



Notions of Connectivity in Overlay Networks

Pierre Fraigniaud, Amos Korman, Shay Kutten, David Peleg, Emek Yuval

► To cite this version:

Pierre Fraigniaud, Amos Korman, Shay Kutten, David Peleg, Emek Yuval. Notions of Connectivity in Overlay Networks. Structural Information and Communication Complexity - 19th International Colloquium, SIROCCO 2012, Jun 2012, Reykjavik, Iceland. 10.1007/978-3-642-31104-8_3. hal-01241098

HAL Id: hal-01241098

<https://inria.hal.science/hal-01241098>

Submitted on 6 Jan 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Notions of Connectivity in Overlay Networks ^{*}

Yuval Emek [†] Pierre Fraigniaud [‡] Amos Korman [§] Shay Kutten [¶]
David Peleg ^{||}

December 9, 2015

Abstract

“How well connected is the network?” This is one of the most fundamental questions one would ask when facing the challenge of designing a communication network. Three major notions of connectivity have been considered in the literature, but in the context of traditional (single-layer) networks, they turn out to be equivalent. This paper introduces a model for studying the three notions of connectivity in *multi-layer* networks. Using this model, it is easy to demonstrate that in multi-layer networks the three notions may differ dramatically. Unfortunately, in contrast to the single-layer case, where the values of the three connectivity notions can be computed efficiently, it has been recently shown in the context of WDM networks (results that can be easily translated to our model) that the values of two of these notions of connectivity are hard to compute or even approximate in multi-layer networks. The current paper shed some positive light into the multi-layer connectivity topic: we show that the value of the third connectivity notion can be computed in polynomial time and develop an approximation for the construction of well connected overlay networks.

^{*}Supported by a France-Israel cooperation grant (“Mutli-Computing” project) from the France Ministry of Science and Israel Ministry of Science. Supported by the ANR projects DISPLEXITY and PROSE, and by the INRIA project GANG

[†]Computer Engineering and Networks Laboratory, ETH Zurich, Zurich, Switzerland. E-mail: yuval.emek@tik.ee.ethz.ch.

[‡]CNRS and University of Paris Diderot, France. E-mail: pierref@liafa.jussieu.fr.

[§]CNRS and University of Paris Diderot, France. E-mail: amos.korman@liafa.jussieu.fr.

[¶]Information Systems Group, Faculty of IE&M, The Technion, Haifa, 32000 Israel. E-mail: kutten@ie.technion.ac.il.

^{||}Department of Computer Science and Applied Mathematics, The Weizmann Institute of Science, Rehovot, 76100 Israel. E-mail: david.peleg@weizmann.ac.il.

1 Introduction

1.1 Background and motivation

The term “connectivity” in networks has more than one meaning, but these meanings are equivalent in “traditional” networks. While the graph theoretic definition of connectivity refers to the ability to reach every node from every other node (a.k.a. *1-connectivity*), the term connectivity is often related also to the *survivability* of a network, namely, the ability to preserve 1-connectivity whenever some links fail.¹ In other words, a network G is said to be *k-connected* if it satisfies the following “connectivity property” (CP):

(CP1) G remains connected whenever up to $k - 1$ links are erased from it.

However, there are also other meanings to connectivity. A network G is also said to be *k-connected* if it satisfies the following “connectivity property”:

(CP2) There exist k pairwise *link-disjoint paths* from s to t for every two nodes s, t in G .

The equivalence of these two properties [12] enables numerous practical applications. For example, one of the applications of the existence of link-disjoint paths (Property (CP2)) is to ensure survivability (Property (CP1)). Often, a backup (link-disjoint) path is prepared in advance, and the traffic is diverted to the backup path whenever a link on the primary path fails. An example is the backup path protection mechanism in SONET networks (see, e.g., [17]).

A third property capturing connectivity is based on the amount of flow that can be shipped in the network between any source and any destination, defining the capacity of a single link to be 1. In other words, a network G is said to be *k-connected* if it satisfies the following “connectivity property”:

(CP3) It is possible to ship k units of flow from s to t for every two nodes s, t in G .

This property too is equivalent to the first two [7], and is also used together with them. For example, routing some flow of information around congestion (which may be possible only if the network satisfies property (CP3) and thus can support this additional flow) uses the second property, i.e., it relies on the existence of multiple link-disjoint paths.

Current networks, however, offer multi-layered structures which yield significant complications when dealing with the notion of connectivity. In particular, the overlay/underlying network dichotomy plays a major role in modeling communication networks, and overlay networks such as peer-to-peer (P2P) networks, MPLS networks, IP/WDM networks, and SDH/SONET-WDM networks, all share the same overall structure: an *overlay* network H is implemented on top of an

¹ The current paper does not deal with *node* connectivity.

underlying network G . This implementation can be abstracted as a *routing scheme* that maps overlay connections to underlying routes. We comment that there are sometimes differences between such a mapping and the common notion of a routing scheme. Still, since the routing scheme often defines the mapping, we shall term this mapping the *routing scheme*.

Often, the underlying network itself is implemented on top of yet another network, thus introducing a multi-layer hierarchy. Typically, the lower level underlying network is “closer” to the physical wires, whereas the higher level network is a traffic network in which edges capture various kinds of connections, depending on the context. For the sake of simplicity, we focus on a pair of consecutive layers G and H . This is sufficient to capture a large class of practical scenarios.

The current paper deals with what happens to the different connectivity properties once we turn to the context of overlay networks. As discussed later on, connectivity has been studied previously in the “overlay network” world under the “survivability” interpretation (CP1), and it has been observed that, in this context, the connectivity parameter changes, i.e., the connectivity of the overlay network may be different from that of the underlying network. Lee et al. [13] demonstrated the significance of this difference by showing that the survivability property is computationally hard and even hard to approximate in the multi-layer case. Since the three aforementioned connectivity parameters may differ in multi-layer networks (see Section 1.2), they also showed a similar result for the disjoint paths connectivity property.

