

Fundamental Limits of Failure Identifiability by Boolean Network Tomography

N. Bartolini, T. He, and H. Khamfroush,

Abstract—Boolean network tomography is a powerful tool to infer the state (working/failed) of individual nodes from path-level measurements obtained by edge-nodes. We consider the problem of optimizing the capability of identifying network failures through the design of monitoring schemes. Finding an optimal solution is NP-hard and a large body of work has been devoted to heuristic approaches providing lower bounds. Unlike previous works, we provide upper bounds on the maximum number of identifiable nodes, given the number of monitoring paths and different constraints on the network topology, the routing scheme, and the maximum path length. The proposed upper bounds represent a fundamental limit on the identifiability of failures via Boolean network tomography. This analysis provides insights on how to design topologies and related monitoring schemes to achieve the maximum identifiability under various network settings. Through analysis and experiments we demonstrate the tightness of the bounds and efficacy of the design insights for engineered as well as real networks.

I. INTRODUCTION AND MOTIVATION

The capability to assess the states of network nodes in the presence of failures is fundamental for many functions in network management, including performance analysis, route selection, and network recovery. In modern networks, the traditional approach of relying on built-in mechanism to detect node failures is no longer sufficient, as bugs and configuration errors in various customer software and network functions often induce “silent failures” that are only detectable from end-to-end connection states [1]. *Boolean network tomography* [2] is a powerful tool to infer the states of individual nodes of a network from binary measurements taken along selected paths. We consider the problem of Boolean network tomography in the framework of group testing [3], [4]. The classic group testing studies the following problem: given a set S of n items, at most d of which are defective, the goal is to identify the defective items through binary measurements taken on subsets $S_i \subseteq S$ of items ($i = 1, \dots, m$). The problem of Boolean network tomography is analogous in that it is also about determining defective items (failed nodes) of a large set S (all the nodes) by performing binary measurements over subsets of items (monitoring paths), where a measurement fails if and only if at least one of the measured nodes has failed.

A significant difference from conventional group testing is that in Boolean network tomography, the subsets used for

tests cannot be designed arbitrarily, but must conform to the structure of the network. In this regard, Cheraghchi et al. have studied graph-constrained group testing in [5], where testing items correspond to either edges or vertices of a graph $\mathcal{G} = (V, E)$. The tests are represented through a so-called *testing matrix* T , which is an $m \times n$ binary matrix, where each row represents a path, each column a vertex/edge, and the (i, j) -th entry an indicator that the j -th vertex/edge belongs to the i -th path. Given the maximum number of defective items, the goal of [5] is to design a testing matrix with the minimum number of rows that guarantees the identification of all the defective items, under the assumption that each monitoring path can span any connected subgraph of \mathcal{G} . In practice, however, the monitoring paths are constrained not only by the network topology, but also by the routing scheme in the network and the endpoints of the probes.

In our work, we tackle the problem of maximizing the number of nodes whose states can be uniquely determined from binary measurements on a given number of monitoring paths. We consider the problem under increasing constraints about the network, from the routing scheme to the endpoints of monitoring paths. Due to the inherent hardness in computing the exact maximum value, we focus on deriving easily computable upper bounds based on the structure of the testing matrix. The bounds allow us to: (i) evaluate the room of improvement for a given monitoring scheme in a specific network setting, and (ii) extract rules for network design to maximize the number of identifiable nodes in a general setting. Although our bounds are derived for the case of a single failure, we show that the bounds are also valid (but looser) for the general case of at most $k \geq 1$ simultaneous failures.

The main contributions of this work are the following:

- We upper-bound the maximum number of identifiable nodes with a given number of monitoring paths, in the following scenarios: (1) paths between arbitrary nodes under arbitrary routing (Theorem IV.1); (2) paths between arbitrary nodes under consistent routing (Theorem IV.2); (3) paths from a single server to multiple clients under consistent routing (Theorem V.1); (4) paths from multiple servers to multiple clients with fixed/flexible assignment under consistent routing (Theorems V.2 and V.3); (5) paths between arbitrary nodes under partially consistent routing (Theorem VI.1).
- We give insights on the design of topologies and monitoring schemes to approximate the bounds, grounded upon the bound analysis.
- We demonstrate the tightness of the upper bounds by comparisons with the results of known heuristics [6] on engineered as well as real network topologies.

N. Bartolini is with the Department of Computer Science, Sapienza University of Rome, Italy. E-mail: bartolini@di.uniroma1.it

T. He and H. Khamfroush are with the Department of Computer Science and Engineering, Pennsylvania State University, USA. E-mail: {tzh58,hxk5299}@cse.psu.edu.

This work was supported by the Defense Threat Reduction Agency under the grant HDTRA1-10-1-0085, and by NATO under the SPS grant G4936 SONICS.

- We compare the bounds in different scenarios to evaluate the impact of the routing scheme, the number of monitoring paths, and the maximum path length on the number of identifiable nodes.

II. RELATED WORK

Pioneered by Duffield [2], Boolean network tomography has direct applications in network failure localization. The early works focused on best-effort inference. For example, Duffield et al. [2], [7] and Kompella et al. [1] aimed at finding the minimum set of failures that can explain the observed measurements, and Nguyen et al. [8] aimed at finding the most likely failure set that explains the observations. Later, the identifiability problem attracted attention. Ma et al. characterized in [9] the maximum number of simultaneous failures that can be uniquely localized, and then extended the results in [10] to characterize the maximum number of failures under which the states of specified nodes can be uniquely identified as well as the number of nodes whose states can be identified under a given number of failures.

The related optimization problems have also been studied. The problem of optimally placing monitors to detect failed nodes via round-trip probing was introduced and proven to be NP-hard by Bejerano et al. in [11]. The work by Cheraghchi et al. [5] aimed at determining the minimum number of monitoring paths to uniquely localize a given number of failures, under the assumption that any path can be monitored. For monitoring paths that start/end at monitors, Ma et al. [12] proposed polynomial time heuristics to deploy a minimum number of monitors to uniquely localize a given number of failures under various routing constraints. When monitoring is performed at the service layer, He et al. [6] proposed service placement algorithms to maximize the number of identifiable nodes by monitoring the paths connecting clients and servers.

Our work also addresses the problem of maximizing the number of identifiable nodes under failures. Unlike previous work, we aim at establishing upper bounds based on general information such as the number/length of monitoring paths, the type of routing scheme, and constraints on the path endpoints. Besides the theoretical value, our results also provide guidelines for network design to facilitate network monitoring.

