

Shared Risk Resource Group : Complexity and Approximability Issues*

D.Coudert¹ and P. Datta² and S. Perennes¹ and H. Rivano¹ and M.-E. Voge¹

1. MASCOTTE Project, I3S(CNRS/UNSA)/INRIA
2004 Route Des Lucioles, B.P. 93
06902 Sophia-Antipolis Cedex – France
{coudert,perennes,rivano,voge}@sophia.inria.fr

2. Computer and Computational Sciences Division
Los Alamos National Laboratory
Los Alamos, NM 87545 – USA
pallab@lanl.gov

Abstract

This article investigates complexity and approximability properties of combinatorial optimization problems yielded by the notion of Shared Risk Resource Group (SRRG). SRRG has been introduced in order to capture network survivability issues where a failure may break a whole set of resources, and has been formalized as *colored graphs*, where a set of resources is represented by a set of edges with same color. We consider here the analogous of classical problems such as determining paths or cuts with the minimum numbers of colors or color disjoint paths. These optimization problems are much more difficult than their counterparts in classical graph theory. In particular standard relationship such as the Max Flow - Min Cut equality do not hold any longer. In this article we identify cases where these problems are polynomial, for example when the edges of a given color form a connected subgraph, and otherwise give hardness and non approximability results for these problems.

Keywords: Survivability, fault tolerance, reliability, Shared Risk Resource Group, colored graphs, complexity, approximability, shortest paths, cuts.

1 Introduction

Network survivability, as an application domain, has fostered a large scope of theoretical and algorithmic developments in the recent years. Although neglected for a long time by survivable network design studies, multiple failures scenarios are now unavoidable with the spread of multilayer networks such as IP/WDM, MPLS networks, P2P or GRID computing overlay structures. For example, in such a multilayer network, multiple and apparently independent virtual links may be paths sharing a link on the underlying network, as illustrated in Figure 1(a) where a single failure

*The research reported in this paper is funded in part by NSF grant ANI-9973102, DARPA/NSA grant N66001-00-1-8949, ACI-SI PRESTO, Région PACA, France Telecom R&D, European projects IST FET AEOLUS and COST 293 Graal.




Figure 1: A multilayer network and its transformation into a colored graph.

of the link (FG) would induce multiple failures on the set of links of the virtual topology. Therefore, the need to take into account correlations between edges has appeared, and the concept of *Shared Risk Link Group*, later generalized to *Shared Risk Resource Group* (SRRG) [4, 17, 19], has been introduced. A SRRG is a set of resources breaking down simultaneously when a given failure occurs.

In addition to their application in telecommunication networks survivability issues for which they have been introduced, Shared Risk Resource Groups are relevant to numerous practical settings. Indeed, they can model correlated traffic jams of road networks, as well as cascading failure scenarios in electrical networks. Related concepts are also relevant in the settings of cryptographic PGP (Pretty Good Privacy) keyrings [13].

Many classical network survivability issues are modeled as graph covering problems in which a minimum cost subset of edges satisfying some given property is to be found. The cost of such solutions is generally the size of the set, or a linear function such as the sum of the cost of the edges. Examples of such problems are minimum cost spanning tree, minimum cut and shortest paths problems. These problems model applications in which edges fail *independently*.

This article focuses on the counterpart of such classical problems in the settings of Shared Risk Resource Group. In the general case, a Shared Risk Resource Group may include links (edges) and nodes (vertices) of the network. However, using straightforward graph transformations, without loss of generality one can consider only sets of edges.

Shared Risk Resource Groups are naturally modeled by associating to each group (or risk) a *color*, and to each edge (or resource) the colors representing the risks affecting it [7, 23]. According to the general case of Shared Risk Resource Groups, an edge may belong to several colors, modeling the fact that a resource may face different and independent risks. In [7], a *colored graph* is thus defined as an undirected graph associated to a set of colors. Each color is a subset of edges and the colors cover the edge set.

However, depending on the context, this property may be interpreted in several ways whether an edge fails when a single of its colors fails or when all its associated colors fail simultaneously.

That is why another definition for colored graphs is presented in [23]. In that case to each edge is associated a single color (the colors then partition the edge set). According to the context, a graph transformation can be applied to a colored graph of the first model to obtain a colored graph of the second model. If in the first model an edge e fails when at least one of its λ colors fails, we replace edge e by a path with λ monochromatic edges. Thus, in the second model, any path containing edge e will fail for any failure of one of its λ colors. Figure 1(b) illustrates this




Figure 2: Examples

transformation. The virtual connection (AH) of Figure 1(a) is routed through the physical links (AF), (FG) and (GH), it is therefore replaced with a path of three monochromatic edges. The colors of these edges are the colors associated to the physical links (AF), (FG) and (GH). The creation of this path requires the creation of two vertices x and y . If in the first model an edge e fails only when all its colors fail, this edge is replaced with λ monochromatic parallel edges.

Therefore, in the following we assume that edges are monochromatic, in other words colors not only cover the edge set but also partition it. This definition is actually encompassed in the definition of [7], consequently all the NP-hardness and inapproximability results proved under the monochromatic assumption extend to the more general settings of [7].

Definition 1.1 (Colored graph) *A colored graph is a triple $G = (V, E, \mathcal{C})$ where (V, E) is an undirected graph and \mathcal{C} is a partition of E .*

In a colored graph, a color containing an edge incident to a vertex is said to be *incident* to this vertex. Two colors are *adjacent* if they are incident to a same vertex. The colored degree of a given vertex is the number of colors which are incident to it.

The span of a color is an important parameter of a colored graph since the complexity and approximability properties of most of the following optimization problems depend on the maximum span of the graph.

Definition 1.2 (span of a color) *The span of a color is the number of connected components of the subgraph induced by the edges of this color.*

An example is given in Figure 2 where color c_1 has span 2 (Fig. 2(b)) and color c_3 has span 1 (Fig. 2(c)).




Figure 3: A minimum color st -path is not necessarily composed of minimum color paths.