Interestingly, the motivation of Lee et al. for addressing the disjoint paths connectivity property was the issue of flow. One of the contributions of the current paper is to directly address this issue, showing that in contrast to the previous two notions of connectivity, the maximum flow supported by an overlay network can actually be computed in polynomial time.

In the specific context of survivability, there has been other papers that have shown that the issue of connectivity in an overlay network is different from that of connectivity in underlying networking. Consider, for example, a situation where several overlay edges (representing connections) of H pass over the same physical link. Then all these overlay edges may be disconnected as a result of a single hardware fault in that link, possibly disconnecting the overlay network. The affected overlay links are said to *share* the *risk* of the failure of the underlying physical link, hence they are referred to in the literature as a *shared risk link group* (SRLG). An SRLG-based model for overlay networks was extensively studied in recent years;² see, e.g., [15] for a useful introduction to this notion and [14] for a discussion of this concept in the context of MPLS. The SRLG model hides the actual structure of the underlying network, in the sense that many different underlying networks can yield the same sets of SRLG. (For certain purposes, this is an advantage of the model.) An even more general notion is that of *Shared Risk Resource Group*. However, sometimes this model abstracts away too much information, making certain computational goals (such as, e.g., flow computations)

² Note that a common underlying link is not the only possible shared risk; overlay links sharing a node may form a shared risk link group too.

harder to achieve.

In contrast to the SRLG model, we present the alternative model of *deep connectivity*, which allows us to simultaneously consider all three components: the overlay network, the underlying network, and the mapping (the routing scheme). Note that all three should be considered: For example, if the underlying network is not connected, then neither can the overlay network be. The routing scheme also affects the connectivity properties as different routing schemes may yield significantly different overlay link dependencies. In some cases, routing is constrained to shortest paths, whereas in other cases it can be very different. In *policy based* routing schemes (see, e.g. [16]), for example, some underlying edges are not allowed to be used for routing from u to v , which may cause the underlying path implementing the overlay link (u, v) to be much longer than the shortest (u, v) -path.

1.2 The deep connectivity model

The underlying network is modeled by a (simple, undirected, connected) graph G whose vertex set $V(G)$ represents the network nodes, and whose edge set $E(G)$ represents the communication links between them. Some nodes of the underlying network are designated as *peers*; the set of peers is denoted by $\mathcal{P} \subseteq V(G)$. The overlay network, modeled by a graph H , spans the peers, i.e., $V(H) = \mathcal{P}$ and $E(H) \subseteq \mathcal{P} \times \mathcal{P}$; H typically represents a “virtual” network, constructed on top of the peers in the underlying communication network G .

An edge (u, v) in the overlay graph H may not directly correspond to an edge in the underlying graph G (that is, $E(H)$ is not necessarily a subset of $E(G)$). Therefore, communication over a (u, v) edge in H should often be routed along some multi-hop path connecting u and v in G . This is the role of a *routing scheme* $\rho : \mathcal{P} \times \mathcal{P} \rightarrow 2^{E(G)}$ that maps each pair (u, v) of peers to some simple path $\rho(u, v)$ connecting u and v in G . A message transmitted over the edge (u, v) in H is physically disseminated along the path $\rho(u, v)$ in G . We then say that (u, v) is *implemented* by $\rho(u, v)$. For the sake of simplicity, the routing scheme ρ is assumed to be symmetric, i.e., $\rho(u, v) = \rho(v, u)$. More involved cases do exist in reality: the routing scheme may be asymmetric, or may map some overlay edge into multiple routes; the simple model given here suffices to show interesting differences between the various connectivity measures.

When a message is routed in H from a peer $s \in \mathcal{P}$ to a non-neighboring peer $t \in \mathcal{P}$ along some multi-hop path $\pi = (x_0, x_1, \dots, x_k)$ with $x_0 = s$, $x_k = t$, and $(x_i, x_{i+1}) \in E(H)$, it is physically routed in G along the concatenated path $\rho(x_0, x_1)\rho(x_1, x_2) \cdots \rho(x_{k-1}, x_k)$. In some cases, when the overlay graph H is known, it will be convenient to define the routing scheme over the edges of H , rather than over all peer pairs.

The notion of *deep connectivity* grasps the connectivity in the overlay graph H , while taking into account its implementation by the underlying paths in G . Specifically, given two peers $s, t \in \mathcal{P}$, we

are interested in three different parameters, each capturing a specific type of connectivity. In order to define these parameters, we extend the domain of ρ from vertex pairs in $\mathcal{P} \times \mathcal{P}$ to collections of such pairs in the natural manner, defining $\rho(F) = \bigcup_{(u,v) \in F} \rho(u,v)$ for every $F \subseteq \mathcal{P} \times \mathcal{P}$. In particular, given an (s,t) -path π in H , $\rho(\pi) = \bigcup_{e \in \pi} \rho(e)$ is the set of underlying edges used in the implementation of the overlay edges along π .

- The *edge-removal deep connectivity* of s and t in H with respect to G and ρ , denoted by $\text{ERDC}_{G,\rho}(s,t,H)$, is defined as the size of the smallest subset $F \subseteq E(G)$ that intersects with $\rho(\pi)$ for every (s,t) -path π in H ; namely, the minimum number of underlying edges that should be removed in order to disconnect s from t in the overlay graph.
- The *path-disjoint deep connectivity* of s and t in H with respect to G and ρ , denoted by $\text{PDDC}_{G,\rho}(s,t,H)$, is defined as the size of the largest collection C of (s,t) -paths in H such that $\rho(\pi) \cap \rho(\pi') = \emptyset$ for every $\pi, \pi' \in C$ with $\pi \neq \pi'$, i.e., the maximum number of overlay paths connecting s to t whose underlying implementations are totally independent of each other.
- The *flow deep connectivity* of s and t in H with respect to G and ρ , denoted by $\text{FDC}_{G,\rho}(s,t,H)$, is defined as the maximum amount of flow³ that can be pushed from s to t in G restricted to the images under ρ of the (s,t) -paths in H , assuming that each edge in $E(G)$ has a unit capacity. Intuitively, if s and t are well connected, then it should be possible to push a large amount of flow between them.

