

FAVE: A fast and efficient network Flow AVailability Estimation method with bounded relative error

Tingwei Liu[†], John C.S. Lui^{*},

Department of Computer Science & Engineering, The Chinese University of Hong Kong, China

Email: [†]twliu@cse.cuhk.edu.hk, ^{*}cslui@cse.cuhk.edu.hk

Abstract—This paper focuses on helping network providers to carry out network capacity planning and sales projection by answering the question: For a given topology and capacity, whether the network can serve current flow demands with high probabilities? We name such probability as “flow availability”, and present the flow availability estimation (FAVE) problem, which is a generalisation of network connectivity or maximum flow reliability estimations. Realistic networks are often large and dynamic, so flow availabilities cannot be evaluated analytically and simulation is often used. However, naive Monte Carlo (MC) or importance sampling (IS) techniques take an excessive amount of time. To quickly estimate flow availabilities, we utilize the correlations among link and flow failures to figure out the importance of roles played by different links in flow failures, and design three “sequential importance sampling” (SIS) methods which achieve “bounded or even vanishing relative error” with linear computational complexities. When applying to a realistic network, our method reduces the flow availability estimation cost by 900 and 130 times compared with MC and baseline IS methods, respectively. Our method can also facilitate capacity planning by providing better flow availability guarantees, compared with traditional methods.

I. INTRODUCTION

Network capacity planning is the process of ensuring sufficient bandwidth is provisioned so that service-level agreement (SLA) objectives like delay, jitter, loss and routing availability can be satisfied [1]. For the purpose of providing a better end user experience and at the same time, keeping the operation cost at an affordable level, effective capacity planning tools are crucial to network providers. Various systems have been built around this problem, such as Cisco’s MATE [2], Facebook’s Prophet [3], and Google’s backbone capacity planning tool [4]. **Motivations of fast flow availability estimations:** The above SLA objectives are called the *demand* of a traffic flow, and the satisfaction probability of demand is defined as the *flow availability*. A key concern of capacity planning is in analyzing the effect of network changes or the arising of new flows on the flow demand satisfaction [2]. To illustrate, consider an example that a network provider needs to serve two flows as shown in Fig.1. In this network, each link i is associated with a failure probability, say $p=0.001$ and a capacity c_i ; also, flow routing follows the “max-min fairness” and “shortest path” policies. Each flow has a bandwidth demand, and an *availability target* (i.e., the lower bound probability that its requested bandwidth needs to be satisfied). The network provider will perform: 1) **Flow availability testing**, i.e., whether flow availability targets can be achieved? E.g., flow 1’s availability target is achieved if it obtains 10 units of bandwidth with a probability no less than




Fig. 1: An example of the FAVE problem

0.9999. 2) **Capacity planning**, e.g. to improve the network, should the provider add more links between node A and B, or add a new node D so to increase the path diversity? 3) **Sales projection**, e.g., to increase profit, can the provider admit a new flow 3? Will the network still provide flow availability guarantees if flow 1 requires 20 additional units of bandwidth? All these cases need flow availabilities to see whether the network (proposed by capacity planning) can serve flow demands (proposed by sales projection). We name the flow availability estimation problem as “FAVE”, and give a formal definition later.

For a large realistic network with complex failure patterns, the flow availability often cannot be evaluated analytically, especially when the traffic scheduling is also considered. Hence, simulation (or sampling) is often used. The Monte Carlo (MC) method [5], which simulates link failures with real link failure probabilities, is the most widely used. Yet, it can be costly for MC to achieve a desired accuracy level, e.g., to estimate flow 2’s availability in Fig.1, by variances collected in Table 1, MC needs at least 3,840 simulation steps to guarantee the 95% confidence interval (CI) width below 10^{-3} . To simulate a large network with many flows, it is crucial to find ways to speed up the simulation process.

High level ideas of our work: The flow availability estimation can be more efficient if flow failures (i.e., flow demands could not be satisfied) happen more frequently by taking a proper distribution to simulate link failures, which is the key idea of “importance sampling (IS)” [6]. We first introduce a baseline IS solution for FAVE using “the correlation between link failures and flow failures” to decide “important links” whose failures are more likely to result in flow failures. We then simulate such links’ failures more often to speed up simulating flow failures. To further improve upon the baseline IS, we show that some link sets’ failures happen more often, and play a role of “root causes” of flow failures. We name such link set as “SEED”. We utilize “the correlation among link failures” captured by SEEDs to decide “important link sets” whose failures are more likely to result in flow failures. We propose three advanced “sequential importance sampling (SIS)” algorithms, among which the SEED-VRE algorithm has the “vanishing relative error (VRE)”

This work is supported in part by the GRF 14200117 and the Huawei Grant.

Table 1: Accuracy, efficiency and computational cost (i.e., simulation steps) comparison

Flow	Theoretical Unavailability	Monte Carlo Method		Number of Simulations	Our method (SEED-VRE)		
		Estimated Unavailability	Empirical Variance		Estimated Unavailability	Empirical Variance	Number of Simulations
1	1.000999×10^{-3}	1.300×10^{-3}	1.29831×10^{-7}	10,000	1.000999×10^{-3}	6.35275×10^{-23}	10
2	1.000000×10^{-3}	8.000000×10^{-4}	7.99360×10^{-8}	10,000	1.000000×10^{-3}	2.11758×10^{-23}	10

and “*linear computational complexity*”. Table 1 illustrates the computational reduction of our methods: Using only **10 simulation steps**, both empirical variance and estimation error are much smaller than that of MC with **10,000 simulation steps**. In other words, not only we can speed up simulations, but *the network operator is able to achieve a higher accuracy of flow availability estimation without incurring an additional cost*.

We emphasize that flow availability estimation can also facilitate capacity planning: for flows with unachievable availability targets, we show that allocating more capacities for the important links determined by our method would bring a better flow availability improvement, compared with traditional methods to allocate more capacities for links with high utilizations. **Contributions:** Our key contribution is in providing efficient and accurate solutions for FAVE. Contributions include:

- We formally define the network Flow Availability Estimation (FAVE) problem, and show it generalizes the classical network reliability estimation problem [7]–[13].
- For FAVE with a single flow, we introduce a novel concept of “SEED” and propose three advanced SIS algorithms which have attractive properties of *bounded relative error (BRE)* or *VRE* with *linear computational complexities*.
- For FAVE with multiple flows, we propose a mixture SIS method which *maintains BRE property and linear computational complexity* when estimating all flows’ availabilities.
- For FAVE with a partial SEED set information, we maintain the BRE and VRE properties of flow availability estimations.
- Extensive results show that our methods greatly speed up the flow availability estimation: 1) Given an illustrative network and full SEED sets or partial SEED sets with good coverage, SEED-VRE achieves *variance reductions of 360,000 and 7,000 times* in the single flow case, and *200,000 and 7 times* for around 80% flows in the multiple flow case, compared with MC and baseline IS. 2) Given a realistic network and partial SEED sets with poor coverage, SEED-VRE *reduces simulation cost by 900 and 130 times* for around 80% flows in the multiple flow case, compared with MC and baseline IS.
- We demonstrate our methods facilitate capacity planning by providing more accurate network reliability evaluations compared with classical methods, and greater flow availability improvements compared with solely using the link capacity utilization information.

Organizations: Section II introduces related work and some preliminaries. Section III formally defines the FAVE problem. In Section IV, we propose a baseline IS solution for FAVE and show its error bounds, then we introduce “SEED” and present SEED based SIS solutions with better error bounds and linear computational complexities. Section V considers more practical issues, e.g., the multiple flows case and partial SEED sets information case. In Section VI, we evaluate our methods on both an illustrative network and a realistic network, i.e., the Abilene network [14]. Section VII shows the utility of our methods in capacity planning. Finally, Section VIII concludes.

II. RELATED WORK & PRELIMINARIES

Our work can be viewed as a generalisation of previous work in estimating *network reliability* [7]–[13] and is related with *importance sampling* [6]. Here, we briefly review previous relevant studies and compare them with our work. We also describe some concepts so readers can gain a better understanding of sampling methods for the network reliability estimation.

A. Network Reliability Estimation

The most relevant literatures to our work focus on evaluating *network reliability*. To design reliable networks, one needs to measure the impact of network failures (e.g., link failures) on network performances [15]. It is known that the exact computation of network reliability is #P-complete, and computational complexities of all known algorithms are exponential increasing with the graph scale [16], which makes the problem intractable even for medium sized networks. Hence, most work on network reliability evaluation considers sampling methods to provide reliability estimations, and they can be classified into “*network connectivity*” based and “*maximum flow*” based.

