FAVE: A Fast and Efficient Network Flow AVailability Estimation Method With Bounded Relative Error

Tingwei Liu and John C. S. Lui, *Fellow, IEEE*

***Abstract*— Capacity planning and sales projection are essential tasks for network operators. This work aims to help network providers to carry out network capacity planning and sales projection by answering: Given topology and capacity, whether the network can serve current flow demands with high probabilities? We name such probability as the "*flow availability*", and present the flow availability estimation (FAVE) problem with generalizing the classical network connectivity based and maximum flow based reliability estimations. To quickly estimate flow availabilities, we utilize correlations among link and flow failures to figure out the importance of roles played by different links in flow failures (i.e., flow demands could not be satisfied). And we design three sequential importance sampling (SIS) estimation methods, which are: (1) *Accurate and efficient*: They achieve a bounded or even vanishing relative error with linear computational complexities. Hence they can provide more accurate estimations in less simulation time. (2) *Robust and scalable:* They maintain such estimation efficiencies even if only a partial SEED set information is available, or when the FAVE problem is extended to the multiple flows case. When applying to a realistic backbone network, our method can reduce the flow availability estimation cost by 900 and 130 times compared with MC and baseline IS methods; and also facilitate capacity planning and sales projection by providing better flow availability guarantees, compared with traditional methods.**

***Index Terms*— Flow availability estimation, sequential importance sampling (SIS), capacity planning.**

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]. To provide a better end-user experience and at the same time, keeping the operating cost at an affordable level, effective capacity planning tools are attracting network providers' attentions. Various systems have been built around this problem, such as Cisco's MATE [2], Facebook's Prophet [3], Juniper's WANDL [4], and Google's backbone network capacity planning tool [5].

SLA objectives mentioned above are called the *demand* of traffic flow, and the demand satisfaction probability is defined as the *flow availability*. One key concern for 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 the example in Fig. 1 where a network provider needs to serve two flows. In the network, each link i is associated with a failure probability 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 specifying the lower bound probability that its requested bandwidth needs to be satisfied. The network provider may want to perform:

- **Flow availability testing**: Whether the flow availability targets can be achieved? For example, flow 1's availability target is achieved if flow 1 obtains 10 units of bandwidth with a probability of no less than 0.9999.
- **Capacity planning**: 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?
- **Sales projection**: To increase the profit, can the provider admit a new flow or upgrade flow 1's bandwidth quota by 20 units, while maintaining flow availability guarantees?

All the above cases need flow availabilities, in order to check whether the network is able to serve the flow demands proposed by sales projection, with the given flow availability targets. We name this flow availability estimation problem as FAVE, and we give a formal definition later.

Realistic networks are often large and with intricate failure patterns, making flow availabilities unable to be evaluated analytically, especially when also taking the traffic scheduling into account. Hence, simulation (or sampling) is often used. Among various sampling methods for FAVE, the Monte Carlo (MC) method [6], which simulates link failures with their nature probabilities, is the most widely used. Yet, it is costly for MC to achieve the desired accuracy: The variance results in Table I implies that, in the example of Fig. 1, to estimate flow 2's availability with guaranteeing the 95% confidence interval (CI) width below 10^{-3}, MC takes at least 3,840 simulation steps. To simulate a large realistic network with many flows, where even a single step is expensive and can take hours, it is important to find ways to reduce the simulation steps.

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 central idea of importance sampling (IS) [7]. In general, designing an efficient IS distribution is highly challenging and problem-dependent. In the case of FAVE, we first introduce a baseline IS solution, where we use "*the correlation between link failures and flow failures*" to decide important links whose failures are more likely to result in flow failures. Specifically, consider the links that frequently fail at the same time when flow fails (e.g., link 1, 4 and 5), we simulate their failures more often to speed up simulating flow failures.

Manuscript received April 1, 2019; revised October 6, 2019 and December 3, 2019; accepted December 16, 2019; approved by IEEE/ACM Transactions on Networking Editor I.-H. Hou. Date of publication February 11, 2020; date of current version April 16, 2020. The work of John C. S. Lui was supported in part by the GRF under Grant 14201819. *(Corresponding author: John C. S. Lui.)*

The authors are with the Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong (e-mail: twliu@cse.cuhk.edu.hk; cslui@cse.cuhk.edu.hk).

Digital Object Identifier 10.1109/TNET.2020.2965161

TABLE I

ACCURACY, EFFICIENCY, AND COMPUTATIONAL COST (I.E., SIMULATION STEPS) COMPARISON

Flow	Theoretical Unavailability	Monte Carlo Method			Our method (SEED-VRE)		
		Estimated Unavailability	Empirical Variance	Number of Simulations	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

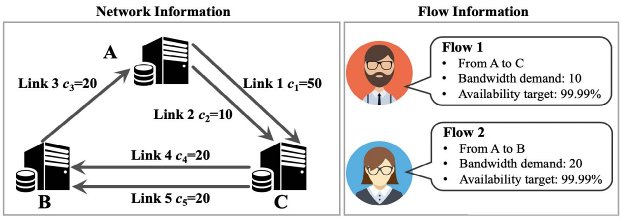

Fig. 1. An example of the FAVE problem.

To further improve upon the baseline IS, we show that some link sets fail more often and play a role of root causes of flow failures. We name such a link set as SEED, and introduce our SEED based sequential importance sampling (SIS) solution, where we use "*the correlation among link failures*" captured by SEEDs to decide important link sets whose failures are more likely to result in flow failures. Specifically, consider the link sets that appear more frequently in observed failures of SEEDs are important. We propose three SEED algorithms, in which the SEED-VRE algorithm has vanishing relative error (VRE) and linear computational complexity. Table I illustrates the computational reduction of our method: With only 10 simulation steps, SEED-VRE achieves much smaller empirical variance and estimation error than MC with 10,000 simulation steps. For a large and complex network, where each simulation step can take hours, this implies the network operator can either significantly speed up simulations, or get more accurate flow availability estimations without incurring additional costs.

We emphasize that the flow availability estimation is crucial to both capacity planning and sales projection: In 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 the classical methods to allocate more capacities for links with high bandwidth utilizations; In sales projection, for flows with flow availabilities far exceeding the availability targets, we allow such flows to request more bandwidth, which results in higher total sales.

Contributions: Our key contribution is in providing efficient and accurate solutions for FAVE. Contributions include:

- We generalize the classical network reliability estimation problem [8]–[14], and formally define the network Flow AVailability Estimation (FAVE) problem.
- For the single flow case, we introduce a novel concept of SEED and propose three advanced SIS algorithms that have attractive properties of bounded relative error (BRE) or VRE with only linear computational complexities.
- For the multiple flow case, we offer a mixture SIS method which maintains BRE property and linear computational complexity when estimating availabilities for all flows.
- For the partial SEED set information case, we maintain BRE and VRE properties of flow availability estimations.
- Extensive results show that our methods greatly speed up the flow availability estimation. Compared with the MC and baseline IS, our SEED-VRE can: 1) Reduce variances by 360,000 and 7,000 times in the single flow case, and 200,000 and 7 times for 80% flows in the multiple flow case, if given an illustrative network and full SEED sets or partial SEED sets with good coverage; 2) Reduce simulation cost by 900 and 130 times for 80% flows in the multiple flow case, if given a realistic network and partial SEED sets with poor coverage.
- We demonstrate our methods facilitate capacity planning by giving more accurate network reliability evaluations compared with classical methods, and higher flow availability improvements compared with solely using the link capacity utilization information. Also, our methods can facilitate sales projection by providing flow availability guarantees even if flows demand more bandwidth, and this results in higher total sales.

Organizations: Section II introduces related work. Section III formally defines the FAVE problem. In Section IV, we offer 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 flow case and partial SEED set information case. In Section VI, we evaluate our methods on both an illustrative network and a realistic network, i.e., the Abilene network [15]. Section VII shows the utility of our methods in capacity planning and sales projection. Finally, Section VIII concludes.

II. RELATED WORK & PRELIMINARIES

We note that our work is a generalisation of previous work in estimating the *network reliability* [8]–[14]. Here, we briefly review previous relevant studies on the network reliability estimation (NRE) and compare them with our work.

A. Network Reliability Estimation

The most relevant literatures to our work focus on evaluating *network reliability*. To design reliable networks, it is necessary to measure the impact of network failures (e.g., link failures) on network performances [16]. It is known that the exact computation of network reliability is #P-complete, and computational complexities of all known algorithms are exponentially increasing with the graph scale [17]. This makes the problem intractable even for medium-sized networks. Hence, most work on NRE considers sampling methods to provide reliability estimations, and they can be classified into "*network connectivity*" based and "*maximum flow*" based.

Network connectivity reliability (NCR): This metric is a classical reliability measure adopted by most work [8]–[11]. The network is modeled 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 [8] take network repair policies (i.e., immediate repair and delay repair) into consideration to model link failures, and estimate NCR with the classical MC. Authors in [9] combine MC with the particle swarm optimization to handle the NCR problem. To improve the efficiency of MC, authors in [10] apply the IS method and use pre-computed "*graph minimal cuts*" to approximate the optimal IS estimator. Authors in [11] extend the NCR problem to the multi-layer networks, and address it by generalizing the traditional minimal cut to the "*cross layer minimal cut*".