III. PROBLEM FORMULATION

We use lower-case letters to denote scalars and vectors and upper-case letters to denote matrices. For a vector p , $p|_i$ denotes the i -th element in the vector. For a matrix M , $M|_{i,j}$ denotes the element in the i -th row and j -th column; moreover, $M|_{i,*}$ denotes the i -th row and $M|_{*,j}$ the j -th column of M .

A. Network Model

We model the network as an undirected graph $\mathcal{G} = (V, E)$, where V is a set of n nodes, and E is the set of links.

Each node may be in normal or failed state. Without loss of generality, we assume that links do not fail, as link failures can be modeled by the failures of logical nodes that represent the links. The set of *all* failed nodes, denoted by $F \subseteq V$, defines the state of a network, and is called *failure set*.

B. Observation Model

We assume that node states cannot be measured directly, but only indirectly via *monitoring paths*. Let $P = \{p_1, p_2, \dots, p_m\}$ be a given set of m monitoring paths. According to the needs of the discussion, each path $p_i \in P$ is represented as either a *set* of nodes p_i , or as an ordered *sequence* of nodes \hat{p}_i , from one endpoint to the other. The state of a path is normal if and only if all traversed nodes (including endpoints) are in normal state. We call the *incident set* of v_i the set of paths affected by the failure of node v_i and denote it with P_{v_i} . We also denote the incident set of paths of a failure set F with $P_F \triangleq \cup_{v_i \in F} P_{v_i}$.

The *testing matrix* T is an $m \times n$ matrix, where $T|_{i,j} = 1$ if $v_j \in p_i$, and zero otherwise. The j -th column of T , denoted with $b(v_j) \triangleq T|_{*,j}$, is the characteristic vector¹ of P_{v_j} . The transpose of $b(v_j)$ is hereby called the *binary encoding* of v_j . Note that multiple nodes may have the same binary encoding.

C. Identifiability

The concept of identifiability refers to the capability of inferring the states of individual nodes from the states of the monitoring paths. Informally, we say that a node v is 1-identifiable, given a set of paths P , if its failure and the failure of any other node w cause the failure of different sets of monitoring paths in P , i.e. v and w have different incident sets. This concept can be extended to the case of concurrent failures of at most k nodes, where a node is k -identifiable in P if any two sets of failures F_1 and F_2 of size at most k , which differ at least in v (i.e., one contains v and the other does not), cause the failures of different monitoring paths in P , i.e. F_1 and F_2 have different incident sets.

He et al. in [6] formalized the concept of k -identifiability that we reformulate as follows:

Definition III.1. *Given a set of monitoring paths P and a node $v_j \in V$, v_j is k -identifiable with respect to (wrt) P if for any failure sets F_1 and F_2 such that $F_1 \cap \{v_j\} \neq F_2 \cap \{v_j\}$, and $|F_i| \leq k$ ($i \in \{1, 2\}$),*

$$\bigvee_{v_i \in F_1} b(v_i) \neq \bigvee_{v_z \in F_2} b(v_z)$$

where with " \bigvee " we refer to the element-wise logical OR.

In the special case of $k = 1$, Definition III.1 implies the following Lemma.

Lemma III.1. *A node v_i is 1-identifiable wrt P if and only if $b(v_i) \neq \mathbf{0}$, and $\forall v_j \neq v_i$, $b(v_j) \neq b(v_i)$, i.e., its binary encoding is not null and not identical with that of any other node.*

D. Bounding Identifiability

The set of monitoring paths P is usually the result of design choices related to topology, monitoring endpoints, routing scheme, etc. Given a collection of candidate path sets² \mathcal{P} under

¹A *characteristic vector* of a subset S of an ordered set of n elements $V = \{v_1, v_2, \dots, v_n\}$ is a binary vector with '1' only in the positions of the elements of V that are included in S .

²For example, \mathcal{P} may be the class of path sets of given cardinality, or paths of a given length.

all possible designs, the question is: how well can we monitor the network using path measurements and which design is the best? Using the notion of k -identifiability, we can measure the monitoring performance by the number of nodes that are k -identifiable wrt P , denoted by $\phi_k(P)$, and formulate this question as an optimization: $\psi_k(P) \triangleq \max_{P \in \mathcal{P}} \phi_k(P)$.

Although extensively studied [11], [5], [12], [6], the optimal solution is hard to obtain due to the (exponentially) large size of \mathcal{P} , and heuristics are used to provide lower bounds. There is, however, a lack of general upper bounds. In this work we establish upper bounds on $\psi_k(P)$ in representative scenarios. Knowledge of these upper bounds is key to understanding the fundamental limits of Boolean network tomography, and gives insights on network design to facilitate network monitoring.

We observe that if a node v_i is k -identifiable wrt P for any $k \geq 1$, then v_i is also 1-identifiable wrt P , which implies that an upper bound on the maximum number of nodes that are 1-identifiable is also an upper bound on the maximum number of nodes that are k -identifiable, as stated below.

Lemma III.2. *For any $k \geq 1$ and any candidate sets of monitoring paths \mathcal{P} , $\psi_1(\mathcal{P}) \geq \psi_k(\mathcal{P})$.*

Proof. Given the optimal choice $P^* \in \mathcal{P}$ achieving $\psi_k(\mathcal{P})$, we have $\psi_1(\mathcal{P}) \geq \phi_1(P^*) \geq \phi_k(P^*) = \psi_k(\mathcal{P})$, where the first inequality is by definition of $\psi_1(\mathcal{P})$ and the second inequality is by Definition III.1. \square

Therefore, in the sequel, we focus on bounding $\psi_1(\mathcal{P})$, simply denoted by $\psi(\mathcal{P})$, where we will replace \mathcal{P} by specific parameters in each network setting. We hereafter shortly call the 1-identifiable nodes ‘‘identifiable’’.

IV. GENERAL NETWORK MONITORING

We initially consider a generic network with a given number of monitoring paths between any nodes. We analyze two cases: (i) arbitrary routing, and (ii) consistent routing.

A. Arbitrary routing

1) *Identifiability Bound:* Under no constraints on the routing scheme, it holds the following bound.

