Detecting a Network Failure

Jon Kleinberg *

Abstract

Measuring the properties of a large, unstructed network can be difficult: one may not have full knowledge of the network topology, and detailed global measurements may be infeasible. A valuable approach to such problems is to take measurements from selected locations within the network and then aggregate them to infer large-scale properties. One sees this notion applied in settings that range from Internet topology discovery tools to remote software agents that estimate the download times of popular Web pages. Some of the most basic questions about this type of approach, however, are largely unresolved at an analytical level. How reliable are the results? How much does the choice of measurement locations affect the aggregate information one infers about the network?

We describe algorithms that yield provable guarantees for a particular problem of this type: detecting a network failure. Suppose we want to detect events of the following form: an adversary destroys up to k nodes or edges, after which two subsets of the nodes, each at least an ε fraction of the network, are disconnected from one another. We call such an event an (ε, k)-partition. One method for detecting such events would be to place "agents" at a set D of nodes, and record a fault whenever two of them become separated from each other. To be a good *detection set*, D should become disconnected whenever there is an (ε, k)-partition; in this way, it "witnesses" all such events.

We show that every graph has a detection set of size polynomial in k and ε^{-1}, and *independent* of the size of the graph itself. Moreover, random sampling provides an effective way to construct such a set. Our analysis establishes a connection between graph separators and the notion of VC-dimension, using techniques based on matchings and disjoint paths.

1 Introduction

Inferring Network Properties. In large, unstructured networks, it can be difficult to answer even the most basic questions — What does the topology look like? How many (large) connected components are there? What is the average distance (or packet round-trip time) between pairs of nodes? Indeed, a vein of recent work has focused on the issue

*Department of Computer Science, Cornell University, Ithaca NY 14853. Email: kleinber@cs.cornell.edu.
Supported in part by a David and Lucile Packard Foundation Fellowship, an ONR Young Investigator Award, and NSF Faculty Early Career Development Award CCR-9701399.

of designing algorithms that operate on networks for which we do not necessarily have an explicit representation, or for which detailed global measurements are not feasible (e.g. [4, 5, 14, 15, 17] and, in a fairly different context, [3, 10]).

Such work has arisen from a variety of different motivations, but much of it — especially in the context of Internet performance metrics — is based on the notion of taking measurements from selected locations within the network, and then aggregating these measurements to infer large-scale network properties. This approach is used, for example, in structural analysis based on *Internet tomography* [4, 5], algorithms for discovering Internet topology [17] and routing structures [15], and mechanisms for estimating a variety of Internet performance metrics [14]. Companies such as Keynote Systems [8, 9] provide information on the average response time of home page requests to high-volume Web sites, by combining data collected from software "agents" distributed through the Internet.

This type of consideration — combining data from measurement agents in a network — forms the crux of our concerns in this paper. How does the choice of locations for such agents affect the aggregate information one infers about the network? How large a subset of locations do we need in order to achieve a desired level of confidence in such information? For most natural performance metrics, these questions lead quickly to open problems.

Network Connectivity. We formulate a concrete inference problem concerning one of the most fundamental network properties — connectivity. Consider an undirected graph G on n nodes, which is initially connected. Suppose, for a parameter $\varepsilon > 0$, we are interested in detecting *ε-partitions*: failures of network elements after which there are two subsets of nodes A and B, each of size at least εn, such that no node in A has a path to any node in B. For a parameter $k > 0$, we wish to be able to detect any ε-partition that is caused by the failure of up to k (adversarially chosen) nodes or edges, and record the occurrence of such an event. Note that increasing k allows an adversary more power, so handling larger values of k represents a sequence of successively more difficult problems.

Here is a general approach for doing this, motivated by the type of analysis discussed above. We place "monitoring agents" at a subset D of the nodes of G, and each of these agents periodically engages in communication with all the others. Now, if at some point in time, there are nodes $u, v \in D$ such that the agents at u and v have no path connecting them, we can record a fault in the network. Such a protocol clearly has the property that a fault is only recorded if the network has actually become disconnected; and we would like to choose D so that a fault is *always* recorded when the deletion of up to k nodes or edges has resulted in an ε-partition. How large a set D do we need in order to achieve this, for an underlying graph G?

We make this question precise as follows. Let $G = (V, E)$ be a connected graph; by an *element* of G, we mean one of its nodes or edges. We say that two sets of nodes A and B are *separated* if there is no path with one end in A and the other in B. We say that a set of elements Z is an (ε, k)-*partitioning set* if $|Z| \leq k$ and $G \setminus Z$ contains disjoint subsets of nodes A and B, each of size at least εn, that are separated. Finally, we say that a set $D \subseteq V$ is an (ε, k)-*detection set* if for every (ε, k)-partitioning set Z, there are two nodes $u, v \in D \setminus Z$ that lie in different connected components of $G \setminus Z$. Such a pair of nodes will no longer be able to communicate after the deletion of the elements in Z, and hence will "witness" the partitioning of the network.

Clearly by taking D to be the entire node set V, we obtain an (ε, k)-detection set. The question is whether there is a much smaller sample with the same property.

