

Optimal Determination of Source-destination Connectivity in Random Graphs

Luoyi Fu, Xinbing Wang
Dept. of Electronic Engineering
Shanghai Jiao Tong University
Shanghai, China
{yiluofu,xwang8}@sjtu.edu.cn

P. R. Kumar
Dept. of Electrical and Computer Engineering
Texas A&M University
College Station, Texas
prk@tamu.edu

ABSTRACT

This paper investigates the problem of optimally determining source-destination connectivity in random graphs. We consider the classic Erdos-Renyi (ER) random graph with n nodes, where an edge independently exists between any two nodes with probability p . The problem examined is that of determining whether a given pair of nodes, a source S and a destination D , are connected by a path. Assuming that at each step one edge can be tested to see if it exists or not, we determine an optimal policy that minimizes the total expected number of steps.

The optimal policy has several interesting features. In order to establish connectivity of S and D , a policy needs to check all edges on some path to see if they all exist, but to establish disconnectivity it has to check all edges on some cut to see if none of them exists. The optimal policy has the following form. At each step it examines the condensation multigraph formed by contracting each known connected component to a single node, and then checks an edge that is simultaneously on a shortest S - D path as well as in a minimum S - D cut. Among such edges, it chooses that which leads to the most opportunities for connection. Interestingly, the optimal strategy does not depend on p or n , even though the entire graph itself undergoes a sharp transition from disconnectivity to connectivity around $p = \ln n/n$. The policy is efficiently implementable, requiring no more than $30 \log_2 n$ operations to determine which edge to test next.

The result also extends to some more general graphs.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: [Computer-Communication Networks]

General Terms

Connectivity, Random Networks, Random Graphs

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
MobiHoc '14, August 11-14, 2014, Philadelphia, PA, USA.
Copyright 2014 ACM 978-1-4503-2620-9/14/08...\$15.00.
<http://dx.doi.org/10.1145/2632951.2632990>.

Keywords

Connectivity; Random graphs; Optimization

1. INTRODUCTION

Connectivity of random graphs has long been a topic of intensive study. A widely studied model, known as the $G(n, p)$ model, was proposed by Gilbert [1] in 1959. It refers to a graph with n nodes, with an edge existing independently between any pair of nodes with probability p . Erdos and Renyi [2,3] showed that this graph is connected with probability approaching one as $n \rightarrow \infty$, if $p(n)$, the probability p as a function of n , satisfies $p(n) > \frac{(1+\epsilon) \ln n}{n}$, and contains isolated nodes with probability approaching one as $n \rightarrow \infty$, if $p(n) < \frac{(1-\epsilon) \ln n}{n}$. We will refer to the $G(n, p)$ model as the ER graph, in conformity with common usage. In a wireless networking context, Gilbert [4], and, more recently, Penrose [5] and Gupta and Kumar [6], have studied geometric random graphs where nodes have random uniform i.i.d. locations in a unit disk, and showed that if the radio transmission range of each node is $r(n) = \sqrt{\frac{\ln n + c(n)}{n\pi^2}}$, then the whole network is connected with probability approaching one as $n \rightarrow \infty$, if and only if $c(n) \rightarrow +\infty$. Ever since, there has been much research effort directed towards studying asymptotic connectivity of randomly distributed wireless networks [7–15].

In this paper we consider ER Graphs. The above works all focus on the connectivity of the entire network in the asymptotic regime where the number of nodes goes to infinity. The focus here is different in the following three ways: (i) It is non-asymptotic, (ii) it is specific to the specific realization of the random network that is being studied, and, (iii) instead of studying the whole network's connectivity, the issue is the connectivity between a specified source S and destination D . We examine the fundamental problem of designing a policy that determines the connectivity of S and D in minimum expected number of steps, where at each step one edge can be chosen and tested to see if it exists. The algorithm we present concludes either with the discovery of a route or the discovery of a cut between S and D . It is designed to reach a conclusion regarding one or the other in the shortest expected time in ER graphs. It should be noted that quickest discovery of a route is very different from discovering the shortest path route, which is a very well understood problem.

We start by considering the classic ER random graph, with a designated source S and destination D . The independence of edge occurrences and their equal probabilities

turn out to give ER graphs a great advantage in terms of solvability over other more complicated graph structures. Subsequently, we consider some slightly more general graphs in Section 5, where there is some additional deterministic structure known a priori about the presence or absence of certain edges.

The question of whether S and D are connected can be resolved if one can either display an S - D path, or an S - D cut. However, we do not know a priori whether the graph is connected or not, thereby making it difficult to know what to do – try to find a path or try to find a cut. The optimal policy has to be dynamic, based on the presence or absence of previously tested edges. This inherent tension in the problem of determining connectivity makes the problem somewhat challenging. Another difficulty is that we do not know the optimal solution for general graphs where some edges are known to exist, some known not to exist, and others existing i.i.d. with probability p , which is what one generally has after some steps of testing. Thus the proof of optimality is not based on dynamic programming-like arguments. Yet another aspect of interest is that the random graph exhibits a phase transition depending on the value of p . For $p > \frac{(1+\epsilon)\ln n}{n}$, the entire graph itself, and not just the particular S - D pair, is connected with high probability for large n , while for $p < \frac{(1-\epsilon)\ln n}{n}$, it is disconnected with high probability. Thus, one may possibly expect that the optimal strategy will depend on the value of p . Very interestingly, however, the optimal policy does not depend on p or n at all!

The main result is the optimality of the following simple testing strategy: At each stage, form the condensation multigraph by contracting each known connected component to a single node. It is a multigraph since the edges between nodes are inherited by the components to which they are contracted. In the set of edges that lie both on a shortest S - D path as well on the minimal S - D cut, test an edge that leads to the most opportunities for connectivity. The policy also has minimal complexity; it requires no more than $30\log_2 n$ operations to determine which edge to test next.

The policy is illustrated for a four-node ER graph in Figure 1. Figure 2 shows the number of steps, averaged over 100 simulations, that the optimal policy takes to establish S - D disconnection/connectivity in an ER graph with 1000 nodes, as a function of p . Determining S - D connectivity at around its phase transition can involve many steps. In between disconnection at very low p (say 10^{-5}) that takes about 999 steps to establish, and connectivity at very high p (around $p = 1$) that can be established in 1 step, the value of p passes through a phase transition around $p = 0.002$ (slightly smaller than the value of $p = \ln n/n \approx 0.006$ under the phase transition to connectivity of the entire network), where it takes a very large number of steps (about 15000) to determine if S and D are connected or not.

