

RESEARCH

Connectivity Problems on Heterogeneous Graphs

Jimmy Wu¹, Alex Khodaverdian², Benjamin Weitz² and Nir Yosef^{2*}

Abstract

Background: Network connectivity problems are abundant in computational biology research, where graphs are used to represent a range of phenomena: from physical interactions between molecules to more abstract relationships such as gene co-expression. One common challenge in studying biological networks is the need to extract meaningful, small subgraphs out of large databases of potential interactions. A useful abstraction for this task turned out to be the Steiner network problems: given a reference “database” graph, find a parsimonious subgraph that satisfies a given set of connectivity demands. While this formulation proved useful in a number of instances, the next challenge is to account for the fact that the reference graph may not be static. This can happen for instance, when studying protein measurements in single cells or at different time points, whereby different subsets of conditions can have different protein milieu.

Results and Discussion: We introduce the *condition* Steiner network problem in which we concomitantly consider a set of distinct biological conditions. Each condition is associated with a set of connectivity demands, as well as a set of edges that are assumed to be present in that condition. The goal of this problem is to find a minimal subgraph that satisfies all the demands through paths that are present in the respective condition. We show that introducing multiple conditions as an additional factor makes this problem much harder to approximate. Specifically, we prove that for C conditions, this new problem is NP-hard to approximate to a factor of $C - \epsilon$, for every $C \geq 2$ and $\epsilon > 0$, and that this bound is tight. Moving beyond the worst case, we explore a special set of instances where the reference graph grows *monotonically* between conditions, and show that this problem admits substantially improved approximation algorithms. We also developed an integer linear programming solver for the general problem and demonstrate its ability to reach optimality with instances from the human protein interaction network.

Conclusion: Our results demonstrate that in contrast to most connectivity problems studied in computational biology, accounting for multiplicity of biological conditions adds considerable complexity, which we propose to address with a new solver. Importantly, our results extend to several network connectivity problems that are commonly used in computational biology, such as Prize-Collecting Steiner Tree, and provide insight into the theoretical guarantees for their applications in a multiple condition setting.

Availability: Our solver for the general *condition* Steiner network problem is available at https://github.com/YosefLab/condition_connectivity_problems

Keywords: Steiner Network; NP Hard; Approximation Algorithm; Protein Protein Interaction

Background

In molecular biology applications, networks are routinely defined over a wide range of basic entities such as proteins, genes, metabolites, or drugs, which serve as nodes. The edges in these networks can have different meanings, depending on the particular context. For instance, in protein-protein interaction (PPI) networks, edges represent physical contact between proteins, either within stable multi-subunit complexes or through

transient causal interactions (i.e., an edge (x, y) means that protein x can cause a change to the molecular structure of protein y and thereby alter its activity). The body of knowledge encapsulated within the human PPI network (tens of thousands of nodes and hundreds of thousands of edges in current databases, curated from thousands of studies [1]) is routinely used by computational biologists to generate hypotheses of how various signals are transduced in eukaryotic cells [2, 3, 4, 5, 6]. The basic premise is that a process that starts with a change to the activity of protein u and ends with the activity of protein v must be propagated

*Correspondence: niryosef@eecs.berkeley.edu

²Department of Electrical Engineering and Computer Science, UC Berkeley, Berkeley, CA USA

Full list of author information is available at the end of the article

through a chain of interactions between u and v . The natural extension regards a process with a certain collection of protein pairs $\{(u_1, v_1), \dots, (u_k, v_k)\}$, where we are looking for a chain of interactions between each u_i and v_i . In another set of applications, the notion of directionality is not directly assumed and instead, one is looking for a parsimonious subgraph that connects together a set S of proteins that are postulated to be active [7, 8].

In most applications, the identity of the so called terminal nodes (i.e., (u_i, v_i) pairs or the set S) is assumed to be known (or inferred from experimental data such as ChIP-seq [7, 8, 5]), while the identity of the intermediate nodes and interactions is unknown. The goal therefore becomes to complete the gap and find a probable subgraph of the PPI network that simultaneously satisfies all the connectivity demands, thereby explaining the overall biological activity. Since the edges in the PPI network can be assigned a probability value (reflecting the credibility of their experimental evidence), by taking the negative log of these values as edge weights, the task becomes minimizing the total edge weight, leading to an instance of the Steiner Network problem. We have previously used this approach to study the propagation of a stabilizing signal in pro-inflammatory T cells, leading to the identification of a new molecular pathway (represented by a sub-graph of the PPI network) that is critical for mounting an auto-immune response, as validated experimentally by perturbation assays and disease models in mice [5]. Tuncbag et al. [8] have utilized the undirected approach using the Prize-Collecting Steiner Tree model, where the input is a network G along with a penalty function, $p(v)$ for each protein (node) in the network (based on their importance; e.g., fold-change across conditions). The goal in this case is to find a probable subtree which contains the majority of the high cost proteins in G , while accounting for penalties paid by both edge usage and missing proteins, in order to capture the biological activity represented in such a network [7, 8].

While these studies contributed to our understanding of signal transduction pathways in living cells, they do not account for a critical aspect of the underlying biological complexity. In reality, proteins (nodes) can become activated or inactivated at different conditions, thereby giving rise to a different set of potential PPIs that might take place [9]. Here, the term *condition* can refer to different points in time [10], different treatments [11], or, more recently, different cells [12]. Indeed, advances in experimental proteomics provide a way to estimate these changes at high throughput, e.g., measuring phosphorylation levels or overall protein abundance, proteome-wide for a limited number


of samples [11]. A complementary line work provides a way to evaluate the abundance of smaller numbers of proteins (typically dozens of them) in hundreds of thousands of single cells [12].

The next challenge is therefore to study connectivity problems that take into account not only the endpoints of each demand, but also the condition in which these demands should be satisfied. This added complication was tackled by Mazza et al. [13], who introduced the “Minimum k -Labeling (MKL)” problem. In this setting, each connectivity demand comes with a label, which represents a certain experimental condition or time point. The task is to label edges in the PPI network so as to satisfy each demand using its respective label, while minimizing the number of edges in the resulting sub-graph and the number of labels used to annotate these edges. While MKL was an important first step, namely introducing the notion of different demands for each condition, the more difficult challenge still remains that of considering variability in the reference graph, namely different sets of proteins that may be active and available for use in each condition.

Summary of Main Contributions

To address this open challenge, here we introduce the Condition Steiner Network (CSN) problem. In this setting, we are given a weighted undirected graph G , a set of C conditions and a set of $k \geq C$ demands, at least one per condition (note that we also cover the case of directed graphs, with similar results). The conditions are specified over a sequence of graphs G_c defined over each condition, where vertices remain the same, but edges are allowed to change across conditions (notably, our results also hold when G_c is defined with changing vertices rather than edges). Furthermore, demands are in the form of “connect node u to node v through a path of nodes that are present in condition c ”. The goal is to find a minimum weight subgraph of G that satisfied all the demands.

We first show that it is NP-hard to find a solution that is better than the trivial one (obtained by solving the problem independently for each condition). This result extends to several types of connectivity problems and provides a theoretical lower bounds to the best-possible approximation guarantee that can be achieved in a multiple condition setting (Table 1). For instance, we can conclude that concomitantly solving the shortest path problem for a set of conditions is hard to approximate, and that the trivial solution (i.e., solving the problem to optimality in each condition) is, theoretically, the best that one can do. Another example, commonly used in PPI analysis, is the Prize-Collecting Steiner tree problem [7, 8]. Here, our results indicate that given a fixed input for this problem (i.e.,



a penalty function $p(v)$ for each vertex), it is NP hard to solve it concomitantly in C conditions, such that the weight of the obtained solution is less than C times that of the optimal solution. Interestingly, a theoretical guarantee of $C \cdot (2 - \frac{2}{|V|})$ ^[1] can be obtained by solving the problem independently for each time point

While these results provide a somewhat pessimistic view, they rely on the assumption that the network frames G_c are arbitrary. In the last part of this paper, we show that for the specific case where the conditions can be ordered such that each condition is a subset of the next (namely, $G_c \subseteq G_{c'}$ for $c \leq c'$) then the CSN problem can be reduced and solved as a standard connectivity problem with a single condition, leading to substantially better theoretical guarantees. Finally, we develop an Integer Linear Programming algorithm for the general CSN problem, and show that provided with real-world input (namely, the human PPI) it is capable of reaching an optimal solution in a reasonable amount of time.

Introduction to Steiner Problems

The Steiner Tree problem, along with its many variants and generalizations, form a core family of NP-hard combinatorial optimization problems. Traditionally, the input to one of these problems is a single (usually weighted) graph, along with requirements about which nodes need to be connected in some way; the goal is to pick a minimum-weight subgraph satisfying the connectivity demands.

^[1] V is the set of nodes in the reference graph G

In this paper, we offer a *multi-condition* perspective; in our setting, multiple graphs over the same vertex set (which one can think of as an initial graph changing over a set of discrete conditions), are all given as input, and the goal is to pick a subgraph satisfying condition-sensitive connectivity requirements. Our study of this problem draws motivation and techniques from several lines of research, which we briefly summarize.

Classic Steiner problems

A basic problem in graph theory is finding the *shortest path* between two nodes; this problem is efficiently solved using, for example, Dijkstra's algorithm.

A natural extension of this is the Steiner Tree problem: given a weighted undirected graph $G = (V, E)$ and a set of terminals $T \subseteq V$, find a minimum-weight subtree that connects all the nodes in T . A further generalization is Steiner Forest: given $G = (V, E)$ and a set of demand pairs $D \subseteq V \times V$, find a subgraph that connects each pair in D . Currently the best known approximation algorithms give a ratio of 1.39 for Steiner Tree [14] and 2 for Steiner Forest [15]. These problems are known to be NP-hard to approximate to within some small constant [16].

For directed graphs, we have the Directed Steiner Network (DSN) problem, in which we are given a weighted directed graph $G = (V, E)$ and k demands $(a_1, b_1), \dots, (a_k, b_k) \in V \times V$, and must find a minimum-weight sub-graph in which each a_i has a path to b_i . When k is fixed, DSN admits a polynomial-time exact algorithm [17]. For general k ,

the best known approximation algorithms have ratio $O(k^{1/2+\epsilon})$ for any fixed $\epsilon > 0$ [18, 19]. On the complexity side, Dodis and Khanna [20] ruled out a polynomial-time $O(2^{\log^{1-\epsilon} n})$ -approximation for unless NP has quasipolynomial-time algorithms^[2]. An important special case of DSN is Directed Steiner Tree, in which all demands have the form (r, b_i) for some root node r . This problem has a $O(k^\epsilon)$ -approximation scheme [21] and a lower bound of $\Omega(\log^{2-\epsilon} n)$ [22].