Example: To illustrate the various definitions, consider the underlying network G depicted in Fig. 1(a), and the overlay network H depicted in Fig. 1(b). The routing scheme ρ assigns each of the 6 overlay edges adjacent to the extreme S and T a simple route consisting of the edge itself. For the remaining three overlay edges, the assigned routes are as follows:

$$\begin{aligned} \rho(U_1, U_4) &= (U_1, M_1, M_2, U_2, U_3, U_4) \\ \rho(M_1, M_4) &= (M_1, M_2, D_2, D_3, M_3, M_4) \\ \rho(D_1, D_4) &= (D_1, D_2, D_3, U_3, U_4, D_4) . \end{aligned}$$

The route $\rho(M_1, M_4)$ is illustrated by the dashed line in Fig. 1(a). Note that in the original (underlying) network G , the connectivity of the extreme nodes S and T is 3 under all three definitions. In contrast, the values of the three connectivity parameters for the extreme nodes S and T in the overlay network H under the routing scheme ρ are as follows:

³ In the setting of undirected graphs, flow may be interpreted in two different ways depending on whether two flows along the same edge in opposite directions cancel each other or add up. Here, we assume the latter interpretation which seems to be more natural in the context of overlay networks.

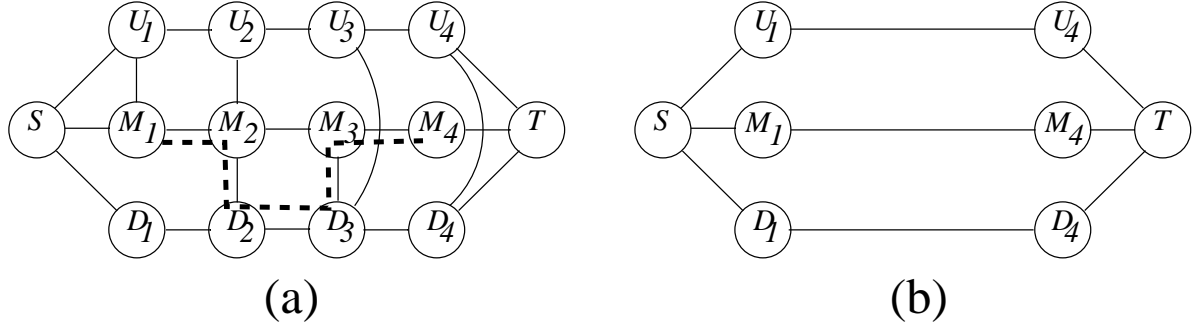


Figure 1: (a) The underlying graph G . (b) The overlay network H .

- The edge-removal deep connectivity of s and t in H w.r.t. G and ρ is $\text{ERDC}_{G,\rho}(s, t, H) = 2$. Indeed, disconnecting the underlying edge (D_2, D_3) plus any edge of the upper underlying route will disconnect S from T .
- The path-disjoint deep connectivity of s and t in H w.r.t. G and ρ is $\text{PDDC}_{G,\rho}(s, t, H) = 1$. Indeed, any two of the three overlay routes connecting S and T share an underlying edge.
- The *flow deep connectivity* of s and t in H w.r.t. G and ρ is $\text{FDC}_{G,\rho}(s, t, H) = 3/2$. This is obtained by pushing $1/2$ flow unit through each of the three overlay routes.

For each deep connectivity (s, t) -parameter $X_{G,\rho}(s, t, H)$, we define the corresponding *all-pairs* variant $X_{G,\rho}(H) = \min_{s,t \in \mathcal{P}} X_{G,\rho}(s, t, H)$. When G and ρ are clear from the context, we may remove them from the corresponding subscripts.

1.3 Our contributions

Our model for overlay networks makes it convenient to explore the discrepancies between the different deep connectivity notions. Classical results from graph theory, e.g., the fundamental min-cut max-flow theorem [7, 6] and Menger’s theorem [12] state that the three connectivity parameters mentioned above are equivalent when a single layer network is considered. Polynomial time algorithms that compute these parameters (in a single layer network) were discovered early [7, 4, 5] and have since become a staple of algorithms textbooks [2, 11]. As mentioned above, previous results [13, 3], when translated to our model, have shown that in multi-layer networks, two of the three connectivity parameters are computationally hard and even hard to approximate. Our first technical contribution is to expand on these negative results by showing that the path-disjoint deep connectivity property cannot be approximated to any finite ratio when attention is restricted to simple paths in the underlying graph.

On the positive side, we show that the flow deep connectivity parameters can be computed in polynomial time (Section 3), thus addressing the motivation of [13] for studying the disjoint paths property in overlay networks. Then, we address the issue of constructing a “good” overlay graph

for a given underlying graph and routing scheme. As opposed to the difficulty of approximating the value of the parameters, we show that the related construction problem can sometimes be well approximated. Specifically, in Section 4, we investigate the problem of constructing 2-edge removal deeply connected overlay graphs with as few as possible (overlay) edges. This problem is shown to be NP-hard, but we show that a logarithmic-approximation for it can be obtained in polynomial-time. We also devise a 2-approximation algorithm for particular, yet natural, instances of the problem.