Theorem IV.1 (Identifiability under arbitrary routing). *Given a network with n nodes and m monitoring paths, the maximum number of identifiable nodes under arbitrary routing satisfies*

$$\psi^{\text{AR}}(m, n) \leq \min\{n; 2^m - 1\}.$$

Proof. By Lemma III.1, every identifiable node v has a different encoding $b(v) \neq \mathbf{0}$. The maximum number of different encodings with m digits, excluding $\mathbf{0}$, is $2^m - 1$. Therefore, the smaller number between $2^m - 1$ and the total number of nodes n bounds the maximum number of identifiable nodes. \square

2) *Tightness of the bound and design insights:* The bound in Theorem IV.1 is tight, as we can construct a topology with m monitoring paths that meets this bound. Given n nodes, we label $\min\{n; 2^m - 1\}$ of them with m -digit binary encodings, excluding $\mathbf{0}$, such that all the encodings are distinct. Then we generate m paths such that path \hat{p}_i ($i = 1, \dots, m$) is one of the

possible sequences of all the nodes v for which $b(v)|_i = 1$. Finally, we connect every two nodes v and w with a link whenever there is a path containing the sub-sequence v, w . An example of this construction is shown in Figure 1, for $m = 3$ and $n = 7$.




Fig. 1. Network meeting the bound of Theorem IV.1 (arbitrary routing).

B. Consistent routing

1) *Identifiability Bound:* In the sequel, we assume that paths satisfy the following property of *routing consistency*.

Definition IV.1. *A set of paths P is consistent if $\forall p, p' \in P$ and any two nodes u and v traversed by both paths (if any), p and p' follow the same sub-path between u and v .*

We remark that routing consistency is satisfied by many practical routing protocols, including but not limited to shortest path routing (where ties are broken arbitrarily but deterministically). Note that this property implies that paths are cycle-free.

We define the *path matrix* of \hat{p}_i as a binary matrix $M(\hat{p}_i)$, in which each row is the binary encoding of a node on the path, and rows are sorted according to the sequence \hat{p}_i . Notice that by definition $M(\hat{p}_i)|_{*,i}$ has only ones, i.e., $M(\hat{p}_i)|_{r,i} = 1, \forall r$.

Lemma IV.1. *Under the assumption of consistent routing, if any two different rows of the matrix $M(\hat{p}_i)$ are equal, then the corresponding nodes are not 1-identifiable.*

Proof. Under consistent routing, the path \hat{p}_i cannot contain any cycle, so every row of $M(\hat{p}_i)$ corresponds to a different node. If two different nodes have the same binary encoding, then by Lemma III.1, the two nodes are not identifiable. \square

Definition IV.2. *A column $M(p)|_{*,k}$ ($k = 1, \dots, m$) of a path matrix $M(p)$ has consecutive ones if all the ‘‘1’’s appear in consecutive rows, i.e., for any two rows i and j ($i < j$), if $M(p)|_{i,k} = M(p)|_{j,k} = 1$, then $M(p)|_{h,k} = 1$ for all $i \leq h \leq j$.*

Lemma IV.2. *Under the assumption of consistent routing, all the columns in all the path matrices have consecutive ones.*

Proof. The assertion is true for $M(\hat{p}_i)|_{*,i}$ since it contains only ones. Let us consider column $M(\hat{p}_i)|_{*,j}$, with $j \neq i$. Assume by contradiction that there are two rows $k_1 < k_2$ s.t. $M(\hat{p}_i)|_{k_1,j} = M(\hat{p}_i)|_{k_2,j} = 1$ but there is a row h with $k_1 < h < k_2$ for which $M(\hat{p}_i)|_{h,j} = 0$. Let v_1, v_2 , and v_h be the nodes with encodings $M(\hat{p}_i)|_{k_1,*}$, $M(\hat{p}_i)|_{k_2,*}$, and $M(\hat{p}_i)|_{h,*}$, respectively. Then the paths \hat{p}_i and \hat{p}_j traverse both nodes v_1 and v_2 following different paths, of which only \hat{p}_i traverses node v_h , in contradiction with consistent routing. \square

Lemma IV.3. *Given $m = |P| > 1$ consistent routing paths, whose length is at most d^* (in number of nodes), the maximum*

number of different encodings in the rows of $M(\hat{p}_i)$ is equal to $\min\{2 \cdot (m - 1), d^*\}$, $\forall p_i \in P$.

Proof. While the number of rows of $M(\hat{p}_i)$ is bounded by the maximum length d^* , the number of different encodings can be lower. Notice that first, column $M(\hat{p}_i)|_{*,i}$ contains only ones, second, for any column $M(\hat{p}_i)|_{*,j}$ with $j \neq i$, it holds, by Lemma IV.2, that it has a consecutive ones.

We say that column k has a *flip* in row r if $M(\hat{p}_i)|_{r-1,k} \neq M(\hat{p}_i)|_{r,k}$. Due to Lemma IV.2 any column of $M(\hat{p}_i)$ can have up to two flips or it would create a fragmented sequence of ones. In order to have a change in the encoding contained in any two successive rows $r-1$ and r of the matrix $M(\hat{p}_i)$, i.e., $M(\hat{p}_i)|_{r-1,*} \neq M(\hat{p}_i)|_{r,*}$, there must be at least a column that flips in r . The number of columns that can flip is $m-1$ and each of them can flip at most two times. The number of different rows that can be observed in $M(\hat{p}_i)$ is therefore upper-bounded by the smallest between the maximum path length d^* and $2 \cdot (m-1)$. \square




Fig. 2. Consistent paths identifying all nodes of the network.

Example: Figure 2 shows a case of consistent routing. The four path matrices have columns with consecutive ones and each column flips at most twice, so the number of different rows is less than $2 \cdot (m-1) = 6$. For instance, $M(\hat{p}_3)$ is:

$$M(\hat{p}_3) = \begin{array}{c|cccc} \text{flips} & b_1 & b_2 & b_3 & b_4 \\ \hline 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 1 & 0 \\ 2 & 1 & 1 & 1 & 0 \\ 3 & 1 & 0 & 0 & 1 \\ 4 & 0 & 0 & 1 & 1 \end{array}$$

We now give an upper bound on the number of identifiable nodes under consistent routing.

Theorem IV.2 (Identifiability with consistent routing). *Given n nodes, and $m > 1$ consistent routing paths of length at most d^* (in number of nodes), the maximum number of identifiable nodes satisfies:*

$$\psi^{\text{CR}}(m, n, d^*) \leq \min \left\{ \sum_{i=1}^{i_{\max}} \binom{m}{i} + \left\lfloor \frac{N_{\max} - \sum_{i=1}^{i_{\max}} i \cdot \binom{m}{i}}{i_{\max} + 1} \right\rfloor; n \right\},$$

where $i_{\max} = \max\{k \mid \sum_{i=1}^k i \cdot \binom{m}{i} \leq N_{\max}\}$, and $N_{\max} = m \cdot \min\{2 \cdot (m-1); d^*\}$.