Our main result is that for every graph G, there is an (ε, k)-detection set whose size is bounded by a function of ε and k only, independent of the number of nodes. Specifically, we show

Theorem 1.1 *Every graph G has an (ε, k)-detection set of size $O(k^3\varepsilon^{-1}\log(k\varepsilon^{-1}))$. Moreover, a set of*

$$O(k^3\varepsilon^{-1}\log(k\varepsilon^{-1}) + \varepsilon^{-1}\log\delta^{-1})$$

nodes chosen uniformly at random is an (ε, k)-detection set with probability at least $1 - \delta$.

Thus, the theorem also shows that uniform random sampling yields a detection set with high probability; in particular, this means that one can construct a small detection set without detailed local knowledge of the network topology.

The heart of the proof is the analysis of this sampling procedure. For this we use the notion of VC-dimension [18]; we define a particular set system over the vertex set of the graph, show that it has small VC-dimension, and then relate detection sets to *ε-nets* [6] for this set system. We discover that the most natural set systems here actually have very large VC-dimension, and so are not useful in our analysis; formulating and analyzing one for which this parameter is bounded raises some interesting graph-theoretic issues, including connections to a theorem of Mader on disjoint paths [12] that generalizes the Tutte-Berge theorem on non-bipartite matchings [2, 11]. We view the connection developed here between VC-dimension and the separators of a graph as one of the contributions of this work, and believe that our style of analysis may be amenable to related problems as well.

The role of node and edge separators here suggests connections to problems such as approximating the failure probability of a network [7]. It is important to note, however, that the issues are quite different at a technical level. Specifically, we are concerned with choosing a fixed set of nodes D from which we can detect any possible (adversarially chosen) failure of a particular type. Also, we allow here for node failures in addition to edge failures; network reliability allowing node failures, on the other hand, is largely an open question [7].

Generalizations. We will actually prove the following strengthening of Theorem 1.1. Suppose that the nodes of our graph are initially partitioned into two classes: *end nodes* V_0, and *internal nodes* V_1. We are only allowed to place monitoring agents at nodes in V_0, and we are interested in detecting ε-partitions of V_0.

Thus, we say that a set of elements Z is an *(ε, k)-partitioning set with respect to V_0* if $|Z| \leq k$ and $G \setminus Z$ contains disjoint subsets $A, B \subseteq V_0$, each of size at least $\varepsilon|V_0|$, that are separated. We say that a set $D \subseteq V_0$ is an *(ε, k)-detection set for V_0* if for every set Z that is (ε, k)-partitioning with respect to V_0, there are two nodes $u, v \in D \setminus Z$ that lie in different connected components of $G \setminus Z$. Our generalization of Theorem 1.1 is now simply the following.

Theorem 1.2 *For every graph $G = (V, E)$, and any set $V_0 \subseteq V$, there is an (ε, k)-detection set for V_0 of size $O(k^3\varepsilon^{-1}\log(k\varepsilon^{-1}))$. Moreover, a set of*

$$O(k^3\varepsilon^{-1}\log(k\varepsilon^{-1}) + \varepsilon^{-1}\log\delta^{-1})$$

nodes chosen uniformly at random from V_0 is an (ε, k)-detection set for V_0 with probability at least $1 - \delta$.

Alternate Notions of Detection. Finally, we note that there are two natural strengthenings of our notion of detection — in which we drop one or the other of the parameters ε and k — but neither of these can be achieved in any non-trivial way.

Specifically, one natural strengthening would be to drop the parameter ε and ask that a set D' have the following property: for any $Z \subseteq V \cup E$ of size at most k, if $G \setminus Z$ is not connected, then some two nodes of $D' \setminus Z$ lie in different components of $G \setminus Z$. But if we consider any 2-connected d-regular graph, and set $k = d$, we see that Z could consist of the d edges incident to any node, and hence we would need to take $D' = V$.

A second natural strengthening would be to drop the parameter k and ask that a set D'' have the following property: for any $Z \subseteq V \cup E$, if $G \setminus Z$ contains separated subsets A and B, each of size at least εn, then some two nodes of $D'' \setminus Z$ lie in different components of $G \setminus Z$. But in this case, consider the star graph $G = K_{1,n-1}$, and suppose $D'' \subseteq V$ has fewer than $(1 - \varepsilon)n$ nodes. Let Z consist of all edges incident to leaves of G not in D''. $G \setminus Z$ contains two separated sets, each of size at least εn, but D'' is completely contained in a single component of $G \setminus Z$.

2 Detecting Edge Failures

We begin by considering the special case in which only edges can fail. We also assume here that $V_0 = V$. Focusing on this special case first provides a good illustration of the way in which we use VC-dimension arguments; it also allows us to work with a particularly natural set system defined over the vertex set for which we can show that our VC-dimension analysis is tight.

Thus, we say that $D \subseteq V$ is an (ε, k)-detection set *with respect to edge failures* if for every (ε, k)-partitioning set Z consisting only of edges, some two nodes of D lie in different components of $G \setminus Z$. We say that a set $S \subseteq V$ is *k-edge-separable* if there exists a set Z of at most k edges such that S is the union of connected components of $G \setminus Z$.

It is not difficult to show that a set which meets every large k-edge-separable set is in fact a detection set.

Lemma 2.1 *If $D \subseteq V$ intersects every k-edge-separable set of size at least εn, then it is an (ε, k)-detection set with respect to edge failures.*

Proof. Let Z be a set of edges of size at most k, and let $U_1, \ldots, U_s$ be the components of $G \setminus Z$. Suppose there exist disjoint sets $A, B \subseteq V$, each of size at least εn, that are separated in $G \setminus Z$. Let us choose such a pair A, B for which $A \cup B$ is maximal subject to this condition.