The proof of optimality of the policy follows from the proofs of optimality of the following three rules, which, when combined together give rise to the policy equivalently described above in terms of the condensation multigraph: Rule 1) The testing starts with a tactic that can lead to early termination by finding an S - D path: it first tests the edges, i.e., one-hop paths, connecting the component containing S and the component containing D . Rule 2) If there are no such edges, then it switches to checking for S - D disconnection by testing the edges on the S - D cut that

contains the minimum number of untested edges. Rule 3) Among the edges in that cut, test the edge that leads to most opportunities for connection. These three rules optimally resolve the tension between checking for paths and cuts.




Figure 1: Illustration of the evolution of the optimal policy on a four-node ER random graph. The line in red shows the edge selected for testing at each step. A bold line indicates that the edge was found to exist after testing, while a dotted line indicates that the edge was found to not exist after testing.

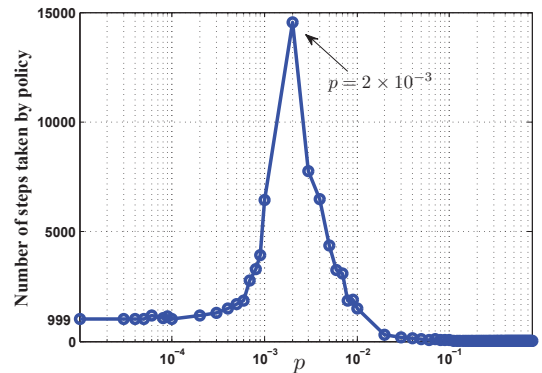


Figure 2: Number of steps needed to determine S - D disconnection/connectivity in a 1000 node ER Graph, as a function of p .

Related Works: We note that there is no other work, to the best of our knowledge, that addresses connectivity between a designated source and destination in random graphs. Rather, as noted earlier, it is asymptotic connectivity of the whole graph that has been intensively studied. Inspired by Erdos and Renyi [2,3], there have been many works [16–21]

dedicated to investigating connectivity of the entire network when the number of nodes is sufficiently large.

Our work also has ties to wireless networks where connectivity is determined by distance. Gilbert [4] initiated the study of random graphs when nodes are randomly distributed on a plane. Penrose [5] and Gupta and Kumar [6] have determined the critical range for establishing overall connectivity with probability approaching one as n goes to infinity. This range results in each node connecting to $\log n$ neighbors on average. Xue and Kumar [8], and, subsequently, [9] and [10], further considered the number of nearest neighbors that nodes need to connect to, in order to ensure the network connectivity. The condition of full connectivity was relaxed by Dousse et. al. [11], which analyzed the scenario where the sink is connected (in a multihop fashion) to a set of nodes that span the entire network. The impact of node degree, density and network dimension on connectivity is also investigated [7], [12]. Due to the emergence of ad hoc networks, mesh networks and sensor networks, there has also been research interest in analysis of k -connectivity [13], [14], [15].

2. MODEL AND DEFINITIONS

Consider an ER Random Graph with n nodes. An edge exists between any pair of nodes with equal probability p , independently of others. At each step, one can test a potential (i.e., untested) edge of the graph to see if it indeed exists. Two specific nodes are identified, labeled S and D . Our objective is to construct a sequential testing strategy terminating in the minimal expected number of steps, to determine if S and D are connected or not. Termination occurs with

- Connectivity, when it has verified the existence of all edges on some S - D path, or
- Disconnectivity, when it has verified the absence of all edges on some S - D cut.

We employ three terms, *known edge*, *known non-edge* and *potential edge* (also sometimes called an “untested edge.”). A known edge is one that has already been tested and found to exist. Similarly, a known non-edge is one that has already been tested but found to not exist. Finally, a potential edge is an edge that is yet to be tested, which upon testing may turn out to exist or not exist. In Table 1, we list other definitions used.

3. POLICY FOR CHECKING S - D CONNECTIVITY IN ER RANDOM GRAPHS

We now present a policy $\pi^* = \{\pi_0^*, \pi_1^*, \dots, \pi_N^*\}$, which we call the alternating policy, for checking S - D connectivity in an ER random graph. We label the nodes other than the designated S and D as $1, 2, \dots, n-2$. Each π_i^* maps \mathcal{G}_t (defined in Table 1) to one of the remaining potential edges. It specifies which potential edge should be tested at time t . Note that $N = n(n-1)/2$ is the upper bound on the number of edges in the graph, and thus the step by which the policy must necessarily have terminated.

The policy π^* has been defined in a brief manner in terms of the condensation multigraph in Section 1. Now we define it in a detailed manner in terms of three rules more suitable for illustration and proof.

Table 1: Notations

Notation	Definition
e_t	The edge tested at step t .
\mathcal{G}_t	The graph state known at step t . It is a list of (i) those potential edges that have been tested and found to exist, i.e., “known edges,” (ii) those potential edges that were tested and found to not exist, called “non-edges,” and (iii) the remaining potential edges that have not yet been tested. Abbreviated to \mathcal{G} if there is no confusion (similarly for the notation below).
$C_{S,t}$ (or $C_{D,t}$, or $C_{i,t}$)	The connected component containing the source (or destination, or the i -th component not containing the source or destination), in the graph comprised of the known edges at step t .
$\mathcal{M}_{\mathcal{G}}$	A <i>minimum cut</i> for a graph state \mathcal{G} , defined as a set which contains the minimum number of potential edges at time t , which if they were found to not exist, would lead to the conclusion that C_S and C_D are not connected by a path. Note that a minimum cut may be non-unique.

3.1 The Alternating Policy π^*

Rule 1: The policy tests a potential edge between C_S and C_D as long as there exists one such direct, one-hop, potential edge between them. Clearly, if such a direct edge is found to exist, then the policy terminates with the finding that S and D are connected.

Rule 2: If there are no potential edges connecting C_S and C_D as specified in Rule 1, then choose an edge from the list \mathcal{L} defined as follows.