The counterparts in colored graphs of the classical graph problems consist of finding a minimum cost set of colors satisfying a given property. For example a **colored path** is defined as a subset of colors whose corresponding edge set contains a usual path (Fig. 2(e)). Note that the edge set corresponding to a given colored path may contain more than one usual path. Nevertheless, a usual path is associated to a single colored path. In the following we use both ways to define a colored path, either as a set of colors, or as a usual path whose edges are colored. We can define a **colored st -cut** (Fig. 2(d)), a **colored cut**, a **colored spanning tree** (Fig. 2(f)) in the same way as we have defined a colored path. Then we define the following optimization problems motivated by network survivability issues.

The **Minimum Color st -Path problem** (MC- st -Path) consists of finding a path between vertices s and t using as few colors as possible and is NP-Hard [23], which is the same as a colored path containing a minimum number of colors. This problem is related to finding a path of minimum failure probability. Furthermore, unlike classical graph shortest path, a subpath of a MC- st -Path is not necessarily a minimum color path. The set of color $\{c_1, c_4\}$ is minimum color st -path in the depicted in Figure 3. There is a single classical path in the subgraph induced by these colors, this path is (s, y, w, x, v, t) . However the color c_3 is a minimum color sw -path unlike the set $\{c_1, c_4\}$.

The **2-Color Disjoint st -Paths problem** (or *diverse routing problem* [1]) consists of finding two paths between a pair of vertices in a graph using distinct color sets, therefore no Shared Risk Resource Group failure cause both paths to fail. Unlike the traditional edge or vertex disjoint paths problem in graph theory, this problem is NP-complete [12]. It follows that minimizing the total number of colors of 2-Color Disjoint st -Paths is difficult. Moreover, it has been shown in [13] for a cryptographic application, that the problem **Maximum Number of Color Disjoint Paths** consisting in finding the maximum number of color disjoint st -paths connecting two vertices is also NP-Hard.

When no two color disjoint st -paths exists, a set of **Minimum Overlapping Color st -Paths**, that is a set of st -paths sharing a minimum number of colors, is of specific interest in survivable network design. It represents protection paths in a failure dependent protection scheme [16]. A variation of this problem, consisting of finding two edge disjoint st -paths minimizing the number of shared colors, has been studied in [23] and shown to be NP-Hard.

The **Minimum Color st -Cut problem** (MC- st -Cut) consists of finding a minimum size set of colors disconnecting two vertices s and t . Similarly, we can define a **Minimum Color Multi-Cut** (MC-Multi-Cut) as a minimum size set of colors whose removal disconnects simultaneously several pairs of given vertices $\{s_1, t_1\} \dots \{s_k, t_k\}$ and a **Minimum Color Cut** (MC-Cut) as a minimum set of colors disconnecting the graph in at least two parts. Note that as in classical graphs, the size of a

MC-Cut is upper bounded by the minimum colored degree which is the minimum number of colors used by the edges adjacent to a vertex. Colored cut problems are related to the vulnerability of multilayer networks since they give tight lower bounds on the number of failures that can disconnect totally or partially a network.

The **Minimum Color Spanning Tree problem** (MC-Spanning Tree) consists of finding a minimum size set of colors containing a connected spanning subgraph, as depicted in Figure 2. A MC-Spanning Tree is a set of Shared Risk Resource Group whose availability guarantees the connectivity of a network. The complexity and approximability properties of this problem are detailed in [3, 15, 21].

In Table 1 we sum up the main results obtained for colored problems. The results with no references are going to be proved in this paper. The weak relations between Color Disjoint st -Paths and MC- st -Cut (analogous to min cut equals max number of edge disjoint paths) are investigated in Section 2. Section 3 is dedicated to cases when the colored problems are polynomial, while Section 4 investigates complexity and approximability of colored problems when the span of colors is bounded by a constant k . Finally Section 5 presents an approximability preserving reduction from the Minimum Label Cover problem to the MC- st -Cut and MC- st -Path problems when there is no restriction on the span of colors.

2 Relations between problems

In this section we show through examples that colored problems differ from classical graph theory, not only by their complexities, but also by their mutual relationships. In particular this is the case for the number of color disjoint st -paths and the MC- st -Cut no longer tightly linked. The max flow-min cut relation being essential in the study of network reliability, this result makes the Color Disjoint st -Paths problem even more thorny and stresses the interest of the Min Overlapping Color st -Paths problem [23].

The sizes of a MC-Spanning Tree and a MC-Cut do not either verify the same inequalities as their classical counterparts [8].

Theorem 2.1 *For each $k \in \mathbb{N}$, there is a colored graph with two vertices s and t such that the MC- st -Cut has value k , while the graph does not contain two color disjoint st -paths.*

Furthermore, this colored graph is also such that the MC-Cut has value k while the graph does not contain two color disjoint spanning trees.

Proof: Let $k \in \mathbb{N}$ be a constant. We construct a colored graph G containing $\binom{2k}{k} + 1$ vertices and $2k$ colors. G is built from a path of length $\binom{2k}{k}$ where each edge is associated with exactly a one subset of k colors. Then we replace each edge of this path by k edges with the k colors of the associated subset, see Figure 4 for $k = 2$. Let s and t be the extremities of the path. A colored spanning tree in G is merely a colored st -path and a MC-Cut is actually a MC- st -Cut.

Assume there is a colored st -path using $l \leq k$ colors. Then there is a subset of at least k colors not used by the path; that implies that the path does not contain any of the multiple edges associated to this subset of colors, and leads to a contradiction. Therefore a colored st -path uses at least $k + 1$ colors. Since there are only $2k$ available colors, Theorem 2.1 follows. \square

Other families of colored graphs exist for the first part of Theorem 2.1, for example we can construct a family of graphs containing k^2 colors and $2k^2$ edges in which a MC- st -Cut uses k colors while no pair of color disjoint st -paths can be found.