Maximum flow reliability (MFR): Another line of work [12]–[14] 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. Authors in [12] assume link capacities are continuous and apply the MC splitting method for the MFR estimation. Authors in [13] 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 according to the link failure distributions. Authors in [14] assume that all graph minimal cut sets are pre-computed and they consider estimating MFR with the order minimal cut sets.

Other reliabilities: Other metrics used include the *connection availability* [18] and *service availability* [19] considering the probability that a connection or service is available. Authors in [18] evaluate the connection availability by computing the connection probability of a small subset of nodes exactly. In [19], the service availability is evaluated by using IS to estimate path availabilities. However, works in [18], [19] are essentially the same with the NCR related work: the problems studied in [18], [19] can be transformed to a problem of determining the connectivity of certain nodes, given the network topology.

Next, we analyse how our work differs from the above classical works.

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 the SLA objective of f.

The flow demand can be, for instance, *bandwidth demand*, *latency demand* or *packet loss demand*, which specifies f's QoS requirement on bandwidth, transmission latency or packet loss when considering different SLA objectives. 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: 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.

Our methods have the following advantages compared with the state-of-the-art methods:

- 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: The success of flow 1 depends both on resource allocation and routing policies and other competing flows.
- 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. We demonstrate this with detailed examples in Section VII.

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., the network topology, capacity and flow information on network performances, and provides more accurate evaluation results. The detailed comparisons can be found in Table II. In the next section, we briefly describe how to address NRE problem with sampling methods and formally define the FAVE problem.

III. Problem Definition

We first introduce the classical NRE problem and discuss how to address it with sampling methods. Then we extend the classical NRE problem and formally define the FAVE problem. We summarize important notations in Table III.

A. Sampling Methods for Network Reliability Estimations

We briefly describe how to apply sampling methods to estimate network reliabilities.

Network reliability estimation problem: Let the network be modelled as a directional multigraph $G{\triangleq}(V,E)$ with N_v nodes in the node set V and N_l links in the 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 fails (succeeds). Let $\boldsymbol{p}{=}\{p_1,\ldots,p_{N_l}\}$, $\boldsymbol{c}{=}\{c_1,\ldots,c_{N_l}\}$ and $\boldsymbol{x}{=}\{x_1,\ldots,x_{N_l}\}$ be the failure probability, capacity and status across all links, respectively.Note that *different links' statuses are statistically independent*. We consider that $\boldsymbol{p}$ and $\boldsymbol{c}$ are known, and the generation of $\boldsymbol{x}$ follows the distribution $p(\boldsymbol{x})$ induced by $\boldsymbol{p}$. There are 2^{N_l} possible realisations of $\boldsymbol{x}$, 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 as the example scenario considered in this work, *the flow demand is unsatisfied*. Given link statuses described by $\boldsymbol{x}$, $\mathcal{R}(\boldsymbol{x}){=}1$ if A is observed and $\mathcal{R}(\boldsymbol{x}){=}0$ vice versa. The network unreliability, i.e., the occurrence probability of A, can be computed via the following integral in *the discrete measure space*:

$$\mu{=}\mathbb{E}_p\left[\mathcal{R}(\boldsymbol{x})\right] = \int \mathcal{R}(\boldsymbol{x})p(\boldsymbol{x})\mathrm{d}\boldsymbol{x}, \tag{1}$$

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

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

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

where N is the number of simulation steps and $\boldsymbol{x}^{(k)}$ is the kth generated link statuses. As MC generates link statuses by true

TABLE II

A Comparison Between Classical Network Reliabilities and Flow Availability

Reliability Measurement	NCR [8]–[11]	MFR [12]–[14]	Flow Availability
Definition	$\mathbb{P}$[node u and v are connected]	$\mathbb{P}$[max flow from u to $v \geq \zeta$]	$\mathbb{P}$[flow f_i's demand is satisfied]
Required Information	Topology Information: $G(V,E)$ and $\boldsymbol{p}$.	Topology Information: $G(V,E)$, $\boldsymbol{p}$ and $\boldsymbol{c}$.	Topology Information: $G(V,E)$, $\boldsymbol{p}$ and $\boldsymbol{c}$; 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 generalisation of network connectivity based and maximum flow based reliabilities.** When link capacity $c_i=\infty$ for $\forall i$, $\mathbb{P}$[flow f_i's demand is satisfied]=$\mathbb{P}$[node u and v are connected]. When flow set F contains only one flow, $\mathbb{P}$[flow f_i's demand is satisfied]=$\mathbb{P}$[max flow from u 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 $\boldsymbol{p}$ and $\boldsymbol{c}$ 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.

TABLE III

Important Notations

Notations	Descriptions
N_l, N_f	The number of transportation links and flows.
e_i	The ith link, where $i \in \{1, \cdots, N_l\}$.
p_i, $p_i(x_i)$, $\boldsymbol{p}$, $p(\boldsymbol{x})$	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. $\boldsymbol{p}=(p_1,...,p_{N_l})$ and $p(\boldsymbol{x})$ is the distribution induced by $\boldsymbol{p}$.
$\boldsymbol{x}$	Link statuses. $\boldsymbol{x}=(x_1,...,x_{N_l})$ and $x_i=1$ ($x_i=0$) means that link e_i fails (succeeds).
f_i	The ith flow, where $i \in \{1, \cdots, N_f\}$.
(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 $\boldsymbol{x}$, $\mathcal{R}(\boldsymbol{x})=1$ ($\mathcal{R}(\boldsymbol{x})=0$) if the flow fails (succeeds).
$L, \mathcal{L}$	$L \subseteq \{1,...,N_l\}$ is a link set. $\mathcal{L}$ is a collection of link sets.
$\Psi(L)$, $\Psi^{-1}(\boldsymbol{x})$	$\Psi(L)$ maps a link set L to link statuses $\boldsymbol{x}$ which satisfies $\forall i \in L$, $x_i=1$; $\forall i \notin L$, $x_i=0$. And $\Psi^{-1}(\boldsymbol{x})$ maps link statuses $\boldsymbol{x}$ to the link set $L=\{i \mid x_i=1, i \in \{1, \cdots, N_l\}\}$.
span(L), span($\mathcal{L}$)	span(L) is the collection of link set L's supersets and span($\mathcal{L}$) is the collection of supersets of all link sets $L \in \mathcal{L}$.
$\mathcal{F}$	$\mathcal{F}=\{L \mid \mathcal{R}(\Psi(L))=1, L \subseteq \{1, \cdots, N_l\}\}$, 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) $\forall L \subsetneq S$, $L \notin \mathcal{F}$; 3) $\forall L \supseteq S$, $L \in \mathcal{F}$. And $\mathcal{S}$ is the collection of all SEEDs.

link failure probabilities $p(\boldsymbol{x})$ (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, which is expensive for an extensive network where even simulating $\mathcal{R}$ on a single sample $\boldsymbol{x}$ can take hours.

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

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

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

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

Estimator efficiency evaluation: The efficiency of an estimator is often measured by its "*variance*". Take the MC estimator as an example, its variance is given by:

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

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

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

Note that the MC estimator is a special case of the IS estimator, given $p(\boldsymbol{x})=q(\boldsymbol{x})$, $\forall \boldsymbol{x}$. Define $\sigma_q^2=\mathbb{V}_q[\mathcal{R}(\boldsymbol{x})\omega(\boldsymbol{x})]$ 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.*

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

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 $q(\boldsymbol{x})=q^*(\boldsymbol{x})$ where:

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

Proof: One can verify this by plugging $q(\boldsymbol{x})=q^*(\boldsymbol{x})$ into the integral form of σ_q^2. The detailed proof is in Appendix A.

Remark: Although the ZV property implies a minimum simulation cost, a procedure to construct such a $q^*(\boldsymbol{x})$ with the ZV property often has a high computational requirement. Hence, the key of designing an efficient IS estimator lies in closely approximating $q^*(\boldsymbol{x})$ with a manageable computational cost. As for different applications and problem definitions, the auxiliary information we can utilize and the way to approximate $q^*(\boldsymbol{x})$ are different, designing efficient IS estimators are highly challenging and problem-dependent.

B. The Flow Availability Estimation (FAVE) Problem

Based on the network reliability estimation problem defined in Section III-A, we now give a formal definition of the FAVE problem. Consider the network $G(V,E)$ with topology information $\boldsymbol{p}$, $\boldsymbol{c}$ and $\boldsymbol{x}$. We consider a flow set F with N_f flows where 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 interested flow fails ($\mathcal{R}=1$) or succeeds ($\mathcal{R}=0$), i.e.,:

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

where $\mathbb{G}$ represents the topology information, including a tuple (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) 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. We assume all information in $(\mathbb{G},\mathbb{F},\mathbb{D},\mathbb{B})$ is known, except the link statuses described by $\boldsymbol{x}$. To simplify the expression, let:

$$\mathcal{R}(\cdot) : \boldsymbol{x} \to \{0, 1\}. \tag{9}$$

Namely, given all the other information, $\mathcal{R}$ only depends on $\boldsymbol{x}$, which follows the link failure distribution $p(\boldsymbol{x})$ induced by $\boldsymbol{p}$.