2 Hardness of approximation

As mentioned earlier, Lee et al. [13] established hardness of approximation results for the problems of computing the parameters $\text{ERDC}_{G,\rho}(s, t, H)$, $\text{PDDC}_{G,\rho}(s, t, H)$, and $\text{ERDC}_{G,\rho}(H)$. For completeness, we observe that the all-pairs variant $\text{PDDC}_{G,\rho}(H)$ of the path-disjoint deep connectivity parameter is also hard to approximate, establishing the following theorem, which is essentially a corollary of Theorem 4 in [13] combined with a result of Håstad [9].

Theorem 2.1. *Unless $NP = ZPP$, the problem of computing the parameter $\text{PDDC}_{G,\rho}(H)$ cannot be approximated to within a ratio of $|E(H)|^{1/2-\epsilon}$ for any $\epsilon > 0$.*

We now turn to show that the following natural variants of the PDDC parameters cannot be approximated to within any finite ratio. Let $\text{SPDDC}_{G,\rho}(s, t, H)$ denote the size of the largest collection C of (s, t) -paths in H such that all paths $\pi \in C$ are implemented by simple paths $\rho(\pi)$ in G and $\rho(\pi) \cap \rho(\pi') = \emptyset$ for every $\pi, \pi' \in C$ with $\pi \neq \pi'$; let $\text{SPDDC}_{G,\rho}(H) = \min_{s,t \in P} \text{SPDDC}_{G,\rho}(s, t, H)$. Note that these parameters are merely a restriction of the PDDC parameters to simple paths (hence the name, which stands for *simple* path-disjoint deep connectivity).

The inapproximability of the $\text{SPDDC}_{G,\rho}(s, t, H)$ parameter is proved by reducing the set packing problem to the problem of distinguishing between the case $\text{SPDDC}_{G,\rho}(s, t, H) = 0$ and the case $\text{SPDDC}_{G,\rho}(s, t, H) \geq 1$. In fact, the vertices $s, t \in V(H)$ that minimize $\text{SPDDC}_{G,\rho}(s, t, H)$ in this reduction are known in advance, thus establishing the impossibility of approximating the all-pairs parameter $\text{SPDDC}_{G,\rho}(H)$ as well. The proofs of the following two theorems are deferred to Appendix A.

Theorem 2.2. *Unless $P = NP$, the problem of computing the parameter $\text{SPDDC}_{G,\rho}(s, t, H)$ cannot be approximated to within any finite ratio.*

Theorem 2.3. *Unless $P = NP$, the problem of computing the parameter $\text{SPDDC}_{G,\rho}(H)$ cannot be approximated to within any finite ratio.*

3 Efficient algorithm for FDC

In this section we develop a polynomial time algorithm that computes the flow deep connectivity parameter $\text{FDC}(s, t, H)$ (which clearly provides an efficient computation of the parameter $\text{FDC}(H)$ as well). Consider some underlying graph G , routing scheme ρ , overlay graph H , and two vertices $s, t \in V(H)$. Let \mathcal{P} denote the collection of all simple (s, t) -paths in H . For each path $p \in \mathcal{P}$ and for each edge $e \in E(G)$, let $\psi(p, e)$ be the number of appearances of the edge e along the image of p under ρ . We begin by representing the parameter $\text{FDC}(s, t, H)$ as the outcome of the following linear program:

$$\begin{aligned} \max \quad & \sum_{p \in \mathcal{P}} x_p \quad \text{s.t.} \\ & \sum_{p \in \mathcal{P}} \psi(p, e) \cdot x_p \leq 1 \quad \forall e \in E(G) \\ & x_p \geq 0 \quad \forall p \in \mathcal{P} \end{aligned}$$

The variable x_p represents the amount of flow pushed along the image under ρ of the path p for every $p \in \mathcal{P}$. The goal is to maximize the total flow pushed along the images of all paths in \mathcal{P} subject to the constraints specifying that the sum of flows pushed through any edge is at most 1. This linear program may exhibit an exponential number of variables, so let us consider its dual program instead:

$$\begin{aligned} \min \quad & \sum_{e \in E(G)} y_e \quad \text{s.t.} \\ & \sum_{e \in E(G)} \psi(p, e) \cdot y_e \geq 1 \quad \forall p \in \mathcal{P} \\ & y_e \geq 0 \quad \forall e \in E(G) \end{aligned}$$

The dual program can be interpreted as fractionally choosing as few as possible edges of G so that the image under ρ of every path p in \mathcal{P} traverses in total at least one edge. We cannot solve the dual program directly as it may have an exponential number of constraints. Fortunately, it admits an efficient separation oracle, hence it can be solved in polynomial time (see, e.g., [8]).

Given some real vector \vec{y} indexed by the edges in $E(G)$, our separation oracle either returns a constraint which is violated by \vec{y} or reports that all the constraints are satisfied and \vec{y} is a feasible solution. Recall that a violated constraint corresponds to some path $p \in \mathcal{P}$ such that $\sum_{e \in E(G)} \psi(p, e) \cdot y_e < 1$. Therefore our goal is to design an efficient algorithm that finds such a path $p \in \mathcal{P}$ if such a path exists.

Let $w(e) = \sum_{e' \in \rho(e)} y_{e'}$ for every edge $e \in E(H)$ and let H' be the weighted graph obtained by assigning weight $w(e)$ to each edge $e \in E(H)$. The key observation in this context is that the (weighted) length of an (s, t) -path p' in H' equals exactly to $\sum_{e \in E(G)} \psi(p, e) \cdot y_e$, where p is the path in H that corresponds to p' in H' . Therefore, our separation oracle is implemented simply by finding a shortest (s, t) -path p^* in H' : if the length of p^* is smaller than 1, then p^* corresponds to a violated constraint; otherwise, the length of all (s, t) -paths in H' is at least 1, hence \vec{y} is a feasible solution. This establishes the following theorem.

Theorem 3.1. *The parameters $\text{FDC}_{G,\rho}(s, t, H)$ and $\text{FDC}_{G,\rho}(H)$ can be computed in polynomial*

time.