Network connectivity reliability (NCR): NCR is a classical reliability measure adopted by most work [7]–[9]. The network is modelled as a graph where links are either failed or operational, and NCR is measured by the probability that a given set of nodes are connected when links fail with given probabilities. Authors in [7] take network repair policies into consideration to model link failures, and estimate NCR with the classical MC. Authors in [8] combine MC with the particle swarm optimization to handle the NCR problem. To improve the efficiency of MC, authors in [9] apply the IS method and use pre-computed “*graph minimal cuts*” to approximate the optimal IS estimator. **Maximum flow reliability (MFR):** Another line of work [11]–[13] generalizes the NCR problem by considering link capacities: link capacities are determined by link statuses, i.e., operational, failed or partially failed, which follow certain probability distributions. Given one source and one sink, MFR is defined as the probability that the maximum flow, i.e., the maximum achievable bandwidth from the source to the sink, is above a given threshold. In [11], link capacities are assumed to be continuous and the MC splitting method is applied for the MFR estimation. Authors in [12] follow the idea of permutating MC and assume all links fail at the beginning and each one of them gets repaired after a random time. Authors in [13] consider estimating MFR with the order minimal cut sets.

Other reliabilities: Some works also study the *connection availability* [17] and *service availability* [18], and consider the probability that a connection or service is available. Authors in [17] evaluate the connection availability by computing the connection probability of a small subset of nodes exactly. Authors in [18] evaluate the service availability by using IS to estimate path availabilities. The problem considered in [17], [18] can be transformed to a problem of determining the connectivity of certain nodes, where the network topology is given. Hence,

Table 2: A comparison between classical network reliabilities and flow availability

Reliability Measurement	NCR [7]–[9], [17], [18]	MFR [11]–[13]	Flow Availability
Definition	$\mathbb{P}[\text{node } u \text{ and } v \text{ are connected}]$	$\mathbb{P}[\text{max flow from } u \text{ to } v \geq \zeta]$	$\mathbb{P}[\text{flow } f_i \text{'s demand is satisfied}]$
Required Information	Topology Information: $G(V, E)$ and \mathbf{p} .	Topology Information: $G(V, E)$, \mathbf{p} and \mathbf{c} .	Topology Information: $G(V, E)$, \mathbf{p} and \mathbf{c} ; Flow Information: F and (s_i, t_i, d_i) for $\forall f_i \in F$.
Application scenarios	1) Topology design evaluation.	1) Topology design evaluation; 2) Capacity planning design evaluation.	1) Topology design evaluation; 2) Capacity planning design evaluation; 3) Sales projection design evaluation.
Relationship	The flow availability is a generalization of network connectivity and maximum flow based reliabilities. When link capacity $c_i = \infty$ for $\forall i$, $\mathbb{P}[\text{flow } f_i \text{'s demand is satisfied}] = \mathbb{P}[\text{node } u \text{ and } v \text{ are connected}]$. When flow set F contains only one flow, $\mathbb{P}[\text{flow } f_i \text{'s demand is satisfied}] = \mathbb{P}[\text{max flow from } u \text{ to } v \geq \zeta]$.		

Note: Consider in network $G(V, E)$, there exist node u, v and flow f_i from u to v with demand ζ . Let vector \mathbf{p} and \mathbf{c} denote information of failure probability and capacity across all links. Consider in flow set F , each flow f_i is associated with a source s_i , a destination t_i and a demand d_i .

[17], [18] are essentially the same with the NCR related work.

B. Comparisons with Classical Reliability Estimation Work

We consider the ‘‘flow availability’’ as our reliability measure. We first give the definition of the flow demand.

Definition 1. The ‘‘demand’’ of flow f is the quality of service (QoS) requirements decided by f 's SLA objective.

Consider different SLA objectives, the flow demand can be, for instance, the *bandwidth demand*, *latency demand* or *packet loss demand*, which specifies f 's QoS requirement on bandwidth, transmission latency or packet loss. To be concrete and so easier to understand, we take the bandwidth demand as an example, and the following analysis works the same for other demands. We define the flow availability as:

Definition 2. For a given topology information, flow information, routing policy and resource allocation policy, the ‘‘flow availability’’ of f is the probability that f 's demand is satisfied.

Comparing with the state-of-the-art methods, our methods have the following advantages:

1) The flow availability can be applied to *evaluate both NCR, given all links have unlimited capacities, and MFR, given the network contains only one flow*. Yet, neither NCR nor MFR can address FAVE. To illustrate, consider the example in Fig.1. If link 1 fails, the network is still connected but neither flow 1 nor 2's demands can be satisfied. Also, the maximum flow from A to B still achieves 10 units, but it does not imply flow 1 succeeds, which depends both on resource allocation and routing policies and other competing flows.

2) Flow availability can be applied to *evaluate not only the reliability of network designs, including topology design and capacity planning, but also the feasibility of sales projection*. In contrast, NCR only applies to the topology design evaluation and MFR only applies to the capacity planning evaluation, for they utilize solely the (partial) topology information.

To summarize, FAVE generalizes the NCR and MFR estimations and considers a more realistic problem setting. Moreover, it can be applied to evaluate impacts of more factors, e.g., network topology, capacity and flow information on network performances, and provides more accurate evaluation results. The detailed comparisons can be found in Table. 2.

C. Sampling Methods for Network Reliability Estimations

Network reliability estimation problem: Let the network be modelled as a directional multigraph $G := (V, E)$ with N_v nodes in node set V and N_l links in link set E . Each link $e_i \in E$ is associated with a tuple (p_i, c_i, x_i) with a small probability p_i to represent e_i 's failure probability, a capacity c_i , and a status x_i , where $x_i = 1$ ($x_i = 0$) means e_i is down (up). Let $\mathbf{p} = \{p_1, \dots, p_{N_l}\}$,

$\mathbf{c} = \{c_1, \dots, c_{N_l}\}$ and $\mathbf{x} = \{x_1, \dots, x_{N_l}\}$ be failure probability, capacity and status across all links respectively. \mathbf{x} , also called the ‘‘failure configuration’’, is generated by the link failure distribution $p(\mathbf{x})$ induced by \mathbf{p} . There are 2^{N_l} possible configurations of \mathbf{x} , which is huge for a large realistic network.

Let \mathcal{R} be the indicator function of some interested event A . According to the reliability definition, A can refer to the event that a subset of nodes are unconnected, or the maximum flow is below the required threshold, or, in our case, the flow demand is unsatisfied. Given link statuses described by \mathbf{x} , $\mathcal{R}(\mathbf{x}) = 1$ if A is observed and $\mathcal{R}(\mathbf{x}) = 0$ vice versa. The network unreliability, which is the occurrence probability of A , can be computed via the following integral in the discrete measure space:

$$\mu = \mathbb{E}_p[\mathcal{R}(\mathbf{x})] = \int \mathcal{R}(\mathbf{x})p(\mathbf{x})d\mathbf{x}, \quad (1)$$

where $\mathbb{E}_p[\cdot]$ means taking the expectation over distribution $p(\mathbf{x})$. Then, the network reliability can be obtained by $1 - \mu$.

Monte Carlo (MC) simulation: The MC simulation draws failure configurations \mathbf{x} independently from $p(\mathbf{x})$ and estimate μ with the following MC estimator:

$$\hat{\mu}_{MC} = \frac{1}{N} \sum_{k=1}^N \mathcal{R}(\mathbf{x}^{(k)}), \quad (2)$$

where N is the number of simulation steps and $\mathbf{x}^{(k)}$ is the k th generated failure configuration. As MC generates link statuses by true link failure probabilities (which can be small), it is rare to observe link failures, and even rarer to observe A . This implies that we need a large N to gain the desired accuracy.

Importance sampling (IS): To improve the efficiency of MC, IS changes the sampling distribution $p(\mathbf{x})$ to increase the occurrence of A , and assigns each sample \mathbf{x} a weight to recover the unbiasedness. Specifically, it replaces Eq.(1) by:

$$\mu = \mathbb{E}_p[\mathcal{R}(\mathbf{x})] = \int \mathcal{R}(\mathbf{x}) \frac{p(\mathbf{x})}{q(\mathbf{x})} q(\mathbf{x}) d\mathbf{x} = \mathbb{E}_q \left[\mathcal{R}(\mathbf{x}) \frac{p(\mathbf{x})}{q(\mathbf{x})} \right], \quad (3)$$

where $q(\mathbf{x})$ is the ‘‘importance distribution’’. For convenience, denote $\omega(\mathbf{x}) = p(\mathbf{x})/q(\mathbf{x})$ as the weight. Therefore, the above expectation is estimated by the following IS estimator:

$$\hat{\mu}_{IS} = \frac{1}{N} \sum_{k=1}^N \mathcal{R}(\mathbf{x}^{(k)}) \frac{p(\mathbf{x}^{(k)})}{q(\mathbf{x}^{(k)})} = \frac{1}{N} \sum_{k=1}^N \mathcal{R}(\mathbf{x}^{(k)}) \omega(\mathbf{x}^{(k)}). \quad (4)$$