Finally, a Steiner variant that has found extensive use in computational biology is the Prize-Collecting Steiner Tree problem, in which the input contains a weighted undirected graph $G = (V, E)$ and penalty function $p : V \rightarrow \mathbb{R}_{\geq 0}$; the goal is to find a subtree which simultaneously minimizes the weights of the edges in the tree and the penalties paid for nodes not included within the tree, i.e. $\text{cost}(T) := \sum_{e \in T} w(e) + \sum_{v \notin T} p(v)$. For this problem, an approximation algorithm with ratio 1.967 is known [23].

Condition Steiner problems

In this paper, we generalize the Shortest Path, Steiner Tree, Steiner Forest, Directed Steiner Network, and Prize-Collecting Steiner Tree problems to the multi-condition setting. In this setting, we have a set of *conditions* $[C] := \{1, \dots, C\}$, and are given a graph for each condition.

Our main object of study is the natural generalization of Steiner Forest (in the undirected case) and Directed Steiner Network (in the directed case), which we call Condition Steiner Network:

Definition 1 (Condition Steiner Network (CSN))

We are given the following inputs:

- 1 A sequence of undirected graphs $G_1 = (V, E_1), G_2 = (V, E_2), \dots, G_C = (V, E_C)$, one for each condition $c \in [C]$. Each edge e in the underlying edge set $E := \bigcup_c E_c$ has a weight $w(e) \geq 0$.
- 2 A set of k connectivity demands $\mathcal{D} \subseteq V \times V \times [C]$. We assume that for every $c \in C$ there exists at least one demand and therefore that $k \geq |C|$.

We call $G = (V, E)$ the underlying graph. We say a subgraph $H \subseteq G$ satisfies demand $(a, b, c) \in \mathcal{D}$ iff H contains an a - b path P along which all edges exist in G_c . The goal is to output a minimum-weight subgraph $H \subseteq G$ that satisfies every demand in \mathcal{D} .

Definition 2 (Directed Condition Steiner Network (DCSN)) *This is the same as CSN except that all the edges are directed, and a demand (a, b, c) must be satisfied by a directed path from a to b in G_c .*

We can also define the analogous generalizations of Shortest Path, (undirected) Steiner Tree, and Prize-Collecting Steiner Tree. We give hardness results and algorithms for these problems by demonstrating reductions to and from CSN and DCSN.

Definition 3 (Condition Shortest Path (CSP), Directed Condition Shortest Path (DCSP)) *These are the special cases of CSN and DCSN in which the demands are precisely $(a, b, 1), \dots, (a, b, C)$ where $a, b \in V$ are common source and target nodes.*

Definition 4 (Condition Steiner Tree (CST)) *We are given a sequence of undirected graphs $G_1 = (V, E_1), \dots, G_C = (V, E_C)$, weights $w(\cdot) \geq 0$ on the edges, and sets of terminal nodes $X_1, \dots, X_C \subseteq V$. We say a subgraph $H \subseteq (V, \bigcup_c E_c)$ satisfies the terminal set X_c iff the nodes in X_c are mutually reachable using edges in H that exist in condition c . The goal is to find a minimum-weight subgraph H that satisfies X_c for every $c \in [C]$.*

Definition 5 (Condition Prize-Collecting Steiner Tree (CPCST)) *We are given a sequence of undirected graph $G_1 = (V, E_1), \dots, G_C = (V, E_C)$, a weight $w(e) \geq 0$ on each $e \in E$, and a penalty $p(v, c) \geq 0$ for each $v \in V, c \in [C]$. The goal is to find a subtree T that minimizes $\sum_{e \in T} w(e) + \sum_{v \notin T, c \in [C]} p(v, c)$.*

Finally, in molecular biology applications, it is often the case that all the demands originate from a common root node. To capture this, we define the following special case of DCSN:

Definition 6 (Single-Source DCSN) *This is the special case of DCSN in which the demands are precisely $(a, b_1, c_1), (a, b_2, c_2), \dots, (a, b_k, c_k)$, for some root $a \in V$. We can assume that $c_1 \leq c_2 \leq \dots \leq c_k$.*

It is also natural to consider variants of these problems in which nodes (rather than edges) vary across the conditions, or in which both nodes and edges vary. In **Problem variants**, we show that all three variants are in fact equivalent; thus we focus on the edge-based formulations.

Our Results

In this work, we perform a systematic study of the condition Steiner problems defined above, from the standpoint of approximation algorithms—that is, algorithms that return subgraphs whose total weights are not too much greater than that of the optimal subgraph—as well as integer linear programming (ILP). Since all of the condition Steiner problems listed

^[2]Throughout this paper, $n := |V|$ denotes the number of nodes in the relevant graph.

Table 1 Approximation bounds for the various Steiner Network Problems in their classic setting and condition setting. For the classic problems, we have indicated the papers in which the bounds are shown. For the condition problems, all the lower bounds are developed in the present work; all the upper bounds are the naive bounds obtained from the “union of shortest paths” heuristic, or from applying the best known approximation algorithm for the appropriate classic Steiner problem to each condition, then taking the union of those solutions.

Problems	Classic		Condition	
	Lower Bound	Upper Bound	Lower Bound(s)	Upper Bound(s)
Steiner Forest	1.01 [16]	2 [15]	$C - \epsilon, k - \epsilon$	$2C, k$
Directed Steiner Network	$k^{1/4-o(1)}$ [24]	$k^{1/2+\epsilon}$ [18, 19]	$C - \epsilon, k - \epsilon$	$C \cdot k^{1/2+\epsilon}, k$
Undirected Shortest Path	N/A	1	$C - \epsilon$	C
Directed Shortest Path	N/A	1	$C - \epsilon$	C
Steiner Tree	1.01 [16]	1.39 [14]	$C - \epsilon$	$1.39C$
Prize-Collecting Steiner Tree	1.01 [16]	1.97 [23]	$C - \epsilon$	$1.97C$

in the previous section turn out to be NP-hard (and in fact all of them except Shortest Path are hard even in the classic single-condition setting) we cannot hope for algorithms that are optimal and run in polynomial time.

First, in [Hardness of condition Steiner problems](#), we show a series of strong negative results, starting with (directed and undirected) Condition Steiner Network:

Theorem 1 (Main Theorem) *CSN and DCSN are NP-hard to approximate to a factor of $C - \epsilon$ as well as $k - \epsilon$ for every fixed $k \geq 2$ and every constant $\epsilon > 0$. For DCSN, this holds even when the underlying graph is acyclic.*

Thus the best approximation ratio one can hope for is C or k ; the latter upper bound is easily achieved by the trivial “union of shortest paths” algorithm: for each demand (a, b, c) , compute the shortest a - b path at condition c ; then take the union of these k paths. This contrasts with the classic Steiner network problems, which have nontrivial approximation algorithms and efficient fixed-parameter algorithms.

Next, we show similar hardness results for the other three condition Steiner problems. This is achieved by a series of simple reductions from CSN and DCSN.

Theorem 2 *Condition Shortest Path, Directed Condition Shortest Path, Condition Steiner Tree, and Condition Prize-Collecting Steiner Tree are all NP-hard to approximate to a factor of $C - \epsilon$ for every fixed $C \geq 2$ and $\epsilon > 0$.*

Note that each of these condition Steiner problems can be naively approximated by applying the best known algorithm for the classic version of that problem in each graph in the input, then taking the union of all those subgraphs. If the corresponding classic Steiner problem can be approximated to a factor of α , then

this process gives an $\alpha \cdot C$ -approximation for the condition version. Thus using known constant-factor approximation algorithms, each of the condition problems in Theorem 2 has an $O(C)$ -approximation algorithm. Our result shows that in the worst case, one cannot do much better.

While these results provide a somewhat pessimistic view, the proofs rely on the assumption that the edge sets in the input networks (that is, E_1, \dots, E_C) do not necessarily bear any relationship to one another. In [Monotonic special cases](#), we move beyond this worst-case assumption by studying a broad class of special cases in which the conditions are *monotonic*: if an edge e exists in some graph G_c , then it exists in all the subsequent graphs $G_{c'}, c' \geq c$. In other words, each graph in the input is a subgraph of the next. For these problems, we prove the following two theorems:

Theorem 3 *Monotonic CSN has a polynomial-time $O(\log k)$ -approximation algorithm. It has no $\Omega(\log \log n)$ -approximation algorithm unless $NP \in DTIME(n^{\log \log \log n})$.*

In the directed case, for monotonic DCSN with a single source (that is, every demand is of the form (r, b, c) for a common root node r), we show the following:

Theorem 4 *Monotonic Single-Source DCSN has a polynomial-time $O(k^\epsilon)$ -approximation algorithm for every $\epsilon > 0$. It has no $\Omega(\log^{2-\epsilon} n)$ -approximation algorithm unless $NP \in ZPTIME(n^{\text{polylog}(n)})$.*

These bounds are proved via approximation-preserving reductions to and from classic Steiner problems, namely Priority Steiner Tree and Directed Steiner Tree. Conceptually, this demonstrates that imposing the monotonicity requirement makes the condition Steiner problems much closer to their classic counterparts, allowing us to obtain algorithms with substantially better approximation guarantees.

Finally in **Application to protein-protein interaction networks**, we show how to model various condition Steiner problems as integer linear programs (ILPs). In experiments on real-world inputs derived from the human PPI network, we find that these ILPs are capable of reaching optimal solutions in a reasonable amount of time.

Table 1 summarizes our results, emphasizing how the known upper and lower bounds change when going from the classic Steiner setting to the condition Steiner setting.

Preliminaries

Note that the formulations of CSN and DCSN in the Introduction involved a fixed vertex set; only the edges change over the conditions. It is also natural to formulate the condition Steiner network problem with nodes changing over condition, or both nodes and edges. However by the following proposition, it is no loss of generality to discuss only the edge-condition variant.

Proposition 1 *The edge, node, and node-and-edge variants of CSN are mutually polynomial-time reducible via strict reductions (i.e. preserving the approximation ratio exactly). Similarly all three variants of DCSN are mutually strictly reducible.*

We defer the precise definitions of the other two variants, as well as the proof of this proposition, to **Problem variants**.