Min. Color ...	st -Path	st -Cut Multi	Cut	2-Disjoint st -Paths	2-Min Overlap. st -Paths	Max. number of Disjoint st -Paths	Spanning Tree
general	complexity	NP-Hard	?	NP-Complete [12]	NP-Hard [23]	NP-Hard [13]	NP-Hard [3]
	non approx	$2^{\log^{1-\delta} C ^{\frac{1}{2}}}$?	—	?	$\forall \varepsilon > 0, V ^{\frac{1}{4}-\varepsilon}$	$o(\log(V))$ [21]
	approx	?	?	—	?	?	$O(\log(V))$ [21]
span k	complexity	NP-Hard	?	NP-Complete	NP-Hard	NP-Hard	NP-Hard [3]
	non approx	$\exists \varepsilon > 0, k^\varepsilon$?	—	?	$\forall \varepsilon > 0, V ^{\frac{1}{4}-\varepsilon}$	$o(\log(V))$ [21]
	approx	k	k	—	?	?	$O(\log(V))$ [21]
span 1	complexity						NP-Hard [3]
	non approx	P [5]	P	P	P	P	$o(\log(V))$ [21]
	approx						$O(\log(V))$ [21]
bounded degree	—	P	P	—	—	—	—

? = open question, $\delta = (\log \log |C|^{\frac{1}{2}})^{-\varepsilon}$ for $\varepsilon < \frac{1}{2}$, — = do not apply.

Table 1: Complexity and approximability properties of colored problems.



Figure 4: Example of colored graph in which the minimum color st -cut is 2 but there is no pair of color disjoint paths.

Note that an upper bound on the ratio *maximum number of color disjoint st -paths*/*MC- st -Cut*, depending on the size of the colored graph, may exist, that is why it may be interesting to find families of colored graphs in which the usual relations do not hold and for which the number of colors and the size of the graphs are known. Note also that when the number of colors of the graph equals the MC- st -Cut value, then there is at least one st -path of each color.

Moreover, another pair of problems, MC- st -Cut and MC-Multi-Cut, has different relationships than their classical counterparts.

Theorem 2.2 *MC-Multi-Cut is equivalent to MC- st -Cut.*

Proof: First MC- st -Cut is a special case of MC-Multi-Cut. Then, we prove that MC-Multi-Cut reduces to MC- st -Cut. To see this, take k copies G_1, \dots, G_k of graph G , one for each pair s_i, t_i to be disconnected in a MC-Multi-Cut instance. Now, merge all vertices $s_i \in G_i$ (resp. t_i) together into a single vertex s (resp. t). A colored st -cut in this new graph is trivially a colored multicut in G . \square

As a consequence of Theorem 2.2, in the following the MC-Multi-Cut problem is not going to be considered.

Finally, as Lemma 2.3 below proves, the MC- st -Cut and MC- st -Path problems are equivalent in the special case of series-parallel graphs. Let us recall that series-parallel graphs are graphs that can be constructed from a single edge by applying a sequence of two operations : we split an edge into two adjacent edges by introducing a vertex in the middle, and we replace an edge by two parallel edges.

Lemma 2.3 *The MC- st -Cut and MC- st -Path problems are equivalent in series-parallel graphs.*

Proof: Let G_1 be a series-parallel graph. G_1 is constructed from a single edge whose extremities are vertices s_1 and t_1 by applying a specific succession of the two operations.

Another series-parallel graph G_2 can be constructed such that the MC- st -Path problem in G_1 is equivalent to the MC- st -Cut problem in G_2 and vice versa in the following way.

Let s_2 and t_2 be the extremities of an isolated edge and construct graph G_2 by applying the complementary succession of operations of the construction of G_1 . That is, when an edge of G_1 is split into two edges, a parallel edge is added to the corresponding edge of G_2 , and when a parallel edge is added to an edge of G_1 , the corresponding edge of G_2 is split into two edges. Give to an edge of G_2 the color given to its corresponding edge in G_1 .

An s_1t_1 -path in G_1 is in bijection with an s_2t_2 -cut in G_2 since a path consists of crossing edges connected in series while a cut consists of cutting "parallel" edges. \square

By consequence of Lemma 2.3, in the rest of the paper, all the complexity and inapproximability results are based on the construction of some specific series-parallel graphs of extremities s and t in which a colored st -path is searched. Moreover, all the NP-hardness and inapproximability proofs can be given for only one of the two problems.

3 Polynomial cases

In this section we present some restrictions under which colored problems are polynomial. We first extend a result of [5] which states that when the edges of each color form a star, the MC-*st*-Path problem is polynomial.

Theorem 3.1 *When each color has span 1, the MC-*st*-Path, MC-*st*-Cut, MC-Cut, 2-Color Disjoint *st*-Paths, 2-Min Overlapping Color *st*-Paths and Maximum Number of Color Disjoint Paths problems are polynomial.*

Proof: We use the following transformation. Given a colored graph $G = (V, E, \mathcal{C})$, we construct the graph $H = (V_H, E_H)$, where each vertex of V_H represents a color of G . There is an edge between two vertices of V_H iff two edges with the corresponding colors are adjacent in G . Now add vertex s (resp. t) and edges between s (resp. t) and any vertex of H representing a color adjacent to s (resp. t) in G .

The MC-*st*-Path, MC-*st*-Cut and MC-Cut problems in G are reduced to their classical counterparts in H , which are respectively : *st*-paths, vertex-*st*-cut and vertex-disjoint *st*-paths. For 2-Min Overlapping Color *st*-Paths, note that any two *st*-paths in H necessarily overlap on vertices that are vertex-*st*-cut of size one.

For MC-Cut, note that a colored graph in which colors have span one can also be seen as an hypergraph in which the vertices spanned by a color form an hyperedge. The MC-Cut problem is then equivalent to the minimum hyperedge cut problem in a general hypergraph which is polynomial [18]. □

The case of MC-Cut is quite particular, since it can be solved in polynomial time when the number of edges of each color is bounded by a given constant k .

Theorem 3.2 *When the number of edges of each color is bounded by a given constant k , MC-Cut can be solved in polynomial time.*

Proof: Let S be the value of the minimum edge cut of the colored graph without taking its colors into account. Then, the value of the optimal color cut is in $[S/k, S]$, and the associated edge cut must have size in $[S, kS]$. In [14], Karger proved that there exist at most $|V|^{2k}$ cuts of size in $[S, kS]$. Moreover, they can be generated in polynomial time.

Hence, the problem can be solved by examining all the k -approximated minimum cuts (about $|V|^{2k}$) and counting the colors they use. □

Note that when the maximum colored degree of the graph is bounded, MC-Cut, MC-*st*-Cut and MC-Multi-Cut are also polynomial since the subsets of colors of size lower than the bound can be enumerated in polynomial time.