Proof. By Lemma III.1, each identifiable node must have a unique encoding. By Lemma IV.3, we can define an upper bound on the number of different node encodings in the path matrices $M(\hat{p}_i)$, $i = 1, \dots, m$ as follows: $N_{\max} = m \cdot \min\{2 \cdot (m-1); d^*\}$. Nevertheless, in the above value of N_{\max} we are counting multiple times the nodes that appear in multiple path matrices. In fact, if an encoding b has k digits equal to 1, then b appears among the rows of k different path matrices.




Fig. 3. An example of semi-grid graph

The number of distinct encodings is maximized when we minimize the number of duplicate encodings and therefore their number of ones. This is achieved when we have $\binom{m}{1}$ different encodings with only one digit equal to 1, $2 \cdot \binom{m}{2}$ with only two digits equal to 1 (appearing in two path matrices), and so forth, until the total number of encodings (counting the duplicates) is N_{\max} .

More formally, let $i_{\max} = \max\{k \mid \sum_{i=1}^k i \cdot \binom{m}{i} \leq N_{\max}\}$. For each $i \leq i_{\max}$, we have $\binom{m}{i}$ encodings containing i digits equal to 1s and appearing in i path matrices. Considering that the remaining $N_{\max} - \sum_{i=1}^{i_{\max}} i \cdot \binom{m}{i}$ encodings will have at least $i_{\max} + 1$ digits equal to 1 and thus appear at least $i_{\max} + 1$ times, the number of distinct encodings out of the N_{\max} encodings is upper bounded by:

$$\psi^{\text{CR}}(m, n, d^*) \leq \sum_{i=1}^{i_{\max}} \binom{m}{i} + \left\lfloor \frac{N_{\max} - \sum_{i=1}^{i_{\max}} i \cdot \binom{m}{i}}{i_{\max} + 1} \right\rfloor.$$

Considering also the bound of n we have the final bound. \square

2) *Tightness of the bound and design insights:* For certain values of m , n and d^* , it is possible to design a network topology that achieves the bound of Theorem IV.2. The construction is suggested by the proof of the theorem, and consists of creating a topology and routing scheme with the maximum number of different binary encodings and the minimum number of 1s. In the following we show such a construction when $n = 36$, $m = 8$, and $d^* = 8$. In this case, $\psi^{\text{CR}} = \binom{8}{1} + \binom{8}{2} + \lfloor \frac{0}{3} \rfloor = 36$. We can obtain this maximum identifiability by intersecting the paths so as to generate the maximum number of dangling nodes traversed by a single path, and all possible intersections of two paths. A topology that meets the above bound is the *half-grid* in Figure 3. All the nodes are identifiable with just 8 monitoring paths with the following pairs of endpoints: (c_1, h_1) , (c_2, h_1) , (c_3, h_2) , (c_4, h_3) , (c_5, h_4) , (c_6, h_5) , (c_7, h_6) , and (c_8, h_7) and routing as in Figure 3. The above construction can be generalized to the case of any m provided that $d^* = m$ and $n = m(m+1)/2$. We show in Figure 10 that in a more general setting, this topology and routing scheme do not always meet the bound of Theorem IV.2, but approximate it closely.

V. SERVICE NETWORK MONITORING

We consider a service network where we monitor paths between clients and servers, under consistent routing in the case of (i) single-server and (ii) multi-server monitoring.

A. Single-Server Monitoring

1) *Identifiability Bound:* Consider the scenario where a single server communicates with multiple clients and we can only monitor the paths in between. The number of paths m

coincides with the number of clients, and all the monitoring paths must share a common endpoint (the server).

We start by showing the special structure of the topology spanned by the monitoring paths.

Lemma V.1. *Under consistent routing, any monitoring paths with a common endpoint r must form a tree rooted at r .*

Proof. We consider any two paths p_i and p_j . Starting from r , the next hops on these paths lead to either a common node or two different nodes. In the latter case, the two paths cannot intersect at any subsequent node v , as otherwise the two path segments from r to v following paths p_i and p_j would violate routing consistency. As this is true for all the paths, the paths must form a tree rooted at r . \square

Given the number of paths m with maximum path length d^* , we define the *optimal monitoring tree* as a tree with m leaves and maximum depth³ $d^* - 1$ that has the maximum number of identifiable nodes when its root-to-leaf paths are monitored.

Lemma V.2. *If the maximum path length d^* satisfies $d^* \geq \lceil \log_2 m \rceil + 1$, the optimal monitoring tree is a full binary tree with m leaves⁴. If $d^* < \lceil \log_2 m \rceil + 1$, then the optimal monitoring tree is a tree composed of $\lfloor \frac{m}{2^{(d^*-2)}} \rfloor$ perfect binary trees⁵ with depth $d^* - 2$, and up to one full binary tree with depth at most $d^* - 2$ and $m \bmod 2^{(d^*-2)}$ leaves, connected to a common root.*

Proof. Let us first consider the case of unbounded path length. By contradiction, assume the existence of an optimal monitoring tree that is not a full binary tree. Such a tree must have at least a node u whose number of children is either (a) strictly greater than two or it is (b) exactly one.

If (a), u has at least three children v_1, v_2 and v_3 . Let p_1, p_2 and p_3 be the paths from these nodes to u , as in Figure 4. We can build a new graph with an additional identifiable node x as in Figure 5, by removing the links between u and v_1, v_2 and adding x as a parent of v_1 and v_2 and children of u . Node x is identifiable as its encoding is different from the encodings of the leaves v_1, v_2 and v_3 and of the root u . If (b), u has




Fig. 4. Three children tree




Fig. 5. Full binary tree

only one child v , as shown in Figure 6. If v is not traversed by any path, or all the paths traversing u also traverse v , then node v is not identifiable, and the removal of v from the tree would not decrease the identifiability. If instead there is a path p_1 traversing both u and v , and a path p_2 traversing u which ends before reaching node v , then path p_2 can be prolonged to traverse a new node x added as a child of node u to increase the identifiability of the topology, as shown in Figure 7.

³The *depth of a tree* is the maximum distance from the root to any leaf, in number of links.

⁴We recall that a *full binary tree* is a binary tree where each node is either a leaf or it has exactly two children.