Then $A = \cup_{i \in I} U_j$ and $B = \cup_{j \in J} U_j$ for disjoint index sets $I, J \subseteq \{1, 2, \ldots, s\}$. Consequently, by assumption, D contains at least one node from each of A and B, and these lie in different components of $G \setminus Z$. ■

We need the following background on the VC-dimension of set systems. Let Ω be a finite set and $\mathcal{X}$ a collection of subsets of Ω. We say that $A \subseteq \Omega$ is *shattered* by $\mathcal{X}$ if for all $B \subseteq A$ there exists an $X \in \mathcal{X}$ such that $B = A \cap X$. The *VC-dimension* of the set system $(\Omega, \mathcal{X})$ is the maximum cardinality of a subset of Ω that is shattered by $\mathcal{X}$. Here is a simple lemma, which we note for later use.

Lemma 2.2 *Let $(\Omega, \mathcal{X})$ be a set system, Ω' a subset of Ω, and $\mathcal{X}|\Omega' = \{X \cap \Omega' : X \in \mathcal{X}\}$. Then the VC-dimension of $(\Omega', \mathcal{X}|\Omega')$ is bounded by the VC-dimension of $(\Omega, \mathcal{X})$.*

We say that $N \subseteq \Omega$ is an *ε-net* with respect to $(\Omega, \mathcal{X})$ if for every $X \in \mathcal{X}$ of cardinality at least $\varepsilon|\Omega|$, the set N contains an element of X.

We now have the following theorem of Haussler and Welzl [6].

Theorem 2.3 ([6]) *There is a function*

$$f(d, \varepsilon, \delta) = O(d\varepsilon^{-1} \log(d\varepsilon^{-1}) + \varepsilon^{-1} \log \delta^{-1})$$

such that the following holds. Let $(\Omega, \mathcal{X})$ be any set system with VC-dimension at most d. Then a random subset of Ω of size $f(d, \varepsilon, \delta)$ is an ε-net with probability at least $1 - \delta$.

The crux of this theorem is the non-obvious point that there exist ε-nets for arbitrary set systems with a size that is polynomial in ε^{-1} and the VC-dimension, and independent of the cardinality of the underlying set Ω.

Taken together, Lemma 2.1 and Theorem 2.3 thus suggest that we focus on the set system $(V, \mathcal{S})$, where $\mathcal{S}$ consists of the k-edge-separable sets. Indeed, Lemma 2.1 states that an ε-net with respect to this set system will be an (ε, k)-detection set.

We bound the VC-dimension of this set system by the following type of argument: Any large set of nodes A can be "paired up" in G via edge-disjoint paths; deleting a set of k edges cannot destroy all these paths; and hence there is no k-edge-separable set whose intersection with A produces exactly one node from each group in our pairing. It then follows that A cannot be shattered.

To show how such a pairing can be achieved we use the following lemma.

Lemma 2.4 *Let $H = (V_H, E_H)$ be a connected graph, and $T \subseteq V_H$ a set of cardinality 2ℓ, for a natural number ℓ. Then there exist mutually edge-disjoint paths $P_1, P_2, \ldots, P_\ell$ such each node in T appears as one end of exactly one of these paths.*

Proof. We define a *T-join* to be a set of edges $F' \subseteq E_H$ such that in the subgraph (V_H, F'), T is the set of nodes incident to an odd number of edges [11]. Let F^* denote a T-join, and let H^* denote the graph obtained from (V_H, F^*) by adding a node v_0 to V_H, with one edge to each node in T. Every node in H^* has even degree, so H^* has an Eulerian tour. Now, if we delete v_0, this tour falls apart at the ℓ places at which it passes through v_0; as a result, we obtain ℓ edge-disjoint paths $P_1, P_2, \ldots, P_\ell$. Each node in T appears as an end of exactly one of these paths, so they satisfy the requirements of the lemma. ■

Lemma 2.5 *The VC-dimension of* $(V, \mathcal{S})$ *is at most* $2k + 1$.

Proof. Let us consider an arbitrary set $A \subseteq V$ of size $2k + 2$; we must show that it is not shattered by $(V, \mathcal{S})$. We apply Lemma 2.4, obtaining a set of mutually edge-disjoint paths $P_1, P_2, \dots, P_{k+1}$ such each node in A appears as one end of exactly one of these paths. Relabeling A if necessary, we will assume that P_i has ends equal to a_i and a_{i+k+1}.

Now we claim that $B = \{a_1, a_2, \dots, a_{k+1}\}$ cannot be written as $A \cap S$ for any k-edge-separable set S. For suppose it could, and let $Z \subseteq E$ have the property that $|Z| \leq k$ and S is a union of connected components of $G \setminus Z$. Since $|Z| < k + 1$, there is some path P_i that contains no edge in Z. But then we have $a_i \in S$, $a_{i+k+1} \notin S$, and a path connecting them in $G \setminus Z$ — this is a contradiction. ■

By Lemma 2.5 and Theorem 2.3, a small random set of nodes of V will form an ε-net of the set system $(V, \mathcal{S})$ with high probability. By Lemma 2.1, such a set system will be a detection set. Thus we have

Theorem 2.6 *A random subset of* V *of size*

$$f(2k + 1, \varepsilon, \delta) = O(k\varepsilon^{-1} \log(k\varepsilon^{-1}) + \varepsilon^{-1} \log \delta^{-1})$$

is a (ε, k)*-detection set with respect to edge failures with probability at least* $1 - \delta$.