Similar to Section III-A, for a specific flow, the unavailability μ can be computed via Eq. (1), and the availability is $1-\mu$. The goal is to evaluate availabilities for (all) flows in F. Yet, the complexity of function $\mathcal{R}$ and the high dimensionality of the topological space make it already expensive to compute $\mathcal{R}(\boldsymbol{x})$ for a single $\boldsymbol{x}$, let alone to evaluate μ analytically. One alternative is to estimate μ via simulations. And we want to reduce the simulation steps, i.e., the times to compute $\mathcal{R}(\boldsymbol{x})$. We name the network flow availability estimation problem as "FAVE". And we show it generalizes the classical network reliability estimation problem [8]–[14] by the following:

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

Proof: The network connectivity based case corresponds to FAVE with $c_i{=}\infty$ for $\forall i$. And the maximum flow based case corresponds to FAVE with $N_f{=}1$. Proofs are in Appendix B. ∎

In Section II-C, we explain how the one-run variance σ^2 reflects the simulation cost and estimation accuracy. In addition to measuring the estimation efficiency with σ^2, we also consider two attractive error bound properties: the *bounded relative error* and *vanishing relative error* [10].

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 $\lim_{\mu\to 0}\varepsilon_{CV}{<}\infty$.

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., $\lim_{\mu\to 0}\varepsilon_{CV}{=}0$.

Remark: Take the MC estimator in Eq. (2) as an example, the variance of MC is $\sigma^2/N{=}(\mu{-}\mu^2)/N$, i.e., $\sigma{=}O(\sqrt{\mu})$. This implies that MC satisfies neither BRE nor VRE property. Also, note that the VRE property is stronger than the BRE property.

In the following, we will discuss how to design estimation methods which have above properties.

IV. Algorithm Design

We first describe our design to address the FAVE problem under "the single flow case". We start with *a baseline IS design* to gain insights for *efficient sampling* with good error bounds. Then we propose our SEED methods.

A. A Baseline Importance Sampling Design

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

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

$$q(\boldsymbol{x}){=}\textstyle\prod_{i=1}^{N_l} q_i(x_i), \tag{10}$$

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

$$q_i(x_i){=}\mathbb{P}\left[x_i|\mathcal{R}(\boldsymbol{x}){=}1\right]. \tag{11}$$

TABLE IV

An Example for the IS Method

$\boldsymbol{x}$	$(1,1,0)$	$(1,1,1)$	$(0,1,1)$
$q^*(\boldsymbol{x}){=}\mathbb{P}\left[\boldsymbol{x}\vert\mathcal{R}(\boldsymbol{x}){=}1\right]$	0.49975	0.00050	0.49975
$q(\boldsymbol{x}){=}\prod_{i=1}^{N_l}\mathbb{P}\left[x_i\vert\mathcal{R}(\boldsymbol{x}){=}1\right]$	0.25000	0.25025	0.25000

Proof: Please see the Appendix C for the detailed proof. ∎

By now, the estimation of $\mathbb{P}\left[x_i|\mathcal{R}(\boldsymbol{x}){=}1\right]$ becomes a new problem. In fact, even given the exact $\mathbb{P}[x_i|\mathcal{R}(\boldsymbol{x}){=}1]$ 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 $q(\boldsymbol{x})$ in Eq. (11) approximates $q^*(\boldsymbol{x})$. To see this, consider the case where a network has 2 nodes connected by 3 parallel links with $\boldsymbol{p}{=}(0.001, 0.2, 0.001)$. The flow fails if link statuses $\boldsymbol{x}$ are (1,1,0), (1,1,1) or (0,1,1). By Theorem 3, one possible importance distribution is $\boldsymbol{q}{=}(0.50025, 1, 0.50025)$. In this case, we show that $q^*(\boldsymbol{x})$ is *not well approximated* by $q(\boldsymbol{x})$ in Table IV.

Remark: The above example illustrates that minimizing KL divergence cannot provide the IS method a performance guarantee: sometimes our chosen $q(\boldsymbol{x})$ can be very different from $q^*(\boldsymbol{x})$. A major reason is that the baseline IS assumes $q(\boldsymbol{x})$ has the product form in Eq. (10), i.e., it considers link failures are independent and ignores correlations among them. We will discuss how to improve the design of $q(\boldsymbol{x})$ 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 $q(\boldsymbol{x})$ in Eq. (11) as its importance distribution, $\mathbb{V}_q[\hat{\mu}_{\mathrm{IS}}]$ is bounded by:

$$\frac{\mu^2\mathrm{KL}(q^*\|q)}{N} {\le} \mathbb{V}_q[\hat{\mu}_{\mathrm{IS}}] {\le} \frac{\mu^2\sqrt{2\log 2\mathrm{KL}(q^*\|q)}}{N\min_{\boldsymbol{x}} q(\boldsymbol{x})}. \tag{12}$$

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

Proof: Please see the Appendix D for the detailed proof. ∎

Remark: Although theorem 4 implies that the variance of baseline IS is upper bounded, this upper bound can be worse than $\mathbb{V}_p[\hat{\mu}_{MC}]$ when $\min_{\boldsymbol{x}} q(\boldsymbol{x})$ is very small: By Eq. (10) and (11), the $\min_{\boldsymbol{x}} q(\boldsymbol{x})$ also depends on μ, and as there are 2^{N_l} realizations of $\boldsymbol{x}$, $\min_{\boldsymbol{x}} q(\boldsymbol{x}){\le}2^{-N_l}$. This motivates us to seek for more efficient sampling methods with better error bounds, which we will discuss in details in the next subsection.

B. Conditions for Efficient Sampling

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

1) ZV sequential importance sampling: To capture the correlation of links' statuses, one possibility is to generate links' statuses x_i "in a sequential manner" (rather than generate x_i independently as the baseline IS). Then each link's status also depends on statuses of previous generated links. We first adapt the ZV importance distribution in Eq. (7) for this sequential design. Let $\boldsymbol{x}_{1:i}{=}(x_1, x_2, \cdots, x_i)$ be statuses of the first ith links.

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

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

Proof: Please see the Appendix E for the detailed proof. ■

Remark: Different from baseline IS, SIS generates links' statuses in a sequential manner, which enables its importance distribution $q_i(\boldsymbol{x})$ to capture the correlation of links' statuses, i.e., *links' importance distributions are dependent*. Note that the generation order will not affect the ZV property of $q^*(\boldsymbol{x})$, but do matter if one wants to approximate $q^*(\boldsymbol{x})$ more accurately. We leave the design of link generation order as a future work and generate links' statuses according to link indexes in this work.

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

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

$$\hat{\mathbb{P}}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i}\right]=O\left(\mathbb{P}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i}\right]\right), \tag{14}$$

and the VRE property if $\forall i=1,\ldots,N_l$, $\hat{\mathbb{P}}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i}\right]$ satisfies:

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

Proof: Please see the Appendix F for the detailed proof. ■

Remark: Theorem 6 gives key guidelines to design sampling methods with BRE and VRE properties, i.e., the estimation $\hat{\mathbb{P}}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i}\right]$ should satisfy conditions in Theorem 6.

In the next subsection, we discuss algorithms which approximate $\mathbb{P}[\mathcal{R}=1|\boldsymbol{x}_{1:i}]$ with BRE and VRE properties.

C. SEED Algorithms

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

Let $\Omega \triangleq \{1,\ldots,N_l\}$. Function $\Psi(\cdot)$ and $\Psi^{-1}(\cdot)$ represent the transformations between link statuses $\boldsymbol{x}$ and failed links $L\subseteq\Omega$:

$$\Psi(\cdot): L\rightarrow\boldsymbol{x}, \text{ where } x_i=\mathbb{1}_{i\in L}, \tag{16}$$
$$\Psi^{-1}(\cdot): \boldsymbol{x}\rightarrow L, \text{ where } L=\{i|x_i=1\}. \tag{17}$$

Given one interested flow, the collection of all link sets, each of which the failure can result in flow failure, is denoted by:

$$\mathcal{F}\triangleq\left\{L|\ \mathcal{R}(\Psi(L))=1, L\subseteq\Omega\right\}. \tag{18}$$

Also, let the collection of all supersets of a single link set L be:

$$\text{span}\left(L\right)\triangleq\left\{L'|L\subseteq L'\subseteq\Omega\right\}. \tag{19}$$

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

$$\text{span}\left(\mathcal{L}\right)\triangleq\bigcup_{L\in\mathcal{L}}\text{span}\left(L\right). \tag{20}$$

The probability that a link set L fails (i.e., all links in L fail) is:

$$\Phi(L)\triangleq\prod_{i\in L}p_i=\sum_{L'\in\text{span}(L)}\mathbb{P}\left[\Psi\left(L'\right)\right]. \tag{21}$$

Definition 6 (SEED): A link set S is called "SEED" if it satisfies the following conditions:

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

The collection of all SEEDs is donoted by $\mathcal{S}$.

Algorithm 1 SEED Based ZV Sampling (SEED-ZV)

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

Output: An importance distribution $q(\boldsymbol{x})$ to achieve the ZV property and a sample of link statuses $\boldsymbol{x}$.

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

Definition 7 (Conditional SEED): Let $\boldsymbol{x}_{1:i}$ be the statuses of the first i links. The *conditional SEED* (*cond-SEED* for short) $S(\boldsymbol{x}_{1:i})$ is a link set which satisfies:

$$(1) S(\boldsymbol{x}_{1:i})\subseteq\{i+1,\cdots,N_l\}; (3)\forall L\subsetneqq S\left(\boldsymbol{x}_{1:i}\right), L\cup\Psi^{-1}(\boldsymbol{x}_{1:i})\notin\mathcal{F};$$
$$(2) S(\boldsymbol{x}_{1:i})\cup\Psi^{-1}(\boldsymbol{x}_{1:i})\in\mathcal{F}; (4)\forall L\supseteq S\left(\boldsymbol{x}_{1:i}\right), L\cup\Psi^{-1}(\boldsymbol{x}_{1:i})\in\mathcal{F}.$$

The collection of all cond-SEEDs is denoted by $\mathcal{S}\left(\boldsymbol{x}_{1:i}\right)$.