4 Sparsest 2-ERDC overlay graphs

In this section we are interested in the following problem, referred to as the *sparsest 2-ERDC overlay graph* problem: given an underlying graph G , a peer set $\mathcal{P} \subseteq V(G)$, and a routing scheme $\rho : \mathcal{P} \times \mathcal{P} \rightarrow 2^{E(G)}$, construct the sparsest overlay graph H for \mathcal{P} (in terms of number of overlay edges) satisfying $\text{ERDC}_{G,\rho}(H) \geq 2$. Of course, one has to make sure that such an overlay graph H exists, so in the context of the sparsest 2-ERDC overlay graph problem we always assume that $\text{ERDC}_{G,\rho}(K_{\mathcal{P}}) \geq 2$, where $K_{\mathcal{P}}$ is the complete graph on \mathcal{P} . This means that a trivial solution with $\binom{n}{2}$ edges, where $n = |\mathcal{P}|$, always exists and the challenge is to construct a sparser one.

4.1 Hardness

We begin our treatment of this problem with a hardness result.

Theorem 4.1. *The sparsest 2-ERDC overlay graph problem is NP-hard.*

Proof. The assertion is proved by a reduction from the *Hamiltonian path* problem. Consider an n -vertex graph G_0 input to the Hamiltonian path problem. Transform it into an instance of the sparsest 2-ERDC overlay graph problem as follows: Construct the underlying graph G by setting $V(G) = V(G_0) \cup \{x, y\}$ and $E(G) = E(G_0) \cup \{(x, y)\} \cup \{(v, x), (v, y) \mid v \in V(G_0)\}$ and let $\mathcal{P} = V(G_0)$. Define the routing scheme ρ by setting

$$\rho(u, v) = \begin{cases} (u, v) & \text{if } (u, v) \in E(G_0) \\ (u, x, y, v) & \text{otherwise.} \end{cases}$$

This transformation is clearly polynomial in n .

We argue that G_0 admits a Hamiltonian path if and only if there exists an overlay graph H for \mathcal{P} so that $|E(H)| = n$ and $\text{ERDC}_{G,\rho}(H) \geq 2$. To that end, suppose that G_0 admits a Hamiltonian path π . If π can be closed to a Hamiltonian cycle (in G_0), then take H to be this Hamiltonian cycle. Otherwise, take H to be the cycle consisting of π and a virtual edge connecting between π 's endpoints. In either case, H clearly has n edges and by the design of ρ , H satisfies $\text{ERDC}_{G,\rho}(H) = 2$.

Conversely, if there exists an overlay graph H for \mathcal{P} so that $|E(H)| = n$ and $\text{ERDC}_{G,\rho}(H) \geq 2$, then H must form a Hamiltonian cycle C in $\mathcal{P} \times \mathcal{P}$. This cycle can contain at most one virtual edge as otherwise, the removal of (x, y) breaks two edges of C which means that $\text{ERDC}_{G,\rho}(H) < 2$. By removing this virtual edge, we are left with a Hamiltonian path in G_0 . \square

4.2 Constructing sparse 2-ERDC overlay graphs

On the positive side, we develop a polynomial time logarithmic approximation algorithm for the sparsest 2-ERDC overlay graph problem. Our algorithm proceeds in two stages: First, we take T to be an arbitrary spanning tree of $\mathcal{P} \times \mathcal{P}$. Subsequently, we aim towards (approximately) solving the following optimization problem, subsequently referred to as the *overlay augmentation* problem: augment T with a minimum number of $\mathcal{P} \times \mathcal{P}$ edges so that the resulting overlay graph H satisfies $\text{ERDC}_{G,\rho}(H) \geq 2$.

We will soon explain how we cope with this optimization problem, but first let us make the following observation. Denote the edges in $\rho(T)$ by $\rho(T) = \{e_1, \dots, e_\ell\}$ and consider some overlay graph H such that $E(H) \supseteq T$ and some $1 \leq i \leq \ell$. Let $F_i(H)$ be the collection of connected components of the graph obtained from H by removing all edges $e \in E(H)$ such that $e_i \in \rho(e)$. Fix

$$\kappa_i(H) = |F_i(H)| - 1 \quad \text{and} \quad \kappa(H) = \sum_{i=1}^{\ell} \kappa_i(H).$$

We think of $\kappa(H)$ as a measure of the distance of the overlay graph H from being a feasible solution to the overlay augmentation problem (i.e., $\text{ERDC}(H) \geq 2$).

Observation. *An overlay graph $H \supseteq T$ satisfies $\text{ERDC}(H) \geq 2$ if and only if $\kappa(H) = 0$.*

Proof. If $\text{ERDC}(H) \geq 2$, then $F_i(H)$ must consist of a single connected component for every $1 \leq i \leq \ell$, thus $\kappa(H) = 0$. Conversely, if $\kappa(H) = 0$, then necessarily $|F_i(H)| = 1$ for every $1 \leq i \leq \ell$, which means that H does not disconnect by the removal of any edge $e_i \in \rho(T)$. It is also clear that the removal of any edge in $E(G) - \rho(T)$ does not disconnect H as the tree T remains intact. Therefore, $\text{ERDC}(H) \geq 2$. \square

Consider some edge $e \in (\mathcal{P} \times \mathcal{P}) - E(H)$ and let $H \cup \{e\}$ denote the overlay graph obtained from H by adding the edge e . By the definition of the parameter κ , we know that $\Delta_i(e, H) = \kappa_i(H) - \kappa_i(H \cup \{e\})$ is either 0 or 1 for any $1 \leq i \leq \ell$. Fixing $\Delta(e, H) = \kappa(H) - \kappa(H \cup \{e\})$, we observe that: $\Delta(e, H) = \sum_{i=1}^{\ell} \Delta_i(e, H)$ referred to as Property (\star) .