4 Bounded color span case

In this section, we assume that the span of each color is bounded by a constant $k \geq 1$. The case $k = 2$ is particularly interesting since if a problem is NP-Hard in this case, it is also NP-Hard for any constant $k \geq 2$.

4.1 2-Color Disjoint st -Paths

The counterpart of the x edge-disjoint st -paths problem consists of finding x color-disjoint st -paths. In the case of paths connecting two nodes (st -paths), the classical problem is easy to solve using single commodity flow techniques, or Menger's augmentations. However, the colored version of the problem turns out to be much harder even when the span of colors is bounded by two. The proof of complexity for the 2-Color Disjoint st -Paths problem presented in [12] can easily be adapted to this case.

Theorem 4.1 *There is a reduction from the Set Splitting problem to the 2-Color Disjoint st -Paths problem restricted to the case of colors of span bounded by 2.*

Proof: Given a collection of subsets $\{S_i | i \in 1, \dots, N\}$ of a finite set S , the Set Splitting problem consists of finding two disjoint subsets of S such that each subset contains at least one element in $S_i, \forall i \in 1, \dots, N$.

We construct a colored graph $G = (V, E, \mathcal{C})$ as follows. V contains $N + 1$ vertices u_1, \dots, u_{N+1} . To each triple s, S_i, S_j such that $s \in S_i \cap S_j$ we associate a color c_{ij}^s .

We connect u_i and u_{i+1} with $|S_i|$ parallel paths, each of them corresponds to an element of S_i . The path associated to s between u_i and u_{i+1} consists of edges of every color c_{ij}^s for each j such that $s \in S_i \cap S_j$.

In this graph colors have span two, and two color disjoint $u_1 u_{N+1}$ -paths correspond to two disjoint subsets of S containing each at least one element in $S_i, \forall i \in 1, \dots, N$ and vice versa. \square

4.2 Maximum Number of Color Disjoint st -Paths

The Maximum Number of Color Disjoint st -Paths Problem has been proved NP-Hard in [13] using a reduction from 3SAT. This reduction fits for the case of colors of span at most two. However the following reduction from the Maximum Independent Set problem implies a stronger inapproximability factor since this problem is not approximable within a factor $n^{\frac{1}{2}-\epsilon}$ for any $\epsilon > 0$ in a graph with n vertices unless $P = NP$ [10].

Theorem 4.2 *There is an approximation preserving reduction from the Maximum Independent Set Problem to the Maximum Number of Color Disjoint st -Paths Problem when colors have span at most two.*

The Maximum Number of Color Disjoint st -Paths Problem is not approximable within a factor $o(n^{\frac{1}{4}-\epsilon})$ for any $\epsilon > 0$ unless $P = NP$.

Proof: To an instance of Maximum Independent Set problem on a graph $G = (V, E)$, we associate the following color disjoint paths instance.

A colored graph H is constructed as follows. A color of H corresponds to an edge of G . Let s and t be two vertices of H . For each vertex v of G an st -path is added in H of length $d(v)$, the colored degree of v in G , and using the $d(v)$ colors adjacent to v in G . An example is given in Figure 5. Note that colors have span 2.

Trivially, two st -paths are color disjoint in H if and only if the two corresponding vertices are not adjacent in G . The maximum number of color disjoint paths is henceforth the size of a maximum independent set of G . In addition, $|V_H| \leq |V|^2$, and so we deduce the inapproximability result of the theorem. \square




Figure 5: Example of a colored graph built from a Max Independent Set instance.

4.3 MC-*st*-Path and MC-*st*-Cut

When the span of each color is at most two (i.e. $k = 2$), the reduction from Min Set Cover of [23] is inadequate to show the NP-hardness of MC-*st*-Path and MC-*st*-Cut since they require that the maximum size of the subsets available to cover the elements of the Min Set Cover instance is bounded by $k = 2$ too. In this case Min Set Cover can be polynomially solved by matching techniques.

The NP-hardness of MC-*st*-Path for $k = 2$ (and so of MC-*st*-Cut, by Lemma 2.3) will be proved by giving a reduction from Maximum 3 Satisfiability. This reduction also proves the inapproximability of these problems.

Theorem 4.3 *The MC-*st*-Path and MC-*st*-Cut problems are NP-Hard when each color contains at most two edges and there is an $\varepsilon > 0$ such that they are hard to approximate within a factor $1 + \varepsilon$.*

Proof: We reduce Maximum 3 Satisfiability to MC-*st*-Path. Consider an instance of Maximum 3 Satisfiability with n variables x_i , $1 \leq i \leq n$, and m clauses C_j , $1 \leq j \leq m$, each being a disjunction of three variables. The Maximum 3 Satisfiability problem consists of finding a truth assignment to the variables which satisfies the maximum number of clauses.

To such an instance we associate a colored graph G as follows. G has $n + m + 1$ vertices u_i $1 \leq i \leq n + m + 1$. Two colors, $T_{i,j}$ (T like true) and $F_{i,j}$ (F like false), are associated to each couple (i, j) such that x_i or \bar{x}_i appears in C_j . Let n_i denote the number of clauses containing either x_i or \bar{x}_i . There are $6m = 2 \sum_{i=1, \dots, n} n_i$ colors.

For each $i \in \{1, \dots, n\}$ two paths of length n_i connect u_i and u_{i+1} , the first one contains one edge of each color $T_{i,j}$ and the second one contains one edge of each color $F_{i,j}$.

For each $j \in \{1, \dots, m\}$, u_{n+j} and u_{n+j+1} are connected with three parallel edges, each one corresponds to a literal y present in the clause C_j . The color of such an edge is either $T_{i,j}$ or $F_{i,j}$ depending on $y = x_i$ or $y = \bar{x}_i$ for some $i \in \{1 \dots n\}$ (see Figure 6). Note that each color is used on exactly two edges.

In the graph G the MC-*st*-Path instance of interest is to find a path from $s = u_1$ to $t = u_{n+m+1}$ using a minimum number of colors.