Examples: We give some examples for the above definitions. Consider the example in Section IV-A, 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 $\boldsymbol{x}_{1:4}=(0,1,0,1)$, there is only one cond-SEED $S(\boldsymbol{x}_{1:4})=\{5\}$.

Next, we use the SEED set to capture the correlation of link failures and approximate $\mathbb{P}[\mathcal{R}=1|\boldsymbol{x}_{1:i}]$ in Theorem 5. We start with considering the full SEED set (i.e., the collection of all SEEDs) as input, and we leave the partial SEED set case and SEEDs collection in the next section.

2) SEED algorithms: Theorem 5 presents the optimal SIS importance distribution $q^*(\boldsymbol{x})$. We propose the SEED-ZV algorithm in Algorithm 1 which computes the exact value of $q^*(\boldsymbol{x})$. When generating x_i, SEED-ZV first assumes x_i=1 and computes the exact value of $\mathbb{P}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i}\right]$ by traversing all combinations of cond-SEEDs $S(\boldsymbol{x}_{1:i})\in\mathcal{S}$ (i.e., lines 3-4). Then it keeps x_i=1 with the probability $q_i(\boldsymbol{x})$, and changes to x_i=0 with the probability $1-q_i(\boldsymbol{x})$ (i.e., line 5). To prove SEED-ZV achieves the ZV property in Theorem 7, we have the following lemma.

Lemma 1: The estimation $\hat{\mathbb{P}}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i}\right]$ *in SEED-ZV satisfies* $\hat{\mathbb{P}}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i}\right]=\mathbb{P}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i}\right]$. *Namely,*

$$\mathbb{P}[\mathcal{R}=1|\boldsymbol{x}_{1:i}]= \sum_{A\neq\emptyset, A\subseteq\mathcal{S}(\boldsymbol{x}_{1:i})}\left((-1)^{|A|-1}\Phi\left[\bigcup_{S(\boldsymbol{x}_{1:i})\in A}S(\boldsymbol{x}_{1:i})\right]\right). \tag{22}$$

Proof: Please see the Appendix G for the detailed proof. ■

Theorem 7: If the link statuses $\boldsymbol{x}$ *are generated by SEED-ZV, the estimator in Eq. (4) has the ZV property.*

Proof: According to Theorem 5 and Lemma 1, the importance distribution q generated by SEED-ZV satisfies $q=q^*$. Namely, SEED-ZV achieves the ZV property. ■

Algorithm 2 SEED Based BRE Sampling (SEED-BRE)

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

Output: An importance distribution $q(\boldsymbol{x})$ to achieve the BRE property and a sample of link statuses $\boldsymbol{x}$.

1: **for** i=1 to N_l **do**
2:   **for** x_i in $\{0,1\}$ **do**
3:     $\hat{\mathbb{P}}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right] \leftarrow \max_{S(\boldsymbol{x}_{1:i})\in\mathcal{S}(\boldsymbol{x}_{1:i})}\Phi(S(\boldsymbol{x}_{1:i}))$;
4:     $q_i(x_i)\leftarrow\frac{\hat{\mathbb{P}}(\mathcal{R}=1|\boldsymbol{x}_{1:i})}{\hat{\mathbb{P}}(\mathcal{R}=1|\boldsymbol{x}_{1:i-1})}p_i(x_i)$;
5:   Normalize $q_i(\boldsymbol{x})$ 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:   $\mathcal{S}(\boldsymbol{x}_{1:i}) \leftarrow$ UPDATECONDSEED$(i, x_i, \mathcal{S}(\boldsymbol{x}_{1:i-1}))$

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

Note that there is a tradeoff between the estimation accuracy and computational complexity in estimating the optimal sampling distribution $q^*(\boldsymbol{x})$. We propose the SEED-BRE algorithm in Algorithm 2, in which we sacrifices some estimation accuracy of $q^*(\boldsymbol{x})$ to achieve a lower linear computational complexity. To generate x_i, SEED-BRE utilizes only *probabilities of the most important cond-SEEDs* for estimating $\hat{\mathbb{P}}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right]$ and $q_i(\boldsymbol{x})$ (i.e., lines 2-4), rather than all combinations of cond-SEEDS (i.e., lines 3-4 in SEED-ZV). To show that SEED-BRE has the BRE property in Theorem 8, we have the following lemma:

Lemma 2: The estimation $\hat{\mathbb{P}}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right]$ in SEED-BRE satisfies $\hat{\mathbb{P}}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right]=O\left(\mathbb{P}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right]\right)$. Namely,

$$\max_{S(\boldsymbol{x}_{1:i})\in\mathcal{S}(\boldsymbol{x}_{1:i})}\Phi\left(S\left(\boldsymbol{x}_{1:i}\right)\right)=O\left(\mathbb{P}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right]\right) \quad (23)$$

Proof: Please see the Appendix I for the detailed proof. ∎

Theorem 8: If link statuses $\boldsymbol{x}$ are generated using SEED-BRE, the estimator in Eq. (4) has the BRE property.

By Lemma 2, the sampling distribution in SEED-BRE satisfies $\hat{\mathbb{P}}[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}]{=}O\big(\mathbb{P}[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}]\big)$. Then, according to Theorem 6, we show that SEED-BRE has the BRE property. Please see the Appendix J for the detailed proof.

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

To further improve estimation accuracy, we then propose the SEED-VRE algorithm in Algorithm 3. The SEED-VRE is similar to SEED-BRE, except that it utilizes *the probability sum of cond-SEEDs* for estimations (i.e., line 3). We next show that SEED-VRE has the VRE property if link failures are rare. We first have the following lemma.

Lemma 3: Assume link failures are rare events and the corresponding failure probabilities have the form $p_i{=}O(\epsilon)$, where $\epsilon{\to}0$, then the estimation $\hat{\mathbb{P}}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right]$ in SEED-VRE satisfies $\hat{\mathbb{P}}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right]=\mathbb{P}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right](1{+}o(1))$. Namely,

$$\sum\nolimits_{S(\boldsymbol{x}_{1:i})\in\mathcal{S}(\boldsymbol{x}_{1:i})}\Phi(S(\boldsymbol{x}_{1:i}))= \mathbb{P}[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}]{+}o(\mathbb{P}[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}]). \quad (24)$$

Proof: Please see the Appendix K for the detailed proof. ∎

Next we show that SEED-VRE achieves the VRE property by the following theorem:

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

Algorithm 3 SEED Based VRE Sampling (SEED-VRE)

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

Output: An importance distribution $q(\boldsymbol{x})$ to achieve the VRE property, and a sample of link statuses $\boldsymbol{x}$.

1: **for** i=1 to N_l **do**
2:   **for** x_i in $\{0,1\}$ **do**
3:     $\hat{\mathbb{P}}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right] \leftarrow \sum_{S(\boldsymbol{x}_{1:i})\in\mathcal{S}(\boldsymbol{x}_{1:i})}\Phi(S(\boldsymbol{x}_{1:i}))$;
4:     $q_i(x_i)\leftarrow\frac{\hat{\mathbb{P}}(\mathcal{R}=1|\boldsymbol{x}_{1:i})}{\hat{\mathbb{P}}(\mathcal{R}=1|\boldsymbol{x}_{1:i-1})}p_i(x_i)$;
5:   Normalize $q_i(\boldsymbol{x})$ 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:   $\mathcal{S}(\boldsymbol{x}_{1:i}) \leftarrow$UPDATECONDSEED$(i, x_i, \mathcal{S}(\boldsymbol{x}_{1:i-1}))$

By Lemma 3, the sampling distribution in SEED-VRE satisfies $\hat{\mathbb{P}}[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}]{=}\mathbb{P}[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}](1{+}o(1))$. ∎

Remark: The computational complexity of SEED-VRE is also $O(N_l|\mathcal{S}|)$, as it needs to traverse all $\mathcal{S}(\boldsymbol{x}_{1:i})$, 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.

V. Practical Consideration

Previous discussions illustrate the effectiveness of our SEED methods in estimating a single flow's availability with the full information of SEED set $\mathcal{S}$. In this section, we take more practical issues into consideration. First, consider the case that one needs to provide flow availability estimation for all flows in the network, of which the amount is at the order of $O(N_v^2)$. It can be 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 of 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. And we propose an algorithm to collect the SEED set information.