We are now ready to complete the description of our approximation algorithm. Starting from $H = T$, the algorithm gradually adds edges to H as long as $\kappa(H) > 0$ according to the following greedy rule. At any intermediate step, we add to H an edge $e \in (\mathcal{P} \times \mathcal{P}) - E(H)$ that yields the maximum $\Delta(e, H)$. When $\kappa(H)$ decreases to zero, the algorithm terminates (recall that this means that $\text{ERDC}(H) \geq 2$).

The analysis of our approximation algorithm relies on the following proposition.

Proposition 4.2. *The parameter κ can be computed in polynomial time. Moreover, for every two overlay graphs H_1, H_2 such that $E(H_1) \subseteq E(H_2)$, we have*

- (1) $\kappa(H_1) \geq \kappa(H_2)$; and

(2) $\Delta(e, H_1) \geq \Delta(e, H_2)$ for every edge $e \in \mathcal{P} \times \mathcal{P}$.

Proof. The fact that κ can be computed efficiently and the fact that $\kappa(H_1) \geq \kappa(H_2)$ are clear from the definition of κ , so our goal is to prove that $\Delta(e, H_1) \geq \Delta(e, H_2)$ for every $e \in \mathcal{P} \times \mathcal{P}$. By Property (\star) , it suffices to show that $\Delta_i(e, H_1) \geq \Delta_i(e, H_2)$ for every $1 \leq i \leq \ell$. If $\Delta_i(e, H_2) = 0$, then this holds vacuously, so suppose that $\Delta_i(e, H_2) = 1$. This means that $e_i \notin \rho(e)$ and the endpoints of e belong to different connected components in $F_i(H_2)$. But since $E(H_1) \subseteq E(H_2)$, it follows that the endpoints of e must also belong to different connected components in $F_i(H_1)$, hence $\Delta_i(e, H_1) = 1$ as well. \square

Proposition 4.2 implies that the overlay augmentation problem falls into the class of *submodular cover* problems (cf. [18, 1]) and our greedy approach is guaranteed to have an approximation ratio of at most $\ln(\kappa(T)) + 1 = O(\log N)$, where $N = |V(G)|$. More formally, letting \hat{H} be a sparsest overlay graph such that $\hat{H} \supseteq T$ and $\text{ERDC}(\hat{H}) \geq 2$, it is guaranteed that the overlay graph H generated by our greedy approach satisfies $|E(H) - T| \leq O(\log N) \cdot |E(\hat{H}) - T|$.

To conclude the analysis, let H^* be an optimal solution to the sparsest 2-ERDC overlay graph problem. Clearly, $|E(H^*)| > n - 1 = |T|$. Moreover, since $E(H^*) \cup T$ is a candidate for \hat{H} , it follows that $|E(H^*) \cup T| \geq |\hat{H}|$, thus $|E(H^*)| \geq |E(\hat{H}) - T|$. Therefore,

$$\begin{aligned} |E(H)| &\leq O(\log N) \cdot |E(\hat{H}) - T| + |T| \\ &\leq O(\log N) \cdot |E(H^*)| + |E(H^*)| \\ &= O(\log(N) \cdot |E(H^*)|) \end{aligned}$$

which establishes the following theorem.

Theorem 4.3. *The sparsest 2-ERDC overlay graph problem admits a polynomial-time logarithmic-approximation.*

4.3 A special case

We now turn to consider instances of the sparsest 2-ERDC overlay graph problem satisfying the following two simplifying assumptions:

- (1) all vertices are peers, namely, $\mathcal{P} = V(G)$; and
- (2) ρ maps every edge in $E(G)$ to itself.

We show that under these assumptions, one can always find a feasible solution with at most $2n - 2$ edges. Since any overlay graph H satisfying $\text{ERDC}(H) \geq 2$ must have at least n edges, this immediately implies a 2-approximation for our problem.

The construction proceeds, once again, in two stages: First, we take H to be an arbitrary spanning tree T of G . Then, for each edge $e \in T$, we add to H some edge $e' \in E(G)$ such that e

belongs to the cycle closed by appending e' to T . Note that such an edge e' exists since G is 2-edge connected (as otherwise, $\text{ERDC}_{G,\rho}(K_{\mathcal{P}}) < 2$). The overlay graph H clearly contains at most $2n - n$ edges. Moreover, assumption (A2) implies that ρ maps every edge in $E(H)$ to itself. Therefore, our construction guarantees that $\text{ERDC}(H) \geq 2$: on one hand, the spanning tree T ensures that the graph obtained from H by removing the edge e is connected for every edge $e \notin T$; on the other hand, H does not disconnect by the removal of any single edge $e \in T$ due to the second stage of the construction.

Theorem 4.4. *Under assumptions (1) and (2), an overlay graph H with $2n - 2$ edges satisfying $\text{ERDC}(H) \geq 2$ can be constructed in polynomial time.*