First list all the paths, comprised of known or potential edges, connecting C_S to C_D with the minimum number of potential edges. One such path is illustrated in Figure 3(a). Such a path must necessarily have potential edges only *between* connected components. If the path traverses only one component, say C_1 , on its path from C_S to C_D , then it must have exactly two potential edges on it, one connecting C_S to C_1 , and another connecting C_1 to C_D , as shown in Figure 3(a). The remaining portion of the path must only consist of known edges *within* connected components C_S , C_1 or C_D . The same rule holds on the remaining potential *shortest* paths, each traversing through only one component, say C_i ($2 \leq i \leq r$), on its path from C_S to C_D .

Fix a minimum cut $\mathcal{M}_{\mathcal{G}}$ for the graph state \mathcal{G} at that time. Figure 3(a) shows two cuts, one separating C_S from $C_1 \cup C_2 \cup \dots \cup C_r$, and another separating C_D from $C_1 \cup C_2 \cup \dots \cup C_r$. (Suppose that, as in the case illustrated, the former is a Minimum Cut, and the latter contains more edges. In the sequel we will show that shortest paths are of length at most two, and minimum cuts will be precisely of this form, as long as the policy keeps testing the edges in the prescribed way.)

Among all the potential edges that are on the aforementioned shortest paths, list only those potential edges that are also in the chosen minimum set $\mathcal{M}_{\mathcal{G}}$. Define this as the list \mathcal{L} . Such a list is shown in Figure 3(b).

For convenience of definition of Rule 3 below, we also list the set of components $\{C_1, C_2, \dots, C_r\}$ ($0 \leq r \leq n-2$) to which the edges belonging to this list \mathcal{L} are connecting. Let us call this set of components as \mathcal{C} , i.e., $\mathcal{C} := \{C_1, C_2, \dots, C_r\}$. Figure 3(b) illustrates this set \mathcal{C} of components too.

Rule 3: In Rule 3 we further sharpen Rule 2 by specifying which particular edge in \mathcal{L} should be tested.

Rule 2 when followed will lead to the following structural property, as we prove in the sequel. All edges in \mathcal{L} are either between C_S and $\cup_{i=1}^r C_i$, or all edges are between C_D and $\cup_{i=1}^r C_i$. (Certainly this is true at the outset for the ER graph, and if an edge is chosen from \mathcal{L} , as in Rule 2, then it retains that property). Suppose all edges in \mathcal{L} connect to C_S , i.e., have one endpoint in C_S . Then, for each C_i , let n_i be the number of direct potential edges between C_i and C_D . Sort the components in \mathcal{C} based on n_i in decreasing order.

Suppose that n_1 is the largest. Test any edge in \mathcal{L} that connects C_S to C_1 . This is illustrated in Figure 3(c).

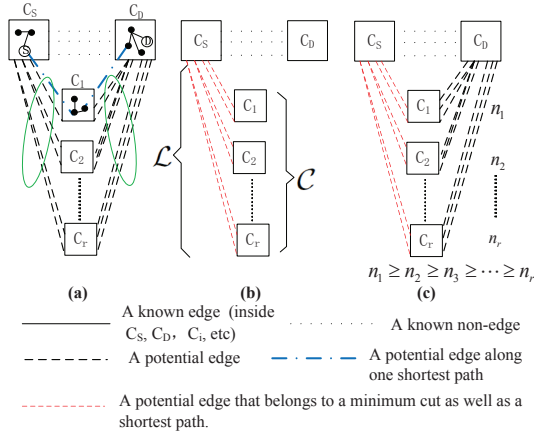


Figure 3: The illustration of Rules 2 and 3.

4. THE PROOF OF OPTIMALITY OF THE ALTERNATING POLICY π^*

We now commence the proof of optimality of the Alternating Policy π^* . It uses a carefully chosen stochastic coupling in several places. We will consider the graph shown in Figure 3(c); the general case proceeds similarly. The edges meeting Rules 2 and 3 are the potential edges between C_S and C_1 . Note that compared to the other C_i 's, C_1 has the greatest chance of subsequently having an edge connected to C_D , if an edge is found between C_1 and C_S .

4.1 Proof of Optimality of Rules 1 and 2

The optimality of Rule 1 follows from the following:

LEMMA 1. *Suppose that there are direct potential edges between C_S and C_D . Then for any policy A that tests an edge other than such a direct potential edge, there is a policy \tilde{A} that does test a direct potential edge and has a lower expected cost than A .*

PROOF. See Appendix A. \square

Next, assuming that Rule 1 is always followed, we prove the optimality of Rule 2. The following lemma establishes a preliminary property.

LEMMA 2. *When the policy follows both Rules 1 and 2, all the edges in the minimum cut at any step will be between $\cup_{i=1}^r C_i$ and C_S , or they will all be between $\cup_{i=1}^r C_i$ and C_D .*

PROOF. A minimum set partitions the connected components into two classes, with one class containing C_S and another class containing C_D . An equivalent claim of the lemma is that at any step t , all the components other than C_S or C_D will lie in one class. Suppose that not all the components lie in the same class. Then, in order to determine the S - D disconnectivity, the edges between components, in particular those that traverse two classes, say C_i and C_j will need to be tested, in addition to those between C_S and the components in C_i , and those between C_D and the components in C_j . Therefore, extra steps are wasted on testing the edges between components. However, those edges will never need to be tested when a cut is chosen to make all the components lie within one class. \square

Therefore, to prove Rule 2 in our scenario we only need to prove Lemma 3 stated as follows.

LEMMA 3. *When there are no direct edges between C_S and C_D , listing all the potential shortest paths and sampling the edges in the minimum set on them will lead to smaller expected cost than sampling any other edge first.*

PROOF. See Appendix B. \square

4.2 Proof of Optimality of Rule 3

Proving Rule 3 is equivalent to proving that in choosing between any two components from the set $\{C_1, \dots, C_r\}$ at step t , the expected cost will be smaller if we first pick the component with a larger number of direct edges to the other side. Without loss of generality, let us suppose that there are only two components, C_1 and C_2 . Based on Lemma 2, all the edges in the minimum cut \mathcal{M}_{G_t} are always on the same side. Let k_{11} and k_{21} be the number of edges from \mathcal{M}_{G_t} that connect to C_1 and C_2 , respectively, as shown in Figure 4(a). Similarly, let k_{12} and k_{22} be the number of direct edges C_1 and C_2 have to C_D . Assuming that $k_{12} \geq k_{22}$ in Figure 4, we have the following lemma regarding Rule 3.