A. Generalization to Multiple Flows Case

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

$$q(\boldsymbol{x}){=}\mathcal{M}(q^{(1)}(\boldsymbol{x}),\ldots,q^{(N_f)}(\boldsymbol{x})). \quad (25)$$

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

$$q(\boldsymbol{x}){=}\sum\nolimits_k w_k q^{(k)}(\boldsymbol{x}),\ \sum\nolimits_k w_k{=}1. \quad (26)$$

Here w_k can be viewed as the probability of taking $q^{(k)}(\boldsymbol{x})$ to generate $\boldsymbol{x}$. Denote $\mu^{(k)}$ as f_k's failure probability and $(\sigma_q^{(k)})^2$ as the one-run variance when taking $q(\boldsymbol{x})$ 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. (26) with pure distributions $q^{(k)}(\boldsymbol{x})$ generated by SEED methods, the IS estimator achieves the BRE property for all flows availabi-lity estimations. Specifically, for flow f_k, the estimator's one-run variance satisfies:

$$\left(\sigma_q^{(k)}\right)^2 \leq \left(\frac{1}{w_k}-1\right)(\mu^{(k)})^2+\frac{1}{w_k}\left(\sigma_{q^{(k)}}^{(k)}\right)^2. \quad (27)$$

Proof: This can be proved by plugging in the mixture distribution $q(\boldsymbol{x})$ in Eq. (26) into the definition of one-run variance $\left(\sigma_q^{(k)}\right)^2$. Please see the Appendix M for the detailed proof. ■

Remark: Theorem 10 states that, if extending to the multiple flow 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

It can be difficult at times 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 $\boldsymbol{x}_{1:i}$ as $\mathcal{S}'(\boldsymbol{x}_{1:i})$. To analyze error bound properties of SEED algorithms, we provide the following lemma.

Lemma 4: Given a partial SEED set $\mathcal{S}'(\boldsymbol{x}_{1:i})$, when estimating $\mathbb{P}\left[\Psi^{-1}(\boldsymbol{x})\in\text{span}(\mathcal{S}'(\boldsymbol{x}_{1:i}))|\boldsymbol{x}_{1:i}\right]$: both 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 has the VRE property.

Proof: The proofs follow the same lines as the proofs in the Theorem 7, Theorem 8 and Theorem 9. ■

Remark: Theorem 6 states that the estimation accuracy depends on how well $\mathbb{P}[\mathcal{R}=1|\boldsymbol{x}_{1:i}]$ is approximated. Given $\mathcal{S}$, SEED methods have good error bound properties for they can well approximate $\mathbb{P}[\Psi^{-1}(\boldsymbol{x})\in\text{span}(\mathcal{S}(\boldsymbol{x}_{1:i}))|\boldsymbol{x}_{1:i}]=\mathbb{P}[\mathcal{R}=1|\boldsymbol{x}_{1:i}]$. Yet, given a partial SEED set $\mathcal{S}'$, the bias between $\mathbb{P}[\mathcal{R}=1|\boldsymbol{x}_{1:i}]$ and $\mathbb{P}\left[\Psi^{-1}(\boldsymbol{x})\in\text{span}(\mathcal{S}'(\boldsymbol{x}_{1:i}))\,|\boldsymbol{x}_{1:i}\right]$ should be considered. Let us consider the following two cases:

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

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

  Here is an example of such a partial SEED set:

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

- **Poor coverage case**. Without prior knowledge of network and flow failures, $\mathcal{S}'$ may have a poor coverage and even $\mathcal{S}'=\emptyset$. If so, the SEED set information can be collected via pre-samplings and updated while simulating flow failures. We provide one possibility to collect SEED set information in the later section.

C. SEED Set Collection

The information of SEED sets can be collected via some pre-samplings or updated adaptive while simulating flow failures. In either case, SEED sets need to be updated according to the observed simulation results. We provide the Algorithm 4 to conduct this update, and we claim that the Algorithm 4 works via the following theorem:

Algorithm 4 SEED Sets Updating (SEED-Updating)

Input: $\{\mathcal{L}^{(k)}\}_{k=1}^{N_f}$, where $\mathcal{L}^{(k)}\subset\mathcal{F}^{(k)}$ and $\mathcal{L}$ is a "sperner family"; N, the number of simulations.
Output: $\{\mathcal{S}'^{(k)}\}_{k=1}^{N_f}$.
for i=1 to N **do**
  Generate topology sample $\boldsymbol{x}$ according to $q(\boldsymbol{x})$;
  for k=1 to N_f **do**
    if $\mathcal{R}^{(k)}(\boldsymbol{x})=1$ and $\Psi^{-1}(\boldsymbol{x})\notin\text{span}\left(\mathcal{L}^{(k)}\right)$ **then**
      $\mathcal{L}^{(k)}\leftarrow\mathcal{L}^{(k)}\cup\{\Psi^{-1}(\boldsymbol{x})\}\backslash\{L|L\in\mathcal{L}^{(k)}, L\supset\Psi^{-1}(\boldsymbol{x})\}$;
$\mathcal{S}'^{(k)}\leftarrow\mathcal{L}^{(k)}$ where $1\leq k\leq N_f$.

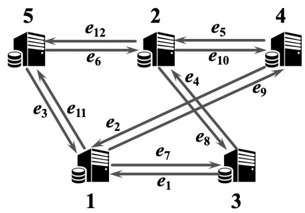

Fig. 2. Topology of a small scale network.

Theorem 11: The SEED-Updating algorithm guarantees that $\text{span}\left(\mathcal{L}^{(k)}\right)\uparrow\text{span}\left(\mathcal{S}^{(k)}\right)$*, i.e.,* $\text{span}\left(\mathcal{L}^{(k)}\right)$ *monotone converges to* $\text{span}\left(\mathcal{S}^{(k)}\right)$*, for* $\forall f_k\in F$*.*

Proof: Please see the Appendix N for the detailed proof. ■

Remark: Theorem 11 implies that SEED-Updating improves the coverage quality of partial SEED set, and it also guarantees the updated partial SEED set converges to the full SEED set.

Note that the sampling distribution $q(\boldsymbol{x})$ is not specified in Algorithm 4. Namely, SEED-Updating can cooperate with any sampling method to collect SEED set information. Besides, the convergence rate depends on the way we generate $\boldsymbol{x}$. We leave the efficient SEED set collection as the future work.

VI. Evaluation of SEED Methods

To evaluate the effectiveness of our methods, we consider both an illustrative small scale network, and a realistic network with the topology and traffic matrices extracted from the Abilene backbone network [15]. 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 $N\geq\frac{4\alpha^2}{\delta^2}\left(\frac{\sigma_q}{\mu}\right)^2$). Hence, we use the *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., $\frac{\sigma^2_{\text{MC}}}{\sigma^2_{\text{SEED}}}$, can imply the simulation cost reduction, i.e., $\frac{N_{\text{MC}}}{N_{\text{SEED}}}$.

A. Experiments on an Illustrative Network

The illustrative network[1] 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 coverages 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. For each link e_i: the link failure probability p_i is uniformly distributed over $[0.5\epsilon, 1.5\epsilon]$ (ϵ is a small positive number); the link capacity c_i is uniformly distributed over $[50, 80, 100, 200]$. The

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

TABLE V

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
ϵ=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}
ϵ=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}
ϵ=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 $\mathcal{S}'$=$\{\{2,5\},\{7,8\}\}$.

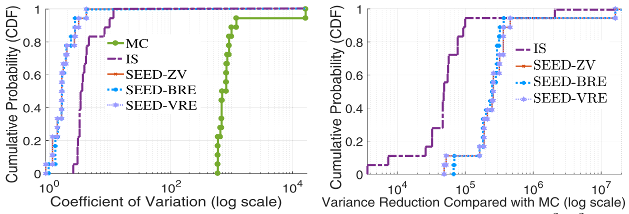

(a) Cdf of coefficient of variation (σ_q/μ) (b) Cdf of variance reduction (σ^2/σ_q^2)

Fig. 3. Performance on the illustrative network (multiple flow case).

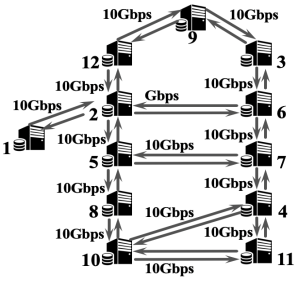

Fig. 4. Topology of the Abilene network.

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 the 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 V. 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 V. Let ϵ=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 ϵ=0.001. We consider an equally weighted sum of flows' pure importance distributions as the mixture SIS distribution. Fig. 3 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 Theorem 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 shows cumulative distributions of the variance reduction compared with MC. With SEED methods, more than 80% flows have *variance reductions* $\sigma^2_{MC}/\sigma^2_{SEED}$>200,000. To better illustrate the efficiency improvement, Table VI summarizes simulation costs to guarantee that for 80% flows, "with 95% confident the relative error is less than 0.01", i.e., $\alpha_{95}\sigma_q/\sqrt{N}{\leq}0.01\hat{\mu}$.

B. Experiments on a Realistic Network

Next, consider a realistic network to show 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 [15], [20] with topology and traffic matrices collected by [21]. The network contains 12 nodes and 30 links. Link capacities are illustrated in Fig. 4. The flow set contains 132 aggregated flows: all flows with the same source and destination are aggregated as a single flow[2]. 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

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

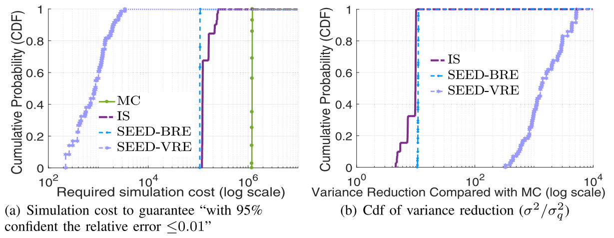

Fig. 5. Performance on the Abilene network (multiple flow case).