Estimator efficiency evaluation: An estimator's efficiency is often measured by its ‘‘variance’’. Take the MC estimator as an example, its variance is given by:

$$\mathbb{V}_p[\hat{\mu}_{MC}] = \frac{1}{N} \mathbb{V}_p[\mathcal{R}(\mathbf{x})] = \frac{1}{N} (\mu - \mu^2), \quad (5)$$

where $\mathbb{V}_p[\cdot]$ means taking the variance over the distribution $p(\mathbf{x})$. The IS estimator's variance can be expressed as:

$$\mathbb{V}_q[\hat{\mu}_{IS}] = \frac{1}{N} \mathbb{V}_q[\mathcal{R}(\mathbf{x})\omega(\mathbf{x})] = \frac{1}{N} (\mathbb{E}_p[\mathcal{R}(\mathbf{x})\omega(\mathbf{x})] - \mu^2). \quad (6)$$

Note that the MC estimator is a special case of the IS estimator if $p(\mathbf{x}) = q(\mathbf{x})$, $\forall \mathbf{x}$. Define $\sigma_q^2 = \mathbb{V}_q[\mathcal{R}(\mathbf{x})\omega(\mathbf{x})]$ as the ‘‘one-run variance’’ of the IS estimator. To achieve a desired estimation accuracy, the CI width should be bounded by a threshold δ , i.e.,

$2\alpha\sigma_q/\sqrt{N}\leq\delta$, and the simulation cost is $N\geq(2\alpha\sigma_q/\delta)^2$, where α is a constant decided by the required confidence level. Thus, a small and bounded σ_q^2 implies an efficient estimator.

Table 3: Important Notations

Notations	Descriptions
N_l, N_f	The number of transportation links and flows.
e_i	The i th link, where $i\in\{1, \dots, N_l\}$.
$p_i, p_i(x_i)$, $\mathbf{p}, p(\mathbf{x})$	p_i and $p_i(x_i)=p_i^{x_i}(1-p_i)^{1-x_i}$ are e_i 's failure probability and distribution. $\mathbf{p}=(p_1, \dots, p_{N_l})$ and $p(\mathbf{x})$ is the distribution induced by \mathbf{p} .
\mathbf{x}	Failure configuration. $\mathbf{x}=(x_1, \dots, x_{N_l})$ and $x_i=1$ ($x_i=0$) means that the link is down (up).
f_i	The i th flow, where $i\in\{1, \dots, N_f\}$.
(s_i, t_i, d_i, o_i)	f_i 's flow information, including source s_i , destination t_i , bandwidth demand d_i and availability target o_i .
\mathcal{R}	Indicator function of traffic engineering simulation. Given \mathbf{x} , $\mathcal{R}(\mathbf{x})=1$ ($\mathcal{R}(\mathbf{x})=0$) if the flow fails (succeeds).
L, \mathcal{L}	$L\subseteq\{1, \dots, N_l\}$ is a link set. \mathcal{L} is a collection of link sets.
$\Psi(L)$, $\Psi^{-1}(\mathbf{x})$	$\Psi(L)$ maps a link set L to the failure configuration \mathbf{x} which satisfies: $\forall i\in L, x_i=1; \forall i\notin L, x_i=0$. $\Psi^{-1}(\mathbf{x})$ maps a failure configuration \mathbf{x} to the link set $L=\{i x_i=1, i\in\{1, \dots, N_l\}\}$.
$\text{span}(L)$, $\text{span}(\mathcal{L})$	$\text{span}(L)$ is the collection of link set L 's supersets and $\text{span}(\mathcal{L})$ is the collection of supersets of all link sets $L\in\mathcal{L}$.
\mathcal{F}	$\mathcal{F}=\{L \mathcal{R}(\Psi(L))=1, L\subseteq\{1, \dots, N_l\}\}$, the collection of all failure link sets which can result in flow failures.
S, \mathcal{S}	A SEED S is a special link set satisfying: 1) $S\in\mathcal{F}$; 2) $\forall L\subseteq S, L\notin\mathcal{F}$; 3) $\forall L\supseteq S, L\in\mathcal{F}$. And \mathcal{S} is the collection of all SEEDs.

Zero-variance (ZV) importance distribution: The following theorem gives an optimal importance distribution $q^*(\mathbf{x})$:

Theorem 1 (Zero-variance importance sampling). *The IS estimator in Eq.(4) can achieve zero-variance, i.e., $\sigma_q^2=0$, if the importance distribution $q(\mathbf{x})=q^*(\mathbf{x})$ where:*

$$q^*(\mathbf{x})=\mathbb{P}[\mathcal{R}(\mathbf{x})=1]. \quad (7)$$

Remark: Although the ZV property implies a minimum simulation cost, it is non-trivial and sometimes even impossible to compute the ZV importance distribution $q^*(\mathbf{x})$. Hence, the key of designing an efficient IS estimator lies in the approximation of $q^*(\mathbf{x})$. As for different applications and problem definitions, the auxiliary information we can utilize and the way to approximate $q^*(\mathbf{x})$ are different, designing efficient IS estimators are highly challenging and problem-dependent.

III. PROBLEM DEFINITION

Consider the network $G(V, E)$ with topology information $\mathbf{p}, \mathbf{c}, \mathbf{x}$, as defined in Section II-C. We consider a flow set F with $N_f=|F|$ flows and each flow $f_i\in F$ is associated with a tuple (s_i, t_i, d_i, o_i) specifying f_i 's source s_i , destination t_i , demand d_i and availability target o_i . We also define the following:

Definition 3. *A flow fails (succeeds) if its demand is unsatisfied (satisfied), e.g., the allocated bandwidth cannot (can) support its bandwidth demand.*

We redefine the function \mathcal{R} to indicate the network routing:

$$\mathcal{R}(\cdot) : (\mathbb{G}, \mathbb{F}, \mathbb{D}, \mathbb{B}) \rightarrow \{0, 1\}, \quad (8)$$

where \mathbb{G} represents the topology information, including a tuple (p_i, c_i, x_i) for every link $e_i\in E$; \mathbb{F} represents the flow information, including a tuple (s_i, t_i, d_i, o_i) for every flow $f_i\in F$; \mathbb{D} and \mathbb{B} represent the underlying routing and resource allocation policies, e.g., shortest path policy and max-min fairness policy. Function \mathcal{R} outputs 1 if the interested flow fails and 0 otherwise. We assume all information in $(\mathbb{G}, \mathbb{F}, \mathbb{D}, \mathbb{B})$ is known, except link statuses described by \mathbf{x} . To simplify the expression, let:

$$\mathcal{R}(\cdot) : \mathbf{x} \rightarrow \{0, 1\}. \quad (9)$$

Namely, given all other information, the routing result indicated

by \mathcal{R} only depends on the failure configuration \mathbf{x} , which is generated by the link failure distribution $p(\mathbf{x})$ induced by \mathbf{p} .

The unavailability μ , also called the flow failure probability, of a specific flow can be computed via the integral in Eq.(1). Our goal is to evaluate availabilities of (all) flows in F . Yet, the complexity of function \mathcal{R} and the high dimensionality of the topological space make it impossible to evaluate the flow availability $1-\mu$ analytically. One alternative is to estimate μ via simulations. We name the network flow availability estimation problem as "FAVE". And we can show it generalizes the classical network reliability estimation problem [7]–[13].

Theorem 2. *The FAVE problem generalizes both the network connectivity based and the maximum flow based network reliability estimation problems.*

Remark: Due to the page limit, we only provide the sketch proof of Thm. 6 in this paper and we leave proofs of all other lemmas and theorems in the technical report [19].

In addition to measuring the estimation efficiency with the variance, we also consider two attractive error bound properties:

Definition 4 (Bounded Relative Error). *An estimator with expectation μ and variance σ^2/N has the bounded relative error (BRE) property if $\sigma=O(\mu)$, i.e., the coefficient of variation (CV) $\varepsilon_{CV}\triangleq\sigma/\mu$ satisfies $\lim_{\mu\rightarrow 0}\varepsilon_{CV}<\infty$.*

Definition 5 (Vanishing Relative Error). *An estimator with expectation μ and variance σ^2/N has the vanishing relative error (VRE) property if $\sigma=o(\mu)$, i.e., $\lim_{\mu\rightarrow 0}\varepsilon_{CV}=0$.*

Remark: Note that the VRE property is stronger than the BRE property. The variance of MC is $\sigma^2/N=(\mu-\mu^2)/N$, i.e., $\sigma=O(\sqrt{\mu})$. This implies that the MC estimator satisfies neither BRE nor VRE property. In the following, we will discuss how to design estimation methods which have above properties.

IV. ALGORITHM DESIGN

We first describe our design for "the single flow case". We start with a baseline IS design to gain insights for efficient estimations. Then we introduce "SEED" and our SEED methods.