To conclude this section, we observe that the bound in Lemma 2.5 is in fact tight.

Proposition 2.7 *There exist graphs* $G = (V, E)$ *for which the VC-dimension of the associated set system* $(V, \mathcal{S})$ *is equal to* $2k + 1$.

Proof. Since Lemma 2.5 provides an upper bound, we need only demonstrate a graph $G = (V, E)$ for which some set of $2k + 1$ nodes is shattered by k-edge-separable sets.

Let G be the star graph $K_{1,2k+1}$, and let A be the leaves of G; we claim that A is shattered by $(V, \mathcal{S})$. Indeed, let B be any subset of A. If $|B| \leq k$, then B is k-edge-separable as we may delete the set Z of all edges incident to nodes in B. Hence, $B = A \cap S$ for a set $S \in \mathcal{S}$. On the other hand, if $|B| > k$, then $|A \setminus B| \leq k$. In this case, we may delete the set Z of all edges incident to nodes in $A \setminus B$; one of the components of the resulting graph is $V \setminus (A \setminus B)$, and $B = A \cap (V \setminus (A \setminus B))$. ■

3 Detecting General Failures

Analyzing the construction of detection sets when both nodes and edges can fail is a much more complicated problem. To see some of the issues involved, we first consider what goes wrong when we try to extend the technique of the previous section in a direct way.

By analogy with the notion of k-edge-separability, we can define a set S to be *k-separable* if there exists a set Z of at most k *elements* (i.e. nodes or edges) such that S is the union of connected components of $G \setminus Z$. Now, a direct analogue of Lemma 2.1 still applies with this new definition: an ε-net with respect to k-separable sets would indeed be an (ε, k)-detection set. Unfortunately, the corresponding analogue of Lemma 2.5 now fails badly: the VC-dimension of the k-separable sets in a graph can be nearly as large as the size of the vertex set itself. For consider the star graph $K_{1,n-1}$: every subset of the leaves is 1-separable, as can be seen by deleting the singleton set Z consisting of the "center" of the star. Consequently, for this graph, the approach of the previous section will yield an upper bound for the size of (ε, k)-detection sets that is actually larger than n, the number of nodes.

However, while the star graph shows that k-separable sets can have large VC-dimension, it does not provide a counter-example to the existence of small detection sets. We simply need an approach that avoids the analysis of k-separable sets. Thus, in this section, we consider a more complicated set system defined over the vertex set V of a graph, establish a bound on the VC-dimension of this set system, and show that ε-nets for this set system yield detection sets.

We begin with a connected graph $G = (V, E)$, and a partition of V into sets V_0 of *end nodes* and V_1 of *internal nodes.* Recall that our detection set D must be a subset of V_0 (i.e. we can only place monitoring agents at end nodes), and that we are trying to detect partitions of V_0. Let $n_0 = |V_0|$ and $n_1 = |V_1|$. We will assume that $k < \frac{1}{3}\varepsilon n_0$; for otherwise, we can choose all of V_0 as our detection set and still remain within the bounds of Theorem 1.2.

We first subdivide each edge of G, adding each of the subdividing nodes to the set V_1. In this way, we may assume that any (ε, k)-partitioning set Z we consider consists only of nodes; for if it contains an edge e, we can replace e with an appropriate subdividing node in V_1, and this new set will induce the same partition of V_0.

Now, what was the underlying problem in using the collection of k-separable sets? Essentially, it allowed too many possible unions of distinct components. We need a way to reduce the combinatorial complexity of the underlying set system, by restricting the way in which we form unions of components, and at the same time keep the set system rich enough that its ε-nets form detection sets. One natural way is to impose a lexicographic order on nodes, and hence on components, and force our unions to respect this ordering. Thus, let us arbitrarily index the vertex set $V = \{v_1, v_2, \ldots, v_n\}$. Given two disjoint subsets A and B, we say that A *precedes* B *lexicographically* if the minimum index of a node in A is less than the minimum index of a node in B. Now, we say that S is a *k-segmental set* if it can be obtained in the following way:

(i) Choose a set Z of at most k nodes. (Recall that we will only consider the deletion of nodes from now on.)

(ii) Let $U_1, U_2, \ldots, U_s$ be the connected components of $G \setminus Z$, indexed in ascending lexicographic order.

(iii) For numbers p and q $(1 \le p \le q \le s)$ define $S = U_p \cup U_{p+1} \cup \cdots \cup U_q$.

Thus, briefly, S is the union of a set of connected components of $G \setminus Z$ that are consecutive in lexicographic order. We will say that Z is the *generator* of S. and that $\{U_p, \ldots, U_q\}$ are the *constituents* of S.

We first establish that the set system of k-segmental sets is related to the construction of detection sets. For a subset of nodes U, we define its *V_0-weight* to be the cardinality of $U \cap V_0$, and we denote this quantity by $\pi_0(U) = |U \cap V_0|$.