References

- [1] J. Bar-Ilan, G. Kortsarz, and D. Peleg. Generalized submodular cover problems and applications. *Theoretical Computer Science* 250:179–200, 2001.
- [2] T.H. Cormen, C.E. Leiserson, R.L. Rivest, and C. Stein. *Introduction to Algorithms*, MIT Press, 2009.
- [3] D. Coudert, P. Datta, S. Perennes, H. Rivano, M.-E. Voge. Shared Risk Resource Group Complexity and Approximability Issues. *Parallel Processing Letters* 17(2): 169-184 (2007)
- [4] E.A. Dinic. Algorithm for solution of a problem of maximum flow in a network with power estimation. *Soviet Math. Doklady (Doklady)*, 11:1277–1280, 1970.
- [5] J. Edmonds and R.M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *J. ACM*, 19(2): 248–264, 1972.
- [6] P. Elias, A. Feinstein, and C.E. Shannon. A note on the maximum flow through a network. *IRE Trans. Inf. Theory*, IT-2(4):117–119, 1956.
- [7] L.R. Ford and D.R. Fulkerson. Maximal flow through a network. *Canadian J. Mathematics*, 8:399–404, 1956.
- [8] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer-Verlag, Berlin, 1993.
- [9] J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182:105–142, 1999.
- [10] R.M. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, R.E. Miller and J.W. Thatcher (Eds), pp. 85–103, 1972.
- [11] J. Kleinberg and E. Tardos. *Algorithm Design*. Addison Wesley, 2005.
- [12] K. Menger. Zur allgemeinen Kurventheorie. *Fund. Math.*, 10:96–115, 1927.
- [13] Kayi Lee, Eytan Modiano, and Hyang Won Lee. Cross-layer Survivability in WDM-Based Networks.” *ACM/IEEE Transactions on Networking*, Vol. 19, No. 4, August 2011.
- [14] MPLS Traffic Engineering: Shared Risk Link Groups (SRLG), Cisco IOS Software Releases 12.0 S, www.cisco.com/en/US/docs/ios/12.0s/feature/guide/fs29srlg.html
- [15] Network protection techniques, network failure recovery, network failure events. Network Protection Website, www.network-protection.net/shared-risk-link-group-srlg/.
- [16] Policy-Based Routing. White Paper, Cisco Systems Inc., 1996. www.cisco.com/warp/public/cc/pd/iosw/tech/policy_wp.pdf.

- [17] S. Ramamurthy and B. Mukherjee. Survivable WDM mesh networks. Part I-Protection. In Proc. INFOCOM, 1999, pp. 744-751, New York.
- [18] L.A. Wolsey. An analysis of the greedy algorithm for the submodular set covering problem. *Combinatorica* 2:385–393, 1982.

APPENDIX

A Hardness of approximation for the simple path variants

Similarly to [13], we base our hardness results on the relationship between path-disjoint deep connectivity and encoding *set systems*. An (m, n) -set system \mathcal{S} is a pair (D, X) , where:

- D is a domain of elements, $|D| = m$; and
- $X \subseteq 2^D$ is a collection of sets in the domain D , $|X| = n$.

It is convenient to represent \mathcal{S} as a Boolean characteristic function, $\mathcal{S} : [m] \times [n] \rightarrow \{0, 1\}$, so that for every $i \in [m]$ and $j \in [n]$, $\mathcal{S}(i, j) = 1$ if and only if the i^{th} element of D is included in the j^{th} set of X .

Lemma A.1. *For every overlay graph H , edge subset $F \subseteq E(H)$ of cardinality $|F| = n$, and (m, n) -set system $\mathcal{S} = (D, X)$, there exist an underlying graph G and a routing scheme $\rho : E(H) \rightarrow 2^{E(G)}$ such that:*

- (1) $|V(G)| = |V(H)| + O(m \cdot n)$;
- (2) $E(G) = (E(H) - F) \cup E_D \cup E_\rho$, where E_D and E_ρ are sets of new edges on $V(G)$, $|E_D| = m$, and $|E_\rho| = O(m \cdot n)$;
- (3) for every $e_j \in F$, $j \in [n]$, and $e_i \in E_D$, $i \in [m]$, it holds that $e_i \in \rho(e_j) \iff \mathcal{S}(i, j) = 1$; and
- (4) for every $e \in E_\rho$, there exists a unique $e_j \in F$, $j \in [n]$, such that $e \in \rho(e_j)$.

Proof. Informally, G and ρ are designed so that the set system $\mathcal{S} = (D, X)$ is encoded in such a way that E_D corresponds to the element domain D and F corresponds to the set collection X . For a more precise description of the design of the underlying graph G and routing scheme ρ , let us first present a simpler construction where G is allowed to be a multigraph (with parallel edges).

The vertex set of G consists of the vertices of H and two additional new vertices v_i^a, v_i^b for every $i \in [m]$. The edge set of G consists of three disjoint subsets. The first subset is just the edges of $E(H) - F$; ρ maps every such H -edge to itself, namely, the edge $(x, y) \in E(H) - F$ is mapped to the path in G that consists of the single edge (x, y) . The second subset is $E_D = \{e_i = (v_i^a, v_i^b) \mid i \in [m]\}$. The routing scheme ρ is designed to guarantee that an edge $e_i \in E_D$ appears in the path $\rho(e_j)$, $e_j \in F$, if and only if $\mathcal{S}(i, j) = 1$. For that purpose, we introduce the third subset E_ρ : for each $j \in [n]$, assuming that the edge $e_j \in F$ connects the vertices x and y in H and that the j^{th} set in \mathcal{S} is $\{i_1, \dots, i_k\} \subseteq [m]$, we add to E_ρ the edge $(x, v_{i_1}^a)$, the edges $(v_{i_\ell}^b, v_{i_{\ell+1}}^a)$ for $\ell = 1, \dots, k-1$, and the edge $(v_{i_k}^b, y)$. It is important to point out that a new copy of those edges are added to E_ρ for each $j \in [n]$, which may create edge multiplicities. Finally, ρ maps the edge e_j to the path $p_j = \langle x, v_{i_1}^a, v_{i_1}^b, \dots, v_{i_k}^a, v_{i_k}^b, y \rangle$. Refer to Figure 2 for an illustration.

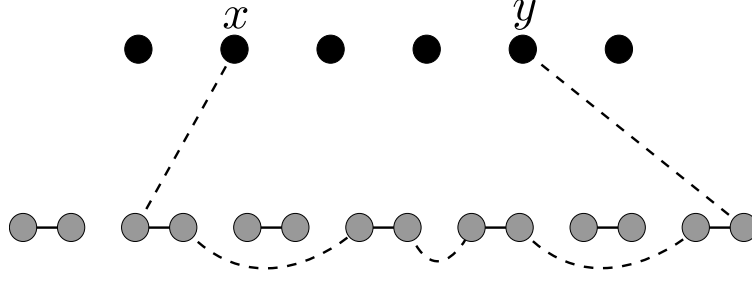


Figure 2: The underlying (multi)graph G . The vertices of H are depicted by the black circles; the vertices in $\{v_i^a, v_i^b \mid i \in [m]\}$ are depicted by the gray circles. The E_D -edges are depicted by the solid segments; the E_ρ -edges along the path p_j to which ρ maps the edge $e_j = (x, y) \in E(H)$ are depicted by the dashed segments.