First we show that to each truth assignment of the variables of the Maximum 3 Satisfiability instance satisfying $\mu \leq m$ clauses corresponds a $u_1 u_{n+m+1}$ -path in the colored graph G with $4m - \mu$ colors.

Consider an assignment of the variables, and construct an $u_1 u_{n+m+1}$ -path as follows. If the variable x_i is set to true (resp. false) the $u_1 u_{n+m+1}$ -path contains between u_i and u_{i+1} the subpath using the colors $T_{i,j}$ (resp. $F_{i,j}$).




Figure 6: Reduction of Maximum 3 Satisfiability to MC- st -Path with colors of span 2

If a clause C_j is satisfied by the assignment, the $u_1 u_{n+m+1}$ -path contains between u_{n+j} and u_{n+j+1} the edge associated to one of the variables satisfying the clause. That is, if the literal $x_i \in C_j$ and the variable x_i is set to true, the path uses the edge of color $T_{i,j}$, or if the literal $\bar{x}_i \in C_j$ and the variable x_i is set to false, the path uses the edge of color $F_{i,j}$.

If the clause C_j is not satisfied by the assignment, then the $u_1 u_{n+m+1}$ -path uses any one of the three edges connecting u_{n+j} and u_{n+j+1} .

The path described above uses $3m + (m - \mu)$ colors. More precisely, $\sum_{i=1, \dots, n} n_i = 3m$ colors are used between u_1 and u_{n+1} and for each of the μ satisfied clauses the color used belongs to these $\sum_{i=1, \dots, n} n_i$ colors. For each of the $m - \mu$ unsatisfied clauses an additional color is needed.

We now prove that a path using $3m + (m - \mu)$ colors represents a truth assignment of the variables of the Maximum 3 Satisfiability instance such that μ clauses are satisfied.

Consider a $u_1 u_{n+m+1}$ -path. The subpath between u_i and u_{i+1} either uses the edges of color $T_{i,j}$ and the variable x_i is set to true, or it uses the edges of color $F_{i,j}$ and the variable x_i is set to false.

Assume that between u_{n+j} and u_{n+j+1} the path uses one of the $3m$ colors already used between u_1 and u_{n+1} , w.l.o.g. let this color be $T_{i,j}$. Then by construction of the colored graph G , the literal x_i belongs to the clause C_j , otherwise the color $T_{i,j}$ would not be available between u_{n+j} and u_{n+j+1} . As the variable x_i is set to true, the clause C_j is satisfied.

As the $u_1 u_{n+m+1}$ -path uses $3m + (m - \mu)$ colors, μ colors among the $3m$ used in the subpath between u_1 and u_{n+1} are used between u_{n+1} and u_{n+m+1} , and thus μ clauses are satisfied. That proves the NP-hardness of MC- st -Path.

In addition, it is hard to decide if the minimum number of clauses unsatisfied is 0 or $\varepsilon_{\text{sat}} m$ for some $\varepsilon_{\text{sat}} > 0$ [20]. Hence it is hard to decide if the optimal number of colors is $3m$ or $3m + \varepsilon_{\text{sat}} m$, which completes the proof.

For the MC- st -Cut problem, note that the constructed graph is a series-parallel graph and thus the result follows by Lemma 2.3. \square

A self-improvement construction similar to the one used in an inapproximability proof for the Maximum Clique problem [11] permits to increase the inapproximability factor of Theorem 4.3 up to the factor given in the following theorem.

Theorem 4.4 *When the span of colors is bounded by a given constant k , there is an $\varepsilon > 0$ independent of k such that MC- st -Path and MC- st -Cut are hard to approximate within a factor k^ε .*

Given a colored graph $G = (V, E, \mathcal{C})$ and two vertices $s, t \in V$, the aforesaid self-improvement construction consists of a *colored graph product* G_{st}^2 (Fig. 7). First we replace each edge (u, v) of G by a copy G_{uv} of the graph G , the vertices u and v are identified with the vertices s and t of the copy. Therefore there are $|E|^2$ edges in G_{st}^2 . Then for an edge (x, y) belonging to a copy G_{uv} , we define its color as (c_{uv}, c_{xy}) , where c_{xy} is the color of edge (x, y) in G . Hence the color set of G_{st}^2 is $\mathcal{C} \times \mathcal{C}$.

Note that two edges belonging to distinct copies of G replacing edges of distinct colors of G are of distinct colors in G_{st}^2 too. In the same way the graph G_{st}^p is built recursively for any $p \in \mathbb{N}$ by replacing each edge of G_{st}^{p-1} by a copy of G .

Two key relations between solutions of MC- st -Path or MC- st -Cut in G_{st}^p and G are stated in Lemmas 4.5 and 4.6. They are based on the decomposition of a path of G_{st}^p into paths of G .

Lemma 4.5 *Let \mathcal{A} be an algorithm for MC- st -Path, then, $\forall p \in \mathbb{N}$, a colored path P_G of G of cost at most $|\mathcal{A}(G_{st}^p)|^{\frac{1}{p}}$ can be deduced from $\mathcal{A}(G_{st}^p)$.*

Proof: The proof is detailed for the case $p = 2$, but we can generalize the used path decomposition to any constant $p \in \mathbb{N}$, there are then p levels of paths instead of only the two following. Consider the path $\mathcal{A}(G_{st}^2)$ returned by algorithm \mathcal{A} for the MC- st -Path in G_{st}^2 . This path can be decomposed into an *external path* P and several *internal subpaths* P_i .

The external path P is the projection on G of the path $\mathcal{A}(G_{st}^2)$ into G , by replacing the subpaths of the copies of G by the edges of G which they correspond to. Let $|C_P|$ the number of colors used by path P in G .

An internal subpath P_i is a subpath of the path $\mathcal{A}(G_{st}^2)$ corresponding to an edge i of G . Let $|C_{P_0}|$ be the minimum number of colors used by an internal subpath among all internal subpaths.

The proof finishes by noticing that $|\mathcal{A}(G_{st}^2)| \geq |C_P| |C_{P_0}| \geq [\min\{|C_P|, |C_{P_0}|\}]^2$. \square

Lemma 4.6 *Let G be a colored graph, s and t two vertices of G and $p \in \mathbb{N}$. Optimal solutions of MC- st -Path in G and in the graph G_{st}^p satisfy $|\text{OPT}(G_{st}^p)| = |\text{OPT}(G)|^p$.*

Proof: As for Lemma 4.5, the proof is detailed for $p = 2$ but it can be easily generalized for any $p \in \mathbb{N}$.

