A Novel Two-Step Approach to Surviving Facility Failures

Chunming Qiao[1], Bingli Guo[2], Shanguo Huang[2], Jianping Wang[3], Ting Wang[4], Wanyi Gu[2]

[1]*Department of Computer Science and Engineering, University at Buffalo (SUNY), New York, USA*
[2]*Key Laboratory of Information Photonics and Optical Communications, Beijing University of Post and Telecommunition, Beijing, China*
[3]*Department of Computer Science, City University of Hong Kong, China*
[4]*NEC Laboratories America, Princeton, NJ, USA*
Email: qiao@buffalo.edu, {bingliguo,shghuang,wyg}@bupt.edu.cn, jianwang@cityu.edu.hk, ting@nec-labs.com,

Abstract: To provide resilient services over an optical network, we first enhance virtual network representing a service request into a *failure dependent and enhanced* VN (FD-EVN), then map FD-EVN to the optical network with shared protection.

OCIS codes: (060.4510) Optical communication; (060.4257) Networks, network survivability

1. Introduction

Future optical networks will provide a rich set of distributed services. Such a service request may be represented using a graph (called *virtual network or VN*) consisting of a number of service nodes and service links. Each service node (link) is associated with some required computing resources (such as processing power and storage) or communication resources (i.e. bandwidth). To provide such a service over a substrate (consisting of *facility* nodes interconnected by an optical network), one needs to find a *distinct* facility node and reserve sufficient computing resources for each service node, and in addition, find a physical path in the optical network and reserve sufficient bandwidth for each service link. This is often referred to as a VN mapping problem.

In this work, we consider the problem of providing resilient services against any single facility node failure which will affect at most one service node of any service request. An entire facilty node may fail due to power outage, virus attacks, disk failures or crashed software etc. Having a survivable optical network alone is *not* sufficient to protect the services against such a failure, as the service node affected by such a failure needs to be migrated to a backup facility node at a geographically different location. There have only been a few approaches to providing resilient services. For example, survivable VN mapping (or SVNM) in [3,6] maps a given N-node (and E-link) VN to *at least* N+1 facility nodes (and > E paths) in the substrate such that after any single facility node failure, there will still be sufficient resources to provide the desired service.

In this paper, we take an alternative, two-step approach: we first *enhance* the VN into an Enhanced VN (EVN) by adding one service node and e≤N additional service links to the VN, and then map the EVN to only N+1 facility nodes and E+e paths in the substrate. Our approach differs from other two-step approaches recently studied in [4,5,7,8] in that we design our EVNs in a *failure-dependent* (FD) fashion whereas the others are *failure–independent* (FI). Our FD-EVN usually requires a fewer virtual resources associated with links and/or service nodes than FI-EVN, and in turn a smaller amount of communication and/or computing resources after mapping such an EVN to the substrate. Note that, our mapping also differs from VN mapping and SVNM studied earlier in that we now have to exploit resources sharing among e.g, the E+e service links, which will not be used concurrently with or without a facility node failure. However, due to space limit, we will focus only on EVN design in the first step in this paper.

This paper is the first that 1) combines the two-step approach with a FD design of EVN and 2) compares two-step FD-EVN and FI-EVN based approaches with SVNM approaches. We show that the proposed approach can improve over all other existing approaches.

2. Design an Enhanced Virtula Network (EVN)

Our EVN design approach can be described based on the concept of *edit grid* and *graph edit distance* [9]. An edit grid is a *completely connected graph,* where each node and link is assumed to have an infinite computing and communication (C&C) resources available. Given an N-node VN, we can obtain a EVN by starting with an (N+1)-node edit grid, selecting an appropriate set of links from the edit grid to connect these N+1 nodes, and reserving a sufficient amount of C&C resources at these nodes and links within the edit grid, respectively. Our design objective is to minimize the total amount of such resources while still guaranteeing that if a node in the edit grid fails (that is, the node and its adjacent links are removed from the edit grid), we can still assign each node/link in VN to a node/link in the edit grid, that have sufficient C&C resources respectively.