⁵We also recall that a *perfect binary tree* is a full binary tree where all leaves are at the same distance from the root.




Fig. 6. One child tree




Fig. 7. Full binary tree

Notice that as long as the maximum path length is $d^* \geq \lceil \log_2 m \rceil + 1$ we can apply the previous discussion and build an optimal full binary tree with up to m leaves and depth $\lceil \log_2 m \rceil + 1$ (maximum distance from the root to the leaves, in number of nodes). If instead $d^* < \lceil \log_2 m \rceil + 1$, the largest number of leaves that can be obtained in a full binary tree topology with depth $d^* - 1$ is 2^{d^*-1} which is lower than the number of paths m . Therefore, in such a case, the maximum identifiability is obtained by creating the maximum number $\lfloor \frac{m}{2^{(d^*-2)}} \rfloor$ of perfect binary trees of depth $d^* - 2$ and up to one full binary tree (not perfect) with depth at most $d^* - 2$, connecting them to a same root, thus ensuring that the number of nodes with either no children or two only children is maximized. \square

Example: Figure 8(a) shows an optimal monitoring tree for $m = 7$ and $d^* = 4$, i.e. a full binary tree. Figure 8(b) considers $m = 7$ but $d^* = 3$, so the optimal monitoring tree is made of 3 perfect binary trees of depth $d^* - 2 = 1$ and a full binary tree of depth at most 1, connected to the same root.




Fig. 8. Optimal monitoring tree: $m = 7$ and $d^* = 4$ (a) or $d^* = 3$ (b).

We derive the following bound for single-server monitoring.

Theorem V.1 (Identifiability for single-server monitoring). *Consider monitoring paths between a server and m clients in a network of n nodes and maximum path length d^* . Then the maximum number of identifiable nodes $\psi^{ss}(m, n, d^*)$ is upper-bounded by:*

$$\begin{cases} \min \{z_{fb}(m), n\} & \text{if } d^* \geq \lceil \log_2 m \rceil + 1, \\ \min \left\{ n; 1 + \left\lfloor \frac{m}{2^{(d^*-2)}} \right\rfloor \cdot z_{fb}(2^{(d^*-2)}) \right. \\ \quad \left. + z_{fb}(m \bmod 2^{(d^*-2)}) \right\} & \text{otherwise,} \end{cases} \quad (1)$$

where $z_{fb}(m) \triangleq \max\{0, 2m - 1\}$ is the number of nodes in a full binary tree with m leaves.

Proof (sketch). First, we show that in the optimal monitoring tree described in Lemma V.2, all nodes are identifiable. Then we show by induction that given a full binary tree with m leaves, the number of nodes $z_{fb}(m)$ is $\max\{0, 2m - 1\}$. Finally, we use this result to count the number of nodes in the optimal monitoring tree. \square

2) *Tightness of the bound and design insights:* Under the constraint that monitoring paths have a common endpoint, for any given number of monitoring paths m , maximum path length d^* , and sufficiently large n , it is possible to design

a network topology according to the structure of an optimal monitoring tree, as described by Lemma V.2, with a number of identifiable nodes equal to the bound in Theorem V.1.

In particular, if $d^* \geq \lceil \log_2 m \rceil + 1$ the topology would be a full binary tree as in the example of Figure 8(a), while if $d^* < \lceil \log_2 m \rceil + 1$ the topology would be the composition of $\lfloor \frac{m}{2^{(d^*-2)}} \rfloor$ perfect binary trees of depth $d^* - 2$, and a full binary tree of depth at most $d^* - 2$, connected to a common root, as in the example of Figure 8(b).

B. Multi-Server Monitoring

1) *Identifiability Bound:* We now consider a multiple server scenario where each server s ($s = 1, \dots, S$) has m_s clients. We analyze two subcases: (i) *fixed client assignment*, where the number of clients m_s for each server is predetermined, and (ii) *flexible client assignment*, where the total number of clients $\sum_{s=1}^S m_s$ is fixed but the distribution across servers can be designed. Following a similar approach as in the proof of Theorem IV.2, we can bound the number of identifiable nodes in each subcase as follows.

Theorem V.2 (Identifiability for multi-server monitoring with fixed client assignment). *Consider the paths between S servers and m_s clients for server s ($s = 1, \dots, S$) in a network of n nodes, with maximum path length d^* . Let $m \triangleq \sum_{s=1}^S m_s$ and $\mathbf{m} \triangleq (m_1, m_2, \dots, m_S)$. Then the maximum number of identifiable nodes $\psi^{MS}(\mathbf{m}, n, d^*)$ is upper-bounded as in Theorem IV.2, except that N_{\max} is specified by $N_{\max} = \min \left\{ m d^*; \sum_{s=1}^S \left[(m_s^2 + 3m_s - 2)/2 + 2m_s(m - m_s) \right] \right\}$.*

Proof (sketch). The proof follows the same arguments of the proof of Theorem IV.2, but considers that some of the monitoring paths share endpoints.

- 1) We observe that, according to Lemma V.1, the monitoring paths form S trees intersecting each other.
- 2) By induction on the value of m_s , we prove a bound on the sum of the maximum number of different binary encodings in the path matrices of a single server s . Considering the only columns of the paths of the same server, let ℓ_k be the maximum number of different binary encodings on the path from a client node v_k to service s , with $k = 1, \dots, m_s$. Let $L_s \triangleq \sum_{k=1}^{m_s} \ell_k$. It holds that $L_s \leq (m_s^2 + 3m_s - 2)/2$.
- 3) We prove that the total number of different binary encodings in all the m paths (including repetitions across different paths) is $N_{\max} \triangleq \sum_{i=1}^S \left[(m_i^2 + 3m_i - 2)/2 + 2m_i \cdot (m - m_i) \right]$. To prove this, we consider the i -th tree individually ($i = 1, \dots, S$). When only considering the m_i columns corresponding to the client-server paths of the i -th service, the path matrices of the i -th tree have a maximum number of different binary encodings $N_{\max_i} = (m_i^2 + 3m_i - 2)/2$. Moreover, in each of these path matrices, there are $m - m_i$ other columns corresponding to client-server paths of the other services. The sequence of bits of each of these columns may flip twice, due to Lemma IV.2, which accounts for $2(m - m_i)$ more column flips in each of the m_i path matrices of the i -th tree. Hence we have

$$N_{\max} \triangleq \sum_{i=1}^S \left[\frac{(m_i^2 + 3m_i - 2)}{2} + 2m_i \cdot (m - m_i) \right]. \quad (2)$$

- 4) We proceed as in the proof of Theorem IV.2, replacing N_{\max} with the righthand side of Equation (2). \square

The following theorem addresses a different case, in which every client can be assigned to any server.