Consider an optimal solution for MC- st -Path in G using color set C_{opt} and remark that the Cartesian product $C_{opt} \times C_{opt}$ is a colored path in G_{st}^2 . Indeed if the path P connects s to t using only colors in C_{opt} in G , the path of G_{st}^2 obtained by replacing each edge of P in G by the path P itself provides a feasible solution in G_{st}^2 and uses only colors in $C_{opt} \times C_{opt}$. Therefore $\text{OPT}(G_{st}^2) \leq \text{OPT}(G)^2$.

By lemma 4.5 if an optimal solution for MC- st -Path in G_{st}^2 uses $\text{OPT}(G_{st}^2)$ colors, then a solution for MC- st -Path in G can be found which uses less than $[\text{OPT}(G_{st}^2)]^{\frac{1}{2}}$. Thus $\text{OPT}(G) \leq [\text{OPT}(G_{st}^2)]^{\frac{1}{2}}$.

\square




Figure 7: A graph G and its square graph G_{st}^2 .

Proof: [Theorem 4.4] From a graph G where each color contains at most two edges, we construct a graph G_{st}^p , an instance of MC- st -Path, for some $p \in \mathbb{N}$. Note that in G_{st}^p each color has span at most 2^p since at most two edges belong to the same color in G and thus the number of edges of a given color is at most twice in G_{st}^p than in G_{st}^{p-1} . By Lemmas 4.5 and 4.6, there is ε_1 such that no algorithm can approximate MC- st -Path within a factor $(1 + \varepsilon_1)^p$, for otherwise it would contradict Theorem 4.3. Since this inapproximability factor is proved for colored graphs of maximum span 2^p for any $p \in \mathbb{N}$, it is still true for any colored graph of maximum span bounded by a constant $k \in \mathbb{N}$. We can write $k = 2^{\lceil \log_2 k \rceil}$. Hence in the case $p = \lceil \log_2 k \rceil$, it is hard to approximate MC- st -Path within a factor $(1 + \varepsilon_1)^{\log_2 k} \leq (1 + \varepsilon_1)^{\lceil \log_2 k \rceil}$. Let $\varepsilon = (\log_2 k)(\ln(1 + \varepsilon_1))/\ln k = \log_2(1 + \varepsilon_1)$, which is independent of k . The previous expression implies $k^\varepsilon = (1 + \varepsilon_1)^{\log_2 k}$. \square

Even if the gap between approximability and inapproximability factors remains large, the following trivial approximation algorithm shows that the MC- st -Path, MC- st -Cut and MC-Cut problems are easier to approximate than the Maximum Number of Color Disjoint Paths problem.

Theorem 4.7 *When the span of each color is bounded by a given constant k , the MC- st -Path, MC- st -Cut, MC-Cut problems can be approximated within a factor k .*

Proof: The transformations presented in the proof of Theorem 3.1 can be generalized to the case of colors of span bounded by a constant k .

Let $G = (V, E, \mathcal{C})$ be a colored graph, then $H = (V_H, E_H)$ is the graph constructed as follows. Each vertex of H represents a connected component of some color of \mathcal{C} . Two vertices of V_H are joined by an edge of E_H if the two associated components have a common vertex in G . Two vertices s and t are added to V_H and an edge of E_H connects s (resp. t) to a vertex v of V_H if s (resp. t) belongs to the component represented by v .

Now assume that Γ is the value of a minimum vertex st -cut in H . Since there are no more than k vertices in H corresponding to the same color of G , and since in the minimum color st -cut of G at most Γ components are cut, the number of colors of a minimum color st -cut in G is between Γ and Γ/k . The same reasoning works for MC- st -Path.

For MC-Cut, construct an hypergraph from the colored graph whose hyperedges are the connected components of each color. Each color generates at most k hyperedges and the same reasoning as for MC- st -Cut and MC- st -Path applies to MC-Cut and the hypergraph cut. \square

5 General Case : MC- st -Path and MC- st -Cut

In the case of colored graphs with generic span of colors, MC- st -Cut and MC- st -Path problems have been shown hard to approximate within a factor $2^{\log^{1-\delta} |\mathcal{C}|^{\frac{1}{4}}}$ with $\delta = (\log \log |\mathcal{C}|^{\frac{1}{4}})^{-\varepsilon}$ for all $\varepsilon < 1/2$ [2, 22].

In this section, we improve this result using a reduction from the Minimum Label Cover problem [6] defined as follows.

Let $B = (U, V, E)$ be a bipartite graph, L_U and L_V two sets of labels that can be assigned to the vertices of U and V , respectively. For each edge $\{u, v\} \in E_B$, with $u \in U$ and $v \in V$, a relationship $\Pi_{uv} \subseteq L_U \times L_V$ defines the pairs of labels that are *admissible* for $\{u, v\}$.

A *labeling* is a pair of functions $f_U : U \rightarrow 2^{L_U}$ and $f_V : V \rightarrow 2^{L_V} - \{\emptyset\}$ defining the assignment of a set of labels to each vertex of B , where 2^L represents the set of all subsets of a given set L .

The cost of a labeling is $\sum_{u \in U} |f_U(u)|$.

An edge $\{u, v\} \in B$, with $u \in U$ and $v \in V$, is covered by a labeling if for *every* label l_v assigned to v , there is a label l_u assigned to u such that the pair (l_u, l_v) is admissible: $(l_u, l_v) \in \Pi_{uv}$. We say

that a label l_v is *eligible* for a vertex $v \in V$ if it belongs to at least one admissible pair $(l_u, l_v) \in \Pi_{uv}$ for *each* neighbor $u \in U$ of v .

The Minimum Label Cover problem consists in finding a minimum cost labeling covering every edge of B . We assume the existence of an eligible label for each vertex $v \in V$. Otherwise, it can be decided in polynomial time that there is no labeling covering every edge of B .