Example: Given the symmetry of the edit grid, how the VN is initially assigned does not matter. Fig.1 (a) shows an example assignment of a 5-node VN, which requires a total of 16 and 24 units of computing resources (shown in

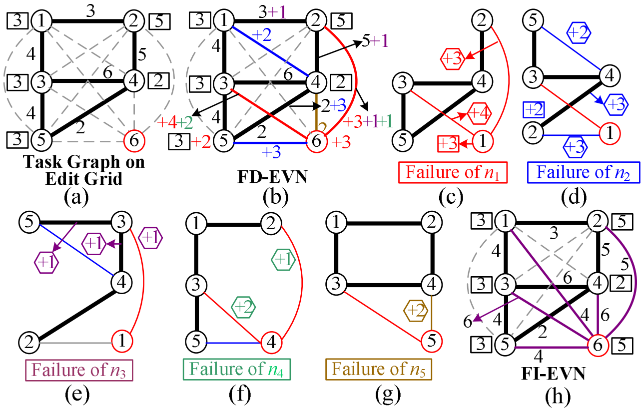

Fig. 1 (a). VN; (b) FD-EVN; (c)-(g). Partially enhanced FD-EVN to handle each node failure scenario; (h) FI-EVN design.

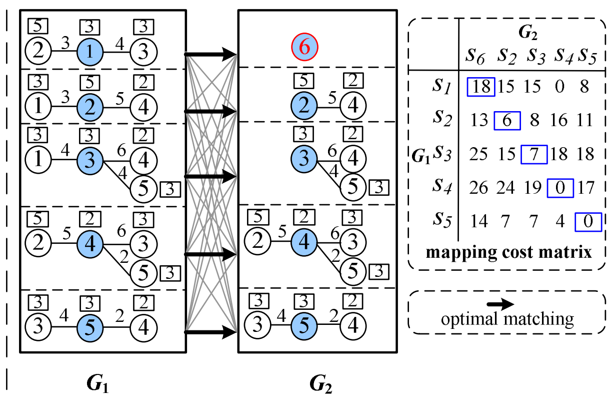

Fig. 2 Graph Transformation and Bipartite Graph Matching.

each square beside a node) and communication resources (shown as a number associated with each link), in a 6-node edit grid. The bold links depict those in the VN while the dashed links depict those in the edit grid. It can be seen that n_i (i=1…5) of the VN is assigned to N_i in the edit grid, and N_6 in the edit grid is not assigned any node in the VN. Fig. 1(b) shows that as a result of our design, an 6-node FD-EVN example with one additional service node (assigned to edit grid node 6), four additional service links (shown in solid but thin lines), and additional C&C resources (shown as the number after the "+" sign). Compared to the VN, such a FD-EVN requires additional 5 and 23 units of C&C resources respectively. These additional resource requirements allow us to assign the original VN to the edit grid even after the removal of any single edit grid node (and its adjacent links) in the edit grid, or equiavelntly, a failure of the assigned service node in the FD-EVN, as shown in Fig. 1 (c) to (g), where the numbers in hexagon and rectangle represent, respectively, additional communication and computing) resources needed by the FD-EVN links and nodes, respectively, for each service node failure scenario. It is worth noting that, to tolerate the failure of edit grid node N_2, service node n_1 (*not* n_2) is assigned to edit grid N_6 as shown in Fig. 1 (d). This failure-dependent assignment *distinguishes* our FD-EVN from a FI-EVN shown in Fig. 1 (h), where a failed service node will *always* be assigned to edit grid node N_6. In addition, our FD-EVN needs one fewer link and two fewer units of communications resources than the FI-EVN (whose design is usually straight-forward).