A. Baseline Importance Sampling Design

1) ZV importance distribution approximation: It seems easy to design an IS estimator if we can well approximate $q^*(\mathbf{x})$ in Eq.(7). Yet, the following discussion shows that this is not an easy task. We use the KL divergence to measure the similarity between $q^*(\mathbf{x})$ and its approximation $q(\mathbf{x})$. We derive $q(\mathbf{x})$ by:

Theorem 3. *Assume the optimal importance distribution $q^*(\mathbf{x})$ in Eq.(7) is approximated by a product form distribution:*

$$q(\mathbf{x})=\prod_{i=1}^{N_l} q_i(x_i), \quad (10)$$

then the KL divergence $\text{KL}(q^*||q)$ is minimized when:

$$q_i(x_i)=\mathbb{P}[x_i|\mathcal{R}(\mathbf{x})=1]. \quad (11)$$

By now, the estimation of $\mathbb{P}[x_i|\mathcal{R}(\mathbf{x})=1]$ becomes a new problem. In fact, even given the exact $\mathbb{P}[x_i|\mathcal{R}(\mathbf{x})=1]$ expressions or values, the performance of this IS method is still not guaranteed: minimizing the KL divergence can only lower bound the estimator's variance, and the lower bound depends on how well $q(\mathbf{x})$ in Eq.(11) can approximate $q^*(\mathbf{x})$. To see this, consider the case where a network has 2 nodes connected by 3 parallel links with $\mathbf{p}=(0.001, 0.2, 0.001)$. The flow fails if link statuses \mathbf{x} are $(1,1,0)$, $(1,1,1)$ or $(0,1,1)$. According to Thm. 3, one possible

importance distribution is $\mathbf{q}=(0.50025, 1, 0.50025)$. In this case, $q^*(\mathbf{x})$ is not well approximated by $q(\mathbf{x})$ in Table 4.

Table 4: An example for the IS method

\mathbf{x}	(1, 1, 0)	(1, 1, 1)	(0, 1, 1)
$q^*(\mathbf{x})=\mathbb{P}[\mathbf{x} \mathcal{R}(\mathbf{x})=1]$	0.49975	0.00050	0.49975
$q(\mathbf{x})=\prod_{i=1}^{N_i} \mathbb{P}[x_i \mathcal{R}(\mathbf{x})=1]$	0.25000	0.25025	0.25000

Remark: The above example illustrates that minimizing KL divergence cannot provide the IS method a performance guarantee: sometimes our chosen $q(\mathbf{x})$ can be very different from $q^*(\mathbf{x})$. A major reason is that the baseline IS assumes $q(\mathbf{x})$ has the product form in Eq.(10), i.e., it considers link failures are independent and ignores the correlations among them. We will discuss how to improve the design of $q(\mathbf{x})$ in later sections.

2) Estimation error bound analysis: The variance bounds of the baseline IS method is given by the following theorem:

Theorem 4. *If the IS estimator in Eq.(4) takes $q(\mathbf{x})$ in Eq.(11) as its importance distribution, then $\mathbb{V}_q[\hat{\mu}_{\text{IS}}]$ is bounded by:*

$$\frac{\mu^2 \text{KL}(q^*||q)}{N} \leq \mathbb{V}_q[\hat{\mu}_{\text{IS}}] \leq \frac{\mu^2 \sqrt{2 \log 2 \text{KL}(q^*||q)}}{N \min_{\mathbf{x}} q(\mathbf{x})}. \quad (12)$$

Here q^* is the optimal importance distribution given by Eq.(7).

Remark: Thm. 4 implies that although the variance of baseline IS is upper bounded, it can be worse than the variance of MC. This motivates us to seek for more efficient sampling methods with better error bounds.

B. Conditions for Efficient Sampling

In the baseline IS design, link importance distributions $q_i(x_i)$ are assumed to be independent. This assumption simplifies the problem, but does not correspond to the reality. Consider the example in Section IV-A, $\mathbb{P}[x_3|\mathcal{R}=1, x_1=1]$ greatly differs from $\mathbb{P}[x_3|\mathcal{R}=1, x_1=0]$, which implies the dependence among $q_i(x_i)$ and the correlation of links' statuses cannot be ignored. Next, we propose our "sequential importance sampling" (SIS) based design and take this correlation into consideration.

1) ZV sequential importance sampling: Let us adapt the ZV importance distribution in Eq.(7) for the SIS method. Denote $\mathbf{x}_{1:i}=(x_1, x_2, \dots, x_i)$ as the status of the first i th links.

Theorem 5. *For the FAVE problem, the IS estimator in Eq.(4) achieves the ZV property if the importance distribution $q(\mathbf{x})=q^*(\mathbf{x})$ for the SIS estimator, where:*

$$q^*(\mathbf{x})=\prod_{i=1}^{N_i} q_i(\mathbf{x}) \quad \text{and} \quad q_i(\mathbf{x})=\frac{\mathbb{P}(\mathcal{R}=1|\mathbf{x}_{1:i})}{\mathbb{P}(\mathcal{R}=1|\mathbf{x}_{1:i-1})} p_i(x_i). \quad (13)$$

Remark: Different from baseline IS, SIS generates links' statuses in a sequential manner, which enables its importance distribution $q_i(\mathbf{x})$ to capture the correlation of links' statuses, i.e., links' importance distributions are dependent.

2) Conditions for good error bounds: To apply the above SIS estimator, we need to estimate (or approximate) $\mathbb{P}[\mathcal{R}=1|\mathbf{x}_{1:i}]$. The following theorem states that the above SIS is robust even if there exists some error when approximating $\mathbb{P}[\mathcal{R}=1|\mathbf{x}_{1:i}]$.

Theorem 6 (Conditions for BRE and VRE properties). *The IS estimator in Eq.(4) has the BRE property if $\forall i=1, \dots, N_i$, $\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i}]$ satisfies:*

$$\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i}]=O(\mathbb{P}[\mathcal{R}=1|\mathbf{x}_{1:i}]), \quad (14)$$

and the VRE property if $\forall i=1, \dots, N_i$, $\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i}]$ satisfies:

$$\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i}]=\mathbb{P}[\mathcal{R}=1|\mathbf{x}_{1:i}]+o(\mathbb{P}[\mathcal{R}=1|\mathbf{x}_{1:i}]). \quad (15)$$

Proof: Assume $\mathbb{P}[\mathcal{R}=1|\mathbf{x}_{1:i}]>0$. Otherwise there is no need to generate the corresponding \mathbf{x} . Then, Eq.(14) is equivalent to:

$$\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i}]=\mathbb{P}[\mathcal{R}=1|\mathbf{x}_{1:i}](\theta_i+o(1)), \quad (16)$$

where θ_i is a constant. Thus, we can derive that:

$$\frac{\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i}]}{\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i-1}]}=\frac{\mathbb{P}[\mathcal{R}=1|\mathbf{x}_{1:i}]}{\mathbb{P}[\mathcal{R}=1|\mathbf{x}_{1:i-1}]} \cdot \left(\frac{\theta_i}{\theta_{i-1}}+o(1)\right). \quad (17)$$

$$\begin{aligned} q(\mathbf{x}) &= \prod_{i=1}^{N_i} \frac{\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i}]}{\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i-1}]} p(x_i) \\ &= \prod_{i=1}^{N_i} \frac{\mathbb{P}[\mathcal{R}=1|\mathbf{x}_{1:i}]}{\mathbb{P}[\mathcal{R}=1|\mathbf{x}_{1:i-1}]} p(x_i) \cdot \prod_{i=1}^{N_i} \left(\frac{\theta_i}{\theta_{i-1}}+o(1)\right) \\ &= q^*(\mathbf{x}) \cdot \left(\prod_{i=1}^{N_i} \frac{\theta_i}{\theta_{i-1}}+o(1)\right). \end{aligned} \quad (18)$$

$$\sigma_q^2=\mu^2 \int \frac{q^*(\mathbf{x})^2}{q(\mathbf{x})} d\mathbf{x}-\mu^2=\mu^2 \left(\prod_{i=1}^{N_i} \frac{\theta_{i-1}}{\theta_i}-1+o(1)\right)=O(\mu^2). \quad (19)$$

Namely, BRE property is achieved. As Eq.(15) is a special case of Eq.(16) by restricting $\theta_i=1$ for all i , it is similar to show that Eq.(15) implies $\sigma_q^2=O(\mu^2)$, i.e., VRE property is achieved. ■

Remark: Thm. 6 provides important guidelines to design sampling methods with BRE and VRE properties: the estimation $\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i}]$ should satisfy the conditions listed in Thm. 6.

C. SEED Algorithms

1) SEED set and related definition: Before introducing how to approximate $\mathbb{P}[\mathcal{R}=1|\mathbf{x}_{1:i}]$ for the SIS estimator, we first present some definitions used in the later discussion.