Next we state the Label Cover problem, which is the starting point of one of our reductions to CSN.

Definition 7 (Label Cover (LC)) *An instance of this problem consists of a bipartite graph $G = (U, V, E)$ and a set of possible labels Σ . The input also includes, for each edge $(u, v) \in E$, projection functions $\pi_u^{(u,v)} : \Sigma \rightarrow C$ and $\pi_v^{(u,v)} : \Sigma \rightarrow C$, where C is a common set of colors; $\Pi = \{\pi_v^e : e \in E, v \in e\}$ is the set of all such functions. A labeling of G is a function $\phi : U \cup V \rightarrow \Sigma$ assigning each node a label. We say a labeling ϕ satisfies an edge $(u, v) \in E$, or (u, v) is consistent under ϕ , iff $\pi_u^{(u,v)}(\phi(u)) = \pi_v^{(u,v)}(\phi(v))$. The task is to find a labeling that satisfies as many edges as possible.*

This slightly generalizes the original definition in [25]. It has the following gap hardness, which follows by combining the PCP theorem [26] with Raz’s parallel repetition theorem [27].

Theorem 5 *For every $\epsilon > 0$, there is a constant $|\Sigma|$ such that the following promise problem is NP-hard: Given a Label Cover instance (G, Σ, Π) , distinguish between the following cases:*

- (YES instance) *There exists a total labeling of G ; i.e. a labeling that satisfies every edge.*
- (NO instance) *There does not exist a labeling of G that satisfies more than $\epsilon|E|$ edges.*

In **Hardness of condition Steiner problems**, we use Label Cover to show $(2 - \epsilon)$ -hardness for 2-CSN and 2-DCSN; that is, when there are only two demands. To prove our main result however, we will actually need a generalization of Label Cover to partite hypergraphs, called k -Partite Hypergraph Label Cover. Out of space considerations we defer the statement of this problem and its gap hardness to **Inapproximability for general C and k** , where the $(2 - \epsilon)$ -hardness result is generalized to show $(C - \epsilon)$ -hardness and $(k - \epsilon)$ -hardness for general number of conditions C and demands k .


Hardness of condition Steiner problems

Overview of the reduction

Here we outline our strategy for reducing Label Cover to the condition Steiner problems. First, we reduce to the CSN problem restricted to having only $C = 2$ conditions and $k = 2$ demands; we call this problem 2-CSN. The directed problem 2-DCSN is defined analogously. Later, we obtain similar hardness for CSN with more conditions or demands by using the same ideas, but reducing from k -Partite Hypergraph Label Cover.

Consider the nodes $u_1, \dots, u_{|U|}$ on the “left” side of the LC instance. We build, for each u_i , a gadget (which is a small sub-graph in the Steiner instance) consisting of multiple parallel directed paths from a source to a sink—one path for each possible label for u_i . We then chain together these gadgets, so that the sink of u_1 ’s gadget is the source of u_2 ’s gadget, and so forth. Finally we create a connectivity demand from the source of u_1 ’s gadget to the sink of $u_{|U|}$ ’s gadget, so that a solution to the Steiner instance must have a path from u_1 ’s gadget, through all the other gadgets, and finally ending at $u_{|U|}$ ’s gadget. This path, depending on which of the parallel paths it takes through each gadget, induces a labeling of the left side of the Label Cover instance. We build an analogous chain of gadgets for the nodes on the right side of the Label Cover instance.

The last piece of the construction is to ensure that the Steiner instance has a low-cost solution if and only if the Label Cover instance has a consistent labeling. This is accomplished by setting all the u_i gadgets to exist only at condition 1 (i.e. in frame G_1), setting the v_j gadgets to exist only in G_2 , and then merging certain edges from the u_i -gadgets with edges from the v_j -gadgets, replacing them with a single, shared edge that exists in both frames. Intuitively, the edges we merge are from paths that correspond to labels that satisfy the Label Cover edge constraints. The result




is that a YES instance of Label Cover (i.e. one with a total labeling) will enable a high degree of overlap between paths in the Steiner instance, so that there is a very low-cost solution. On the other hand, a NO instance of LC will not result in much overlap between the Steiner gadgets, so every solution will be costly.

Let us define some of the building blocks of the reduction we just sketched:


- A *bundle* is a graph gadget consisting of a source node b_1 , sink node b_2 , and parallel, disjoint *strands* (defined shortly) from b_1 to b_2 .
- A *chain* of bundles is a sequence of bundles, with the sink of one bundle serving as the source of another.
- A *simple strand* is a directed path of the form $b_1 \rightarrow c_1 \rightarrow c_2 \rightarrow b_2$.
- In a simple strand, we say that (c_1, c_2) is the *contact edge*. Contact edges have weight 1; all other edges in our construction have zero weight.
- More generally, a *strand* can be made more complicated, by replacing a contact edge with another bundle (or even a chain of them). In this way, bundles can be nested, as shown in Figure 1.
- We can *merge* two or more simple strands from different bundles by setting their contact edges to be the same edge, and making that edge existent at the union of all conditions when the original edges existed (Figure 1).

Before formally giving the reduction, we illustrate a simple example of its construction.

Example 1 Consider a toy Label Cover instance whose bipartite graph is a single edge, label set is $\Sigma = \{1, 2\}$, and projection functions are shown:



Our reduction outputs this corresponding 2-CSN instance:



G_1 comprises the set of blue edges; G_2 is green. The demands are $(u_1^S, u_2^S, 1)$ and $(v_1^S, v_2^S, 2)$. The gadget for the Label Cover node u (the blue sub-graph) consists of two strands, one for each possible label. In the v -gadget (green sub-graph), the strand corresponding to a labeling of ‘2’ branches further, with one simple strand for each agreeing labeling of u . Finally, strands (more precisely, their contact edges) whose labels map to the same color are merged.

The input is a YES instance of Label Cover whose optimal labelings (u gets either label 1 or 2, v gets label 2) correspond to 2-CSN solutions of cost 1 (both gadgets traverse the $(u, 1, v, 2)$ -path, or both traverse the $(u, 2, v, 2)$ -path). If this were a NO instance and edge e could not be satisfied, then the resulting 2-CSN gadgets would have no overlap.

Inapproximability for two demands

We now formalize the reduction in the case of two conditions and two demands; later, we extend this to general C and k .

Theorem 6 2-CSN and 2-DCSN are NP-hard to approximate to within a factor of $2 - \epsilon$ for every constant $\epsilon > 0$. For 2-DCSN, this holds even when the underlying graph is acyclic.

Proof We describe a reduction from Label Cover (LC) to 2-DCSN with an acyclic graph. Given the LC instance $(G = (U, V, E), \Sigma, \Pi)$, construct a 2-DCSN instance $(\mathcal{G} = (G_1, G_2))$, along with two connectivity demands) as follows. Create nodes $u_1^S, \dots, u_{|U|+1}^S$ and

$v_1^S, \dots, v_{|V|+1}^S$. Let there be a bundle from each u_i^S to u_{i+1}^S ; we call this the u_i -bundle, since a choice of path from u_i^S to u_{i+1}^S in \mathcal{G} will indicate a labeling of u_i in G .

The u_i -bundle has a strand for each possible label $\ell \in \Sigma$. Each of these ℓ -strands consists of a chain of bundles—one for each edge $(u_i, v) \in E$. Finally, each such (u_i, v) -bundle has a simple strand for each label $r \in \Sigma$ such that $\pi_{u_i}^{(u_i, v)}(\ell) = \pi_v^{(u_i, v)}(r)$; call this the (u_i, ℓ, v, r) -path. In other words, there is ultimately a simple strand for each possible labeling of u_i 's neighbor v such that the two nodes are in agreement under their mutual edge constraint. If there are no such consistent labels r , then the (u_i, v) -bundle consists of just one simple strand, which is not associated with any r . Note that every minimal $u_1^S \rightarrow u_{|U|+1}^S$ path (that is, one that proceeds from one bundle to the next) has total weight exactly $|E|$.

Similarly, create a v_j -bundle from each v_j^S to v_{j+1}^S , whose r -strands (for $r \in \Sigma$) are each a chain of bundles, one for each $(u, v_j) \in E$. Each (u, v_j) -bundle has a (u, ℓ, v_j, r) -path for each agreeing labeling ℓ of the neighbor u , or a simple strand if there are no such labelings.

Set all the edges in the u_i -bundles to exist in G_1 only. Similarly the v_j -bundles exist solely in G_2 . Now, for each (u, ℓ, v, r) -path in G_1 , merge it with the (u, ℓ, v, r) -path in G_2 , if it exists. The demands are $\mathcal{D} = \left\{ \left(u_1^S, u_{|U|+1}^S, 1 \right), \left(v_1^S, v_{|V|+1}^S, 2 \right) \right\}$.

We now analyze the reduction. The main idea is that any $u_i^S \rightarrow u_{i+1}^S$ path induces a labeling of u_i ; thus the demand $\left(u_1^S, u_{|U|+1}^S, 1 \right)$ ensures that any 2-DCSN solution indicates a labeling of all of U . Similarly, $\left(v_1^S, v_{|V|+1}^S, 2 \right)$ forces an induced labeling of V . In the case of a YES instance of Label Cover, these two connectivity demands can be satisfied by taking two paths with a large amount of overlap, resulting in a low-cost 2-DCSN solution. In contrast when we start with a NO instance of Label Cover, any two paths we can choose to satisfy the 2-DCSN demands will be almost completely disjoint, resulting in a costly solution. We now fill in the details.

Suppose the Label Cover instance is a YES instance, so that there exists a labeling ℓ_u^* to each $u \in U$, and r_v^* to each $v \in V$, such that for all edges $(u, v) \in E$, $\pi_{u_i}^{(u, v)}(\ell_u^*) = \pi_v^{(u, v)}(r_v^*)$. The following is an optimal solution \mathcal{H}^* to the constructed 2-DCSN instance:

- To satisfy the demand at condition 1, for each u -bundle, take a path through the ℓ_u^* -strand. In particular for each (u, v) -bundle in that strand, traverse the (u, ℓ_u^*, v, r_v^*) -path.

- To satisfy the demand at condition 2, for each v -bundle, take a path through the r_v^* -strand. In particular for each (u, v) -bundle in that strand, traverse the (u, ℓ_u^*, v, r_v^*) -path.