Lemma 3.1 *If $D \subseteq V$ intersects every k-segmental set of V_0-weight at least $\frac{1}{3}\varepsilon n_0$, then it is* *an (ε, k)-detection set for V_0.*

Proof. Let Z be a set of nodes of size $k' \leq k$, and let $U_1, \ldots, U_s$ be the components of $G \setminus Z$ in ascending lexicographic order. Suppose there exist disjoint sets $A, B \subseteq V \setminus Z$, each of V_0-weight at least εn_0, that are separated in $G \setminus Z$. In this case, it follows that no component U_i has V_0-weight greater than $(1-\varepsilon)n_0$.

Now, let j be the minimum index such that $\pi_0(U_1 \cup \cdots \cup U_j) > \frac{1}{3}\varepsilon n_0$; then we must have

$$\begin{aligned}\pi_0(U_1 \cup \cdots \cup U_j) &\leq \frac{1}{3}\varepsilon n_0 + (1-\varepsilon)n_0 \\ &\leq (1-\frac{2}{3}\varepsilon)n_0.\end{aligned}$$

Thus,

$$\begin{aligned}\pi_0(U_{j+1} \cup \cdots \cup U_s) &\geq n_0 - k' - (1-\frac{2}{3}\varepsilon)n_0 \\ &= \frac{2}{3}\varepsilon n_0 - k' \geq \frac{1}{3}\varepsilon n_0.\end{aligned}$$

Now, $U_1 \cup \cdots \cup U_j$ and $U_{j+1} \cup \cdots \cup U_s$ are both k-segmental sets of V_0-weight at least $\frac{1}{3}\varepsilon n_0$, and so D contains at least one node in each; and these lie in different components of $G \setminus Z$, as required. ■

Let $\mathcal{E}$ denote the collection of all k-segmental subsets of V. We wish to show that $(V, \mathcal{E})$ has bounded VC-dimension; for then Lemma 2.2 will imply that $(V_0, \mathcal{E}|V_0)$ has bounded VC-dimension as well, allowing us to use Theorem 2.3 in conjunction with Lemma 3.1. In the analogous situation in the previous section, we used an argument based on the fact that every subset of nodes of even cardinality could be partitioned into two subsets linked by edge-disjoint paths. But this approach cannot be directly applied here: in the example of $K_{1,n-1}$ we see that a graph can have a large set of nodes (the leaves) for which no two large disjoint subsets can be paired up via node-disjoint paths.

Thus linkages via node-disjoint paths behave very differently from linkages via edge-disjoint paths. To understand the obstacles to linkages via node-disjoint paths, we use a deep min-max theorem of Mader [12], in an equivalent formulation described by Robertson, Seymour, and Thomas [16]. Consider a graph $H = (V_H, E_H)$, let $L_1, L_2, \ldots, L_t$ be subsets of V_H, and let $\mathcal{L} = \{L_1, \ldots, L_t\}$. An *$\mathcal{L}$-path* is a path in H whose ends u, v have the property that $u \in L_i$ and $v \in L_j$ for indices $i \neq j$. An *$\mathcal{L}$-linkage* is a set of mutually node-disjoint $\mathcal{L}$-paths. An *$\mathcal{L}$-blocker* is a partition of V_H into sets $W, Y_1, Y_2, \ldots, Y_t$, together with sets $X_i \subseteq Y_i$, such that the following conditions hold.

(i) For each j, $L_j \subseteq W \cup (\cup_i X_i)$.

(ii) If $v \in Y_i \setminus X_i$, then all neighbors of v lie in $W \cup Y_i$.

(iii) For every $\mathcal{L}$-path P in H that does not meet W, there is an i so that some edge of P has both ends in Y_i.

Theorem 3.2 ([12, 16]) *Let $H = (V_H, E_H)$ be a graph, and $\mathcal{L} = \{L_1, \ldots, L_t\}$ with each $L_i \subseteq V_H$. There exists an $\mathcal{L}$-linkage of size ℓ if and only if there does not exist an $\mathcal{L}$-blocker* *consisting of sets $W, Y_1, Y_2, \ldots, Y_t,\ X_1, X_2, \ldots, X_t$ such that*

$$|W| + \sum_{i=1}^{t} \left\lfloor \frac{|X_i|}{2} \right\rfloor < \ell.$$

We first derive a consequence of this theorem that is slightly weaker, but easier to use in the arguments to follow.

Lemma 3.3 *Let $H = (V_H, E_H)$ be a graph, and $A \subseteq V_H$. Suppose that for a number ℓ, there do not exist $\ell + 1$ mutually node-disjoint paths with mutually distinct ends in A. Then there is a set $W' \subseteq V_H$ of size at most 3ℓ such that each node in $A \setminus W'$ belongs to a different component of $H \setminus W'$.*

Proof. Let $s = |A|$. We write $A = \{a_1, \ldots, a_s\}$, $L_i = \{a_i\}$, and $\mathcal{L} = \{L_1, L_2, \ldots, L_s\}$. Then by hypothesis, there is no $\mathcal{L}$-linkage of size $\ell + 1$; so by Theorem 3.2, there is an $\mathcal{L}$-blocker consisting of sets $W, Y_1, Y_2, \ldots, Y_t,\ X_1, X_2, \ldots, X_t$ satisfying

$$|W| + \sum_{i=1}^{t} \left\lfloor \frac{|X_i|}{2} \right\rfloor \leq \ell.$$

We say that a non-empty set X_i is *large* if $|X_i| \geq 2$, and we let W^* denote the union of all the large sets X_i. Note that there can be at most ℓ large sets X_i, and

$$\begin{aligned}
|W \cup W^*| &= |W| + \sum_{X_i \text{ large}} |X_i| \\
&\leq |W| + \left[\sum_{X_i \text{ large}} 1 + 2\left\lfloor \frac{|X_i|}{2} \right\rfloor \right] \\
&\leq \ell + 2\left[|W| + \sum_{i=1}^{t} \left\lfloor \frac{|X_i|}{2} \right\rfloor \right] \\
&\leq 3\ell.
\end{aligned}$$