TABLE VI
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}

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 methods[3]. 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 *13 times* via SEED-BRE, compared with MC. Fig. 5 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 AND SALES PROJECTION

Now, we demonstrate the utility of our methods in the capacity planning, sales projection and topology planning. By testing flow availabilities for the planning proposals of capacity, sales or topology, we can not only *determine which proposal allows the network to provide better flow availability guarantees* even if flows demand for more bandwidth resources, but also *utilize flow availability feedbacks for further refinements of infeasible proposals*. We use Abilene as an example to show how does our method work in details, and compare with the NCR and MFR based methods.

A. Applications in the Sales Projection

Consider the case that a network provider wants to increase sales via, e.g, oversubscribing flow demands, providing higher flow availability guarantees or admitting new flows. For a new sales projection proposal $\mathcal{SP}$, we will:

- **Test flow availabilities**. Collect flow availability feedback for each flow f_i, including a flow availability estimation $\hat{\mu}_i$ with upper confidence bound $\overline{\mu_i}$ and lower bound confidence $\underline{\mu_i}$ derived from the empirical variance.
- **Check feasibility**. We run enough simulations to make sure $o_i \notin [\underline{\mu_i}, \overline{\mu_i}]$. f_i's availability target o_i is *achieved* if

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

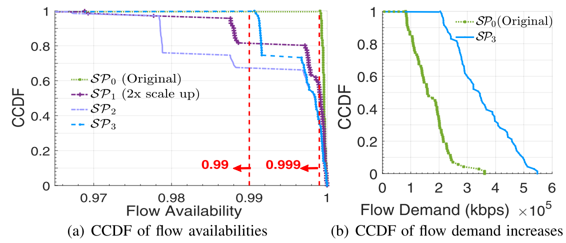

Fig. 6. An example of applying FAVE to the sales planning.

$\underline{\mu_i} \geq o_i$ and *unreached* if $\overline{\mu_i} < o_i$. $\mathcal{SP}$ is *feasible* if all flow availability targets are achieved; and infeasible otherwise.
- **For a feasible proposal** $\mathcal{SP}$, we either apply it directly or refine it using availability feedbacks. Specifically, we can improve f_i's experiences by supporting its current demand d_i with a higher availability target o_i' or supporting a higher demand d_i' with its current availability target o_i when $\underline{\mu_i} > o_i$. We can also admit a new flow f_j.
- **For a infeasible proposal** $\mathcal{SP}$, we can take a reasonable compromise between $\mathcal{SP}$ and a feasible $\mathcal{SP}'$, e.g., for flow f_i with d_i, o_i in $\mathcal{SP}$ and d_i', o_j' in $\mathcal{SP}'$, we take $(d_i + d_i')/2$ and $(o_i + o_i')/2$ as the new demand and flow availability target, respectively.

To see the achievable gains by doing, we take the Abilene network as an example. Let flow availability targets $o_i = 99\%$ and take raw demands in Section VI-B as the original proposal $\mathcal{SP}_0$. First, as $\mathcal{SP}_0$ is feasible and flow availabilities far surpass availability targets, we double the demands in $\mathcal{SP}_0$ as our new proposal $\mathcal{SP}_1$. Then, as $\mathcal{SP}_1$ is infeasible, we take a compromise between $\mathcal{SP}_0$ and $\mathcal{SP}_1$ and obtain the proposal $\mathcal{SP}_2$. Next, as $\mathcal{SP}_2$ is still infeasible, we continue the refinements until we obtain a feasible proposal $\mathcal{SP}_3$. Fig. 6(a) summarizes the complementary cumulative probabilities (CCDF) of flow availabilities achieved by applying different proposals. One can observe that, by refining the sales projection proposals, we obtain a feasible oversubscribe factor for each flow demand such that the network can still support these oversubscribed flow demands with given flow availability targets. Fig. 6(b) shows the significant flow demand improvements achieved by applying our method, which implies a *90.5% sales increase*.

B. Applications in the Capacity Planning

Consider the case where a network provider builds new links with certain capacities to achieve all flow availability targets. For a capacity planning proposal $\mathcal{CP}$, similar to the sales projection case, we can *evaluate its feasibility and adopt it if feasible*. We are more concern about how the availability feedbacks can be applied to refine infeasible proposals. Specifically, we refine an infeasible proposal with the following information:

- **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.
- **Maximum flow based**: The maximum flow value is a widely adopted network reliability measure [12]–[14]. We take the increase 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 tar-

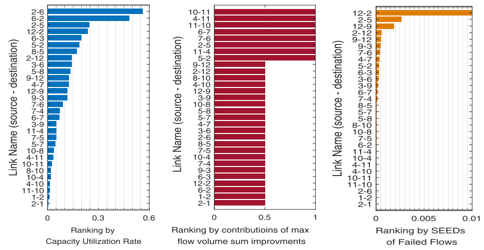

(a) (b) (c)

Fig. 7. Link importance ranking.

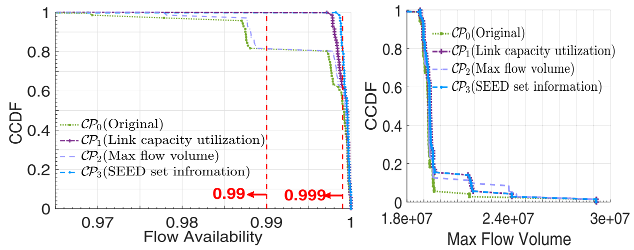

(a) CCDF of flow availability (b) CCDF of max flow volume

Fig. 8. An example of applying FAVE to the sales planning.

gets. The number of failed flows in $\mathcal{F}'$ when e_j fails ($\sum_{f_i \in \mathcal{F}'} \mathbb{P}\left[\mathcal{R}_i{=}1|x_j{=}1\right]$) can be estimated using SEED algorithms and can imply the importance of links.[4]

We take the current capacity design of Abilene as the original proposal $\mathcal{CP}_0$, and we double the raw demands as flow demands so that $\mathcal{CP}_0$ is infeasible. We rank links using the above metrics and summarize rankings in Fig. 7. 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. 7, we have three proposals, i.e., $\mathcal{CP}_1$, $\mathcal{CP}_2$ and $\mathcal{CP}_3$ by taking the top four links in Fig. 7(a), 7(b) and 7(c). Fig.8(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%. Our method also provides *more accurate evaluations to determine better capacity proposals*. According to Fig. 8(a), we can easily determine a better proposal by comparing $\mathcal{CP}_1$, $\mathcal{CP}_2$ and $\mathcal{CP}_3$ over the flow availability. However, as shown in Fig. 8(b), when comparing these three proposals over maximum flow volumes, it is hard to tell which one is better.

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 high capacity utilization if many flows' shortest paths go through it. Yet, if it is easy to find an 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.

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

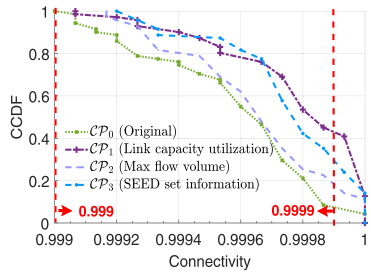

Fig. 9. An example of topology design.

C. Applications in the Topology Design

Consider the case where the network provider wants to strengthen the network by building some new links among some already directly connected nodes. We can use the method in capacity planning to determine important links. If a network provider plans to add some new nodes and build links to connect them, our method can help to check feasibilities and compare the flow availability guarantees of different topology designs.

We consider the example in the capacity planning case. Traditional methods prefer topology designs with better network connectivities. We compare the evaluation results on both the flow availabilities in Fig. 8(a) and the achieved network connectivities in Fig. 9(a), for the above four proposals. One can observe that: our method can easily determine that the proposal $\mathcal{CP}_3$ results in the best flow demands satisfaction; on the contrary, when comparing these proposals over the network connectivity, $\mathcal{CP}_2$ has the highest network connectivity and even the infeasible proposals $\mathcal{CP}_0$ and $\mathcal{CP}_1$ achieve high network connectivities of 99.9%. This implies that the network connectivity is not a proper metric for evaluating the satisfaction of flow demands.

VIII. Conclusion

In this paper, we generalize the classical network reliability problem and consider the flow availability estimation (FAVE) problem, where a flow availability is defined as the satisfaction probability of flow demands. 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 the 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 around 900 and 130 times compared with the MC and baseline IS methods on the Abilene network. We demonstrate that our methods can facilitate the capacity planning and sales projection: For capacity planning, our methods provide more accurate network reliability estimations compared with classical methods, and greater flow availability improvements compared with solely using link capacity utilizations; for sales projection, our methods provide flow availability guarantees even if flows demand more bandwidth, and so increases the total sales.

Appendix

We give the detailed proofs for some theorems and lemmas appeared in this paper in the following:

A. The Proof of Theorem 1

By Eq. (6), the estimator in Eq. (4) has the ZV property if:

$$\mathbb{E}_p\left[\mathcal{R}(\boldsymbol{x})\omega(\boldsymbol{x})\right] = \mu^2. \quad (28)$$

Next, extend Eq. (28) as the integral form and plug in $q{=}q^*$:

$$\begin{aligned}\mathbb{E}_p\left[\mathcal{R}(\boldsymbol{x})\omega(\boldsymbol{x})\right] &= \int \frac{\mathcal{R}(\boldsymbol{x})p^2(\boldsymbol{x})}{\mathbb{P}\left[\boldsymbol{x}|\mathcal{R}(\boldsymbol{x}){=}1\right]}\mathrm{d}\boldsymbol{x} {=} \int \frac{\left(\mathcal{R}(\boldsymbol{x})p(\boldsymbol{x})\right)^2}{\mathbb{P}\left[\boldsymbol{x}|\mathcal{R}(\boldsymbol{x}){=}1\right]}\mathrm{d}\boldsymbol{x} \\ &= \int \frac{\mathbb{P}\left[\boldsymbol{x},\mathcal{R}(\boldsymbol{x}){=}1\right]}{\mathbb{P}\left[\boldsymbol{x}|\mathcal{R}(\boldsymbol{x}){=}1\right]}\mathcal{R}(\boldsymbol{x})p(\boldsymbol{x})\mathrm{d}\boldsymbol{x} \\ &= \mu \int \mathcal{R}(\boldsymbol{x})p(\boldsymbol{x})\mathrm{d}\boldsymbol{x} {=} \mu^2. \end{aligned} \quad (29)$$

Hence, q^* is the optimal importance distribution. The proof of Theorem 1 is completed.