In tallying the total edge cost, $\mathcal{H}^* \cap G_1$ (i.e. the subgraph at condition 1) incurs a cost of $|E|$, since one contact edge in \mathcal{G} is encountered for each edge in G . $\mathcal{H}^* \cap G_2$ accounts for no additional cost, since all contact edges correspond to a label which agrees with some neighbor's label, and hence were merged with the agreeing contact edge in $\mathcal{H}^* \cap G_1$. Clearly a solution of cost $|E|$ is the best possible, since every $u_1^S \rightarrow u_{|U|+1}^S$ path in G_1 (and every $v_1^S \rightarrow v_{|V|+1}^S$ path in G_2) contains at least $|E|$ contact edges.

Conversely suppose we started with a NO instance of Label Cover, so that for any labeling ℓ_u^* to u and r_v^* to v , for at least $(1 - \epsilon)|E|$ of the edges $(u, v) \in E$, we have $\pi_{u_i}^{(u, v)}(\ell_u^*) \neq \pi_v^{(u, v)}(r_v^*)$. By definition, any solution to the constructed 2-DCSN instance contains a simple $u_1^S \rightarrow u_{|U|+1}^S$ path $P_1 \in G_1$ and a simple $v_1^S \rightarrow v_{|V|+1}^S$ path $P_2 \in G_2$. P_1 alone incurs a cost of exactly $|E|$, since one contact edge in \mathcal{G} is traversed for each edge in G . However, P_1 and P_2 share at most $\epsilon|E|$ contact edges (otherwise, by the merging process, this implies that more than $\epsilon|E|$ edges could be consistently labeled, which is a contradiction). Thus the solution has a total cost of at least $(2 - \epsilon)|E|$.

The underlying directed graph we constructed is acyclic, as every edge points “to the right” as in Example 1. It follows from the gap between the YES and NO cases that 2-DCSN is NP-hard to approximate to within a factor of $2 - \epsilon$ for every $\epsilon > 0$, even on acyclic graphs. Finally, note that the same analysis holds for 2-CSN, by simply making every edge undirected; however in this case the graph is clearly not acyclic. \square

Inapproximability for general C and k

Theorem 1 (Main Theorem) *CSN and DCSN are NP-hard to approximate to a factor of $C - \epsilon$ as well as $k - \epsilon$ for every fixed $k \geq 2$ and every constant $\epsilon > 0$. For DCSN, this holds even when the underlying graph is acyclic.*

Proof We perform a reduction from k -Partite Hypergraph Label Cover, a generalization of Label Cover to hypergraphs, to CSN, or DCSN with an acyclic graph. Using the same ideas as in the $C = k = 2$ case, we design k demands composed of parallel paths corresponding to labelings, and merge edges so that a good global labeling corresponds to lots of overlaps between those paths. The full proof is left to [Inapproximability for general \$C\$ and \$k\$](#) . \square

Note that a k -approximation algorithm is to simply choose $\mathcal{H} = \bigcup_c \tilde{P}_c$, where \tilde{P}_c is the shortest $a_c \rightarrow b_c$ path in G_c . Thus by Theorem 1, essentially no better approximation is possible in terms of k alone. In contrast, most classic Steiner problems have good approximation algorithms, or are even exactly solvable for constant k .

Inapproximability for Steiner variants

We take advantage of our previous hardness of approximation results in Theorem 1 and show, via a series of reductions, that CSP, CSN, and CPCST are also hard to approximate.

Theorem 2 *Condition Shortest Path, Directed Condition Shortest Path, Condition Steiner Tree, and Condition Prize-Collecting Steiner Tree are all NP-hard to approximate to a factor of $C - \epsilon$ for every fixed $C \geq 2$ and $\epsilon > 0$.*

Proof We first reduce from CSN to CSP (and DCSN to DCSP). Suppose we are given an instance of CSN with graph sequence $\mathcal{G} = (G_1, \dots, G_C)$, underlying graph $G = (V, E)$, and demands $\mathcal{D} = \{(a_i, b_i, c_i) : i \in [k]\}$. We build a new instance $(\mathcal{G}' = (G'_1, \dots, G'_k), G' = (V', E'), \mathcal{D}')$ as follows.

Initialize G' to G . Add to G' the new nodes a and b , which exist at all conditions/in all frames G'_i . For all $e \in E$ and $i \in [k]$, if $e \in G_{c_i}$, then let e exist in G'_i as well. For each $(a_i, b_i, c_i) \in \mathcal{D}$,

- 1 Create new nodes x_i, y_i . Create zero-weight edges (a, x_i) , (x_i, a_i) , (b_i, y_i) , and (y_i, b) .
- 2 Let (a, x_i) , (x_i, a_i) , (b_i, y_i) , and (y_i, b) exist only in frame G'_i .

Lastly, the demands are $\mathcal{D}' = \{(a, b, i) : i \in [k]\}$.

Given a solution $H' \subseteq G'$ containing an $a \rightarrow b$ path at every condition $i \in [k]$, we can simply exclude nodes $a, b, \{x_i\}$, and $\{y_i\}$ to obtain a solution $H \subseteq G$ to the original instance, which contains an $a_i \rightarrow b_i$ path in G_{c_i} for all $i \in [k]$, and has the same cost. The converse is also true by including these nodes.

Observe that essentially the same procedure shows that DCSN reduces to DCSP; simply ensure that the edges added by the reduction are directed rather than undirected.

Next, we reduce CSP to CST. Suppose we are given an instance of CSP with graph sequence $\mathcal{G} = (G_1, \dots, G_C)$, underlying graph $G = (V, E)$, and demands $\mathcal{D} = \{(a, b, i) : i \in [C]\}$. We build a new instance of CST as follows:

$(\mathcal{G}' = (G'_1, \dots, G'_C), G' = (V', E'), \mathcal{X} = (X_1, \dots, X_C))$. Set \mathcal{G}' to \mathcal{G} , and G' to G . Take the set of terminals in each condition to be $X_i = \{a, b\}$. We note that a solution $H' \subseteq G'$ to the CST instance is trivially

a solution the CSP instance with the same cost, and vice-versa.

Finally, we reduce CST to CPCST. We do this by making an appropriate assignment of the penalties $p(v, c)$. Suppose we are given an instance of CST with graph sequence $\mathcal{G} = (G_1, \dots, G_C)$, underlying graph $G = (V, E)$, and terminal sets $\mathcal{X} = (X_1, \dots, X_C)$. We build a new instance of CPCST, $(\mathcal{G}' = (G'_1, \dots, G'_C), G' = (V', E'), p(v, c))$. In particular, set \mathcal{G}' to \mathcal{G} , and G' to G . Set $p(v, c)$ as follows:

$$p(v, c) = \begin{cases} \infty, & v \in X_c \\ 0, & \text{otherwise} \end{cases}$$

Consider any solution $H \subseteq G$ to the original CST instance. Since H spans the terminals X_1, \dots, X_C (thus avoiding any infinite penalties), and since the non-terminal vertices have zero cost, the overall cost of H remains the same cost in the constructed CPCST instance. Conversely, suppose we are given a solution $H' \subseteq G'$ to the constructed CPCST instance. If the cost of H' is ∞ , then there exists no solution that can span all the X_c 's simultaneously, and thus no solution exists for the CST instance. On the other hand if H' has finite cost, then H' is also a solution for the CST instance, with the same cost.

To summarize: in the first reduction from CSN to CSP, the number of demands, k , in the CSN instance is the same as the number of the conditions, C , in the CSP instance; we conclude that CSP is NP-hard to approximate to a factor of $C - \epsilon$ for every fixed $C \geq 2$ and $\epsilon > 0$. Since C remains the same in the two subsequent reductions, we also have that CST and CPCST are NP-hard to approximate to a factor of $C - \epsilon$. \square

Monotonic special cases

In light of the strong lower bounds in the previous theorems, in this section we consider more tractable special cases of the condition Steiner problems. A natural restriction is that the changes over conditions are *monotonic*:

Definition 8 (Monotonic {CSN, DCSN, CSP, DCSP, CST, CPCST}) *In this special case (of any of the condition Steiner problems), we have that for each $e \in E$ and $c \in [C]$, if $e \in G_c$, then $e \in G_{c'}$ for all $c' \geq c$.*

We now examine the effect of monotonicity on the complexity of the condition Steiner problems.

Monotonicity in the undirected case

In the undirected case, we show that monotonicity has a simple effect: it makes CSN equivalent to the following well-studied problem:

Definition 9 (Priority Steiner Tree [28]) *The input is a weighted undirected multigraph $G = (V, E, w)$, a priority level $p(e)$ for each $e \in E$, and a set of k demands (a_i, b_i) , each with priority $p(a_i, b_i)$. The output is a minimum-weight forest $F \subseteq G$ that contains, between each a_i and b_i , a path in which every edge e has priority $p(e) \leq p(a_i, b_i)$.*

Priority Steiner Tree was introduced by Charikar, Naor, and Schieber [28], who gave a $O(\log k)$ approximation algorithm. Moreover, it cannot be approximated to within a factor of $\Omega(\log \log n)$ assuming $\text{NP} \notin \text{DTIME}(n^{\log \log \log n})$ [29]. We now show that the same bounds apply to Monotonic CSN.

Lemma 1 *Priority Steiner Tree and Monotonic CSN have the same approximability.*

Proof We transform an instance of Priority Steiner Tree into an instance of Monotonic CSN as follows: the set of priorities becomes the set of conditions; if an edge e has priority $p(e)$, it now exists at all conditions $t \geq p(e)$; if a demand (a_i, b_i) has priority $p(a_i, b_i)$, it now becomes $(a_i, b_i, p(a_i, b_i))$. If there are parallel multiedges, break up each such edge into two edges of half the original weight, joined by a new node. Given a solution $H \subseteq G$ to this CSN instance, contracting any edges that were originally multiedges gives a Priority Steiner Tree solution of the same cost. This reduction also works in the opposite direction (in this case there are no multiedges), which shows the equivalence. \square

Furthermore, the $O(\log k)$ upper bound applies to CST (We note that Monotonic CSP admits a trivial algorithm, namely take the subgraph induced by running Dijkstra's Algorithm on G_1).