Let $\Omega \triangleq \{1, \dots, N_i\}$ be the index set of all links. We use function $\Psi(\cdot)$ and $\Psi^{-1}(\cdot)$ for transformations between the failure configuration \mathbf{x} and the set of failure links $L \subseteq \Omega$:

$$\Psi(\cdot): L \rightarrow \mathbf{x}, \quad \text{where } \forall i \in L, x_i=1; \forall i \notin L, x_i=0. \quad (20)$$

$$\Psi^{-1}(\cdot): \mathbf{x} \rightarrow L, \quad \text{where } L=\{i|x_i=1\}. \quad (21)$$

For one interested flow, denote the collection of all link sets whose failures can result in the flow failure by:

$$\mathcal{F} \triangleq \{L | \mathcal{R}(\Psi(L))=1, L \subseteq \Omega\}. \quad (22)$$

Also, denote the collection of all supersets of link set L by:

$$\text{span}(L) \triangleq \{L' | L \subseteq L' \subseteq \Omega\}. \quad (23)$$

Accordingly, for a collection of link sets $\mathcal{L}=\{L\}$,

$$\text{span}(\mathcal{L}) \triangleq \bigcup_{L \in \mathcal{L}} \text{span}(L). \quad (24)$$

The probability of the event that "all links in L fail" is:

$$\Phi(L) \triangleq \prod_{i \in L} p_i = \sum_{L' \in \text{span}(L)} \mathbb{P}[\Psi(L')]. \quad (25)$$

Definition 6 (SEED). *Define the SEED S as a special link set which satisfies the following conditions:*

$$(1) S \in \mathcal{F}; \quad (2) \forall L \subseteq S, L \notin \mathcal{F}; \quad (3) \forall L \supseteq S, L \in \mathcal{F}.$$

We also denote the collection of all SEEDs by \mathcal{S} .

Definition 7 (Conditional SEED). *Consider the statuses of some links are specified by $\mathbf{x}_{1:i}$, define the conditional SEED (cond-SEED for short) $S(\mathbf{x}_{1:i})$ as a special link set satisfying:*

$$(1) S(\mathbf{x}_{1:i}) \subseteq \{i+1, \dots, N_i\}; \quad (3) \forall L \subseteq S(\mathbf{x}_{1:i}), L \cup \Psi^{-1}(\mathbf{x}_{1:i}) \notin \mathcal{F}; \\ (2) S(\mathbf{x}_{1:i}) \cup \Psi^{-1}(\mathbf{x}_{1:i}) \in \mathcal{F}; \quad (4) \forall L \supseteq S(\mathbf{x}_{1:i}), L \cup \Psi^{-1}(\mathbf{x}_{1:i}) \in \mathcal{F}.$$

Also, we denote the collection of all cond-SEEDs by $\mathcal{S}(\mathbf{x}_{1:i})$.

Examples: We give some examples for the above definitions. Consider the example in Section I, we have $\mathcal{S}=\{\{1\}, \{4,5\}\}$. By the definition of SEED, the failure of any subset of a SEED, e.g., $L=\{4\}$ or $L=\{5\}$, will not result in flow 1's failure; and the failure of any superset of a SEED, e.g., $L=\{1,2\}$, will result in flow 1's failure. If given the first four links' statuses by $\mathbf{x}_{1:4}=(0, 1, 0, 1)$, there is only one cond-SEED $S(\mathbf{x}_{1:4})=\{5\}$.

Algorithm 1 SEED Based ZV Sampling (SEED-ZV)

Input: The collection of all SEEDs \mathcal{S}

Output: An importance distribution $q(\mathbf{x})$ to achieve the ZV property and a sample of failure configuration \mathbf{x} .

```
1: for  $i=1$  to  $N_l$  do
2:    $x_i \leftarrow 1$  and  $\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i}] \leftarrow 0$ ;
3:   for all  $\emptyset \neq A \subseteq \mathcal{S}(\mathbf{x}_{1:i})$  do
4:      $\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i}] \leftarrow \hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i}] + (-1)^{|A|-1} \Phi\left(\bigcup_{S(\mathbf{x}_{1:i}) \in A} S(\mathbf{x}_{1:i})\right)$ ;
5:   Keep  $x_i=1$  with  $q_i(\mathbf{x}) = \frac{\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i}]}{\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i-1}]} p_i$ ;
6:    $\mathcal{S}(\mathbf{x}_{1:i}) \leftarrow \text{UPDATECONDSEED}(i, x_i, \mathcal{S}(\mathbf{x}_{1:i-1}))$ 
7:   function  $\text{UPDATECONDSEED}(i, x_i, \mathcal{S}(\mathbf{x}_{1:i-1}))$ 
8:     if  $x_i = 1$  then
9:       return  $\mathcal{S}(\mathbf{x}_{1:i}) \leftarrow \{L \setminus \{i\} \mid L \in \mathcal{S}(\mathbf{x}_{1:i-1})\}$ 
10:    else
11:    return  $\mathcal{S}(\mathbf{x}_{1:i}) \leftarrow \{L \mid L \in \mathcal{S}(\mathbf{x}_{1:i-1}), i \notin L\}$ 
```

Next, we use ‘‘SEED’’ and ‘‘SEED set’’ to capture the correlation of link failures and approximate $\mathbb{P}[\mathcal{R}=1|\mathbf{x}_{1:i}]$ in Thm. 5.

2) SEED based IS algorithms: Thm. 5 presents the optimal importance distribution $q^*(\mathbf{x})$ for the SIS method. We next show that $q^*(\mathbf{x})$ can be computed exactly via Algorithm 1.

Theorem 7. *If failure configurations \mathbf{x} are generated using SEED-ZV, the estimator in Eq. (4) has the ZV property.*

Remark: Though SEED-ZV has the ZV property, it is computationally expensive: as the need for traversing all combinations of cond-SEEDs $S(\mathbf{x}_{1:i}) \in \mathcal{S}$ for each link e_i , the computational complexity is $O(N_l \cdot 2^{|\mathcal{S}|})$.

Note that there is a tradeoff between estimation accuracy and computational complexity to estimate $q^*(\mathbf{x})$. Next, we consider sacrificing some estimation accuracy of $q^*(\mathbf{x})$, i.e., *utilize probabilities of the most important cond-SEEDs* for estimations (in line 3 of Algorithm 2), and propose SEED-BRE, which has a lower linear computational complexity. We show SEED-BRE has the BRE property by the following theorem.

Theorem 8. *If failure configurations \mathbf{x} are generated using SEED-BRE, the estimator in Eq. (4) has the BRE property.*

Remark: The computational complexity of SEED-BRE is only $O(N_l |\mathcal{S}|)$, as it needs to traverse all cond-SEEDs in $\mathcal{S}(\mathbf{x}_{1:i})$, for each e_i . The size of $\mathcal{S}(\mathbf{x}_{1:i})$ decreases when i increases.

To further improve estimation accuracy, we *utilize the probability sum of cond-SEEDs for estimations* (in line 3 of Algorithm 3), and propose the SEED-VRE algorithm. We next show that SEED-VRE has the VRE property if link failures are rare.

Theorem 9. *When link failure probabilities are small and in the form of $p_i = O(\epsilon)$, $\forall i \in \Omega$. If failure configurations \mathbf{x} are generated by SEED-VRE, the estimator in Eq. (4) has the VRE property.*

Remark: The computational complexity of SEED-VRE is also $O(N_l |\mathcal{S}|)$, as it needs to traverse all the cond-SEED sets in $\mathcal{S}(\mathbf{x}_{1:i})$, for each e_i .

By now, we have three SEED algorithms, i.e., SEED-ZV, SEED-BRE and SEED-VRE, to compute the importance distributions of SIS estimator. They can achieve ZV, BRE and VRE properties respectively, and with the computational complexities $O(N_l \cdot 2^{|\mathcal{S}|})$, $O(N_l |\mathcal{S}|)$ and $O(N_l |\mathcal{S}|)$ respectively. However, all above discussions focus on the ‘‘single flow case’’. Next, we will generalize our methods to handle multiple flows and take other practical issues into consideration.