We define $W' = W \cup W^*$ and claim that no two nodes of $A \setminus W'$ belong to the same component of $H \setminus W'$. For suppose there were a path P in $H \setminus W'$ with ends equal to $a_i, a_j \in A$ (with $i \neq j$), and assume without loss of generality that P does not contain any node more than once. Consider any set Y_i that intersects P. The only nodes in Y_i with neighbors outside $Y_i \cup W'$ are in X_i, and so we must $X_i \neq \phi$; but since $X_i \not\subseteq W^*$, we have $|X_i| = 1$. Let us write $X_i = \{x_i\}$.

Thus, in the order of traversal from a_i to a_j, the first node in Y_i met by P is x_i, and the last node in Y_i met by P is x_i. It follows that $|P \cap Y_i| = 1$. But this applies to any set Y_i that intersects P, and this violates condition (iii) in the definition of an $\mathcal{L}$-blocker. ■

We now analyze the VC-dimension of $(V, \mathcal{E})$ using a two-pronged argument: for any large set A, either G contains many node-disjoint paths with mutually distinct ends in A, or it contains a disconnecting set W' as in Lemma 3.3. In both cases, we argue that A cannot be shattered.

Lemma 3.4 *The VC-dimension of* $(V, \mathcal{E})$ *is at most* $O(k^3)$.

Proof. Recall that we have fixed a lexicographic ordering of V, as $v_1, v_2, \ldots, v_n$. We define $r = 3(3k+1)(k+1)^2 + 3k$ and suppose by way of contradiction that there is a set A, of size greater than r, that is shattered by $(V, \mathcal{E})$.

We first claim

(1) *There are not* $k+1$ *mutually node-disjoint paths with mutually distinct ends in* A.

For suppose there were such paths $P_1, \ldots, P_{k+1}$, such that P_i has ends $a_i, a'_i \in A$. Now, by our assumption that A is shattered, there is a k-segmental set S such that $B = \{a_1, \ldots, a_{k+1}\} = A \cap S$. Let Z be a generator of S. Then since $|Z| < k+1$, there is some path P_i that contains no node in Z. But this is a contradiction, since $a_i \in S$, $a'_i \notin S$, and P_i is a path connecting them in $G \setminus Z$.

By (1), and Lemma 3.3, there is a set W' of size at most $3k$ such that every node in $A \setminus W'$ belongs to a different component of $G \setminus W'$. For a node $v \in V \setminus W'$, we use $C(v)$ to denote the component of $G \setminus W'$ containing v, and $N(v)$ to denote the set of nodes in W' that are neighbors of at least one node in $C(v)$. We write $A' = A \setminus W' = \{a'_1, \ldots, a'_s\}$, where $s \geq r - 3k$ and the indexing need not be related to the lexicographic order on V.

The remainder of the proof is guided by the following basic idea. We will consider a subset of A' (called B'' below) such that the components of $G \setminus W'$ containing nodes of B'' are "interlaced" in lexicographic order with components containing nodes of $A' \setminus B''$. Now, for any supposed k-segmental set S so that $A \cap S = B''$, we will show that several of the components of $G \setminus W'$ are in fact constituents of S. Because of the lexicographic interlacing of the components of $G \setminus W'$, and because S must consist of components that are contiguous in lexicographic order, we will find that in order to include all nodes of B'', S must include at least one node from $A' \setminus B''$. This will contradict our assumption that $A \cap S = B''$, and hence show that A cannot be shattered.

We define a parameter $\gamma = 3(k+1)^2$. We define the *weight* w_v of a node $v \in W'$ to be the number of sets $N(a'_i)$, $a'_i \in A'$, that contain v. We say $v \in W'$ is *heavy* if $w_v \geq \gamma$; otherwise, we say it is *light*. Let W_h denote the heavy nodes in W', and W_ℓ denote the light nodes. We say a node $a'_i \in A'$ is *typical* with respect to W_h and W_ℓ if $N(a'_i) \subseteq W_h$; we say it is *atypical* otherwise.

Consider performing the following procedure. We delete all the sets $C(a'_i)$ for which a'_i is atypical; update A', the weights $\{w_v : v \in W'\}$, and the sets W_h and W_ℓ; and then iterate. We stop when there are no atypical nodes remaining. We claim

(2) *At most* $3k\gamma$ *nodes of* A' *are deleted by the above procedure.*

We prove this by the following counting argument. First, observe that for any node $v \in W'$, there is at most one iteration in which it is light, but has positive weight. Indeed, the weight of v never increases; and in the first iteration in which it is declared light, all sets $C(a'_i)$ for which $v \in N(a'_i)$ are deleted, and so the weight of v is reduced to zero. Now, consider any node $a'_i \in A'$ that is deleted. In the iteration in which it is deleted, there is some $v \in W_\ell \cap N(a'_i)$; we *charge* a'_i to v. A node $v \in W'$ can only be charged in the single iteration when it is light but has positive weight; and in this iteration, at most γ nodes in A' can be charged to it. Since $|W'| \leq 3k$, (2) follows.