Lemma 2 *If Monotonic CSN can be approximated to a factor of $f(k)$ for some function f , then Monotonic CST can also be approximated to within $f(k)$.*

Proof We now show a reduction from CST to CSN. Suppose we are given a CST instance on graphs $\mathcal{G} = (G_1, \dots, G_C)$ and terminal sets $\mathcal{X} = (X_1, \dots, X_C)$. Our CSN instance has precisely the same graphs, and has the following demands: for each terminal set X_c , pick any terminal $a \in X_c$ and create a demand (a, b, c) for each $b \neq a \in X_c$. A solution to the original CST instance is a solution to the constructed CSN instance

with the same cost, and vice-versa; moreover, if the CST instance is monotonic, then so is the constructed CSN instance. Observe that if the total number of CST terminals is k , then the number of constructed demands is $k - C$, and therefore an $f(k)$ -approximation for CSN implies an $f(k - C) \leq f(k)$ -approximation for CST, as required. \square

Monotonicity in the directed case

In the directed case, we give an approximation-preserving reduction from a single-source special case of DCSN to the Directed Steiner Tree (DST) problem, then apply a known algorithm for DST. Recall the definition of Single-Source DCSN:

Definition 6 (Single-Source DCSN) *This is the special case of DCSN in which the demands are precisely $(a, b_1, c_1), (a, b_2, c_2), \dots, (a, b_k, c_k)$, for some root $a \in V$. We can assume that $c_1 \leq c_2 \leq \dots \leq c_k$.*

Lemma 3 *Monotonic Single-Source DCSN and Directed Steiner Tree have the same approximability.*

For the remainder of this section, we refer to Monotonic Single-Source DCSN as simply DCSN. Towards proving the theorem, we now describe a reduction from DCSN to DST. Given a DCSN instance $(G_1 = (V, E_1), G_2 = (V, E_2), \dots, G_C = (V, E_C), \mathcal{D})$ with underlying graph $\mathcal{G} = (V, E)$, we construct a DST instance $(G' = (V', E'), D')$ as follows:

- G' contains a vertex v^i for each $v \in V$ and each $i \in [k]$. It contains an edge (u^i, v^i) with weight $w(u, v)$ for each $(u, v) \in E_i$. Additionally, it contains a zero-weight edge (v^i, v^{i+1}) for each $v \in V$ and each $i \in [k]$.
- D' contains a demand $(a^1, b_i^{c_i})$ for each $(a, b_i, c_i) \in \mathcal{D}$.

Now consider the DST instance (G', D') .

Lemma 4 *If the DCSN instance $(G_1, \dots, G_C, \mathcal{D})$ has a solution of cost C^* , then the constructed DST instance (G', D') has a solution of cost at most C^* .*

Proof Let $\mathcal{H} \subseteq \mathcal{G}$ be a DCSN solution having cost C^* . For any edge $(u, v) \in E(\mathcal{H})$, define the *earliest necessary condition* of (u, v) to be the minimum c_i such that removing (u, v) would cause \mathcal{H} not to satisfy demand (a, b_i, c_i) .

Claim 1 *There exists a solution $\mathcal{C} \subseteq \mathcal{H}$ that is a directed tree and has cost at most C^* . Moreover for every path P_i in \mathcal{C} from the root a to some target b_i , as we traverse P_i from a to b_i , the earliest necessary conditions of the edges are non-decreasing.*

Proof of Claim 1 Consider a partition of \mathcal{H} into edge-disjoint sub-graphs $\mathcal{H}_1, \dots, \mathcal{H}_k$, where \mathcal{H}_i is the sub-graph whose edges have earliest necessary condition c_i . Clearly each \mathcal{H}_i is a single component.

If there is a directed cycle or parallel paths in the first sub-graph \mathcal{H}_1 , then there is an edge $e \in E(\mathcal{H}_1)$ whose removal does not cause \mathcal{H}_1 to satisfy fewer demands at condition c_1 . Moreover by monotonicity, removing e also does not cause \mathcal{H} to satisfy fewer demands at any future conditions. Hence there exists a directed tree $\mathcal{T}_1 \subseteq \mathcal{H}_1$ such that $\mathcal{T}_1 \cup \left(\bigcup_{i=2}^k \mathcal{H}_i\right)$ has cost at most C^* and still satisfies \mathcal{T} .

Now suppose by induction that for some $j \in [k-1]$, $\bigcup_{i=1}^j \mathcal{T}_i$ is a tree such that $\left(\bigcup_{i=1}^j \mathcal{T}_i\right) \cup \left(\bigcup_{i=j+1}^k \mathcal{H}_i\right)$ has cost at most C^* and satisfies \mathcal{D} . Consider the partial solution $\left(\bigcup_{i=1}^j \mathcal{T}_i\right) \cup \mathcal{H}_{j+1}$; if this sub-graph is not a directed tree, then there must be an edge $(u, v) \in E(\mathcal{H}_{j+1})$ such that v has another in-edge in the sub-graph. However by monotonicity, (u, v) does not help satisfy any new demands, as v is already reached by some other path from the root. Hence by removing all such redundant edges, we have $\mathcal{T}_{j+1} \subseteq \mathcal{H}_{j+1}$ such that $\left(\bigcup_{i=1}^{j+1} \mathcal{T}_i\right) \cup \left(\bigcup_{i=j+2}^k \mathcal{H}_i\right)$ has cost at most C^* and satisfies \mathcal{D} , which completes the inductive step.

We conclude that $\mathcal{T} := \bigcup_{i=1}^k \mathcal{T}_i \subseteq \mathcal{H}$ is a tree of cost at most C^* satisfying \mathcal{D} . Observe also that by construction, \mathcal{T} has the property that if we traverse any $a \rightarrow b_i$ path, the earliest necessary conditions of the edges never decrease. \square

Now let \mathcal{T} be the DCSN solution guaranteed to exist by Claim 1. Consider the sub-graph $H' \subseteq G'$ formed by adding, for each $(u, v) \in E(\mathcal{T})$, the edge $(u^c, v^c) \in E'$ where c is the earliest necessary condition of (u, v) in $E(\mathcal{H})$. In addition, add all the free edges (v^i, v^{i+1}) . Since $w(u^c, v^c) = w(u, v)$ by construction, $\text{cost}(H') \leq \text{cost}(\mathcal{T}) \leq C^*$.

To see that H' is a valid solution, consider any demand $(a^1, b_i^{c_i})$. Recall that \mathcal{T} has a unique $a \rightarrow b_i$ path P_i along which the earliest necessary conditions are nondecreasing. We added to H' each of these edges at the level corresponding to its earliest necessary condition; moreover, whenever there are adjacent edges $(u, v), (v, x) \in P_i$ with earliest necessary conditions c and $c' \geq c$ respectively, there exist in H' free edges $(v^t, v^{c+1}), \dots, (v^{c'-1}, v^{c'})$. Thus H' contains an $a^1 \rightarrow b_i^{c_i}$ path, which completes the proof. \square

Lemma 5 *If the constructed DST instance (G', D') has a solution of cost C^* , then the original DCSN instance $(G_1, \dots, G_C, \mathcal{D})$ has a solution of cost at most C^* .*

Proof First note that any DST solution ought to be a tree; let $T' \subseteq G'$ be such a solution of cost C . For each $(u, v) \in G$, T' might as well use at most one edge of the form (u^i, v^i) , since if it uses more, it can be improved by using only the one with minimum i , then taking the free edges (v^i, v^{i+1}) as needed. We create a DCSN solution $\mathcal{T} \subseteq \mathcal{G}$ as follows: for each $(u^i, v^i) \in E(T')$, add (u, v) to \mathcal{T} . Since $w(u, v) = w(u^i, v^i)$ by design, we have $\text{cost}(\mathcal{T}) \leq \text{cost}(T') \leq C$. Finally, since each $a^1 \rightarrow b_i^{c_i}$ path in G' has a corresponding path in \mathcal{G} by construction, \mathcal{T} satisfies all the demands. \square

Lemma 3 follows from Lemma 4 and Lemma 5. Finally we can obtain the main result of this subsection:

Theorem 4 *Monotonic Single-Source DCSN has a polynomial-time $O(k^\epsilon)$ -approximation algorithm for every $\epsilon > 0$. It has no $\Omega(\log^{2-\epsilon} n)$ -approximation algorithm unless $\text{NP} \in \text{ZPTIME}(n^{\text{polylog}(n)})$.*

Proof The upper bound follows by composing the reduction (from Monotonic Single-Source DCSN to Directed Steiner Tree) with the algorithm of Charikar et al. [21] for Directed Steiner Tree, which achieves ratio $O(k^\epsilon)$ for every $\epsilon > 0$. More precisely they give a $i^2(i-1)k^{1/i}$ -approximation for any integer $i \geq 1$, in time $O(n^i k^{2i})$. The lower bound follows by composing the reduction (in the opposite direction) with a hardness result of Halperin and Krauthgamer [22], who show the same bound for Directed Steiner Tree. \square

In **Explicit algorithm for Monotonic Single-Source DCSN**, we show how to modify the algorithm of Charikar et al. to arrive at a simple, explicit algorithm for Monotonic Single-Source DCSN achieving the same guarantee.

Application to protein-protein interaction networks

Methods such as Directed Condition Steiner Network can be key in identifying underlying structure in biological processes. As a result, it is important to assess the runtime feasibility of solving for an optimal solution. We show via simulation on human protein-protein interaction networks, that our algorithm on single-source instances is able to quickly and accurately infer maximum likelihood subgraphs for a certain biological process.

Building the protein-protein interaction network

We represent the human PPI network as a weighted directed graph, where proteins serve as nodes, and interactions serve as edges. The network was formed by aggregating information from four sources of interaction

data, including Netpath [30], Phosphosite [31], HPRD [32], and InWeb [33], altogether, covering 16222 nodes and 437888 edges. Edge directions are assigned where these annotations were available (primarily in Phosphosite and NetPath). The remaining edges are represented by two directed edges between the proteins involved. Edge weights were assigned by taking the negative logarithm of the associated confidence score, indicating that finding the optimal Steiner Network would be the same as finding the most confident solution (assuming independence between edges). Confidence data was available for the largest of the data sets (InWeb). For HPRD edges that are not in InWeb, we used the minimum nonzero confidence value by default. For the smaller and highly curated data-sets, Phosphosite and NetPath, we used the maximal confidence level.