Algorithm 2 SEED Based BRE Sampling (SEED-BRE)

Input: The collection of all SEEDs \mathcal{S}

Output: An importance distribution $q(\mathbf{x})$ to achieve the BRE property and a sample of failure configuration \mathbf{x} .

```
1: for  $i=1$  to  $N_l$  do
2:   for  $x_i$  in  $\{0, 1\}$  do
3:      $\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i}] \leftarrow \max_{S(\mathbf{x}_{1:i}) \in \mathcal{S}(\mathbf{x}_{1:i})} \Phi(S(\mathbf{x}_{1:i}))$ ;
4:      $q_i(x_i) \leftarrow \frac{\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i}]}{\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i-1}]} p_i(x_i)$ ;
5:   Normalize  $q_i(\mathbf{x})$  by  $q_i(1) \leftarrow \frac{q_i(1)}{q_i(1)+q_i(0)}$ ,  $q_i(0) \leftarrow 1 - q_i(1)$ ;
6:   Set  $x_i$  as 1 with probability  $q_i(1)$  and 0 with probability  $q_i(0)$ ;
7:    $\mathcal{S}(\mathbf{x}_{1:i}) \leftarrow \text{UPDATECONDSEED}(i, x_i, \mathcal{S}(\mathbf{x}_{1:i-1}))$ 
```

Algorithm 3 SEED Based VRE Sampling (SEED-VRE)

Input: The collection of all SEEDs \mathcal{S}

Output: An importance distribution $q(\mathbf{x})$ to achieve the VRE property and a sample of failure configuration \mathbf{x} .

```
1: for  $i=1$  to  $N_l$  do
2:   for  $x_i$  in  $\{0, 1\}$  do
3:      $\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i}] \leftarrow \sum_{S(\mathbf{x}_{1:i}) \in \mathcal{S}(\mathbf{x}_{1:i})} \Phi(S(\mathbf{x}_{1:i}))$ ;
4:      $q_i(x_i) \leftarrow \frac{\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i}]}{\hat{\mathbb{P}}[\mathcal{R}=1|\mathbf{x}_{1:i-1}]} p_i(x_i)$ ;
5:   Normalize  $q_i(\mathbf{x})$  by  $q_i(1) \leftarrow \frac{q_i(1)}{q_i(1)+q_i(0)}$ ,  $q_i(0) \leftarrow 1 - q_i(1)$ ;
6:   Set  $x_i$  as 1 with probability  $q_i(1)$  and 0 with probability  $q_i(0)$ ;
7:    $\mathcal{S}(\mathbf{x}_{1:i}) \leftarrow \text{UPDATECONDSEED}(i, x_i, \mathcal{S}(\mathbf{x}_{1:i-1}))$ 
```

V. PRACTICAL CONSIDERATION

Previous discussions illustrate that SEED methods work well in estimating a single flow’s availability with the full information of \mathcal{S} . Next, we consider more practical issues. First, as there are $O(N_v^2)$ flows in the network, it is costly to design a ‘‘customized’’ estimator for each flow and individually estimate their availabilities. If the designed estimator works for a group of flows, the computational cost can be reduced significantly. Furthermore, as SEED methods rely on the SEED set which may be difficult to obtain the full information at times, we consider the case that only a partial information of \mathcal{S} is available, e.g., we only know some frequently observed SEEDs.

A. Generalization to Multiple Flows Case

To provide efficient and accurate availability estimations for a group of flows at the same time, one possibility is to utilize SEED methods to design a *pure importance distribution* $q^{(k)}(\mathbf{x})$ for each flow $f_k \in F$, then take a mixture of these pure distributions with a strategy \mathcal{M} to simulate link failures:

$$q(\mathbf{x}) = \mathcal{M}(q^{(1)}(\mathbf{x}), \dots, q^{(N_f)}(\mathbf{x})). \quad (26)$$

To derive such a *mixture importance distribution*, we take a weighted sum of these pure distributions:

$$q(\mathbf{x}) = \sum_k w_k q^{(k)}(\mathbf{x}), \sum_k w_k = 1. \quad (27)$$

Here w_k can be viewed as the probability of taking $q^{(k)}(\mathbf{x})$ to generate \mathbf{x} . Denote $\mu^{(k)}$ as f_k ’s failure probability and $(\sigma_q^{(k)})^2$ as the one-run variance when taking $q(\mathbf{x})$ to estimate f_k ’s failure probability. Next we analyze error bounds of this mixture sampling strategy, when applying to the multiple flows case.

Theorem 10. *Using the mixture sampling strategy in Eq. (27) with pure distributions $q^{(k)}(\mathbf{x})$ generated by SEED methods, the IS estimator achieves the BRE property for all flows availability estimations. Specifically, for flow f_k , the estimator’s one-run variance satisfies:*

$$\left(\sigma_q^{(k)}\right)^2 \leq \left(\frac{1}{w_k} - 1\right) \left(\mu^{(k)}\right)^2 + \frac{1}{w_k} \left(\sigma_{q^{(k)}}^{(k)}\right)^2. \quad (28)$$

Remark: Thm. 10 states that, when extending to the multiple flows case, our methods guarantee the estimation efficiency for all flows. Designing proper or even optimal weights $\{w_k\}$ is challenging. Online learning is a good approach to find a more efficient weight setting, and we leave this as a future work.

B. Partial Seed Set Information

At times, it may be difficult to obtain a “full” SEED set \mathcal{S} , especially when the network is large and flow failures are rare. To provide robust estimations, consider the case that we have only a partial information of \mathcal{S} , e.g., limited historical data of flow failures which gives $\mathcal{S}' \subset \mathcal{S}$. Denote the cond-SEED set induced by \mathcal{S}' and $\mathbf{x}_{1:i}$ as $\mathcal{S}'(\mathbf{x}_{1:i})$. To analyze error bound properties of SEED algorithms, we provide the following lemma.

Lemma 1. *Given a partial SEED set $\mathcal{S}'(\mathbf{x}_{1:i})$, when estimating $\mathbb{P}[\Psi^{-1}(\mathbf{x}) \in \text{span}(\mathcal{S}'(\mathbf{x}_{1:i})) | \mathbf{x}_{1:i}]$: SEED-ZV and SEED-BRE have ZV and BRE properties respectively; assume link failure probabilities are small and follow the form of $p_i = O(\epsilon)$, $\forall i \in \Omega$, SEED-VRE have the VRE property.*

Remark: Thm.5 states that the estimation accuracy depends on how well $\mathbb{P}[\mathcal{R}=1 | \mathbf{x}_{1:i}]$ is approximated. Given \mathcal{S} , SEED methods have good error bound properties for they can well approximate $\mathbb{P}[\Psi^{-1}(\mathbf{x}) \in \text{span}(\mathcal{S}(\mathbf{x}_{1:i})) | \mathbf{x}_{1:i}] = \mathbb{P}[\mathcal{R}=1 | \mathbf{x}_{1:i}]$. However, given a partial SEED set \mathcal{S}' , the bias between $\mathbb{P}[\mathcal{R}=1 | \mathbf{x}_{1:i}]$ and $\mathbb{P}[\Psi^{-1}(\mathbf{x}) \in \text{span}(\mathcal{S}'(\mathbf{x}_{1:i})) | \mathbf{x}_{1:i}]$ should be considered. Let us consider the following two cases:

- 1) **Good coverage case.** SEED methods maintain good error bound properties if \mathcal{S}' has a *good coverage*, which is defined formally as:

$$\mathbb{P}[\Psi^{-1}(\mathbf{x}) \in \text{span}(\mathcal{S}')] = \mathbb{P}[\Psi^{-1}(\mathbf{x}) \in \text{span}(\mathcal{S})] (1 + o(1)). \quad (29)$$

Here is an example of such a partial SEED set:

$$\mathcal{S}' = \{S | S \in \mathcal{S}, |S| \leq \min_{S_j \in \mathcal{S}} |S_j| + k\}, \text{ where } p_i = O(\epsilon), k \geq 0.$$

- 2) **Poor coverage case.** \mathcal{S}' may not have a good coverage, e.g., without prior knowledge of network and flow failures, $\mathcal{S}' = \emptyset$. In such case, the SEED set information can be collected via pre-samplings and updated while simulating flow failures.

VI. EVALUATION OF SEED METHODS

We evaluate our methods on both an illustrative small scale network and a realistic network with topology and traffic matrices extracted from the Abilene network [14]. The simulation cost to guarantee the estimation error (or relative error) below a constant δ is $N \geq \frac{4\alpha^2}{\delta^2} \sigma_q^2$ (or $N \geq \frac{4\alpha^2}{\delta^2} (\frac{\sigma_q}{\mu})^2$). Hence, we use *one-run variance* σ_q^2 and *coefficients of variation* (CV) $\varepsilon_{CV} = \frac{\sigma_q}{\mu}$ to quantify the estimation efficiency. The variance reduction, i.e., $\frac{\sigma_{MC}^2}{\sigma_{SEED}^2}$, can imply the simulation cost reduction, i.e., $\frac{N_{MC}}{N_{SEED}}$.

A. Experiments on an Illustrative Network

The illustrative network¹ demonstrates the “best achievable theoretical improvements” using our method, compared with MC and baseline IS. We start with the single flow case, where full SEED sets or partial SEED sets with good coverage are provided. Then we extend it to the multiple flows case.

Experiment setting: The network is modelled as a directional multigraph G with five nodes and 12 links as depicted in Fig.2.

¹This is provided so that readers can simulate and validate our methods.

For each link e_i : p_i is uniformly distributed over $[0.5\epsilon, 1.5\epsilon]$ (ϵ is a small positive number); c_i is uniformly distributed over $\{50, 80, 100, 200\}$. The flow set F contains 18 flows. For each flow f_k : the source and destination are randomly selected; d_k is the bandwidth demand and uniformly distributed over $[5, 25]$. We consider traffic engineering follows the shortest path and max-min fairness policies. Note that the above setting provides an instance of network routing function $\mathcal{R}(\cdot)$ in Eq.(8).