Thus, when this deletion procedure terminates, we have a final partition of W' into heavy and light nodes, which we will refer to as W_h and W_ℓ for the remainder of the proof. We define $A'' \subseteq A'$ to be the set of nodes not deleted from A' by the deletion procedure; by (2) we know that there are at least $s - 3k\gamma$ such nodes. We also know that all nodes in A'' are typical with respect to W_h and W_ℓ, and that for each node $v \in W_h$, there are at least γ nodes $a'' \in A''$ for which $v \in N(a'')$.

We now partition A'' into sets B and B' as follows. For each node $v \in W_h$, we select nodes $b_{v,1}, \ldots, b_{v,k+1} \in A''$ with the property that $v \in N(b_{v,i})$, and place them in B. We define $B' = A'' \setminus B$. Since $|B| \leq 3k(k+1)$, $|B'| = t \geq s - 3k\gamma - 3k(k+1) = 3(k+1)$. We index $B' = \{b'_1, b'_2, \ldots, b'_t\}$ in ascending lexicographic order of the sets $\{C(b'_i)\}$.

Now, partition B' into $B'_1 = \{b'_1, b'_2, \ldots, b'_{k+1}\}$, $B'_2 = \{b'_{k+2}, \ldots, b'_{2k+2}\}$, and $B'_3 = \{b'_{2k+3}, \ldots, b'_t\}$. Note that each of B'_1, B'_2, and B'_3 has at least $k+1$ elements. We define $B'' = B'_1 \cup B'_3$, and claim

(3) *There is no k-segmental set $S \subseteq V$ so that $B'' = A \cap S$.*

This will yield a contradiction to our assumption that A is shattered, and hence prove the lemma.

To prove (3), suppose there were such a set S, with generator Z. Let $U_1, U_2, \ldots, U_s$ denote the components of $G \setminus Z$ in lexicographic order; for some indices p and q, we have $S = U_p \cup U_{p+1} \cup \cdots \cup U_q$. First we claim

(3.1) *Each node of B'' belongs to a different constituent of S.*

For suppose not, and let $b'_i, b'_j \in B''$ belong to the same constituent U_ℓ. Then there is a path $P \subseteq U_\ell$ with ends equal to b'_i and b'_j. Since b'_i and b'_j belong to distinct components of $G \setminus W'$, there is a node $v \in W' \cap P$, and hence in $W' \cap U_\ell$. Now consider the nodes $b_{v,1}, \ldots, b_{v,k+1}$ that were placed in B because $v \in N(b_{v,i})$ for $i = 1, 2, \ldots, k+1$. For each such node $b_{v,i}$, there is a path Q_i with ends equal to v and $b_{v,i}$ such that $W' \cap Q_i = \{v\}$. But $v \notin Z$, and $|Z| < k+1$; thus, there is some Q_i that does not meet Z. It follows that the corresponding node $b_{v,i}$ belongs to U_ℓ, and hence belongs to S. But $b_{v,i} \notin B''$, contradicting our assumption that $B'' = A \cap S$. This proves (3.1).

Next, we claim

(3.2) $W_h \subseteq Z$.

Again, suppose not, and say that $v \in W_h \setminus Z$. Now, A'' contains at least γ nodes a'' for which $v \in N(a'')$. We have $|B| \leq 3k(k+1)$ and $|B'_2| = k+1$; thus $B'' = A'' \setminus (B \cup B'_2)$ contains at least $\gamma - (3k+1)(k+1) \geq k+2$ nodes b'' for which $v \in N(b'')$. Let B''_v denote this set of nodes in B''. Since $|Z| \leq |B''_v| - 2$, there exist two nodes $b'_v, b''_v \in B''_v$ such that $Z \cap (C(b'_v) \cup C(b''_v)) = \phi$. But this implies there is a path in $G \setminus Z$ with ends equal to b'_v and b''_v, contradicting (3.1); this proves (3.2)

Since A'' consists entirely of typical nodes, the following is a direct consequence of (3.2).

(3.3) *For each $a'' \in A''$, we have $N(a'') \subseteq Z$.*

Now, since B'_1, B'_2, and B'_3 each have at least $k+1$ nodes, there exist nodes $b'_i \in B'_1$, $b'_j \in B'_2$, and $b'_\ell \in B'_3$ such that $Z \cap (C(b'_i) \cup C(b'_j) \cup C(b'_k)) = \phi$. Combined with (3.3), this implies

(3.4) *Each of $C(b_i')$, $C(b_j')$, and $C(b_\ell')$ is a component of $G \setminus Z$.*

Finally, recall that we indexed $B' = \{b_1', b_2', \dots, b_t'\}$ in ascending lexicographic order of the sets $C(b_1'), \dots, C(b_t')$. But since $b_i', b_\ell' \in B''$, (3.4) implies that $C(b_i')$ and $C(b_\ell')$ are constituents of S. Now, $C(b_j')$ is a component of $G \setminus Z$ that comes between $C(b_i')$ and $C(b_\ell')$ lexicographically; since S is k-segmental, it follows that $C(b_j') \subseteq S$ and hence $b_j' \in S$. But $b_j' \notin B''$, contradicting our assumption that $B'' = A \cap S$. This contradiction establishes (3), and hence the lemma. ■

Finally, we prove the theorem described initially.

