Pseudo-Separation for Assessment of Structural Vulnerability of a Network

Alan Kuhnle, Tianyi Pan, Victoria G. Crawford, Md Abdul Alim, My T. Thai

Department of Computer and Information Science at the University of Florida

Introduction

Many studies of network vulnerability, or the degree to which the functionality of a network may be disrupted by failures, have incorporated connectivity as a fundamental measure of network functionality. Whatever functionality a network may provide to a pair of nodes is usually absent if the pair is disconnected. In this work, we generalize the idea of network impairment resulting from disconnection of network elements to impairment upon sufficient separation in the network according to a given metric; that is, if $d(s,t) > T$ for two vertices s, t, distance metric d, and threshold T. **The full version of this paper is available at `https://arxiv.org/abs/1704.04555`**.

Our Contributions

Based upon this idea of functionality impairment under sufficient separation, we:

- Introduce novel T-separation analogues to the the min-cut and multi-cut problem, T-PCUT and T-MULTI-PCUT, respectively, and analyze their computational complexity.
- Present GEN, an $O(\log n)$-approximation algorithm, FEN, a $(T+1)$-approximation algorithm, and GEST, an efficient, randomized algorithm with probabilistic performance guarantee.
- Provide an experimental evaluation of these algorithms.

Computational Complexity

- **Proposition 1** *For $T \leq 3$, T-PCUT with uniform lengths and costs is solvable in polynomial time.*
- **Proposition 2** *Let D be a constant, T-PCUT$(G,(s,t))$ be an instance of T-PCUT for some constant T with uniform lengths and uniform costs. If the maximum degree δ in G satisfies $\delta \leq D$, then the optimal solution W is computable in polynomial time.*
- **Theorem 3** *Consider the decision version of 1-PCUT with uniform costs and arbitrary lengths; that is, given problem instance 1-PCUT$(G,(s,t))$ with uniform costs and arbitrary lengths, and given constant $D > 0$, determine if a solution $W \subset V$ exists with $|W| \leq D$. This problem is NP-complete.*
- **Theorem 4** *Unless $P = NP$, there is no polynomial-time approximation to uniform length, cost T-MULTI-PCUT within a factor of 1.3606, for $T \geq 1$.*

Approximation Algorithms

We present three approximation algorithms for arbitrary vertex cost T-MULTI-PCUT, when the length function on edges is bounded below: $d(e) > q_{min}$ for some constant $q_{min} > 0$. Let $T_0 = T/q_{min}$. Then, all paths in $\mathcal{P}$ can be enumerated in $O(n^{T_0})$. Furthermore, an optimal solution S of vertices must intersect every path in $\mathcal{P}$. Thus, T-MULTI-PCUT may be considered as a special case of the set covering problem, where the paths correspond to the elements to be covered and vertices correspond to sets.

- **The greedy algorithm (GEN)**, choosing at each iteration the vertex covering the most paths in $\mathcal{P}$.
- **The frequency rounding algorithm (FEN)** based upon the optimal solution to the linear program corresponding to the set cover instance.
- **The probabilistic approximation algorithm, GEST**, which is similar to GEN except that GEST estimates which node intersects the most paths by sampling paths.

We have the following results for the approximation algorithms:

- **Theorem 5** *GEN achieves a performance guarantee of $O(\log n)$ with respect to the optimal solution with running time bounded by $O(kn^{T_0})$. Furthermore, for each n, there exists an instance of the single pair PCUT problem where GEN returns a solution of cost greater than a factor $\Omega(\log n)$ of the optimal.*
- **Theorem 6** *FEN achieves a performance guarantee of $T_0 + 1$ with respect to the optimal solution.*
- **Theorem 7** *Given an instance $(G, c, d, \mathcal{S})$ of uniform vertex cost T-MULTI-PCUT whose length function d is bounded below, let δ be the maximum degree in G, and let $\alpha \in (0,1)$. With probability at least $1 - 1/n$, GEST returns a feasible solution W with cost within ratio $O\left(\alpha\delta^{T_0} + \log|\mathcal{S}|\right)$ of optimal. The running time of GEST is $O(k^3 n \log(2n^2)/2\alpha^2)$.*

Problem Definitions

Let T be an arbitrary but fixed constant. The problems will take as input a triple (G, c, d), where G is a directed graph $G = (V, E)$; $c : V \to \mathbf{R}^+$ is a cost function on vertices representing the difficulty of removing each node; and $d : E \to \mathbf{R}^+$ is a length function on edges.

- **Minimum T-pseudocut (T-PCUT)** *Given triple (G, c, d) and a pair (s, t) of vertices of G, determine a minimum cost set $W \subset V \backslash \{s, t\}$ of vertices such that $d(s, t) > T$ after the removal of W from G.*
- **Minimum T-multi-pseudocut (T-MULTI-PCUT)** *Given triple (G, c, d), and a target set of pairs of vertices of G, $\mathcal{S} = \{(s_1, t_1), (s_2, t_2), \ldots, (s_k, t_k)\}$, determine a minimum cost set W of vertices such that $d(s_i, t_i) > T$ for all i after the removal of W from G.*

In the above two formulations, we emphasize that the threshold T is a fixed constant independent of the input.