B. The Proof of Theorem 2

One can easily verify that:

- When there is no link capacity constraints, i.e., $c_i{=}\infty$, the satisfaction of flow demands only depends on the connectivity between the source and the destination. Namely, the FAVE problem becomes a NCR problem.
- When $N_f{=}1$, i.e., the flow set $\mathbb{F}$ contains only a single flow, the satisfaction of flow demands only depends on the proba- bility that the maximum achievable bandwidth from the sourceto the destination exceeds the flow demand value. Namely, the FAVE problem becomes a MFR problem.

Therefore, the proof of Theorem 2 is completed.

C. The Proof of Theorem 3

Assume q^* is approximated via a product form distribution:

$$q(\boldsymbol{x}){=}\prod_{i=1}^{N_l} q(x_i), \quad (30)$$

by minimizing their KL divergence:

$$\begin{aligned} KL(q^*\|q) &= \int q^*(\boldsymbol{x}) \log \frac{q^*(\boldsymbol{x})}{q(\boldsymbol{x})}\mathrm{d}\boldsymbol{x} \\ &= \int q^*(\boldsymbol{x}) \log q^*(\boldsymbol{x})\mathrm{d}\boldsymbol{x} - \int q^*(\boldsymbol{x}) \log q(\boldsymbol{x})\mathrm{d}\boldsymbol{x} \end{aligned} \quad (31)$$

This is equivalent to maximize the following:

$$\begin{aligned} &\int q^*(\boldsymbol{x}) \log q(\boldsymbol{x})\mathrm{d}\boldsymbol{x} \\ &= \int \mathbb{P}\left[\boldsymbol{x}|\mathcal{R}(\boldsymbol{x}){=}1\right] \log \prod_{i=1}^{N_l} q(x_i)\mathrm{d}\boldsymbol{x} \\ &= \sum_{i=1}^{N_l} \int \mathbb{P}\left[\boldsymbol{x}|\mathcal{R}(\boldsymbol{x}){=}1\right] \log q(x_i)\mathrm{d}\boldsymbol{x} \\ &= \sum_{i=1}^{N_l} \int \left(\int \mathbb{P}\left[\boldsymbol{x}|\mathcal{R}(\boldsymbol{x}){=}1\right]\mathrm{d}\boldsymbol{x}_{-i}\right) \log q(x_i)\mathrm{d}x_i \\ &= \sum_{i=1}^{N_l} \int \left(\mathbb{P}\left[x_i|\mathcal{R}(\boldsymbol{x}){=}1\right]\right) \log q(x_i)\mathrm{d}x_i \\ &= \sum_{i=1}^{N_l} (-KL(\mathbb{P}\left[x_i|\mathcal{R}(\boldsymbol{x}){=}1\right] \| q(x_i)) \\ &\quad + \int \mathbb{P}\left[x_i|\mathcal{R}(\boldsymbol{x}){=}1\right] \log \mathbb{P}\left[x_i|\mathcal{R}(\boldsymbol{x}){=}1\right]\mathrm{d}x_i). \end{aligned} \quad (32)$$

where $\boldsymbol{x}_{-i}{=}(x_1,\cdots,x_{i-1},x_{i+1},\cdots,x_{N_l})$. It is easy to show that Eq. (32) is maximized when:

$$q(x_i) = \mathbb{P}\left[x_i|\mathcal{R}(\boldsymbol{x}){=}1\right]. \quad (33)$$

Hence, taking the importance distribution in Eq. (33), the KL divergence in Eq. (31) is minimized and:

$$KL(q^*\|q){=}\int \mathbb{P}\left[\boldsymbol{x}|\mathcal{R}(\boldsymbol{x}){=}1\right] \log \frac{\mathbb{P}\left[\boldsymbol{x}|\mathcal{R}(\boldsymbol{x}){=}1\right]}{\prod_{i=1}^{N_l}\mathbb{P}\left[x_i|\mathcal{R}(\boldsymbol{x}){=}1\right]}\mathrm{d}\boldsymbol{x} \quad (34)$$

The proof of Theorem 3 is completed.

D. The Proof of Theorem 4

According to Eq. (6), we have

$$\begin{aligned} \mathbb{V}_q\left[\hat{\mu}_{IS}\right] &= \frac{\mathbb{V}_q\left[\mathcal{R}(\boldsymbol{x})\omega(\boldsymbol{x})\right]}{N} {=} \frac{1}{N}\left(\int \frac{\left(\mathcal{R}(\boldsymbol{x})p(\boldsymbol{x})\right)^2}{q(\boldsymbol{x})}\mathrm{d}\boldsymbol{x} {-} \mu^2\right) \\ &= \frac{1}{N}\left(\int \frac{\mathbb{P}^2\left[\boldsymbol{x},\mathcal{R}(\boldsymbol{x}){=}1\right]}{q(\boldsymbol{x})}\mathrm{d}\boldsymbol{x} {-} \mu^2\right) \\ &= \frac{\mu^2}{N}\int \frac{q^*(\boldsymbol{x})^2}{q(\boldsymbol{x})}\mathrm{d}\boldsymbol{x} {-} \frac{\mu^2}{N} {=} \frac{\mu^2}{N}\chi^2(q^*\|q). \end{aligned} \quad (35)$$

$$\chi^2(q^*||q) \geq \log\left[\chi^2(q^*||q)+1\right] \geq KL(q^*||q). \quad (36)$$

$$\chi^2(q^*||q) \leq \frac{1}{min_{\boldsymbol{x}} q(\boldsymbol{x})}\sqrt{2\log 2KL(q^*||q)}. \quad (37)$$

Thus, Eq. (12) is proved. The proof of Theorem 4 is completed.

E. The Proof of Theorem 5

We first derive the sequential form of the previous ZV importance distribution:

$$\begin{aligned} q^*(\boldsymbol{x}) &= \mathbb{P}\left[\boldsymbol{x}|\mathcal{R}(\boldsymbol{x}){=}1\right] \\ &= \mathbb{P}\left[x_1|\mathcal{R}{=}1\right]\cdot\mathbb{P}\left[x_2|x_1,\mathcal{R}{=}1\right]\cdots\mathbb{P}\left[x_{N_l}|\boldsymbol{x}_{1:N_l},\mathcal{R}{=}1\right] \end{aligned} \quad (38)$$

$$\begin{aligned} \mathbb{P}\left[x_i|\boldsymbol{x}_{1:i-1},\mathcal{R}{=}1\right] &{=} \frac{\mathbb{P}\left[\boldsymbol{x}_{1:i},\mathcal{R}{=}1\right]}{\mathbb{P}\left[\boldsymbol{x}_{1:i-1},\mathcal{R}{=}1\right]}, \\ &= \frac{\mathbb{P}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right]\mathbb{P}\left[\boldsymbol{x}_{1:i}\right]}{\mathbb{P}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i-1}\mathbb{P}\left[\boldsymbol{x}_{1:i-1}\right]\right]} {=} \frac{\mathbb{P}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right]}{\mathbb{P}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i-1}\right]}p_i(x_i). \end{aligned} \quad (39)$$

Hence, the proof of Theorem 5 is completed.

F. The Proof of Theorem 6

Without loss of generality, we assume $\mathbb{P}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right]{>}0$ (otherwise it is not necessary to generate the corresponding sample $\boldsymbol{x}$). Then, Eq. (14) is equivalent to:

$$\begin{aligned} \hat{\mathbb{P}}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right] &= \theta_i\mathbb{P}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right] {+} o\left(\mathbb{P}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right]\right) \\ &= \mathbb{P}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right]\left(\theta_i {+} o(1)\right), \end{aligned} \quad (40)$$

where θ_i is a constant. Thus,

$$\begin{aligned} \frac{\hat{\mathbb{P}}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right]}{\hat{\mathbb{P}}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i-1}\right]} &= \frac{\mathbb{P}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right]}{\mathbb{P}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i-1}\right]}\cdot\frac{(\theta_i+o(1))}{(\theta_{i-1}+o(1))} \\ &= \frac{\mathbb{P}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right]}{\mathbb{P}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i-1}\right]}\cdot\left(\frac{\theta_i}{\theta_{i-1}}{+}o(1)\right). \end{aligned} \quad (41)$$

Therefore,

$$\begin{aligned} q(\boldsymbol{x}) &{=} \prod_{i=1}^{N_l}\frac{\hat{\mathbb{P}}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right]}{\hat{\mathbb{P}}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i-1}\right]}p(x_i) \\ &{=} \prod_{i=1}^{N_l}\frac{\mathbb{P}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i}\right]}{\mathbb{P}\left[\mathcal{R}{=}1|\boldsymbol{x}_{1:i-1}\right]}p(x_i)\cdot\prod_{i=1}^{N_l}\left(\frac{\theta_i}{\theta_{i-1}}{+}o(1)\right) \quad (42) \\ &{=} q^*(\boldsymbol{x})\cdot\left(\prod_{i=1}^{N_l}\frac{\theta_i}{\theta_{i-1}}{+}o(1)\right). \quad (43) \end{aligned}$$

$$\begin{aligned} \mathbb{V}_q\left[\mathcal{R}(\boldsymbol{x})\omega(\boldsymbol{x})\right] &= \mu^2\int\frac{q^*(\boldsymbol{x})^2}{q(\boldsymbol{x})}\mathrm{d}\boldsymbol{x} - \mu^2 \\ &= \mu^2\left(\prod_{i=1}^{N_l}\frac{\theta_{i-1}}{\theta_i}{-}1{+}o(1)\right){=}O(\mu^2). \end{aligned} \quad (44)$$

Namely, BRE property is achieved.