Single flow analysis: We start with the single flow case, and select one particular flow f_{i^*} from the 18 flows. The detailed information of f_{i^*} , together with the SEED set information, are introduced in notes of Table 5. We compare our SEED methods with MC and baseline IS. The comparison result, including the expectation μ , theoretical one-run variance σ_q^2 and CV ε_{CV} , are summarized in Table 5. Let $\epsilon = 0.05$: given a full SEED set \mathcal{S} , SEED-BRE and SEED-VRE achieve variance reductions of around 2,000 and 360,000 times compared with MC, and around 45 and 7,000 times compared with baseline IS; given a partial SEED set \mathcal{S}' , our SEED methods estimate flow availabilities with very small biases and much smaller variances, i.e., with a small simulation cost, the estimation can be very close to the theoretical value. We also reduce ϵ from 0.05 to 0.001, to validate the vanishing property of SEED methods. While μ reduces with the decreasing ϵ : ε_{CV} of MC increases significantly as we have discussed in Section III; ε_{CV} of baseline IS is relatively stable; ε_{CV} of SEED methods reduces significantly, and ε_{CV} of SEED-VRE even achieves a 300 times reduction.

Multiple flows analysis: Next, we take all 18 flows into consideration. Let $\epsilon = 0.001$. We consider an equally weighted sum of flows’ pure importance distributions as the mixture SIS distribution. Fig.3(a) illustrates cumulative distributions of ε_{CV} if pure distributions are generated by different methods. With SEED methods, for around 80% flows $\varepsilon_{CV} \leq 2$, which is smaller than the best case of ε_{CV} of baseline IS. This demonstrates the BRE property of SEED methods as stated in Thm. 10. Furthermore, both SEED methods and baseline IS, their ε_{CV} are 1,000 times smaller than that of MC. To depict the variance reduction compared with MC much clearer, Fig.3(b) shows cumulative distributions of the variance reduction compared with MC. With SEED methods, more than 80% flows have *variance reductions* $\sigma_{MC}^2 / \sigma_{SEED}^2 > 200,000$. To better illustrate the efficiency improvement, Table 6 summarizes simulation costs to guarantee that for 80% flows, “with 95% confident the relative error is less than 0.01”, i.e., $\alpha_{95} \sigma_q / \sqrt{N} \leq 0.01 \hat{\mu}$.

Table 6: Summary of simulation costs (N steps)

	MC	IS	SEED-BRE
Simulation cost (N)	3.11×10^{10}	7.78×10^5	2.01×10^5

B. Experiments on a Realistic Network

Next, consider a realistic network to demonstrate the “improvements in practice” by using our methods. As it is hard to obtain the full SEED sets information in the complex realistic case, simulations on the realistic network can validate the efficiency of SEED methods when estimating all flows’ availabilities, given partial SEED sets with poor coverage property.

Experiment setting: We use the Abilene network [14], [20] with topology and traffic matrices collected by [21]. The net-

Table 5: Flow unavailability analysis results for the network in Fig.2

		Full SEED Set					Partial SEED Set		
		MC	IS	SEED-ZV	SEED-BRE	SEED-VRE	SEED-ZV	SEED-BRE	SEED-VRE
$\epsilon=0.05$	μ	3.808×10^{-3}	3.808×10^{-3}	3.808×10^{-3}	3.808×10^{-3}	3.808×10^{-3}	3.520×10^{-3}	3.520×10^{-3}	3.520×10^{-3}
	σ_q^2	3.793×10^{-3}	7.549×10^{-5}	0	1.663×10^{-6}	1.059×10^{-8}	0	1.145×10^{-7}	1.277×10^{-9}
	σ_q/μ	1.617×10^1	2.281	0	3.387×10^{-1}	2.702×10^{-2}	0	8.886×10^{-2}	9.384×10^{-3}
$\epsilon=0.01$	μ	1.419×10^{-4}	1.419×10^{-4}	1.419×10^{-4}	1.419×10^{-4}	1.419×10^{-4}	1.393×10^{-4}	1.393×10^{-4}	1.393×10^{-4}
	σ_q^2	1.418×10^{-4}	1.323×10^{-7}	0	6.586×10^{-10}	1.523×10^{-13}	0	4.282×10^{-11}	1.743×10^{-14}
	σ_q/μ	8.395×10^1	2.564	0	1.809×10^{-1}	2.751×10^{-3}	0	4.697×10^{-2}	9.477×10^{-4}
$\epsilon=0.001$	μ	1.392×10^{-6}	1.392×10^{-6}	1.392×10^{-6}	1.392×10^{-6}	1.392×10^{-6}	1.389×10^{-6}	1.389×10^{-6}	1.389×10^{-6}
	σ_q^2	1.392×10^{-6}	1.348×10^{-11}	0	6.903×10^{-15}	1.563×10^{-20}	0	4.431×10^{-16}	1.768×10^{-21}
	σ_q/μ	8.48×10^2	2.637	0	5.969×10^{-2}	8.982×10^{-5}	0	1.515×10^{-2}	3.026×10^{-5}

Note: Information of flow f_{i^*} : 1) the tuple of source, destination and demand $(s_{i^*}, t_{i^*}, d_{i^*}) = (4, 3, 20.25)$; 2) the full SEED set $\mathcal{S} = \{\{2, 5\}, \{7, 8\}, \{2, 3, 8\}, \{2, 8, 12\}, \{5, 6, 7\}, \{5, 7, 11\}\}$; 3) the partial SEED set with good coverage $\mathcal{S}' = \{\{2, 5\}, \{7, 8\}\}$.





Fig. 3: Performance comparison for the illustrative network (multiple flow case).

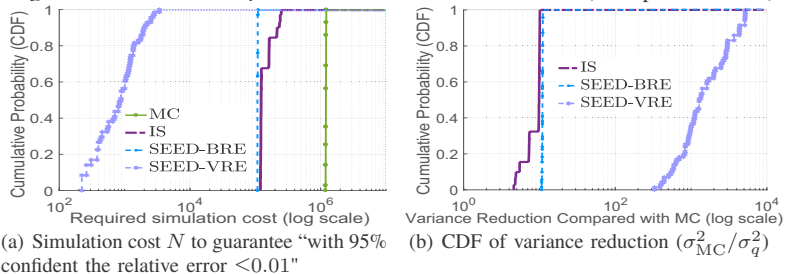


Fig. 4: Topology of the Abilene network.

Fig. 5: Performance comparison for the Abilene network (multiple flow case).

Fig. 4: Topology of the Abilene network.

work contains 12 nodes and 30 links. The flow set contains 132 aggregated flows: all flows with the same source and destination are aggregated as a single flow². We take each flow's peak (99 percentile) throughput [22] as its raw demand. As Abilene has a sufficient capacity to serve raw demands, we double raw demands to see whether the network can still support oversubscribed demands. The routing follows the shortest path policy. The capacity allocation follows the max-min fairness policy, which is also adopted by Google's B4 backbone network [23]. **Multiple flows analysis:** We consider all aggregated flows and estimate their availabilities at the same time. Due to the high dimensionality of FAVE in this realistic network, it is costly to obtain theoretical variances of different methods. Thus, we run each method 10,000 times and use the *empirical variance* $\hat{\sigma}_q^2$ [24] to estimate the one-run variance σ_q^2 and compute the simulation cost N to guarantee that "with 95% confident the relative error is below 0.01". Fig.5(a) shows cumulative distributions of N by taking the mixture of pure distributions generated by different methods³. With SEED-VRE, we find that to achieve the desired accuracy level, for around 60% flows the required simulation costs $N \leq 1,000$, and for around 80% flows $N \leq 1,400$. Simulation costs for SEED-BRE, baseline IS and MC methods to guarantee 80% flows to achieve the accuracy target are 100,000, 180,000 and 1,260,000, respectively. So the efficiency is improved by around 900 times via SEED-VRE and

²We take the Abilene network as an example and consider aggregated flows due to limitations of the accessible realistic traffic data.

³Due to the exponential complexity, SEED-ZV is not applied in this case.

13 times via SEED-BRE, compared with MC. Fig.5(b) illustrates cumulative distributions of the variance reduction compared with MC. With SEED methods, 80% of the flows have variance reductions larger than 900 times.

VII. APPLICATIONS IN CAPACITY PLANNING

Now, we demonstrate the utility of our method in capacity planning on the Abilene network.

Consider the case where a network provider needs to build new links with certain capacities so as to achieve all flow availability targets. For a capacity planning proposal \mathcal{CP} , we can *evaluate its feasibility and adopt it if feasible*, as follows:

- We obtain *availability feedbacks* of each flow f_i , including a flow availability estimation $\hat{\mu}_i$ with upper (lower) confidence bound $\overline{\mu}_i$ ($\underline{\mu}_i$) computed via empirical variances.
- f_i 's availability target o_i is *achieved* if $\mu_i \geq o_i$ and *unreached* if $\overline{\mu}_i < o_i$.⁴ \mathcal{CP} is *feasible* if all flows' availability targets are achieved; and *infeasible* otherwise.