Theorem V.3 (Identifiability for multi-server monitoring with flexible client assignment). *Consider monitoring the paths between S servers and m clients with arbitrary client-server assignment in a network of n nodes, with maximum path length d^* . Then the maximum number of identifiable nodes $\psi^{MS}(m, S, n, d^*)$ is upper-bounded as in Theorem IV.2, except that N_{\max} is specified by $N_{\max} = \min \left\{ m \cdot d^*; m^2 \left(2 - \frac{3}{2S} \right) + 3m/2 - S \right\}$.*

Proof. Let \mathcal{A} be the set of possible assignments of m clients to S servers: $\mathcal{A} = \{ \mathbf{m} \in \mathbb{N} | m_s \geq 0, \text{ and } \sum_{s=1}^S m_s = m \}$.

The bound on the number of identifiable nodes in the case of S servers and undistinguished clients can be formulated as in Theorems IV.2 and V.2, where $N_{\max} = \min \left\{ m d^*; \max_{\mathbf{m} \in \mathcal{A}} \sum_{s=1}^S \left[\frac{m_s^2 + 3m_s - 2}{2} + 2m_s \cdot (m - m_s) \right] \right\}$. In order to calculate N_{\max} we address the optimization, in the integer variables m_s , of the objective function $f(\mathbf{m}) = \sum_{s=1}^S \left[(m_s^2 + 3m_s - 2)/2 + 2m_s \cdot (m - m_s) \right] = 2m^2 + 3m/2 - S - 3/2 \sum_{s=1}^S m_s^2$ (obtained by replacing $\sum_{s=1}^S m_s$ with m where possible), under the constraint that $\mathbf{m} \in \mathcal{A}$. A relaxation of this problem leads to the following solution: $m_s = m/S, \forall s = 1, \dots, S$, and an objective value of $m^2 \left(2 - \frac{3}{2S} \right) + 3m/2 - S$, which is an upper bound to $f(\mathbf{m})$, from which we derive the assertion of the theorem. \square

2) *Design insights:* In a setting in which the m monitoring paths connect a given number of servers to their clients, the maximum identifiability is obtained by letting the branches of several server-rooted optimal monitoring trees intersect with each other, while satisfying the consistent routing assumption and the constraint on the maximum path length d^* .

While in the case of fixed client assignment to servers, the number of leaves of each tree is predetermined, in the case of flexible client assignment, the proof of Theorem V.3 suggests that the highest identifiability is obtained through a uniform assignment of clients to servers. In terms of topology design this implies that the maximally identifiable topology would require uniformly sized monitoring trees.

VI. DATA-CENTER NETWORK MONITORING

The last scenario is a data-center network where we monitor paths between end-hosts. Data-center networks have unique topology and routing properties that require dedicated analysis.

A. Half-consistent Routing

Typical data-center topologies are based on two or three layers of homogeneous k -port switches arranged into tree-like topologies. A common topology built of commodity Ethernet switches is the *fat-tree* topology [13]. Recent works on data-center design and optimization propose the use of fat-tree topologies to deliver high bandwidth to hosts at the leaves of the fat-tree, which requires spreading the pod's outgoing traffic uniformly to the core switches. A special instance of a k -ary fat-tree together with a related addressing and




Fig. 9. Routing inconsistency in a fat-tree

routing scheme is described in the work of Al-Fares et al. in [14]. An example with 3 layers and $k = 4$ is shown in Figure 9. The authors of [14] propose the use of a joint routing and addressing scheme which violates the consistent routing assumption in two aspects: (1) routes between different source-destination pairs may not be consistent, (2) routes in different directions between the same source-destination pair may not be consistent either.

An example of two paths violating routing consistency is given in Figure 9. The example highlights routing inconsistency, as the path between the aggregation switches 10.1.3.1 and 10.3.3.1 can be different depending on the source and the destination hosts.

In this scenario, while Theorem IV.1 is still valid, Theorem IV.2 cannot be applied due to the lack of routing consistency and symmetry. To address the lack of routing symmetry, in the following we consider one-way monitoring paths, while we address the lack of consistency by introducing the following concept of half-consistency.

Definition VI.1. *If a routing scheme guarantees that any path $p_i \in P$ can be divided into two segments $s_1(p_i)$ and $s_2(p_i)$, such that the property of routing consistency (Definition IV.1) holds for the set $P' = \cup_{p_i \in P} \{s_1(p_i), s_2(p_i)\}$, then the routing scheme is called half-consistent.*

Observation VI.1. *Any shortest-path routing scheme on a fat-tree is half-consistent.*

Proof. Let us call $u_s(p)$ and $u_t(p)$ the source and the destination endpoints of p , and let us call the *upper node* $u_m(p)$ the node of p that is the farthest from the endpoints. Due to the structure of the fat-tree, there is only a unique path $s_1(p)$ from $u_s(p)$ to $u_m(p)$, and a unique other path $s_2(p)$ from $u_m(p)$ to $u_t(p)$. Therefore, for any two intermediate nodes on $s_i(p)$ ($i = 1, 2$), there cannot be any alternative path between them, and the routing of these path segments is consistent. \square

B. Identifiability Bound

In the case of half-consistent routing (including routing in fat-trees), the following bound holds.

Theorem VI.1 (Identifiability under half-consistent routing). *Given m ($m > 1$) monitoring paths of maximum length d^* in an n -node network, the maximum number of identifiable nodes under the assumption of half-consistent routing $\psi^{HR}(m, n, d^*)$ is upper-bounded as in Theorem IV.2, except that N_{\max} is specified by $N_{\max} = m \cdot \min\{2^{m-1}, 4 \cdot (m-1); d^*\}$.*

Proof (sketch). The proof derives from the same arguments of the proof of Theorem IV.2, but replaces the concept of consistency with half-consistency as follows.