Mathematic Formulation and Heuristics: Based on the discussion above, we mathematically formulate the optimal FD-EVN design problem with the objective being to minimize the total cost (or in terms of units) of C&C resources. Due to space limit, we omit the detailed description but sufficie it to say that since the optimal FD-EVN design problem is known to be NP-hard, the mathematical formulation is intractable for a large N.

Accordingly, we also develop efficient heuristic algorithms. One basic idea of our heuristics is that, given an initial VN assignment, we consider all possible N failures (one for each edit grid node) *one at a time* according to a random order. An alternative to this *Random* heuristic is that we will consider the failure of the edit grid node requiring the *minimum amount of C&C resources first* (called the *MinFrst* heuristic hereafter). Our design objective is to find *an incrementally optimal partial enhancement* requiring a *minimum amount of additional* C&C resources. Below, we illustrate this idea assuming that we consider the failure of edit grid node N_1 first.

Graph Transformation/Decompsition and Bipartite Graph Matching: To simplify the task of VN enhancement, Fig. 2 shows how the 5-node VN in Fig. 1 can be decomposed into a set of five (5) star structures (shown as G_1), each containing a primary service node (which is highlighted), its adjacent links and its neighboring nodes. Similarly, the latest (partially enhanced) version of the FD-EVN in the edit grid (which at first is the same as the VN), with N_1 and its adjacent links removed, can also be decomposed (shown as G_2). For each star-structure s_i (where $1\leq i\leq 5$) in G_1, we calculate the cost, in terms of additional C&C resources, of enhancing s_j (where $1\leq j\leq 5$) in G_2 with additional nodes and links such that the latter contains the s_i in G_1 as a subgraph. Note that the constraints are 1) the primay service node x of s_i in G_1 has to be assigned to the primary edit grid node of s_j in G_2, (after which the latter is renumbered to x), and 2) no other edit grid nodes can be renumbered although new edit grid nodes can be created, along with new edit grid links with additional amount of C&C resources. This cost is shown as the (i, j)-th entry in an $N\times N$ mapping cost matrix in Fig. 2.

By adding a weighted edge between the two sets of star-structures, we derive a bipartite graph, whose optimal solution (obtainable in polynomial time [10]) will lead to the partial enhancement requiring the minimum additional C&C resources, which can handle the failure of N_1 in edit grid. Finally, we repeat the procedure of graph transformation/decomposition and the bipartite graph matching for each following node failure to construct the final

FD-EVN. It is worth noting that, every step of partial enhancement is based on the latest version of edit grid including the additional C&C resources needed for the pervious node failure, which explains why the sequence of node failure to be considered would impact the *total amount of additional* C&C resources needed.

3. Numerica Results

We have compared our FD-EVN designs based on our mathematic formulation and the *MinFst* and *Random* heuristics for VNs whose N≤4, and found that the two heuristics result in close to optimal FD-EVN designs, with *MinFst* slightly better than *Random*. We have also implemented our two-step approaches that use a heuristic similar to that in [2] to map such FD-EVNs during the second step, and compared their performance with the following two approaches: a SVNM approach whereby a VN is directly mapped to the substrate based on FD protection using a heuristic similar to that in [3,6] (hereafter shown as FD-IOCM), and a two-step approach that builds a FI-EVN first, and then maps it to the subsrate using a heuristic similar to that in [7] (shown as FI-EVN). We have used the GT-ITM tool [1] to generate a 100-node substrate and assumed a uniformly distributed C&C resources at its nodes and links. In each simulation run of 50,000 time units, a failure occurs at a randomly selected facility location every 2,000 units. We have also randomly generated about 5000 service requests in each run following a Poissona arrival process (at a rate of 0.1) and having an exponentially distributed service duration (with a mean of 1000 time units). Each service request has a randomly generated size (N) and topology as well as C&C resource requirements.