LEMMA 4. *Among all the potential edges that C_1 and C_2 connect to in \mathcal{M}_{G_t} , it incurs smaller expected cost to first test the ones that C_1 is connected to.*

PROOF. We prove this by induction on the number of potential edges, in the graph of the type that results at each step from following Rules 1 and 2. Based on Lemma 2, such type of graph contains singleton components with the edges in the minimum cut all being between C_S and singletons, or all being between C_D and singletons. Clearly, Rules 1, 2 and 3 are all true for the graph with 2 potential edges. Assume that Rule 3 is true for graphs with k potential edges. We now consider the case where the graph has $k+1$ edges. Similar to the proof of Rule 1, we introduce two policies $\tilde{\mathcal{A}}$ and \mathcal{A} , both of which always follow Rules 1 and 2. $\tilde{\mathcal{A}}$ first tests an edge connecting C_S to C_1 , as prescribed by Rule 3. If it finds an edge, then by Rule 1 it subsequently tests edges from C_1 to C_D . If not, according to Rule 2, the same cut continues to be minimum, and by induction, since the number of potential edges is reduced by one after the first test, it continues to test the remaining edges between C_1 and C_S . In contrast, the policy \mathcal{A} violates Rule 3 on the first




Figure 4: The illustration of the four cases $(\tilde{\mathcal{E}}, \mathcal{E})$, $(\tilde{\mathcal{E}}^c, \mathcal{E}^c)$, $(\tilde{\mathcal{E}}, \mathcal{E}^c)$ and $(\tilde{\mathcal{E}}^c, \mathcal{E})$.

test by testing an edge connecting to C_2 , and then follows the optimal policy for the graph with k edges. We have four possible cases, which are listed as follows:

$\tilde{\mathcal{E}}$: $\tilde{\mathcal{A}}$ finds an edge between C_S and C_1 , and subsequently no edges between C_1 and C_D , and then the minimum cut turns out to be on the same side as C_S ; or $\tilde{\mathcal{A}}$ finds an edge between C_S and C_1 and subsequently an edge between C_1 and C_D ; or $\tilde{\mathcal{A}}$ finds no edges between C_S and C_1 .

$\tilde{\mathcal{E}}^c$: $\tilde{\mathcal{A}}$ finds an edge between C_S and C_1 , and subsequently no edges between C_1 and C_D , and the minimum cut therefore subsequently switches to the side of C_D .

\mathcal{E} : \mathcal{A} finds an edge between C_S and C_2 , and subsequently no edges between C_2 and C_D , and then the minimum cut turns out to be on the same side as C_S ; or \mathcal{A} finds an edge between C_S and C_2 and subsequently an edge between C_2 and C_D ; or \mathcal{A} finds no edges between C_S and C_2 .

\mathcal{E}^c : \mathcal{A} finds an edge between C_S and C_2 , and subsequently no edges between C_2 and C_D , and the minimum cut therefore subsequently switches to the side of C_D .

The four cases, i.e., $(\tilde{\mathcal{E}}, \mathcal{E})$, $(\tilde{\mathcal{E}}^c, \mathcal{E}^c)$, $(\tilde{\mathcal{E}}, \mathcal{E}^c)$ and $(\tilde{\mathcal{E}}^c, \mathcal{E})$ are illustrated in Figures 4(b), (c), (d) and (e), respectively. Consider Figure 4(b), which is the case $(\tilde{\mathcal{E}}, \mathcal{E})$. We will prove using a stochastic coupling argument that testing C_1 first leads to smaller expected cost. We couple the edges labeled with the same symbol (e.g., β) tested under $\tilde{\mathcal{A}}$ and \mathcal{A} , as shown in Figures 5 (a), (b), (c) and (d). Note in particular that, edges in α and β under policy $\tilde{\mathcal{A}}$ are switched under policy \mathcal{A} , and edges in θ and ε under policy $\tilde{\mathcal{A}}$ are also

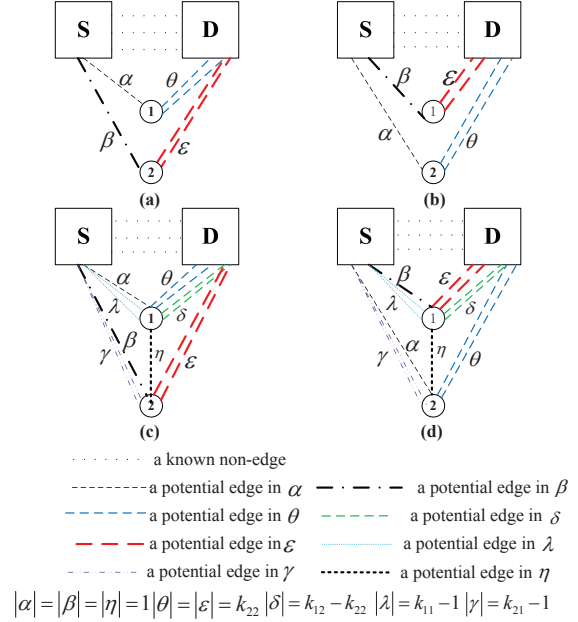


Figure 5: The stochastic coupling of the edges under policy $\tilde{\mathcal{A}}$ (shown in (a)) and policy \mathcal{A} (shown in (b)).

switched under policy \mathcal{A} , as shown in Figures 5 (a) and (b). The edges in δ , λ , γ and η under policy $\tilde{\mathcal{A}}$ are the same as those in δ , λ , γ and η under policy \mathcal{A} , as shown in Figures 5 (c) and (d). The edges with different labels are shown in Figures 5 (c) and (d). Since testing each potential edge has two possible outcomes depending on whether it exists or not, the corresponding sample paths generated under $\tilde{\mathcal{A}}$ and \mathcal{A} are illustrated in Figures 6 (a) and (b), respectively. The nodes in the figures represent the tested edges, while the outcome of each tested edge is indicated by a label, 1 if it exists, or 0 otherwise. Let $P_{D,i}$ ($i = 1, 2, \dots$) be the i th path where C_S and C_D are found to be disconnected under $\tilde{\mathcal{A}}$ (\mathcal{A}), and let $P_{C,i}$ be ($i = 1, 2, \dots$) the i th path where C_S and C_D are found to be connected under $\tilde{\mathcal{A}}$ (\mathcal{A}). Note that the paths with the same label have the same probability. As an example, consider paths labeled as $P_{D,4}$ in both Figures 6 (a) and (b). We can see that the outcome 0 from node α in Figure 6(a) has the same probability as the outcome 0 from nodes β in Figure 6(b). The outcome 1 from node β in Figure 6(a) has the same probability as the outcome 1 from node α in Figure 6(b). Nodes α and β have the same number of edges, and the outcomes of the edges from nodes θ , δ , ε and λ are the same in both figures. Another example is the paths labeled with $P_{C,2}$ in both figures. It can be seen that the outcome 1 from node θ in Figure 6(a) has the same probability as the outcome 1 from node ε in Figure 6(b). Nodes θ and ε have the same number of edges, and the outcomes of the edges emanating from nodes α and λ are the same in both figures. Similar relationships hold for the remaining paths with the same labels in Figures 4 (a) and (b), and we therefore omit their explanations. As for the paths $\tilde{P}_{D,1}$ and $\tilde{P}_{D,2}$ in Figure 6(a) and paths $\hat{P}_{D,1}$ and $\hat{P}_{D,2}$ in Figure 6(b), it is trivially true that any terminating time belonging to the range $[k_{12}+2, 1+k_{11}+k_{12}]$