Solving DCSN to optimality

Definition 6 (Single-Source DCSN) *This is the special case of DCSN in which the demands are precisely $(a, b_1, c_1), (a, b_2, c_2), \dots, (a, b_k, c_k)$, for some root $a \in V$. We can assume that $c_1 \leq c_2 \leq \dots \leq c_k$.*

We can derive a natural integer linear program for the Single-Source Directed Condition Steiner Network in terms of network flows, with each demand being met by a flow from source to target:

$$\begin{aligned} & \text{minimize} && \sum_{(u,v) \in E} d_{uv} \cdot w(u,v) \\ & \text{subject to} && \\ & && k \cdot d_{uv} \geq d_{uvc} \quad \forall c, (u,v) \in E_c \\ & && \sum_{(v,w) \in E_c} d_{vwc} = \sum_{(u,v) \in E_c} d_{uvc} + \delta_{vc} \quad \forall c, v \in V \\ & && d_{uvc} \in \{0, 1, \dots, k\} \quad \forall c, (u,v) \in E_c \\ & && d_{uv} \in \{0, 1\} \quad \forall (u,v) \in E \end{aligned}$$

Figure 3 Integer Linear Program for Single-Source Condition Steiner Network. δ_{vc} = number of demands at condition c in which v is a source, minus the number in which v is a target.

Each variable d_{uvc} denotes the flow through edge (u, v) at condition c , if it exists; each variable d_{uv} denotes whether (u, v) is ultimately in the chosen solution sub-graph. The first constraint ensures that if an edge is used at any condition, it is chosen as part of the solution. The second constraint enforces flow conservation, and hence that the demands are satisfied, at all nodes and all conditions.

We note that DCSN easily reduces DCSP, as outlined in Theorem 2. However, DCSP is a special case

of Single-Source DCSN. Therefore, the Integer Linear Program defined above can be applied to any DCSN instance with a transformation of the instance to DCSP.

Performance Analysis of Integer Linear Programming

Given the protein protein interaction network G , we sample an instance of the node-variant Single-Source DCSN as so^[3]:

- Instantiate a source node a .
- Independently sample β nodes reachable from a , for each of the γ conditions, giving us $\{b_{1,1}, \dots, b_{\gamma,\beta}\}$.
- For each node $v \in V$, include $v \in V_c$ if v lies on the shortest path from a to one of $\{b_{1,c}, \dots, b_{\gamma,c}\}$
- For all other nodes $v \in V_c$ for all c , include $v \in V_c$ with probability p .

Using a workstation running an Intel Xeon E5-2690 processor and 250GB of RAM, **optimal solutions** to instances of modest size (generated using the procedure just described) were within reach:

Table 2 ILP solve times for some random instances generated by our random model.

β	γ	p	Time to solve
10	1	.25	1m \pm 10s
100	1	.25	1m \pm 10s
10	1	.75	1m \pm 10s
100	1	.75	1m \pm 10s
10	10	.25	9m \pm 30s
10	10	.75	11m \pm 30s
100	10	.25	12m \pm 30s
100	10	.75	17m \pm 2m
10	100	.75	2h 30m \pm 12m
100	100	.75	4h \pm 40m

We notice that our primary runtime constraint comes from γ , the number of conditions. In practice, the number of conditions does not exceed 100. Therefore, we present a model which can also easily translate and find **optimal solutions** to real world biological problems with practical runtime.

Conclusion and Discussion

In this paper we introduced the Condition Steiner Network (CSN) problem and its directed variant, in which the goal is to find a minimal subgraph satisfying a set of k condition-sensitive connectivity demands. We show, in contrast to known results for traditional Steiner problems, that this problem is NP-hard to approximate to a factor of $C - \epsilon$, as well as $k - \epsilon$, for every $C, k \geq 2$ and $\epsilon > 0$. We then explored a special case, in which the conditions/graphs satisfy

^[3]As previously mentioned, this variant reduces to the edge variant via reduction, and vice versa

a *monotonicity* property. For such instances we proposed algorithms significantly beating the pessimistic lower bound for the general problem; this was accomplished by reducing the problem to certain traditional Steiner problems. Lastly, we developed and applied an integer programming-based exact algorithm on simulated instances built over the human protein-protein interaction network, and reported feasible runtimes for real-world problem instances.

Importantly, along the way we showed implications of these results for CSN on other network connectivity problems that are commonly used in PPI analysis—such as Shortest Path, Steiner Tree, Prize-Collecting Steiner Tree—when conditions are added. We showed that for each of these problems, we cannot guarantee (in polynomial time) a solution with a value below $C - \epsilon$ times the optimal value. These lower bounds are quite strict, in the sense that naively approximating the problem separately in every condition, and taking the union of those solutions, already gives an approximation ratio of $O(C)$. At the same time, by relating the various condition Steiner problems to one another, we also obtained some positive results: the condition versions of Shortest Path and Steiner Tree admit good approximations when the conditions are monotonic. Moreover, all of the condition problems (with the exception of Prize-Collecting Steiner Tree) can be solved using a natural integer programming framework that works well in practice.

Proofs of Main Theorems

Problem variants

There are several natural ways to formulate the condition Steiner network problem, depending on whether the edges are changing over condition, or the nodes, or both.

Definition 10 (Condition Steiner Network (edge variant)) *This is the formulation described in the Introduction: the inputs are $G_1 = (V, E_1), \dots, G_C = (V, E_C)$, $w(\cdot)$, and $\mathcal{D} = \{(a_i, b_i, c_i)\}$. The task is to find a minimum-weight sub-graph $\mathcal{H} \subseteq \mathcal{G}$ that satisfies all of the demands.*

Definition 11 (Condition Steiner Network (node variant)) *Let the underlying graph be $\mathcal{G} = (V, E)$. The inputs are $G_1 = (V_1, E(V_1)), \dots, G_C = (V_C, E(V_C))$, $w(\cdot)$, and \mathcal{D} . Here, $E(V_c) \subseteq E$ denotes the edges induced by $V_c \subseteq V$. A path satisfies a demand at condition t iff all edges along that path exist in G_c .*

Definition 12 (Condition Steiner Network (node and edge variant)) *The inputs are precisely $G_1 = (V_1, E_1), \dots, G_C = (V_C, E_C)$, $w(\cdot)$, and \mathcal{D} . This is the same as the node variant except that each E_c can be any subset of $E(V_c)$.*

Similarly, define the corresponding directed problem Directed Condition Steiner Network (DCSN) with the same three variants. The only difference is that the edges are directed, and a demand (a, b, c) must be satisfied by a directed $a \rightarrow b$ path in G_c .

The following observation enables all our results to apply to all problem variants.

Proposition 2 *The edge, node, and node-and-edge variants of CSN are mutually polynomial-time reducible via strict reductions (i.e. preserving the approximation ratio exactly). Similarly all three variants of DCSN are mutually strictly reducible.*

Proof The following statements shall hold for both undirected and directed versions. Clearly the node-and-edge variant generalizes the other two. It suffices to show two more directions:

- (Node-and-edge reduces to node) Let (u, v) be an edge existent at a set of conditions $\tau(u, v)$, whose endpoints exist at conditions $\tau(u)$ and $\tau(v)$. To make this a node-condition instance, create an intermediate node $x_{(u,v)}$ existent at conditions $\tau(u, v)$, an edge $(u, x_{(u,v)})$ with the original weight $w(u, v)$, and an edge $(x_{(u,v)}, v)$ with zero weight. A solution of cost W in the node-and-edge instance corresponds to a node-condition solution of cost W , and vice-versa.
- (Node reduces to edge) Let (u, v) be an edge whose endpoints exist at conditions $\tau(u)$ and $\tau(v)$. To make this an edge-condition instance, let (u, v) exist at conditions $\tau(u, v) := \tau(u) \cap \tau(v)$. Let every node exist at all conditions; let the edges retain their original weights. A solution of cost W in the node-condition instance corresponds to an edge-condition solution of cost W , and vice-versa. \square

Inapproximability for general C and k

Here we prove our main theorem, showing optimal hardness for any number of demands. To do this, we introduce a generalization of Label Cover to partite hypergraphs:

Definition 13 (k -Partite Hypergraph Label Cover (k-PHLC)) *An instance of this problem consists of a k -partite, k -regular hypergraph $G = (V_1, \dots, V_k, E)$ (that is, each edge contains exactly one vertex from each of the k parts) and a set of possible labels Σ . The input also includes, for each hyperedge $e \in E$, a projection function $\pi_e^v : \Sigma \rightarrow C$ for each $v \in e$; Π is the set of all such functions. A labeling of G is a function $\phi : \bigcup_{i=1}^k V_i \rightarrow \Sigma$ assigning each node a label. There are two notions of edge satisfaction under a labeling ϕ :*

- ϕ strongly satisfies a hyperedge $e = (v_1, \dots, v_k)$ iff the labels of all its vertices are mapped to the same color, i.e. $\pi_{v_i}^e(\phi(v_i)) = \pi_{v_j}^e(\phi(v_j))$ for all $i, j \in [k]$.
- ϕ weakly satisfies a hyperedge $e = (v_1, \dots, v_k)$ iff there exists some pair of vertices v_i, v_j whose labels are mapped to the same color, i.e. $\pi_{v_i}^e(\phi(v_i)) = \pi_{v_j}^e(\phi(v_j))$ for some $i \neq j \in [k]$.

Theorem 7 For every $\epsilon > 0$ and every fixed integer $k \geq 2$, there is a constant $|\Sigma|$ such that the following promise problem is NP-hard: Given a k -Partite Hypergraph Label Cover instance (G, Σ, Π) , distinguish between the following cases:

- (YES instance) There exists a labeling of G that strongly satisfies every edge.
- (NO instance) Every labeling of G weakly satisfies at most $\epsilon|E|$ edges.

Theorem 7 follows from Feige's k -prover system [34] by taking the number of repetitions to be (a constant depending on k and ϵ) large enough so that the error probability drops below ϵ .

The proof of $(C - \epsilon)$ -hardness and $(k - \epsilon)$ -hardness follows the same outline as the $C = k = 2$ case (Theorem 6).

Theorem 8 (Main Theorem) CSN and DCSN are NP-hard to approximate to a factor of $C - \epsilon$ as well as $k - \epsilon$ for every fixed $k \geq 2$ and every constant $\epsilon > 0$. For DCSN, this holds even when the underlying graph is acyclic.