Theorem 1.2 *For every graph $G = (V, E)$, and any set $V_0 \subseteq V$, there is an (ε, k)-detection set for V_0 of size $O(k^3\varepsilon^{-1}\log(k\varepsilon^{-1}))$. Moreover, a set of*

$$O(k^3\varepsilon^{-1}\log(k\varepsilon^{-1}) + \varepsilon^{-1}\log\delta^{-1})$$

nodes chosen uniformly at random from V_0 is an (ε, k)-detection set for V_0 with probability at least $1 - \delta$.

Proof. By Lemmas 2.2 and 3.4, the VC-dimension of $(V_0, \mathcal{E}|V_0)$ is $O(k^3)$. Thus, by Theorem 2.3, a random set of $O(k^3\varepsilon^{-1}\log(k\varepsilon^{-1}) + \varepsilon^{-1}\log\delta^{-1})$ nodes chosen uniformly at random from V_0 is a $\frac{1}{3}\varepsilon$-net for this set system with probability at least $1 - \delta$.

By definition, such a $\frac{1}{3}\varepsilon$-net D has the property that for any $S \in \mathcal{E}$, if $|S \cap V_0| \geq \frac{1}{3}\varepsilon n_0$, then $S \cap V_0$ contains an element of D. Thus, D intersects every k-segmental set of V_0-weight at least $\frac{1}{3}\varepsilon n_0$, and so by Lemma 3.1, it is an (ε, k)-detection set for V_0. ■

4 Further Extensions

Using our analysis of detection sets, we can extend our results to a setting in which nodes and edges can fail according to an unknown (adversarially chosen) probability distribution.

We consider a model in which an adversary assigns a failure probability p_j to each element j of G, subject only to the condition that the mean number of failures $\sum_{j \in V \cup E} p_j$ is bounded by a number k. Note that since an adversary could assign a failure probability of 1 to k elements, and 0 to all others, this model clearly includes our initial adversarial model. Using Markov's inequality in conjunction with Theorem 1.2, we can show that a random subset D of V_0 of size $O(k^3\varepsilon^{-1}\delta^{-3}\log(k\varepsilon^{-1}\delta^{-1}))$ has the following property with high probability: If elements fail according to any assignment of probabilities $\{p_j\}$ as above, the probability that two subsets of V_0 of size $\varepsilon|V_0|$ become separated but all nodes of D remain connected is at most δ. In other words, regardless of the distribution of failures in this model (assuming only that they have bounded mean), there is only a small probability that our set D will fail to "witness" any significant partition of the end nodes.

References

[1] N. Alon, J. Spencer, *The Probabilistic Method*, Wiley, 1992.

[2] C. Berge, "Sur le couplage maximum d'un graphe," *C.R. Acad. Sci. Paris* 247(1958), pp. 258–259.

[3] K. Bharat, A. Broder, "A technique for measuring the relative size and overlap of public Web search engines," *Proc. 7th Intl. WWW Conference*, 1998.

[4] F. Chung, M. Garrett, R. Graham, D. Shallcross, "Distance realization problems with applications to Internet tomography," preprint, http://www.math.ucsd.edu/∼fan.

[5] K. Claffy, T. Monk, D. McRobb, "Internet Tomography," *Nature, Web Matters*, Jan. 7, 1999.

[6] D. Haussler, E. Welzl, "ε-nets and simplex range queries," *Discrete and Computational Geom.* 2(1987), pp. 127–151.

[7] D. Karger, "A randomized fully polynomial time approximation scheme for the all-terminal network reliability problem," *Proc. ACM Symp. on Theory of Computing*, 1995.

[8] Keynote Systems. http://www.keynote.com.

[9] Keynote Systems. "Data Accuracy and Statistical Methodology." White paper at www.keynote.com/downloads/wp_statmethod.doc.

[10] S. Lawrence, C.L. Giles, "Accessibility of Information on the Web," *Nature* 400(1999), pp. 107–109.

[11] L. Lovász, M. Plummer, *Matching Theory.* North-Holland, 1986.

[12] W. Mader, "Über die Maximalzahl kreuzungsfreier H-Wege," *Arch. Math.* 31(1978), pp. 387–402.

[13] J. Mahdavi, V. Paxson, "IPPM Metrics for Measuring Connectivity." RFC 2678, Network Working Group, Sept. 1999.

[14] V. Paxson, "Towards a Framework for Defining Internet Performance Metrics. *Proc. INET* 1996, at http://www.aciri.org/vern/papers.html.

[15] S. Ratnasamy and S. McCanne, "Inference of Multicast Routing Trees and Bottleneck Bandwidths using End-to-end Measurements," *Proc. IEEE INFOCOM*, 1999.

[16] N. Robertson, P.D. Seymour, R. Thomas, "Hadwiger's conjecture for K_6-free graphs," *Combinatorica* 13(1993), pp. 279–361.

[17] R. Siamwalla, R. Sharma, and S. Keshav, "Discovering Internet Topology," July 1998, at http://www.cs.cornell.edu/skeshav/papers.html.

[18] V.N. Vapnik, A.Y. Chervonenkis, "On the uniform convergence of relative frequencies of events to their probabilities," *Theory of Prob. App.*, 16(1971), pp. 264–280.