By evaluating flow availabilities, we can not only *determine which proposal allows the network to provide better flow availability guarantees*, but also *utilize flow availability feedbacks for further refinements of infeasible proposals*. We consider the following information to refine an infeasible proposal:

- **Link capacity utilization based:** The utilization metric is the primary metric of interest in capacity planning [1]. Hence, link capacity utilizations can imply the importance of links.

⁴We run enough simulations to guarantee $o_i \notin [\underline{\mu}_i, \overline{\mu}_i]$.

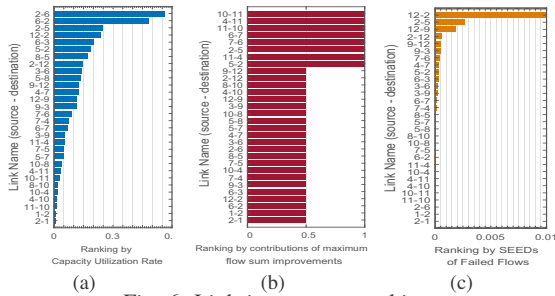


Fig. 6: Link importance ranking

- **Maximum flow based:** The maximum flow value is a widely adopted network reliability measure [11]–[13]. We take the increased of the sum of maximum flows brought by increasing one unit link capacity to measure the importance of links.
- **SEED based:** By testing flow availabilities, we get \mathcal{F}' , the set of flows with unsatisfied availability targets. The number of failed flows in \mathcal{F}' when e_j fails ($\sum_{f_i \in \mathcal{F}'} \mathbb{P}[\mathcal{R}_i=1|x_j=1]$) can be estimated using SEED algorithms and can imply the importance of links.⁵

Let the current capacity design of Abilene be the original proposal \mathcal{CP}_0 , and twice the raw demands be flow demands, such that \mathcal{CP}_0 is infeasible. We rank links using the above metrics and summarize rankings in Fig.6. Assume the provider has a budget and only affords to build four new links, each has a capacity of 2.5Gbps and a failure probability of 0.01. Based on rankings in Fig.6, we have three proposals, i.e., \mathcal{CP}_1 , \mathcal{CP}_2 and \mathcal{CP}_3 by taking the top four links in Fig.6(a), 6(b) and 6(c). Fig.7(a) shows flow availability evaluation results. As our method selects links with the largest impact on flow failures, it achieves *greater flow availability improvements*: in \mathcal{CP}_3 , flow availabilities of around 80% of the flows reach 99.9%.

Insights 1: *Improper capacity planning offers little help on improving flow availabilities.* E.g., although \mathcal{CP}_2 maximizes the sum of maximum flows, it does not consider the distribution of traffic demands across the network, and thus only brings little improvements on flow availabilities.

Insights 2: *Link utilization is not always the best indicator of capacity planning.* With the shortest path routing policy, link e_i has a high capacity utilization if many flows' shortest paths go through it. Yet, if it is easy to find some alternate link when e_i fails, e_i 's failure will not result in flow failures and so e_i is not the most important if aiming at improving flow availabilities.

VIII. CONCLUSION

In this paper, we propose fast and accurate methods in solving the FAVE problem. We introduce the concept of "SEED" to determine the importance of roles played by different links in flow failures, and propose three SEED based SIS methods which achieve BRE and VRE properties with linear computational complexities. To provide robust and scalable estimations, we extend FAVE to the multiple flows case and partial SEED set case, and our methods maintain the estimation efficiency. We apply our methods on both an illustrative network and a realistic network, and our methods reduce the simulation cost by

⁵As $\mathbb{P}[\mathcal{R}_i=1|x_{1:k}]$, the probability that flow f_i fails given the first k th links' statuses $\mathbf{x}_{1:k}$, can be estimated using SEED algorithms, $\mathbb{P}[\mathcal{R}_i=1|x_j=1]$ can also be estimated by taking e_j as the first link and let $k=1$, $\mathbf{x}_{1:k}=1$.

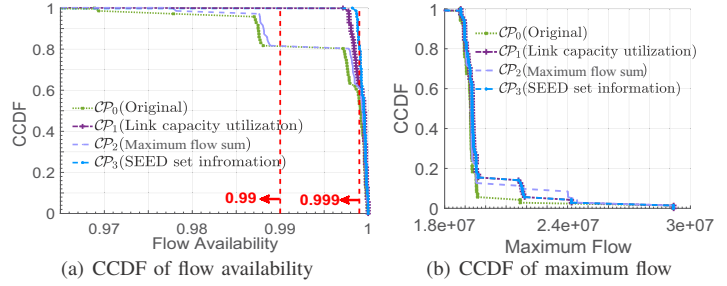


Fig. 7: An example of applying FAVE to the capacity planning.

around 900 and 130 times compared with MC and baseline IS methods on the Abilene network. We show that our method facilitates capacity planning by providing more accurate network reliability estimations compared with classical methods, and greater flow availability improvements compared with solely using link capacity utilizations for capacity planning.

REFERENCES

- [1] Cisco.Com, "Best practices in core network capacity planning," <https://communities.cisco.com/docs/DOC-35673>, 2013.
- [2] —, "Planning and designing networks with the Cisco MATE portfolio," <https://communities.cisco.com/docs/DOC-36973>, 2013.
- [3] S. J. Taylor and B. Letham, "Prophet: forecasting at scale," <https://research.fb.com/prophet-forecasting-at-scale/>, 2017.
- [4] A. K. Bangla, A. Ghaffarkhah *et al.*, "Capacity planning for the Google backbone network," in *Proc. ISMP*, 2015.
- [5] J. M. Hammersley and D. C. Handscomb, "The general nature of Monte Carlo methods," in *Monte Carlo Methods*, 1964, pp. 1–9.
- [6] P. W. Glynn and D. L. Iglehart, "Importance sampling for stochastic simulations," *Manag. Sci.*, 1989.
- [7] Y. Jiang, R. Li *et al.*, "The method of network reliability and availability simulation based on Monte Carlo," in *Proc. IEEE ICQR2MSE*, 2012.
- [8] W. C. Yeh, Y. C. Lin *et al.*, "A particle swarm optimization approach based on Monte Carlo simulation for solving the complex network reliability problem," *IEEE Trans. Rel.*, 2010.
- [9] P. L'Ecuyer, G. Rubino *et al.*, "Approximate zero-variance importance sampling for static network reliability estimation," *IEEE Trans. Rel.*, 2011.
- [10] K. Lee, H.-W. Lee, and E. Modiano, "Reliability in layered networks with random link failures," in *Proc. IEEE INFOCOM*, 2010.
- [11] Z. I. Botev, S. Vaisman *et al.*, "Reliability of stochastic flow networks with continuous link capacities," in *Proc. IEEE WSC*, 2014.
- [12] I. Gertsbakh, R. Rubinstein *et al.*, "Permutational methods for performance analysis of stochastic flow networks," *Probab. Eng. Information Sci.*, 2014.
- [13] E. Datta and N. K. Goyal, "Reliability estimation of stochastic flow networks using pre-ordered minimal cuts," in *Proc. IEEE MicroCom*, 2016.
- [14] "Abilene," <https://www.internet2.edu/>.
- [15] J. Xiong and W. Gong, "A novel algorithm on network reliability estimation," *Mathematical and computer modelling*, 2005.
- [16] M. O. Ball, "Computational complexity of network reliability analysis: An overview," *IEEE Trans. Rel.*, 1986.
- [17] A. M. Shooman, "Algorithms for network reliability and connection availability analysis," in *Proc. IEEE ELECTRO*, 1995.
- [18] A. E. Conway, "Fast simulation of service availability in mesh networks with dynamic path restoration," *IEEE TON*, 2011.
- [19] "Technical report," <https://dl.dropboxusercontent.com/s/5tjzn90wsys81h1/FAVE.pdf?dl=0>.
- [20] W. Jiang, R. Zhang-Shen *et al.*, "Cooperative content distribution and traffic engineering in ISP network," in *Proc. ACM SIGMETRICS*, 2009.
- [21] Y. Zhang, "Six months of Abilene traffic matrices," <http://www.cs.utexas.edu/~yzhang/>, 2014.
- [22] I. Raicu, I. Foster, and A. Szalay, "Harnessing grid resources to enable the dynamic analysis of large astronomy datasets," in *Proc. IEEE Supercomputing*, 2006.
- [23] S. Jain, A. Kumar *et al.*, "B4: Experience with a globally-deployed software defined WAN," in *Proc. ACM SIGCOMM*, 2013.
- [24] A. Müller, "Improved variance reduced Monte-Carlo simulation of in-the-money options," *J. Math. Finance*, 2016.