Proof Given the k -PHLC instance in the form $(G = (V_1, \dots, V_k, E), \Sigma, \Pi)$, and letting $v_{c,i}$ denote the i -th node in V_c , construct a DCSN instance $(\mathcal{G} = (G_1, \dots, G_k)$, along with k demands) as follows. For every $c \in [k]$, create nodes $v_{c,1}^S, \dots, v_{t,|V_c|+1}^S$. Create a $v_{c,i}$ -bundle from each $v_{c,i}^S$ to $v_{c,i+1}^S$, whose ℓ -strands (for $\ell \in \Sigma$) are each a chain of bundles, one for each incident hyperedge $e = (v_{1,i_1}, \dots, v_{c,i}, \dots, v_{k,i_k}) \in E$. Each $(v_{1,i_1}, \dots, v_{c,i}, \dots, v_{k,i_k})$ -bundle has a $(v_{1,i_1}, \ell_1, \dots, v_{c,i}, \ell_c, \dots, v_{k,i_k}, \ell_k)$ -path for each agreeing combination of labels—that is, every k -tuple $(\ell_1, \dots, \ell_c, \dots, \ell_k)$ such that: $\pi_{v_{1,i_1}}^e(\ell_1) = \dots = \pi_{v_{c,i}}^e(\ell_c) = \dots = \pi_{v_{k,i_k}}^e(\ell_k)$, where e is the shared edge. If there are no such combinations, then the e -bundle is a single simple strand.

For $c \in [k]$, set all the edges in the $v_{c,i}$ -bundles to exist in G_c only. Now, for each $(v_{1,i_1}, \ell_1, \dots, v_{k,i_k}, \ell_k)$, merge together the $(v_{1,i_1}, \ell_1, \dots, v_{k,i_k}, \ell_k)$ -paths across all G_c that have such a strand. Finally, the connectivity demands are $\mathcal{D} = \left\{ (v_{c,1}^S, v_{c,|V_c|+1}^S, c) : c \in [k] \right\}$.

The analysis follows the $k = 2$ case. Suppose we have a YES instance of k -PHLC, with optimal labeling ℓ_v^* to each node $v \in \bigcup_{t=1}^k V_c$. Then an optimal solution \mathcal{H}^* to the constructed DCSN instance is to traverse, at each condition c and for each $v_{c,i}$ -bundle, the path through the $\ell_{v_{c,i}}^*$ -strand. In particular for each $(v_{1,i_1}, \dots, v_{k,i_k})$ -bundle in that strand, traverse the $(v_{1,i_1}, \ell_1^*, \dots, v_{k,i_k}, \ell_k^*)$ -path.

In tallying the total edge cost, $\mathcal{H}^* \cap G_1$ (the sub-graph at condition 1) incurs a cost of $|E|$, one for each contact edge. The sub-graphs of \mathcal{H}^* at conditions $2, \dots, k$ account for no additional cost, since all contact edges correspond to a label which agrees with all its neighbors' labels, and hence were merged with the agreeing contact edges in the other sub-graphs.

Conversely suppose we have a NO instance of k -PHLC, so that for any labeling ℓ_v^* , for at least $(1 - \epsilon)|E|$ hyperedges e , the projection functions of all nodes in e disagree. By definition, any solution to the constructed DCSN instance contains a simple $v_{t,1}^S \rightarrow v_{t,|V_c|+1}^S$ path P_c at each condition c . As before, P_1 alone incurs a cost of exactly $|E|$. However, at least $(1 - \epsilon)|E|$ of the hyperedges in G cannot be weakly satisfied; for these hyperedges e , for every pair of neighbors $v_{c,i_c}, v_{c',i_{c'}} \in e$, there is no path through the e -bundle in v_{t,i_c} 's $\ell_{v_{c,i_c}}^*$ -strand that is merged with any of the paths through the e -bundle in $v_{c',i_{c'}}$'s $\ell_{v_{c',i_{c'}}}^*$ -strand (for otherwise, it would indicate a labeling that weakly satisfies e in the k -PHLC instance). Therefore paths P_2, \dots, P_k each contribute at least $(1 - \epsilon)|E|$ additional cost, so the solution has total cost at least $(1 - \epsilon)|E| \cdot k$.

It follows from the gap between the YES and NO cases that DCSN is NP-hard to approximate to within a factor of $k - \epsilon$ for every constant $\epsilon > 0$; and since $C = k$ in our construction, it is also NP-hard for $C - \epsilon$. Moreover since The directed condition graph we constructed is acyclic, this result holds even on DAGs. As before, the same analysis holds for the undirected problem CSN by undirecting the edges. \square

Explicit algorithm for Monotonic Single-Source DCSN

We provide a modified version of the approximation algorithm presented in Charikar et al. [21] for Directed Steiner Tree (DST), which achieves the same approximation ratio for our problem Monotonic Single-Source DCSN.

We provide a similar explanation as of that presented in Charikar et al. Consider a trivial approximation algorithm, where we take the shortest path from the

source to each individual target. Consider the example where there are edges of cost $C - \epsilon$ to each target, and a vertex v with distance C from the source, and with distance 0 to each target. In such a case, this trivial approximation algorithm will achieve only an $\Omega(k)$ -approximation. Consider instead an algorithm which found, from the root, an intermediary vertex v , which was connected to all the targets via shortest path. In the case of the above example, this would find us the optimal sub-graph. The algorithm below generalizes this process, by progressively finding optimal substructures with good cost relative to the number of targets connected. We show that this algorithm provides a good approximation ratio.

Definition 14 (Metric closure of a condition graph) *For a directed condition graph $\mathcal{G} = (G_1 = (V, E_1), G_2 = (V, E_2), \dots, G_C = (V, E_C))$, define its metric closure to be $G = (V, E, \tilde{w})$ where $E = \bigcup_c E_c$ and $\tilde{w}(u, v, c)$ is the length of the shortest $u \rightarrow v$ path in G_c (note that in contrast with w , \tilde{w} takes three arguments).*

Definition 15 ($V(T)$) *Let T be a tree with root r . We say a demand of the form (r, b, c) is satisfied by T if there is a path in T from r to b at condition c . $V(T)$ is then the set of demands satisfied by T .*

Definition 16 ($D(T)$) *The density of a tree T is $D(T) = \frac{\text{cost}(T)}{|V(T)|}$, where $\text{cost}(T)$ is the sum of edge weights of T .*

```

1: function  $A_i(\text{transitive closure } G = (V, E, w), r, c, k, \mathcal{D} \subseteq V \times [C])$ 
2:   if  $(r, b_i, c_i)$  does not exist for least  $k(b_i, c_i) \in \mathcal{D}, c_i \geq c$ 
   then return NO SOLUTION
3:    $T \leftarrow \emptyset$ 
4:   while  $k > 0$  do
5:      $T_{\text{best}} \leftarrow \emptyset$ 
6:     for all  $(v, t') \in V \times [C], c' \geq c$  and  $k', 1 \geq k' \geq k$  do
7:        $T' \leftarrow A_{i-1}(G, v, c', k', \mathcal{D}) \cup (r, v, c')$ 
8:       if  $d(T_{\text{BEST}}) > d(T')$  then  $T_{\text{BEST}} \leftarrow T'$ 
       Demand  $i$  satisfied only if edge to  $b_i$  at  $c_i$  (ie  $(x, b_i, c_i)$  for some  $x$ )
9:    $T \leftarrow T \cup T_{\text{BEST}}; k \leftarrow k - |D \cap V(T_{\text{BEST}})|; X \leftarrow X - V(T_{\text{BEST}})$ 
10:  return  $T$ 

```

The way we will prove the approximation ratio of this algorithm is to show that it behaves precisely as the algorithm of Charikar et al. does, when given as input the DST instance produced by our reduction from Monotonic Single Source DCSN (Lemma 3).

Proposition 3 *The algorithm above is equivalent to the algorithm of Charikar et al., when applied to the DST instance output by the reduction of Lemma 3.*

Proof To see this, note that in our reduced instance, we see a collection of vertices, $v^1, \dots, v^{|C|}$. Therefore, the only equivalent modifications needed to the original algorithm are:

- In the input, rather than keeping track of the current root as some vertex v^i , keep track of v at the current condition instead, i.e. (v, i) .
- The distance from some v^i to $x^j, j \geq i$ is simply the distance from v to x at condition j , i.e. $\tilde{w}(v, x, j)$.
- Instead of looping through all vertices in the form $v^1, \dots, v^{|C|}$, we instead loop through all vertices, and all conditions.

Therefore this algorithm guarantees the same approximation ratio for Monotonic Single Source DCSN as the original algorithm achieved for DST. In particular for all $i > 1$, $A_i(G, a, 0, k, D)$ provides an $i^2(i-1)k^{1/i}$ approximation to DCSN, in time $O(n^i k^{2i})$ [21, 35]^[4]. \square

List of Abbreviations

- CPCST: Condition Prize-Collecting Steiner Tree
- CSN: Condition Steiner Network
- CST: Condition Steiner Tree
- CSP: Condition Shortest Path
- DSN: Directed Steiner Network
- DST: Directed Steiner Tree
- DCSN: Directed Condition Steiner Network
- DCSP: Directed Condition Shortest Path
- k -PHLC: k -Partite Hypergraph Label Cover
- MKL: Minimum k -Labeling
- PPI: Protein Protein Interaction

Declarations

Ethics approval and consent to participate

Not applicable

Consent for publication

Not applicable

Availability of data and material

Our simulation code can be found at the following URL:

https://github.com/YosefLab/condition_connectivity_problems

Author's contributions

All authors conceived and designed the study. JW and BW derived the main hardness results. AK derived the monotonic hardness results and approximation algorithm. NY was the PI and oversaw the project.

Competing interests

The authors declare that they have no competing interests.

Funding and Acknowledgements

This work was partially supported by the National Science Foundation Graduate Research Fellowship Program award DGE 1106400, NIH grants U01HG007910 and U01MH105979, and the U.S.-Israel Binational Science Foundation.

Author details

¹Department of Computer Science, Stanford University, Stanford, CA USA.

²Department of Electrical Engineering and Computer Science, UC Berkeley, Berkeley, CA USA.

^[4]The first paper [21] incorrectly claims a bound of $i(i-1)k^{1/i}$; this was corrected in [35].