- 1) We prove that each column of any path matrix $M(\hat{p}_i)$ has up to two sequences of consecutive ones. This is due to the half-consistency property which implies that each column has two parts, each meeting the consecutive ones property.
- 2) We prove that the maximum number of different encodings in a path matrix $M(\hat{p}_i)$ is $\min\{2^{m-1}, d^*, 4(m-1)\}$. In fact, the number of different encodings of $M(\hat{p}_i)$ is bounded by the maximum number of nodes d^* , and by the maximum number of encodings of $m-1$ digits (recall that the i -th digit is always equal to 1). Additionally, by considering each column of $M(\hat{p}_i)$ separately we notice that besides the i -th column which contains only 1s, any other column can flip its value up to 4 times, two for each sequence of consecutive ones. In order to have a change in the encoding in any two successive rows $r-1$ and r there must be at least a column that flips in r . Notice that the number of columns that can flip is $m-1$. As a consequence, the number of different rows of $M(\hat{p}_i)$ is upper-bounded by $4(m-1)$.
- 3) We observe that the maximum number of different encodings in m path matrices, is bounded by

$$N_{\max} = m \cdot \min\{2^{m-1}, 4 \cdot (m-1); d^*\}. \quad (3)$$
- 4) We proceed as in the proof of Theorem IV.2, replacing the value of N_{\max} with the value of Equation (3). \square

C. Design insights

The identifiability of a fat-tree depends on the topology parameters k , ℓ and the number of paths m . In Section VII-D we show that for a small number of layers ℓ , half-consistency does not imply a notable increase in the upper bound on identifiability compared with consistency. In contrast, in determining a lower bound, we notice that an empirical choice of paths that achieves good identifiability makes ample use of the half-consistency property to enable the identifiability of the routers at the upper layers of the fat-tree. Results are shown in Figure 17. General guidelines on how to optimize identifiability via path selection in such topologies are left to future work.

VII. PERFORMANCE EVALUATION

To evaluate the tightness of the proposed upper bounds, we compare them with lower bounds obtained by known heuristics on synthetic and real network topologies. Since the bound in Theorem IV.1 is achievable under arbitrary routing (see Section IV-A2) and loose otherwise, we show it once in Figure 10 and we omit it in the rest of the evaluation.

A. Consistent Routing

We analyze the tightness of the upper bound in Theorem IV.2 under consistent routing. Figure 10 considers a network having a half-grid topology as in Figure 3, with 78 nodes arranged into 12 lines and 12 columns. It shows the upper bound (UB) of Theorems IV.1 and IV.2 as well as a lower bound (LB) obtained by placing monitoring endpoints as in Section IV-B2. We vary the number of paths while fixing the maximum length at $d^* = 12$. We show the bound in Theorem IV.1 to highlight the potential in improving identifiability by relaxing routing consistency. For instance, regardless of the topology and the monitoring endpoints, we see that if




Fig. 10. Bounds in a half-grid with $n = 78$, varying m , and $d^* = 12$.




Fig. 11. Bound of Th. IV.2, varying m and d^* (network of Fig. 10).

the number of paths is 6, the routing is consistent, and the maximum path length is 12, then the number of identifiable node is no greater than 30, much smaller than the bound of $\min\{2^6 - 1, 78\} = 63$ obtained by Theorem IV.1 without the constraints on routing and path length.


In Figure 11 we show, for the same network, how the bound of Theorem IV.2 varies with the number of monitoring paths m and the maximum path length d^* . For small values of d^* the bound has an almost linear growth with m . For larger values of d^* the bound shows two regions: an initial super-linear growth for small values of m , and a linear growth for large values of m . The figure also shows that while the number of paths m has a major impact on the number of identifiable nodes, the length of the monitoring paths has a significant impact only when d^* is small, and diminishing impact otherwise.

B. Single-Server Monitoring


Figure 12 shows two scenarios with different topologies. The first scenario is a network of 95 nodes, connected as a full binary tree with 48 leaves, with $d^* = 7$ (in number of nodes). The figure shows the increase of the optimal number of identifiable nodes by varying the number of monitoring paths having a common endpoint. By using 48 paths of maximum length $d^* = 7$ from the leaves to the root, it is possible to identify all the network nodes. Notice that the optimal number of identifiable nodes that can be obtained by varying server location and placement of clients coincide with the value of the bound of Theorem V.1. Lemma V.2 shows in fact that for such a topology, the optimal identifiability is achieved by placing the endpoints of the m different monitoring paths one in the root of the tree and the others in a way that the paths form a full binary tree topology.

For the second scenario we consider a stricter limit on the path length $d^* = 3$. We consider a tree topology where a common root is connected to 24 binary trees of depth 1, for a total of 48 leaves, and 73 nodes (this topology is constructed extending the case of Figure 8(b) to connect 24 subtrees). In this topology, by using 48 paths of maximum length $d^* = 3$, from the leaves to the root, it is possible to identify all the nodes. Also in this case, the bound of Theorem V.1 is tight, and coincides with the optimal, which is a tree of paths where $\lceil m/2 \rceil$ binary trees of depth 1 descend from a common root.

The Figure also shows that the values of the bound obtained with Theorem IV.2, are considerably looser than those of Theorem V.1. This is because the former considers any m paths generated with any consistent routing scheme, while the latter considers the additional requirement that the monitoring paths share a unique common endpoint.



(a)



(b)

Fig. 12. Bound for single-server monitoring (Th. V.1) - full binary tree for $d^* = 7$ (a), multiple binary trees connected to a single root for $d^* = 3$ (b).




Fig. 13. UB of Th. V.1 and LB of random placement, AT&T topology, $S = 1$, $d^* = 4$, varying m .




Fig. 14. UB of Th. V.2 and LB of GD [6], AT&T topology, $d^* = 4$, varying m and S (3 clients per server).

Figure 13 illustrates an experiment on an existing AT&T topology mapped with Rocketfuel [15], with 108 nodes and 141 links. We consider a single server and a random placement of m clients. We obtained a lower bound, called "Random", by running 100 trials for each value of m and using the largest number of nodes identified by client-server paths under consistent shortest path routing. We then compare this value to the upper bound given by Theorem V.1. As the figure shows, the lower bound is not as close to the upper bound as in the case of the engineered topologies in Figure 12.

C. Multi-Server Monitoring

In these experiments we also consider the AT&T topology with 108 nodes and 141 links. We analyze the case of multiple servers, each serving 3 clients. We increase the number of servers and vary the number of clients accordingly. Figure 14 shows the upper bound of Theorem V.2 compared to a lower bound obtained with the heuristic *greedy distinguishability maximization (GD)*⁶ proposed in [6]. Notice that this heuristic finds a good approximation to the optimal number of identifiable nodes in this problem setting. Although the heuristic only optimizes server placement, while Theorem V.2 considers the optimal placement of servers as well as clients, the experiment shows a good approximation of the upper and the lower bounds when m is sufficiently small.