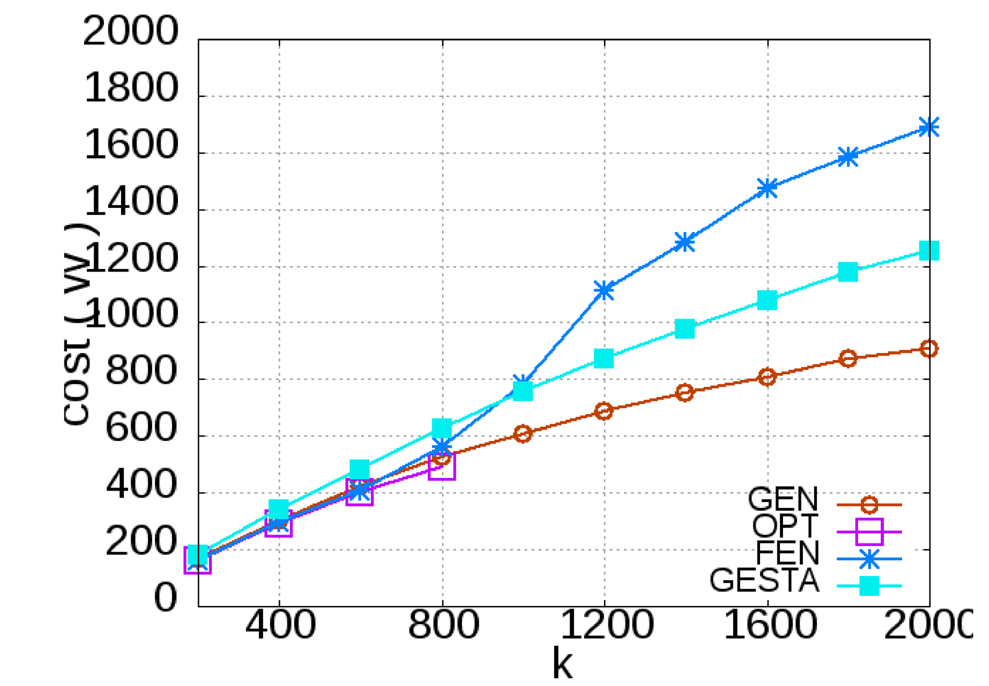

Figure 1: Performance vs. number of pairs on a synthesized internet topology, with $T = 6$. Edge weights are uniformly distributed on $[1, 10]$ to simulate an additive Quality-of-Service metric.

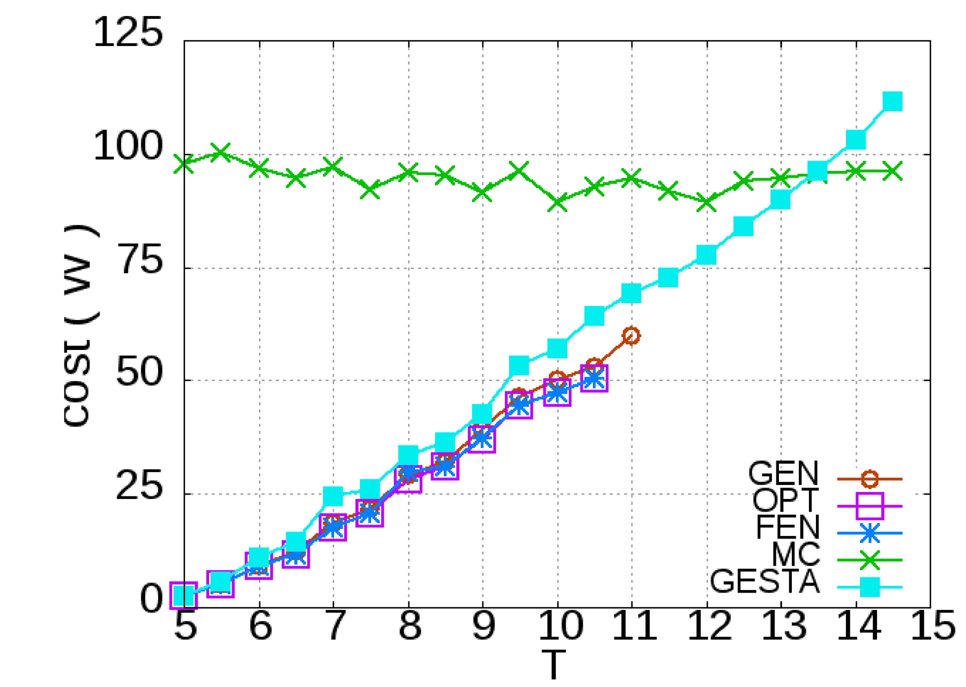

Figure 2: Performance vs. threshold T, for $k = 1$. MC is mininum-cut algorithm. All results are averaged over 10 choices of $\mathcal{S}$. This experiment is on an Erdos-Renyi random graph with $|V| = 1000, |E| = 49995$.

Acknowledgements

This work is supported in part by NSF grant CNS-1443905 and DTRA grant HDTRA1-14-1-0055.

Contact Information

- Learn more about the Optima Network Science group at the University of Florida by visiting www.cise.ufl.edu/research/OptimaNetSci/
- For further information, contact Alan Kuhnle at kuhnle@cise.ufl.edu