the likelihood of each scenario, it could be presented to the DSS, which would adjust the widths of the corresponding rectangles accordingly, (keeping the total of the widths equal to 1.0). This would also allow the system to ignore or deemphasize scenarios that are deemed unlikely or unimportant. Planners performing "what-if" analyses would note that, for a fixed monetary budget, A_2 is somewhat stable. The planner will likely have to trade off low values of A_1 (representing the potential severity of some catastrophic loss) with higher values for A_3,(the probability of some loss in the demand). Systems having algorithms to generate planning solutions would provide the user with the ability to set some lower bound constraint on the level of protection provided to each demand, based on one (ore more) of these measures. For example, one planner requiring 90% survivability may state that A_1 must be less than 10% for all of the demands, while another would base the constraints on A_2.

Scenario-Based Measures

An additional opportunity is afforded to the planner that wishes to measure (or constrain) the severity of any particular failure scenario *s*. The following diagram is developed by ordering the point-to-point demands *k* by increasing value of α_{ks}.

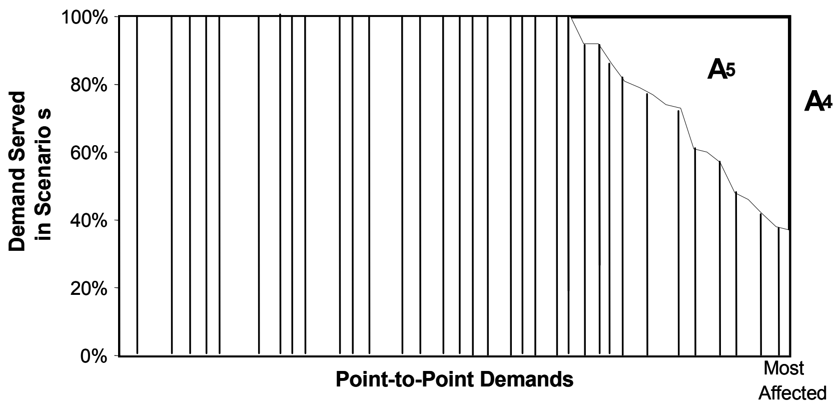

Figure 3. The Impact of Network Failure Scenario s

In this diagram, the bars have widths that are proportional to the number of units in the demand, (where the total width is set to 1.0). The diagrams associated with the more likely failure scenarios identified by the planner could be presented more prominently by the DSS. This would help the planner improve her "what-if" analyses and help her identify locations where additional capacity purchases would be most effective. The area represented by A_5 measures the total impact of the node or link failure scenario *s*, as a percentage of the total demand. (1-A_4) represents the maximum percentage of restoration guaranteed to the demands if failure scenario *s* should occur. A high value for A_4 would mean that some demand or demands are particularly sensitive to this particular failure. It may also indicate when the plan for rerouting demands after the failure is not equitable, particularly if the other demands have a smaller value for α. (Note that a major challenge in the design of DSS algorithms for partial survivability is to establish a mechanism that matches a user's preferences when determining which demands to restore and which demands to leave unserved after a failure). A high value for either of the measures would indicate to the planner that survivability would be improved if there were a greater number or more capacious paths that can serve as a backup to the affected link or node.

Aggregate Measurements

In addition to (or instead of) placing constraints on individual point-to-point demands or individual network failure scenarios, the DSS may allow the planner to measure or control some more broad (network-based) measure for survivability. This would allow the planner to rank plans based on their survivability or to evaluate the trade-off between some single survivability measure and the cost of the plan. (While we strongly doubt that any DSS could provide a credible regression model that would quantify the unit cost of a percent of survivability, we believe that some "efficient frontier" representing the trade-offs can be determined with the help of a DSS). The most obvious choices for a single measure would be the largest of the α values or some weighted average of the α values, (using either the relative demand volumes or the relative scenario likelihoods as weights).