Figure 15 shows a comparison of the three bounds of Theorems IV.2 (arbitrary sources/destinations), V.2 (fixed client assignment) and V.3 (flexible client assignment) for the same topology, where we vary the numbers of services and clients, with a maximum path length $d^* = 20$. In the figure, the bound of Theorem IV.2 represents the special case of one client

⁶Note that GD requires client locations to be predetermined. Here we place clients on some of the 78 dangling nodes, and then use GD to place servers.




Fig. 15. UB of Theorems IV.2, V.2 and V.3, AT&T topology, $d^*=20$, S servers, m clients - even (a) and uneven (b) distribution of clients to servers.




Fig. 16. UB of Theorems IV.2 and VI.1 - 100 nodes, varying d^* . Fig. 17. UB of Th. VI.1 and LB for a 4-ary fat-tree with 3 layers.

per server. We calculate the bound of Theorem V.2 assuming first a uniform assignment of clients to servers, as shown in Figure 15(a), and then an uneven assignment, which is shown in Figure 15(b). For uneven assignment: in the case of two servers, one server is assigned to 4/5 of the clients, while the other to the rest 1/5; in the case of three servers, one server is assigned to 3/4 of the clients, the second server to 3/16, and the third server to 1/16. It can be seen that in the case of even assignment of clients to servers, the two bounds of Theorems V.2 (fixed client assignment) and V.3 (flexible client assignment) give the same values. By contrast, in the case of uneven distribution of clients to servers, Theorem V.2 gives a considerably smaller bound than Theorem V.3, which assumes an even distribution of clients to servers.

D. Data-Center Network Monitoring

Before considering the case of a specific fat-tree topology we highlight that, given the small diameter (and consequently small path length) that typically characterizes these topologies, we do not expect to see a considerable difference by introducing half-consistent routing. We see that only with a high number of layers, routing half-consistency plays a role in optimizing identifiability. To this purpose Figure 16 evidences the difference in the upper bounds of the case of a more flexible half-consistent routing scheme considered in Theorem VI.1, with respect to the case of consistent routing considered in Theorem IV.2. It considers a general network with 100 nodes. The difference of identifiability between consistent and half-consistent routing grows by increasing the maximum length of monitoring paths as $d^* = 5, 15, 25$, which in a fat-tree would correspond to values of $\ell = 2, 7, 12$.

In conclusion, we can affirm that for topologies with very short diameter, such as in the case of fat-trees, having a higher degree of freedom in routing (half-consistent routing) has a significant impact on the identifiability of the network only for a high number of layers.

We now consider the case in which monitoring is performed along paths between hosts of a data-center network with a fat-tree topology and the routing scheme proposed in [14]. In Figure 17 we consider a 4-ary fat-tree with three layers and study the tightness of the bound of Theorem VI.1. Due to the high complexity in selecting the optimal monitoring paths, we resort to an empirical selection of paths that give us a lower bound on the number of identifiable nodes. It is interesting to see that with only 16 monitoring paths we are able to monitor all the 36 nodes of this fat-tree.

VIII. CONCLUSION

We consider the problem of maximizing the number of nodes whose states can be identified via Boolean network tomography. We formulate the problem in terms of graph-based group testing and exploit the combinatorial structure of the testing matrix to derive upper bounds on the number of identifiable nodes under different assumptions, including: arbitrary routing, consistent routing, monitoring through client-server paths with one or multiple servers (and even or uneven distribution of clients), and half-consistent routing. These bounds show the fundamental limits of Boolean network tomography in both real and engineered networks. Besides the theoretical value of this analysis, we use the bounds to derive insights for the design of topologies and monitoring schemes with high identifiability in different network scenarios. Through analysis and experiments we evaluate the tightness of the bounds and demonstrate the efficacy of the design insights for engineered as well as real networks.

REFERENCES

- [1] R. R. Kompella, J. Yates, A. Greenberg, and A. Snoeren, "Detection and localization of network black holes," *IEEE INFOCOM*, 2007.
- [2] N. Duffield, "Simple network performance tomography," in *ACM IMC*, 2003.
- [3] R. Dorfman, "The detection of defective members of large populations," *The Annals of Mathematical Statistics*, 1943.
- [4] G. K. Atia and V. Saligrama, "Boolean compressed sensing and noisy group testing," *IEEE Trans. on Inf. Theory*, vol. 58, no. 3, March 2012.
- [5] M. Cheraghchi, A. Karbasi, S. Mohajer, and V. Saligrama, "Graph-constrained group testing," in *IEEE Trans. on Inf. Theory*, no. 1, 2012.
- [6] T. He, N. Bartolini, H. Khamfroush, I. Kim, L. Ma, and T. La Porta, "Service placement for detecting and localizing failures using end-to-end observations," in *IEEE ICDCS*, 2016.
- [7] N. Duffield, "Network tomography of binary network performance characteristics," *IEEE Trans. on Inf. Theory*, vol. 52, 2006.
- [8] H. Nguyen and P. Thiran, "The Boolean solution to the congested IP link location problem: Theory and practice," in *IEEE INFOCOM*, 2007.
- [9] L. Ma, T. He, A. Swami, D. Towsley, K. K. Leung, and J. Lowe, "Node Failure Localization via Network Tomography," in *ACM IMC*, 2014.
- [10] L. Ma, T. He, A. Swami, D. Towsley, and K. K. Leung, "Network capability in localizing node failures via end-to-end path measurements," *IEEE/ACM Transactions on Networking*, June 2016.
- [11] Y. Bejerano and R. Rastogi, "Robust monitoring of link delays and faults in IP networks," in *IEEE INFOCOM*, 2003.
- [12] L. Ma, T. He, A. Swami, D. Towsley, and K. Leung, "On optimal monitor placement for localizing node failures via network tomography," *Elsevier Performance Evaluation*, vol. 91, pp. 16–37, September 2015.
- [13] C. E. Leiserson, "Fat-trees: Universal networks for hardware-efficient supercomputing," *IEEE Trans. on Computers*, vol. 34, no. 10, 1985.
- [14] M. Al-Fares, A. Loukissas, and A. Vahdat, "A Scalable, Commodity Data Center Network Architecture," in *ACM SIGCOMM*, 2008.
- [15] N. Spring, R. Mahajan, and D. Wetheral, "Measuring isp topologies with rocketfuel," in *ACM SIGCOMM*, August 2002.