References

- Chatr-aryamontri, A., Breitkreutz, B.-J., Oughtred, R., Boucher, L., Heinicke, S., Chen, D., Stark, C., Breitkreutz, A., Kolas, N., O'Donnell, L., Reguly, T., Nixon, J., Ramage, L., Winter, A., Sellam, A., Chang, C., Hirschman, J., Theesfeld, C., Rust, J., Livstone, M.S., Dolinski, K., Tyers, M.: The biogrid interaction database: 2015 update. *Nucleic Acids Research* **43**(Database issue), 470–478 (2015). doi:[10.1093/nar/gku1204](https://doi.org/10.1093/nar/gku1204)
- Ben-Shitrit, T., Yosef, N., Shemesh, K., Sharan, R., Ruppin, E., Kupiec, M.: Systematic identification of gene annotation errors in the widely used yeast mutation collections. *Nat Meth* **9**(4), 373–378 (2012)
- Huang, S.-s.C., Fraenkel, E.: Integration of proteomic, transcriptional, and interactome data reveals hidden signaling components. *Science signaling* **2**(81), 40 (2009)
- Scott, M.S., Perkins, T., Bunnell, S., Pepin, F., Thomas, D.Y., Hallett, M.: Identifying regulatory subnetworks for a set of genes. *Molecular & Cellular Proteomics* **4**(5), 683–692 (2005)
- Wu, C., Yosef, N., Thalhamer, T., Zhu, C., Xiao, S., Kishi, Y., Regev, A., Kuchroo, V.K.: Induction of pathogenic th17 cells by inducible salt-sensing kinase sgk1. *Nature* **496**(7446), 513–517 (2013)
- Yosef, N., Ungar, L., Zalckvar, E., Kimchi, A., Kupiec, M., Ruppin, E., Sharan, R.: Toward accurate reconstruction of functional protein networks. *Molecular systems biology* **5**(1), 248 (2009)
- Huang, S.-s.C., Fraenkel, E.: Integrating proteomic, transcriptional, and interactome data reveals hidden components of signaling and regulatory networks. *Science Signaling* **2**(81), 40–40 (2009). doi:[10.1126/scisignal.2000350](https://doi.org/10.1126/scisignal.2000350). <http://stke.sciencemag.org/content/2/81/ra40.full.pdf>
- Tuncbag, N., Gosline, S.J.C., Kedaigle, A., Soltis, A.R., Gitter, A., Fraenkel, E.: Network-based interpretation of diverse high-throughput datasets through the omics integrator software package. *PLOS Computational Biology* **12**(4), 1–18 (2016). doi:[10.1371/journal.pcbi.1004879](https://doi.org/10.1371/journal.pcbi.1004879)
- Przytycka, T.M., Singh, M., Slonim, D.K.: Toward the dynamic interactome: it's about time. *Briefings in Bioinformatics* **11**(1), 15–29 (2010). doi:[10.1093/bib/bbp057](https://doi.org/10.1093/bib/bbp057)
- Mertins, P., Przybylski, D., Yosef, N., Qiao, J., Clauser, K., Raychowdhury, R., Eisenhaure, T.M., Maritzen, T., Haucke, V., Satoh, T., Akira, S., Carr, S.A., Regev, A., Hacohen, N., Chevrier, N.: An integrative framework reveals signaling-to-transcription events in toll-like receptor signaling. *Cell Reports* **19**(13), 2853–2866 (XXXX). doi:[10.1016/j.celrep.2017.06.016](https://doi.org/10.1016/j.celrep.2017.06.016)
- Kanshin, E., Bergeron-Sandoval, L.-P., Isik, S.S., Thibault, P., Michnick, S.W.: A cell-signaling network temporally resolves specific versus promiscuous phosphorylation. *Cell Reports* **10**(7), 1202–1214 (2015). doi:[10.1016/j.celrep.2015.01.052](https://doi.org/10.1016/j.celrep.2015.01.052)
- Bendall, S.C., Simonds, E.F., Qiu, P., Amir, E.-a.D., Krutzik, P.O., Finck, R., Bruggner, R.V., Melamed, R., Trejo, A., Ornatsky, O.I., Balderas, R.S., Plevritis, S.K., Sachs, K., Pe'er, D., Tanner, S.D., Nolan, G.P.: Single-cell mass cytometry of differential immune and drug responses across a human hematopoietic continuum. *Science* **332**(6030), 687–696 (2011). doi:[10.1126/science.1198704](https://doi.org/10.1126/science.1198704). 21551058[pmid]
- Mazza, A., Gat-Viks, I., Farhan, H., Sharan, R.: A minimum-labeling approach for reconstructing protein networks across multiple conditions. *Algorithms for Molecular Biology* **9**(1), 1 (2014)
- Byrka, J., Grandoni, F., Rothvoß, T., Sanità, L.: An improved lp-based approximation for steiner tree. In: *Proceedings of the Forty-second ACM Symposium on Theory of Computing*, pp. 583–592 (2010). ACM
- Agrawal, A., Klein, P., Ravi, R.: When trees collide: An approximation algorithm for the generalized steiner problem on networks. *SIAM Journal on Computing* **24**(3), 440–456 (1995)
- Chlebík, M., Chlebíková, J.: The steiner tree problem on graphs: Inapproximability results. *Theoretical Computer Science* **406**(3), 207–214 (2008)
- Feldman, J., Ruhl, M.: The directed steiner network problem is tractable for a constant number of terminals. In: *Foundations of Computer Science, 1999. 40th Annual Symposium On*, pp. 299–308 (1999). IEEE
- Feldman, M., Kortsarz, G., Nutov, Z.: Improved approximating algorithms for directed steiner forest. In: *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 922–931 (2009). Society for Industrial and Applied Mathematics
- Chekuri, C., Even, G., Gupta, A., Segev, D.: Set connectivity problems in undirected graphs and the directed steiner network problem. *ACM Transactions on Algorithms (TALG)* **7**(2), 18 (2011)
- Dodis, Y., Khanna, S.: Designing networks with bounded pairwise distance. In: *Proceedings of the Thirty-first Annual ACM Symposium on Theory of Computing*, pp. 750–759 (1999). ACM
- Charikar, M., Chekuri, C., Cheung, T.-y., Dai, Z., Goel, A., Guha, S., Li, M.: Approximation algorithms for directed steiner problems. *Journal of Algorithms* **33**(1), 73–91 (1999)
- Halperin, E., Krauthgamer, R.: Polylogarithmic inapproximability. In: *Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing*, pp. 585–594 (2003). ACM
- Archer, A., Bateni, M., Hajiaghayi, M., Karloff, H.: Improved approximation algorithms for prize-collecting steiner tree and tsp. *SIAM journal on computing* **40**(2), 309–332 (2011)
- Dinur, I., Manurangsi, P.: Eth-hardness of approximating 2-csp and directed steiner network. In: *LIPICs-Leibniz International Proceedings in Informatics*, vol. 94 (2018). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik
- Arora, S., Babai, L., Stern, J., Sweedy, Z.: The hardness of approximate optima in lattices, codes, and systems of linear equations. In: *Foundations of Computer Science, 1993. Proceedings., 34th Annual Symposium On*, pp. 724–733 (1993). IEEE
- Arora, S., Lund, C., Motwani, R., Sudan, M., Szegedy, M.: Proof verification and the hardness of approximation problems. *Journal of the ACM (JACM)* **45**(3), 501–555 (1998)
- Raz, R.: A parallel repetition theorem. *SIAM Journal on Computing* **27**(3), 763–803 (1998)
- Charikar, M., Naor, J., Schieber, B.: Resource optimization in qos multicast routing of real-time multimedia. *IEEE/ACM Transactions on Networking* **12**(2), 340–348 (2004)
- Chuzhoy, J., Gupta, A., Naor, J.S., Sinha, A.: On the approximability of some network design problems. *ACM Transactions on Algorithms (TALG)* **4**(2), 23 (2008)
- Kandasamy, K., Mohan, S.S., Raju, R., Keerthikumar, S., Kumar, G.S.S., Venugopal, A.K., Telikicherla, D., Navarro, J.D., Mathivanan, S., Pecquet, C., Gollapudi, S.K., Tattikota, S.G., Mohan, S., Padhukasahasram, H., Subbannayya, Y., Goel, R., Jacob, H.K., Zhong, J., Sekhar, R., Nanjappa, V., Balakrishnan, L., Subbaiah, R., Ramachandra, Y., Rahiman, B.A., Prasad, T.K., Lin, J.-X., Houtman, J.C., Desiderio, S., Renaud, J.-C., Constantinescu, S.N., Ohara, O., Hirano, T., Kubo, M., Singh, S., Khatri, P., Draghici, S., Bader, G.D., Sander, C., Leonard, W.J., Pandey, A.: Netpath: a public resource of curated signal transduction pathways. *Genome Biology* **11**(1), 3 (2010). doi:[10.1186/gb-2010-11-1-r3](https://doi.org/10.1186/gb-2010-11-1-r3)
- Hornbeck, P.V., Zhang, B., Murray, B., Kornhauser, J.M., Latham, V., Skrzypek, E.: Phosphositeplus, 2014: mutations, ptms and recalibrations. *Nucleic Acids Res* **43**(Database issue), 512–520 (2015). doi:[10.1093/nar/gku1267](https://doi.org/10.1093/nar/gku1267)
- Keshava Prasad, T.S., Goel, R., Kandasamy, K., Keerthikumar, S., Kumar, S., Mathivanan, S., Telikicherla, D., Raju, R., Shafreen, B., Venugopal, A., Balakrishnan, L., Marimuthu, A., Banerjee, S., Somanathan, D.S., Sebastian, A., Rani, S., Ray, S., Harrys Kishore, C.J., Kanth, S., Ahmed, M., Kashyap, M.K., Mohmood, R., Ramachandra, Y.L., Krishna, V., Rahiman, B.A., Mohan, S., Ranganathan, P., Ramabadran, S., Chaerkady, R., Pandey, A.: Human protein reference database—2009 update. *Nucleic Acids Research* **37**(suppl.1), 767–772 (2009). doi:[10.1093/nar/gkn892](https://doi.org/10.1093/nar/gkn892)
- Li, T., Wernersson, R., B Hansen, R., Horn, H., Mercer, J., Slodkiewicz, G., T Workman, C., Rigina, O., Rapacki, K., Stærfeldt, H., Brunak, S., S Jensen, T., Lage, K.: A scored human protein–protein interaction network to catalyze genomic interpretation **14** (2016)
- Feige, U.: A threshold of $\ln n$ for approximating set cover. *Journal of the ACM (JACM)* **45**(4), 634–652 (1998)
- Helvig, C.S., Robins, G., Zelikovsky, A.: An improved approximation scheme for the group steiner problem. *Networks* **37**(1), 8–20 (2001)