Alternatively, a diagram, similar in structure to the one in Figure 2, can be developed by taking the summary value of the choice from each of the scenario-based diagrams, (e.g., A_5), and then tracking these values (in increasing order) across all likely scenarios. Then one of the three summaries would be calculated to provide a single measure for the whole network, (e.g., the measure selected in Myung, et. al. (1999) is the maximum over all scenarios of the A_5 from each scenario). Similarly, by taking a summary value from each of the demand-based diagrams (e.g., A_1), and tracking the values across all demands, one can develop a diagram, similar in structure to the one in Figure 3. Then some summary from these values can be calculated, (e.g., the "average worst case demand loss").

Conclusions

In this paper, we have described a variety of means by which a network planning DSS can be improved to measure the (partial) survivability provided in a plan. These data can be presented to a planner in forms that demonstrate how well protected the demands in a given plan are. They also provide a language where the planner can express specific survivability constraints to DSS algorithms that generate potential planning solutions. Based on the prevalent objectives and the network application under consideration, the planner may want to limit the worst-case severity of a network failure instance. In another case, she may want to direct the impact of a failure onto a select set of demands while keeping the remainder of the plan robust. Alternatively she may decide that any loss of demand is unacceptable; the DSS would have to find the most cost effective network plan that provides 100% survivability. The measures described in this document allow any of these objectives to be represented.

The development of descriptive, flexible models for measuring survivability is just one part of the process of improving the available mechanized decision support for robust network planning. Follow up activities include designing additional DSS features that poll the planner to determine the best way to present the information so that the more preferred measurements are more prominently presented. (In addition, any appropriate (linear) combination of measurements can also be calculated, based on the information obtained about the planner's preferences). Also helpful would be features that demonstrate how sensitive a network plan is to the constraints set by the planner. The DSS could demonstrate the impacts of marginal changes to help the planner learn how to set values for each measure and exercise better control over the final plan. Another enhancement would be features that test the consistency of planner inputs to identify when some mix of constraints is impossible to achieve.

The constraints for survivability complement other settings and parameters made available to the user of a network planning DSS to control the solutions it generates. Experienced planners can convey information about issues like how much embedded capacity to use, which demands to protect, how diversely to set the initial routings, which specific elements to include into the plan, etc. For the novice, who may be unfamiliar with the network planning domain, a DSS can employ a knowledge-based approach, (see, e.g., Binbasioglu and Jarke (1986)), to explicate the various model components and interactively and adaptively guide the planner through the process of determining appropriate parameters to effectively control the quality of the plan. Based on the information provided herein, survivability parameters can be one of the first components so implemented.

Another set of activities to enhance network planning decision support is the development of algorithms for generating (cost-effective) solutions that are consistent with the varied survivability constraints ultimately set by the (expert or novice) planner. Existing methodologies are somewhat capable of determining capacity levels that meet planner constraints on individual demands or on individual failure scenarios, but are less effective in satisfying average case or overall network aggregate targets. We anticipate that the optimization problems arising in this domain would be sufficiently complex to attract the interest of researchers for years to come!

References

Binbasioglu, M., Jarke, M., "Domain Specific DSS Tools for Knowledge-Based Model Building", Decision Support Systems (2), pp. 213-223, 1986.

Cosares, S., Deutsch, D., Saniee, I., and Wasem, O., "SONET Toolkit: A Decision Support System for the Design of Robust and Cost-Effective Fiber-Optic Networks", Interfaces (25), pp. 20-40, 1995.

Demeester, P., Wu, T. H., and Yoshikai, N., eds. IEEE Communications Magazine, Special Issue on Survivable Communication Networks, (37), August 1999.

Mikhail, M., Shallcross, D., Dean, N., and Mostrel, M., "A Commercial Application of Survivable Network Design: ITP/INPLANS CCS Network Topology Analyzer", Symposium on Discrete Algorithms 1996, pp. 279-287.

Myung, Y., Kim, H., and Tcha, D., "Design of Communication Networks with Survivability Constraints", Management Science, (45), pp. 238-252, 1999.

"RSoft Design Group's MetroWAND Makes Metro Mesh Migration Easy", Press Release, September 2002.

Sun, L. and Liu, K., "A method for interactive articulation of information requirements for strategic decision support", Information and Software Technology (43), pp. 247-263, 2001.

Wu, T. H., Fiber Network Service Survivability: Architectures, Technologies, and Designs, Artech House, 1992.