under the paths $\tilde{P}_{D,1}$ and $\tilde{P}_{D,2}$ can also be found under the paths $\hat{P}_{D,1}$ and $\hat{P}_{D,2}$. Therefore, the only paths that have different probabilities are shown in the form of bold lines under both \tilde{A} and A (specifically, path $\alpha\theta^c\delta$ under \tilde{A} , and paths $\alpha\theta^c\beta\delta$, $\alpha\theta^c\beta^c\lambda\delta$ and $\alpha^c\lambda^c\beta\varepsilon^c\delta$ under A). Obviously, $\alpha\theta^c\delta$ dominates the paths $\alpha\theta^c\beta\delta$, $\alpha\theta^c\beta^c\lambda\delta$ and $\alpha^c\lambda^c\beta\varepsilon^c\delta$, which indicates that testing C_1 first will lead to smaller expected cost. This completes the proof of the case in Figure 4(b).

Table 2: Definitions of the edges for coupling \tilde{A} and A

Notation	Definition
α	The first potential edge that policy \tilde{A} (A) tests.
β	The first potential edge that policy \tilde{A} (A) tests between C_S and C_2 (C_1).
θ	The first k_{22} potential edges that policy \tilde{A} (A) tests after an edge is found between C_1 (C_2) and C_S . Recall that $k_{12} \geq k_{22}$.
ε	k_{22} potential edges to be tested by policy \tilde{A} (A) after β is found to exist.
δ	The $k_{12} - k_{22}$ remaining potential edges yet to be tested by policy \tilde{A} (A) between C_D and C_1 , after no edges are found to exist in θ (ε).
λ	The $k_{11} - 1$ remaining potential edges yet to be tested by policy \tilde{A} (A) between C_S and C_1 , after the first edge between C_S and C_1 is found to not exist.
γ	The $k_{21} - 1$ remaining potential edges between C_S and C_2 .
η	The potential edge between C_1 and C_2 .

Next we turn to the case $(\tilde{\mathcal{E}}^c, \mathcal{E}^c)$ shown in Figure 4(c). The paths generated under \tilde{A} and A using coupling are illustrated in Figures 7 (a) and (b), respectively. Again, the paths with the same label have the same probability, and they can be checked as in the aforementioned case $(\tilde{\mathcal{E}}, \mathcal{E})$. As for the paths $\tilde{P}_{D,1}$ and $\tilde{P}_{D,2}$ in Figure 7(a) and paths $\hat{P}_{D,1}$ and $\hat{P}_{D,2}$ in Figure 7(b), it is trivially true that any terminating time belonging to the range $[k_{12} + k_{21} + k_{22} + 2, k_{11} + k_{12} + k_{21} + k_{22} + 1]$ under the paths $\tilde{P}_{D,1}$ and $\tilde{P}_{D,2}$ can also be found under the paths $\hat{P}_{D,1}$ and $\hat{P}_{D,2}$. The same conclusion also holds for the paths $\tilde{P}_{C,1}$ and $\tilde{P}_{C,2}$ in Figure 7(a), and paths $\hat{P}_{C,1}$ and $\hat{P}_{C,2}$ in Figure 7(b). Hence, we highlight the paths that will terminate in different time with bold lines under both \tilde{A} and A (specifically path $\alpha\theta^c\delta$ under \tilde{A} and paths $\alpha\theta^c\delta\beta$, $\alpha\theta^c\delta\beta^c\lambda$ and $\alpha^c\lambda^c\beta\varepsilon^c\delta$ under A). Obviously, $\alpha\theta^c\delta$ terminates earlier than $\alpha\theta^c\delta\beta$, $\alpha\theta^c\delta\beta^c\lambda$ as well as $\alpha^c\lambda^c\beta\varepsilon^c\delta$, which indicates that testing C_1 first will lead to smaller expected cost. This completes the proof of the case in Figure 4(c).

Now we proceed to prove cases $(\tilde{\mathcal{E}}, \mathcal{E}^c)$ and $(\tilde{\mathcal{E}}^c, \mathcal{E})$, shown in Figures 4(d) and 4(e), respectively. We first consider case $(\tilde{\mathcal{E}}, \mathcal{E}^c)$. This case implies that the minimum cut turns out to be on the same side of C_S , after an edge is found between C_1 and C_S but no edges are found between C_1 and C_D by \tilde{A} . Due to Rule 2, it leads to smaller expected cost for \tilde{A} to subsequently test the edges between C_2 and C_S rather than those between C_2 and C_D . This means the expected cost of case $(\tilde{\mathcal{E}}, \mathcal{E}^c)$ is smaller than that of case $(\tilde{\mathcal{E}}^c, \mathcal{E}^c)$. Moreover, we have already proved that the expected cost is smaller




Figure 6: Paths generated by the coupling of \tilde{A} (a) and A (b) for the case shown in Figure 4(b).

under policy \tilde{A} in case $(\tilde{\mathcal{E}}^c, \mathcal{E}^c)$. By transitivity, we conclude that \tilde{A} leads to smaller cost than A does in case $(\tilde{\mathcal{E}}, \mathcal{E}^c)$. Now consider the last case, $(\tilde{\mathcal{E}}^c, \mathcal{E})$. This case implies that the minimum cut switches to the side of C_D , after an edge is found between C_1 and C_S but no edges are found between C_1 and C_D by \tilde{A} . According to Rule 2, it incurs smaller expected cost if \tilde{A} subsequently tests the edges between C_2 and C_D rather than those between C_2 and C_S . Since we have already proved that the cost is smaller under policy \tilde{A} in case $(\tilde{\mathcal{E}}, \mathcal{E})$, it follows that by transitivity that \tilde{A} leads to smaller expected cost than A in case $(\tilde{\mathcal{E}}^c, \mathcal{E})$. \square

The low complexity of the policy, of no more than $30 \log_2 n$ operations per step, follows from the fact that except for C_S and C_D , all other components are singletons, and moreover are only of three kinds, either having a known non-edge to S , or a known non-edge to D , or neither. Hence all computations involving these components are extremely simple.