The construction of G immediately implies that $|V(G)| = |V(H)| + 2m$ and that $|E_D| = m$. To see that $|E_\rho| = O(n \cdot m)$, observe that each E_ρ -edge belongs to some path p_j , $j \in [n]$, and that each such path admits at most two hops for every $i \in [m]$. It remains to show that for every $j \in [n]$ and $i \in [m]$, $e_i \in \rho(e_j)$ if and only if $\mathcal{S}(i, j) = 1$, which follows directly from the design of ρ as the path p_j goes through an edge $e_i \in E_D$ if and only if $\mathcal{S}(i, j) = 1$.

Finally, note that edge multiplicities in the constructed multigraph G may occur only among the edges in E_ρ . Hence, by subdividing each edge $e = (u, v) \in E_\rho$, replacing it with the edges (u, z_e) and (v, z_e) , where z_e is a new vertex, we can turn the multigraph G into a simple graph at the cost of adding $O(m \cdot n)$ extra vertices. The lemma follows. \square

The hardness results established in this section are based on Lemma A.1 by a reduction from the *set packing* problem: Given an (m, n) -set system $\mathcal{S} = (D, X)$ and a positive integer k , the set packing problem asks whether there exist k pairwise disjoint sets in X . The problem is known to be NP-complete; in fact, it is among the original 21 NP-complete problems listed by Karp [10].

The inapproximability of the $\text{SPDDC}_{G,\rho}(s, t, H)$ parameter is proved by reducing the set packing problem to the problem of distinguishing between the case $\text{SPDDC}_{G,\rho}(s, t, H) = 0$ and the case $\text{SPDDC}_{G,\rho}(s, t, H) \geq 1$. Given as input to the set packing problem some (m, n) -set system \mathcal{S} with set collection $X = \{S^1, \dots, S^n\}$ and positive integer k , we first construct the $(m, k \cdot n)$ -set system \mathcal{S}' obtained from \mathcal{S} by creating k identical copies S_1^j, \dots, S_k^j of each set S^j in \mathcal{S} . Clearly, \mathcal{S}' admits a set packing of size k if and only if \mathcal{S} admits a set packing of size k . Then, we construct the overlay graph H as illustrated in Figure 3. Let $\{F = (u_{\ell-1}, v_\ell^j) \mid 1 \leq \ell \leq k, j \in [n]\}$ and take G and ρ to be the underlying graph and routing scheme promised by Lemma A.1 when applied to H , F , and \mathcal{S}' , where ρ is organized so that edge $(u_{\ell-1}, v_\ell^j)$ in F corresponds to set S_ℓ^j in \mathcal{S}' .

Assume first that $\text{SPDDC}_{G,\rho}(s, t, H) \geq 1$ and let π be an (s, t) -path in H such that its image under ρ is a simple path in G . Let $(u_0, v_1^{j_1}), \dots, (u_{k-1}, v_k^{j_k})$ be the edges traversed by π on odd hops. By the definition of \mathcal{S}' and the design of ρ , it must be that $j_\ell \neq j_{\ell'}$ for every $1 \leq \ell < \ell' \leq k$.

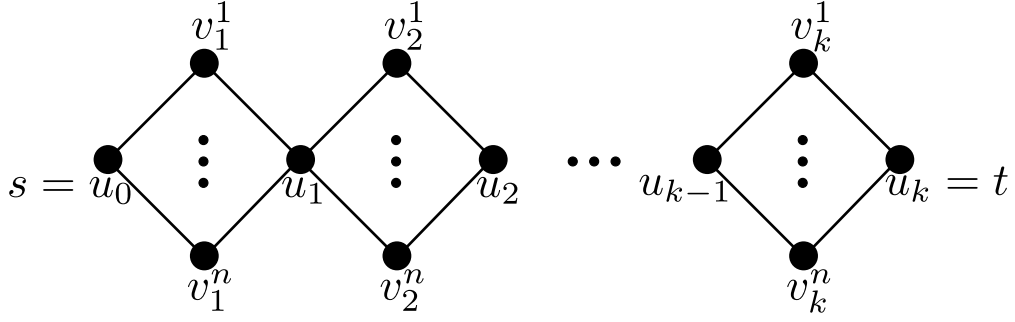


Figure 3: The overlay graph H used in the hardness proof of $\text{SPDDC}_{G,\rho}(s, t, H)$.

as otherwise, $\rho(\pi)$ does not form a simple path in G . Lemma A.1 then guarantees that the sets S^{j_1}, \dots, S^{j_k} in \mathcal{S} are pairwise disjoint. In the converse direction, assume that the sets S^{j_1}, \dots, S^{j_k} in \mathcal{S} are pairwise disjoint. By the definition of \mathcal{S}' , the sets $S_1^{j_1}, \dots, S_k^{j_k}$ in \mathcal{S}' are also pairwise disjoint. Lemma A.1 then guarantees that the image under ρ of the path $\langle u_0, v_1^{j_1}, u_1, v_2^{j_2}, \dots, u_{k-1}, v_k^{j_k}, u_k \rangle$ in H is simple, hence $\text{SPDDC}_{G,\rho}(s, t, H) \geq 1$, which establishes Theorem 2.2. Theorem 2.3 follows by observing that in the aforementioned construction, $\text{SPDDC}_{G,\rho}(x, y, H)$ is minimized by taking $x = s$ and $y = t$.