In the following, we describe a reduction from Minimum Label Cover to MC-*st*-Cut and MC-*st*-Path problems with a tight control of the dimensions of the instances that are built. This allows us to increase the order of magnitude of the inapproximability factor.

Theorem 5.1 *MC-*st*-Path and MC-*st*-Cut problems are NP-Hard to approximate within a factor $2^{\log^{1-\delta} |\mathcal{C}|^{\frac{1}{2}}}$ where $\delta = (\log \log |\mathcal{C}|^{\frac{1}{2}})^{-\varepsilon}$ for any $\varepsilon < 1/2$.*

Proof: Let $B = (U, V, E), L_U, L_V, \{\Pi_{uv}\}$ be an instance of Minimum Label Cover. We built a colored graph $G = (V_{\mathcal{C}}, E_{\mathcal{C}}, \mathcal{C})$ and an instance of MC-*st*-Path in G as follows.

1. For each pair $(l, u) \in L_U \times U$, there is a color c_{lu} in \mathcal{C} . Note that $|\mathcal{C}| = |L_U||U|$.
2. For each vertex $v_i \in V$, there is a vertex x_i in $V_{\mathcal{C}}$.
3. There is an additional vertex $x_{|V|+1}$ in $V_{\mathcal{C}}$.
4. For each label $l_v \in L_V$ which is eligible for vertex $v_i \in V$, there is a path connecting x_i to x_{i+1} . For each neighbor $u \in U$ of v_i , there is a corresponding edge in the path. This edge can be multiple since it must contain an edge with color c_{lu} for each label $l \in L_U$ such that (l, l_v) is an admissible pair of labels for the edge $\{u, v_i\} : (l, l_v) \in \Pi_{uv_i}$. Thus, if there is a single label $l \in L_U$ such that (l, l_v) is an admissible pair of labels for the edge $\{u, v_i\}$, the corresponding edge is simple. However, if there are more than one such label, the corresponding edge of the path is indeed a multiple edge.

Figure 8(b) depicts the MC-*st*-Path instance that corresponds to the Minimum Label Cover instance of Figure 8(a). Note that in this case, the label 2 is not eligible for vertex e : even if it belongs to the relation Π_{ae} , it does not belong to Π_{ce} . Assigning this label to vertex e would therefore prevent from covering the edge $\{c, e\}$.

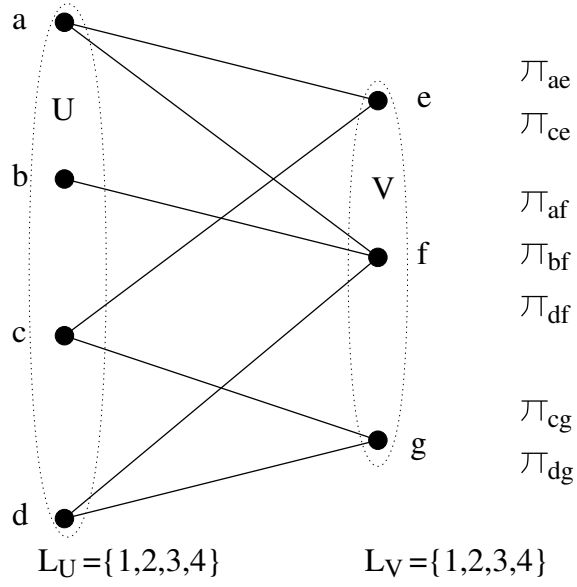
In the following we consider a $x_1 x_{|V|+1}$ -colored path \mathcal{P} with μ colors, and built a corresponding labeling covering all edges of E with cost μ .

The subpath of \mathcal{P} between vertices x_i and x_{i+1} is associated to a label $l_v \in L_V$ that is eligible for v_i . Let assign l_v to v_i in the labeling: $f_V(v_i) = \{l_v\}$.


By construction, for each u neighbor of v_i , the subpath contains an edge of color $c_{l_u u}$, where l_u belongs to an admissible pair of labels for the edge $\{u, v_i\}$. Let assign l_u to u in the labeling: $l_u \in f_U(u)$.

Consequently, the labeling constructed that way covers all edges. Its cost is μ since for each color c_{lu} in the path, there is a label that is assigned to a vertex of U , and two distinct colors correspond to distinct pairs (label, vertex).

For example the path $\{c_{a4}, c_{c2}, c_{b2}, c_{d1}\}$ in Figure 8(b) induces the assignment of label 4 to a since the color c_{a4} is used, the assignment of 2 to b and c and 1 to d . The label 1 is also assigned to e since between vertices x_1 and x_2 the path uses the subpath corresponding to label 1. In the same way, label 4 is assigned to f and label 2 to g .



(a) Instance of Minimum Label Cover



(b) Instance of MC-*st*-Path

Figure 8: Illustration of the reduction from Minimum Label Cover to MC-*st*-Path.

Conversely, we prove in the following that a labeling covering all the edges of E with cost μ induces a $x_1 x_{|V|+1}$ -colored path using μ colors.

Consider the set of colors $C = \{c_{lu} | l \in f_U(u) \subseteq L_U, u \in U\}$ with all colors corresponding to a (label, vertex) pair of the labeling. This set is of size μ . Moreover, for each vertex $v_i \in V$, each edge $\{u, v_i\}$ adjacent to v_i is covered by a label l_u . Therefore, the colors in C contains at least the edges of a $x_i x_{i+1}$ -path, as we claimed.

Henceforth, an algorithm that would yield a ρ -approximation for MC-*st*-Path would also yield a ρ -approximation for Minimum Label Cover. As a matter of fact, Minimum Label Cover cannot be approximated within $2^{\log^{1-\delta} |V|}$, where $\delta = (\log \log |V|)^{-\varepsilon}$ and $\varepsilon < \frac{1}{2}$ [6]. This results holds for instances where $|L_U| \leq 2^{\log^{1-\delta} |V|} \leq |V|^{\frac{1}{2}}$ and $|U| \leq |V|(\log \log |V|)^{\frac{1}{2}} \leq |V|^{\frac{3}{2}}$. As far as the size of the color set of G is $|C| = |U||L_U| \leq |V|^2$, MC-*st*-Path cannot be approximated within $2^{\log^{1-\delta} |C|}$ where $\delta = (\log \log |V|)^{-\varepsilon}$ and $\varepsilon < 1/2$.