5. EXTENSION TO GENERAL GRAPHS

In this section, we indicate extensions of the optimality results, without proofs, to certain more general classes of random graphs.




Figure 7: Paths generated by the coupling of \tilde{A} (a) and A (b) for the case shown in Figure 4(c).

5.1 (1,p) Random Graphs

By a $(1,p)$ random graph, we mean a graph that initially contains two types of edges, “1” edges and “ p ” edges. The former are edges that are known to exist, i.e., exist with probability one, while the latter are potential edges that independently exist with probability p . The Alternating Policy remains optimal under the condition that the size of each component, excluding C_S and C_D , is the same.

5.2 (1,0,p) Random Graphs

By a $(1,0,p)$ random graph, we mean a graph that also initially contains “0” edges. A “0” edge is simply an edge that we already know to not exist. Unfortunately, determining the optimal policy for all types of $(1,0,p)$ random graphs appears to be intractable. The Alternating Policy remains optimal for certain types of $(1,0,p)$ graph patterns – series graphs, parallel graphs, SP graphs, PS graphs, series of parallel of series (SPS) graphs and parallel of series of parallel (PSP) graphs.

5.2.1 Optimal policy for Series of Parallel (SP) Graphs

An SP graph consists of several parallel graphs arranged in series, as shown in Figure 8. It consists of n parallel graphs labeled $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_n$, with m_i potential edges (i.e.,

existing independently with equal probability p) in the i th parallel graph.




Figure 8: A series of parallel (SP) graph.

THEOREM 1. The optimal policy is to test the parallel subgraph with the fewest number of potential edges.

5.2.2 Optimal Policy for Parallel of Series (PS) Graphs

A PS graph consists of a series graphs arranged in parallel, as shown in Figure 9. It consists of n series graphs labeled $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_n$, with m_i potential edges in the i th series graph.




Figure 9: A parallel of series (PS) graph.

THEOREM 2. The optimal policy is to test any edge on the series subgraph that contains the fewest number of potential edges.

5.2.3 The Optimal Policy for SPS Graphs

An SPS graph consists of several PS graphs arranged in series, as shown in Figure 10. It consists of n PS graphs labeled $\mathcal{PS}_1, \mathcal{PS}_2, \dots, \mathcal{PS}_n$, with m_i potential edges in the i th ($1 \leq i \leq n$) PS graph.




Figure 10: A series of parallel of series (SPS) graph.

THEOREM 3. The optimal policy is to test any edge in that series graph which has the fewest number of edges in the PS graph that contains the minimum number of series graphs.

5.2.4 The Optimal Policy for PSP Graphs

A PSP graph consists of several SP graphs arranged in parallel, as shown in Figure 11. It consists of n SP graphs labeled as $\mathcal{SP}_1, \mathcal{SP}_2, \dots, \mathcal{SP}_n$, with m_i potential edges in the i th ($1 \leq i \leq n$) SP graph.

THEOREM 4. The optimal policy is to test any edge in that parallel graph which contains the fewest number of edges in the SP graph that contains the minimum number of parallel graphs.




Figure 11: A parallel of series of parallel (PSP) graph.

6. CONCLUSION

We have studied the problem of optimally determining S - D connectivity of random graphs. We have considered the class of classic ER Random Graphs with a finite number n of nodes, for any value of n . Assuming that each testing of an edge has the same unit cost, we have determined a policy for establishing whether the designated source and destination are connected with minimum expected cost. Interestingly, though ER graphs exhibit a sharp transition between disconnectivity of the entire graph and connectivity of the entire graph around $p = \ln n/n$, the optimal policy turns out to not depend on the specific value of p , or even of which side of the phase transition p lies. The policy simply contracts each known connected component to a single super node at each step, and in that condensation multigraph it simply tests an edge that is both on the shortest path containing the super nodes containing S and D , as well as on a minimum S - D cut, giving higher priority to sampling the edges that will lead to more opportunities for S - D connectivity. The policy is also extendable to certain types of more generalized graph structures consisting of additional deterministic structure.

Acknowledgement

This paper is partially based on work supported by NSF under Contract Nos. CNS-1302182 and CCF-0939370, and AFOSR under Contract No. FA-9550-13-1-0008, and NSF China (No. 61325012, 61271219), and National Key Project (2012ZX03002009-004).

7. REFERENCES

- [1] E. N. Gilbert. Random graphs. *Annals of Mathematical Statistics*, 30(4):1141–1144, 1959.
- [2] P. Erdos and A. Renyi. On random graphs. I. *Publicationes Mathematicae* 6, pages 290–297, 1959.
- [3] P. Erdos and A. Renyi. On the evolution of random graphs. *the Mathematical Institute of the Hungarian Academy of Sciences* 5, 17-61, 1960.
- [4] E. N. Gilbert. Random plane networks. *J. SIAM*, 9:533–543, Dec. 1961.
- [5] M. D. Penrose. The longest edge of the random minimal spanning tree. *The Annals of Probability*, 7:340–361, 1997.
- [6] P. Gupta and P. R. Kumar. Critical power for asymptotic connectivity in wireless networks. *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of W.H. Fleming, W.M. McEneaney, G. Yin, and Q. Zhang*, pages 547–566, 1998.