The proposed approaches outperform the other representative approaches in terms of total amount of C&C resources needed and service request acceptance ratio as shown in Fig. 3 and 4, but as a tradeoff, require more service nodes to migrate after failures as shown in Fig. 5. In particular, the proposed FD-EVN approaches save approximately 25% to 30% resources when compared to FD-IOCM and FI-EVN. Also, a 20% improvement in the acceptance ratio is achieved by FD-EVN. As a tradeoff, however, in FD-EVN approaches, a service node has a higher (up to 4 times higher when $N \geq 8$) probability to migrate (in order to cope with a facility node failure) than FI-EVN and FD-IOCM.

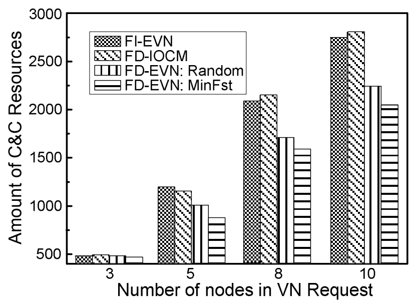

Fig. 3 The amount of C&C resources needed

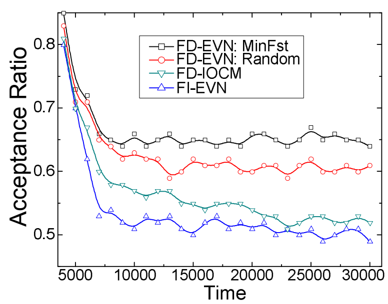

Fig. 4 Acceptance Ratio

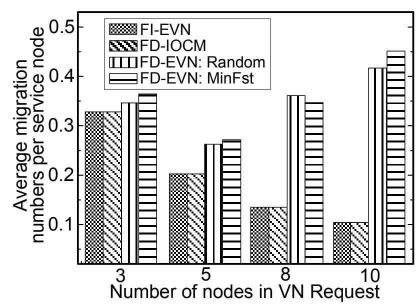

Fig. 5 Averag migration numbers per service node

4. Concluding Remarks

In this work, we have taken a holistic approach to providing reliable services by first enhancing a service request's resilience with built-in redundancy in a proactive way, and then mapping it. Our novel failure-dependent enhancement is the first of its kind and has shown to be resource efficient and in particular, outperform other approaches. The proposed heuristics that use graph transformation, decomposition and bipartite matching may also be useful for finding minimum graph edit distance in pattern recognition and biological network structure analysis.

References

[1] N. M. Mosharaf at al., "Virtual Network Embedding with Coordinated Node and Link Mapping", IEEE INFOCOM, 2009
[2] M. Yu, Y. Yi, J. Rexford, and M. Chiang, "Rethinking virtual network embedding: Substrate support for path splitting and migration," ACM SIGCOM MCCR, vol.38, no.2, pp.17 29, 2008.
[3] H. Yu, C. Qiao, et al, "Survivable Virtual Infrastructure Mapping in a Federated Computing and Networking System under Single Regional failures", accepted by IEEE Globecom 2010.
[4] X.Liu, C.Qiao, et al., "Robust Application Specific and Agile Private (ASAP) Networks Withstanding Multi-layer Failures", OFC/NFOEC, 2009.
[5] X. Yu, C. Qiao, et al., "Survivable logical topology design for distributed computing in WDM networks", OFC/NFOEC, 2009,
[6] H. Yu, C. Qiao, et al. "On the Survivable Virtual Infrastructure Mapping Problem", IEEE ICCCN, Aug. 2010.
[7] H. Yu et al, "Cost efficient design of survivable virtual infrastructure to recover from facility node failure", submitted to ICC 2011.
[8] Yeow, Wai-Leong et al., "Designing and embedding reliable virtual infrastructures", VISA '10, 2010, 33-40.
[9] Derek Justice, Alfred Hero, "A binary linear programming formulation of the graph edit distance", IEEE Transactions on Pattern Analysis and Machine Intelligence. 2006, vol. 28, p.1214.
[10] H. W. Kuhn. "The Hungarian method for the assignment problem. Naval Research Logistics", 2: 83–97, 1955.