For the MC-*st*-Cut problem, note that G is a series-parallel graph and that the inapproximability factor depends only on the number of colors. Thus by Lemma 2.3 the theorem is also proved for MC-*st*-Cut. \square

Remark that the MC-*st*-Path problem is a sub-case of the colored problem consisting of finding a Steiner tree[9] using the minimum number of colors. The latest problem is therefore at least as hard to approximate as MC-*st*-Path.

6 Conclusion

In this paper, we have investigated complexity and approximability properties of optimization problems yielded by the notion of Shared Risk Resource Group. This work has been conducted in the combinatorial framework given by the notion of colored graphs.

We have shown that standard relationship between classical graph theory problems cannot be imported in the colored graph settings. In particular no min/max relation analogous to the classical max flow-min cut theorem holds, while *st*-cut and multicut problems are equivalent. Nevertheless the existence of a relation between colored problems depending on the size of the colored graph has to be investigated as well as many other open questions.

We have identified an important parameter in the complexity and approximability of some problems : the span of a color. When the maximum span of the colored graph is 1, most problems are polynomial, while when the maximum span is bounded by a constant $k > 1$ they become NP-Hard or NP-Complete. These problems are even harder to approximate when there is no assumption on the span of the colors.

Concerning generic settings, *global* problems (MC-Cut and MC-Spanning Tree) and *local* problems, i.e. concerning two vertices such as MC-*st*-Path and MC-*st*-Cut, appear to belong to different complexity classes (Table 1). Global problems seem to be easier than local ones.

However, the complexity of MC-Cut remains a conjecture in the general case. Note that in different settings, the problem has been proven to be NP-Hard [7].

We proved that MC-*st*-Path and MC-*st*-Cut problems belong to the class III defined in particular in [11], i.e., that they are not approximable within a factor $2^{\log^{1-\delta} |C|^{\frac{1}{2}}}$ for some $\delta > 0$. But inapproximability within a factor $|V|^\varepsilon$ or $|C|^\varepsilon$ is still an open problem. However, we conjecture that one can find a *booster* construction for MC-*st*-Path and MC-*st*-Cut that can close the question, in a similar way as for the Maximum Clique problem [11].

References

- [1] R. Bhandari. *Survivable Networks: Algorithms for Diverse Routing*. Kluwer Academic, 1999.
- [2] R.D. Carr, S. Doddi, G. Konjevod, and M. Marathe. On the red-blue set cover problem. In *Proceedings of the eleventh annual ACM-SIAM symposium on Discrete algorithms (SODA)*, pages 345–353. Society for Industrial and Applied Mathematics, 2000.
- [3] R. Chang and S. Leu. The minimum labeling spanning trees. *Information Processing Letters*, 63(5):277–282, 1997.
- [4] A. L. Chiu, J. Strand, and R. Tkach. Issues for routing in the optical layer. *IEEE Communications Magazine*, 39(2):81–87, 2001.
- [5] P. Datta and A.K. Somani. Diverse routing for shared risk resource groups (SRRG) failures in WDM optical networks. In *Proceedings of IEEE BroadNets*, pages 120–129, October 2004.
- [6] I. Dinur and S. Safra. On the hardness of approximating label-cover. *Information Processing Letters*, 89(5):247–254, 2004.
- [7] A. Faragó. A graph theoretic model for complex network failure scenarios. In *Proceedings of the Eighth INFORMS Telecommunications Conference*, Dallas, Texas, March 2006.
- [8] A. Frank. *Connectivity and Network Flows*, volume 1 of *Handbook of combinatorics*, chapter 2. Elsevier Science, R.L. Graham, M. Grötschel, and L. Lovász edition, 1995.
- [9] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the theory of NP-completeness*. Freeman NY, 1979.
- [10] J. Hastad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182:105–142, 1999.
- [11] D.S. Hochbaum, editor. *Approximation Algorithms for NP-Hard Problems*. PWS Publishing Company, 1997.
- [12] J.Q. Hu. Diverse routing in mesh optical networks. *IEEE Transactions on Communications*, 51(3):489–494, 2003.
- [13] Q. Jiang, D.S. Reeves, and P. Ning. Improving robustness of PGP keyrings by conflict detection. In *Proceedings of the RSA Conference Cryptographers’ Track (CT-RSA2004)*, pages 194–207. LNCS 2964, February 2004.
- [14] D.R. Karger. Global min-cuts in rnc and other ramifications of a simple mincut algorithm. In *Proceedings of the fourth annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 21–30, Austin, Texas, United States, January 1993.
- [15] S.O. Krumke and H.-C. Wirth. On the minimum label spanning tree problem. *Information Processing Letters*, 66(2):81–85, 1998.
- [16] G. Mohan and C.S.R. Murthy. Lightpath restoration in WDM optical networks. *IEEE Network*, 14(6):24–32, 2000.
- [17] D. Papadimitriou, F. Poppe, J. Jones, S. Venkatachalam, S. Dharanikota, R. Jain, R. Hartani, and D. Griffith. Inference of shared risk link groups. IETF Draft, OIF Contribution, OIF 2001-066.

- [18] R. Rizzi. On minimizing symmetric set functions. Technical report, University of Trento, 1999.
- [19] P. Sebos, J. Yates, G. Hjlmsson, and A. Greenberg. Auto-discovery of shared risk link groups. In *Proceedings of IEEE/OSA OFC*, volume 3, pages WDD3–1 – WDD3–3, 2001.
- [20] V.V. Vazirani. *Approximation Algorithms*. Springer, 2002.
- [21] Y. Wan, G. Chen, and Y. Xu. A note on the minimum label spanning tree. *Information Processing Letters*, 84(2):99–101, 2002.
- [22] H.-C. Wirth. *Multicriteria Approximation of Network Design and Network Upgrade Problems*. PhD thesis, Bayerische Julius-Maximilians-Universität Würzburg, 2001.
- [23] S. Yuan, S. Varma, and J.P. Jue. Minimum-color path problems for reliability in mesh networks. In *Proceedings of IEEE InfoCom*, volume 4, pages 2658–2669, 2005.