- [7] C. Bettstetter. On the minimum node degree and connectivity of a wireless multihop network. *Proc. of ACM MobiHoc*, 2002.
- [8] F. Xue and P. R. Kumar. The number of neighbors needed for connectivity of wireless networks. *Wireless Networks*, 10:169–181, Mar. 2004.
- [9] F. Xue and P. R. Kumar. On the theta-coverage and connectivity of large random networks. *Joint Special Issue of IEEE Trans. on Inform. Theory and IEEE/ACM Trans.s on Netw. and Inform. Theory*, 52:2289–2399, June 2006.
- [10] P. Balister, B. Bollobas, A. Sarkar, and M. Walters. Connectivity of random k -nearest-neighbour graphs. *Advances in Applied Probability*, 37:1–24.
- [11] O. Dousse, P. Mannersalo, and P. Thiran. Latency of wireless sensor networks with uncoordinated power saving mechanisms. *Proc. of ACM MobiHoc’04*, May 2004.
- [12] H. M. Ammari and S. K. Das. Asymptotic critical transmission radius for k -connectivity in wireless ad hoc networks. *IEEE Trans. Para. and Dis. System*, 20:872–885, 2009.
- [13] P. J. Wan and C. W. Yi. Critical density for coverage and connectivity in three-dimensional wireless sensor networks using continuum percolation. *IEEE/ACM Transactions on Networks and Information Theory*, 56:2867–2874, May 2010.
- [14] X. L. Bai, C. Zhang, D. Xuan, and W. J. Jia. Full-coverage and k -connectivity ($k = 14, 6$) three dimensional networks. *Proc. of IEEE INFOCOM*, July 2009.
- [15] X. Wang, X. Lin, Q. Wang, and W. Luan. Mobility increases the connectivity of wireless networks. *IEEE/ACM Trans. on Netw.*, 21:440–454, Apr. 2012.
- [16] P. Erdos, E. M. Palmer, and R. W. Robinson. Local connectivity of a random graph. *J. of Graph Theory*, 7(4):411–417, 1983.
- [17] B. Bollobas, S. Janson, and O. Riordan. The phase transition in inhomogeneous random graphs. *J. Random Structures and Algorithms*, 31(1):3–122, 2007.
- [18] G. Han and A. M. Makowski. One-dimensional geometric random graphs with nonvanishing densities: Part I: A strong zero-one law for connectivity. *J. IEEE. Trans. on Inform. Theory*, 55(12):5832–5839, 2009.
- [19] M. D. Penrose. On k -connectivity for a geometric random graph. *Random Structure and Algorithms*, 15(2):145–164, Sept. 1999.
- [20] C. Hirsch, D. Neuhauser, and V. Schmidt. Connectivity of random geometric graphs related to minimal spanning forests. *Advances in Applied Probability*, 45:1–36, 2013.
- [21] L. Devroye and N. Fraiman. Connectivity of inhomogeneous random graphs. *J. Random Structures and Algorithms*, 2013.

Appendix

Appendix A: Proof of Lemma 1

The proof uses stochastic coupling. We need to consider a general hybrid graph that may be the state at some

intermediate time. Such a graph will contain C_S , C_D , and some other components with different sizes, as well as some edges that we already know to not exist. Consider such an initial graph state, and consider any policy \mathcal{A} which does not first test a potential direct edge (denoted by d) between C_S and C_D at step 0. (We note that all direct potential edges are identical.). Let $U_i(\omega)$ denote the choice made by \mathcal{A} at a subsequent step i , with ω representing the sample path, with an outcome of sampling being either 1 (edge existing) or 0 (not existing). Denote by $T(\omega)$ the termination time of \mathcal{A} , and by $n_d(\omega)$ the first time that \mathcal{A} tests d , i.e., $n_d(\omega) = \inf\{t \geq 1 : U(t) = d\}$. Since \mathcal{A} may never get to check d , we say that $n_d(\omega) = +\infty$ if $U_1(\omega) \neq d, U_2(\omega) \neq d, \dots, U_{T(\omega)}(\omega) \neq d$.

The sequence of edges tested under \mathcal{A} till it terminates can be divided into three cases, as follows:

- Case 1: $\{U_1(\omega), U_2(\omega), \dots, U_{T(\omega)}(\omega)\}$ where $U_i(\omega) \neq d$, $i = 1, \dots, T(\omega)$. In this case, we have $T(\omega) < n_d(\omega) = +\infty$.
- Case 2: $\{U_1(\omega), U_2(\omega), \dots, U_{n_d(\omega)-1}(\omega), U_{n_d(\omega)}(\omega) = d\}$. Obviously, in this case we have $T(\omega) = n_d(\omega)$.
- Case 3: $\{U_1(\omega), U_2(\omega), \dots, U_{n_d(\omega)-1}(\omega), U_{n_d(\omega)}(\omega) = d, U_{n_d(\omega)+1}(\omega), \dots, U_{T(\omega)}(\omega)\}$ where we have $T(\omega) > n_d(\omega)$.

Now we construct a family of policies $\{\mathcal{A}^i\}$, where \mathcal{A}^i follows \mathcal{A} , except that if \mathcal{A} has not tested d by step i , then \mathcal{A}^i does test d at step i , and subsequently continues to test the same edges tested by \mathcal{A} except delayed by one step, and catches up with \mathcal{A} when \mathcal{A} tests d . We note that each \mathcal{A}^i is a feasible policy, i.e., it only uses information that \mathcal{A}^i will provide (even though it is supposed to emulate \mathcal{A} after step i). In particular, let us consider \mathcal{A}^1 , where the d is tested at the first step, and which subsequently follows \mathcal{A} (delayed by one step below). Correspondingly, we denote T^i as the termination time of \mathcal{A}^i . The sequence of resulting edges tested by \mathcal{A}^1 till it terminates has three possible cases.

- Case 1: $\{d\}$ if $T^1 = n_d(\omega)$ or $\{T^1 < n_d(\omega), e_d = 1\}$.
- Case 2: $\{d, U_1(\omega), U_2(\omega), \dots, U_{T^1}(\omega)\}$ if $T^1 < n_d(\omega)$ and $e_d = 0$.
- Case 3: $\{d, U_1(\omega), U_2(\omega), \dots, U_{n_d(\omega)-1}(\omega), U_{n_d(\omega)}(\omega), U_{n_d(\omega)+1}(\omega), \dots, U_{T^1}(\omega)\}$ if $T^1 > n_d(\omega)$ and $e_d = 0$.

Based on this, T^1 can be expressed as

$$\begin{aligned} |T^1| &= \mathbf{1}(T^1 = n_d(\omega)) + \mathbf{1}(T^1 < n_d(\omega), e_d = 1) \\ &\quad + \mathbf{1}(T^1 < n_d(\omega), e_d = 0) [|T^1| + 1] \\ &\quad + \mathbf{1}(T^1 > n_d(\omega), e_d = 0) |T^1|. \end{aligned}$$

Based on this, we have

$$\begin{aligned} T^1 - T^i &= \mathbf{1}(T^i = n_d(\omega)) [1 - T^i] \\ &\quad + \mathbf{1}(T^i < n_d(\omega), d = 1) [1 - T^i] \\ &\quad + \mathbf{1}(T^i < n_d(\omega), d = 0) 1. \end{aligned}$$




Figure 12: The edges coupled under policies $\tilde{\mathcal{A}}$ (a) and \mathcal{A} (b).

Now we consider the policy \mathcal{A}^2 , and define T^2 similarly. Then we have

$$\begin{aligned} \mathbb{E}[T^1 - T^2] &= \mathbb{E}[\mathbf{1}(T^2 = 2)(1 - 2) + \mathbf{1}(T^2 < 2, e_d = 1)(1 - 1) \\ &\quad + \mathbf{1}(T^2 < 2, e_d = 0)|T^1|] \\ &= \mathbb{E}[\mathbf{1}(T^2 = 2)(-1) + \mathbf{1}(T^2 = 1, e_d = 0)] \\ &= -p + 0 = -p. \end{aligned}$$

Hence \mathcal{A}^1 has lower cost than \mathcal{A}^2 . Similarly, we can prove that each \mathcal{A}^i has lower cost than \mathcal{A}^{i+1} . Noting that $\mathcal{A}^N = \mathcal{A}$, we thus conclude that \mathcal{A}^1 has lower cost than \mathcal{A} , proving the lemma. \square

Appendix B: Proof of Lemma 3

We prove by induction on the number of steps. Clearly, by Lemma 1, the conclusion holds at step 1. Suppose it is true till a certain step t . From Lemma 2, it follows that all the edges in the minimum cut at step t are either all between C_S and $\cup_{i=1}^r C_i$, or they are all between C_D and $\cup_{i=1}^r C_i$. Furthermore, since all those edges are also along all the potential shortest paths, they also belong to the set \mathcal{L} . Without loss of generality, suppose that all the edges in \mathcal{L} at step t are between C_S and $\cup_{i=1}^r C_i$. Now we consider step $t + 1$.

We introduce two policies $\tilde{\mathcal{A}}$ and \mathcal{A} , where $\tilde{\mathcal{A}}$ always follows Rule 2. Specially, at step $t + 1$, $\tilde{\mathcal{A}}$ will test the edges in \mathcal{M}_{G_t} , starting from an edge between some component, say C_1 , and C_S . If an edge is found between C_1 and C_S , then, due to Rule 1, $\tilde{\mathcal{A}}$ subsequently tests all the edges between C_1 and C_D . If not, then the same cut continues to be the minimum cut, and $\tilde{\mathcal{A}}$ continues to test the remaining edges between C_1 and C_S . In contrast, \mathcal{A} violates Rule 2 at step $t + 1$ by testing an edge between C_1 and then follows $\tilde{\mathcal{A}}$.

We prove that testing the edges between C_1 and C_S leads to smaller expected cost, by again employing a stochastic coupling. Note that the number of edges between C_1 and C_S can be either larger or smaller than the number of the edges between C_1 and C_D . We only consider the former case; the latter is proved similarly. We couple the edges tested under $\tilde{\mathcal{A}}$ and \mathcal{A} , as shown in Figures 12. Table 3 lists the permutations of the nodes labels involved in the coupling. The paths generated under both $\tilde{\mathcal{A}}$ and \mathcal{A} for the cases in Figures 12(a) and (b) are shown in Figure 13. The nodes in the figures represent the edges tested while the outcomes of the tested edges are indicated using lines with 1 (or 0)

Table 3: Definitions of the edges coupled under $\tilde{\mathcal{A}}$ and \mathcal{A}

Notation	Definition
α	The first potential edge that policy $\tilde{\mathcal{A}}$ (\mathcal{A}) tests.
β	The first potential edge that policy $\tilde{\mathcal{A}}$ (\mathcal{A}) tests between C_1 and C_D (C_S).
θ	The $k_{11} - 1$ potential edges that policy $\tilde{\mathcal{A}}$ (\mathcal{A}) tests between C_1 and C_D (C_S) after β is found not to exist between C_1 and C_D (C_S).
δ	The remaining $k_{12} - k_{11}$ ($k_{11} - k_{12}$) potential edges to be tested by policy $\tilde{\mathcal{A}}$ (\mathcal{A}) between C_D and C_1 .
ε	The remaining $k_{11} - 1$ potential edges between C_1 and C_S to be tested by policy $\tilde{\mathcal{A}}$ and the $k_{11} - 1$ potential edges between C_1 and C_D to be tested by policy \mathcal{A} .




Figure 13: Paths generated by the coupling on $\tilde{\mathcal{A}}$ (a) and \mathcal{A} (b) for the case shown in Figures 12 (a)-(b).

meaning that the edges turn out to exist (or not exist). The paths that terminate at the same time are marked with the same labels under $\tilde{\mathcal{A}}$ and \mathcal{A} . The two paths highlighted in bold lines in Figure 13(b), i.e., $\hat{P}_{D,1}$ and $\hat{P}_{D,2}$, represent the graph states where none of the edges between C_1 and C_S exist while one edge between C_1 and C_D is found to exist in $\hat{P}_{D,2}$ and not to exist in $\hat{P}_{D,1}$ under policy \mathcal{A} . The path in bold line in Figure 13(a) represents the graph state where none of the edges between C_1 and C_S exist under policy $\tilde{\mathcal{A}}$, with the circle containing dots being the possible edges that $\tilde{\mathcal{A}}$ will subsequently test. Notice that if the circle represents an edge between C_1 and C_D , i.e., if $\tilde{\mathcal{A}}$ chooses to test an edge between C_1 and C_D , then both Figures 13(a) and (b) will have the same termination probability. However, the expected cost will be smaller if \mathcal{A} tests any edge between C_S (C_D) and one of the remaining components. Similarly, we can prove that the expected cost can be even smaller if the edge tested belongs to the remaining edges in \mathcal{L} . The proof is completed by checking all the components. \square