Eq. (15) is a special case of Eq. (40) by restricting θ_i=1 for all i. Similar to the above proof, we can show that $\mathbb{V}_q[\mathcal{R}(\boldsymbol{x})\omega(\boldsymbol{x})]=o(\mu^2)$. Namely, the VRE property is achieved. The proof of Theorem 6 is completed.

G. The Proof of Lemma 1

Consider the statuses of some links are fixed by $\boldsymbol{x}_{1:i}$. We use $L\subseteq\{i+1,\ldots,N_l\}$ to denote the set of failed links in the remaining dimensions. The event "flow fails and all links in the cond-SEED $S(\boldsymbol{x}_{1:i})$ fail" has the following probability:

$$\mathbb{P}\left[\Psi^{-1}(\boldsymbol{x}_{i+1:N_l})\in\left\{L|L\cup\Psi^{-1}(\boldsymbol{x}_{1:i})\in\mathcal{F}, L\in\text{span}\left(S\left(\boldsymbol{x}_{1:i}\right)\right)\right\}\right] = \mathbb{P}\left[\Psi^{-1}(\boldsymbol{x}_{i+1:N_l})\in\text{span}\left(S\left(\boldsymbol{x}_{1:i}\right)\right)\right]=\Phi\left(S\left(\boldsymbol{x}_{1:i}\right)\right). \quad (45)$$

If $L\cup\Psi^{-1}(\boldsymbol{x}_{1:i})\in\mathcal{F}$, there must exist $S(\boldsymbol{x}_{1:i})\in\mathcal{S}(\boldsymbol{x}_{1:i})$ such that $L\in\text{span}\left(S\left(\boldsymbol{x}_{1:i}\right)\right)$. Also, according to the definition of the cond-SEED, if there exists a cond-SEED $S(\boldsymbol{x}_{1:i})$ which makes $L\in\text{span}(S(\boldsymbol{x}_{1:i}))$, then $L\cup\Psi^{-1}(\boldsymbol{x}_{1:i})\in\mathcal{F}$. Therefore,

$$\{L|L\cup\Psi^{-1}(\boldsymbol{x}_{1:i})\in\mathcal{F}, L\subseteq\{i+1,\ldots,N_l\}\} = \bigcup\nolimits_{S(\boldsymbol{x}_{1:i})\in\mathcal{S}(\boldsymbol{x}_{1:i})}\text{span}\left(S\left(\boldsymbol{x}_{1:i}\right)\right). \quad (46)$$

$$\mathbb{P}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i}\right] =\mathbb{P}\Big[\Psi^{-1}(\boldsymbol{x}_{i+1:N_l})\in\{L|L\cup\Psi^{-1}(\boldsymbol{x}_{1:i})\in\mathcal{F}, L\subseteq\{i+1,\ldots,N_l\}\}\Big] = \mathbb{P}\Big[\Psi^{-1}(\boldsymbol{x}_{i+1:N_l})\in\bigcup\nolimits_{S(\boldsymbol{x}_{1:i})\in\mathcal{S}(\boldsymbol{x}_{1:i})}\text{span}\left(S\left(\boldsymbol{x}_{1:i}\right)\right)\Big]. \quad (47)$$

Hence, according to the "inclusion−exclusion principle", we have Eq. (22). The proof of Lemma 1 is completed.

H. The Proof of Theorem 7

According to Theorem 5 and Lemma 1, the importance distribution q generated in the SEED-ZV algorithm satisfies q=q^*. Namely, SEED-ZV can achieve the ZV property. The proof of Theorem 7.

I. The Proof of Lemma 2

Denote $\Phi\big(S^*(\boldsymbol{x}_{1:i})\big)=\max_{S(\boldsymbol{x}_{1:i})\in\mathcal{S}(\boldsymbol{x}_{1:i})}\Phi\big(S(\boldsymbol{x}_{1:i})\big)$. According to Eq. (47), we have:

$$\mathbb{P}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i}\right]\geq\Phi\left(S^*\left(\boldsymbol{x}_{1:i}\right)\right) \quad (48)$$

$$\mathbb{P}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i}\right] \leq \sum\nolimits_{S(\boldsymbol{x}_{1:i})\in\mathcal{S}(\boldsymbol{x}_{1:i})}\mathbb{P}\left[\Psi^{-1}(\boldsymbol{x}_{i+1:N_l})\in\text{span}\left(S\left(\boldsymbol{x}_{1:i}\right)\right)\right] = \sum\nolimits_{S(\boldsymbol{x}_{1:i})\in\mathcal{S}(\boldsymbol{x}_{1:i})}\Phi\left(S\left(\boldsymbol{x}_{1:i}\right)\right)\leq\left|\mathcal{S}\left(\boldsymbol{x}_{1:i}\right)\right|\Phi\left(S^*\left(\boldsymbol{x}_{1:i}\right)\right) \quad (49)$$

The proof of Lemma 2 is completed.

J. The Proof of Theorem 8

According to Lemma 2, we have:

$$\hat{\mathbb{P}}[\mathcal{R}=1|\boldsymbol{x}_{1:i-1},x_i=1]p_i+\hat{\mathbb{P}}[\mathcal{R}=1|\boldsymbol{x}_{1:i-1},x_i=0](1-p_i) =O(\mathbb{P}[\mathcal{R}=1|\boldsymbol{x}_{1:i-1},x_i=1])p_i+O(\mathbb{P}[\mathcal{R}=1|\boldsymbol{x}_{1:i-1},x_i=0])(1-p_i) = O(\mathbb{P}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i-1}\right]) \quad (50)$$

As there is a normalizing step in Algorithm 2, $q_i(1)$ is normalized by the following:

$$q_i(1)/(q_i(1)+q_i(0)) = \frac{p_i\hat{\mathbb{P}}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i-1},x_i=1\right]}{p_i\hat{\mathbb{P}}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i-1},x_i=1\right]+(1-p_i)\hat{\mathbb{P}}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i-1},x_i=0\right]} = \frac{O\left(\mathbb{P}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i-1},x_i=1\right]\right)}{O\left(\mathbb{P}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i-1}\right]\right)}p_i. \quad (51)$$

Similar conclusion holds for $q_i(0)$. Therefore, the sampling distribution satisfies:

$$q(x_i|\boldsymbol{x}_{1:i-1})=\frac{O(\mathbb{P}(\mathcal{R}=1|\boldsymbol{x}_{1:i}))}{O(\mathbb{P}(\mathcal{R}=1|\boldsymbol{x}_{1:i-1}))}p_i(x_i), \quad (52)$$

and it matches the definition in Theorem 5. Hence, according to the proof of Theorem 6, SEED-BRE can achieve bounded relative error. The proof of Theorem 8 is completed.

K. The Proof of Lemma 3

Note that the cond-SEED $\mathcal{S}_{1:i}$ is a "*sperner family*". Hence, there is no inclusion among all members of $\mathcal{S}_{1:i}$. For $A\subseteq\mathcal{S}_{1:i}$, consider the case that $|A|\geq 2$,

$$\Big|\bigcup\nolimits_{S(\boldsymbol{x}_{1:i})\in A}S(\boldsymbol{x}_{1:i})\Big|\geq\max\nolimits_{S(\boldsymbol{x}_{1:i})\in A}\big|S(\boldsymbol{x}_{1:i})\big|+1. \quad (53)$$

Therefore, $\forall S_j(\boldsymbol{x}_{1:i})\in A$

$$\Phi\Big(\bigcup\nolimits_{S(\boldsymbol{x}_{1:i})\in A}S(\boldsymbol{x}_{1:i})\Big)=O\big(\Phi\left(S_j(\boldsymbol{x}_{1:i})\right)\epsilon\big), \quad (54)$$

$$\sum_{\emptyset\neq A\subseteq\mathcal{S}(\boldsymbol{x}_{1:i})}\Big((-1)^{|A|-1}\Phi\Big[\bigcup\nolimits_{S(\boldsymbol{x}_{1:i})\in A}S\left(\boldsymbol{x}_{1:i}\right)\Big]\Big) = \sum_{S(\boldsymbol{x}_{i:i})\in\mathcal{S}(\boldsymbol{x}_{i:i})}\Phi\big(S(\boldsymbol{x}_{1:i})\big)+O\Big(\epsilon\sum\nolimits_{S(\boldsymbol{x}_{i:i})\in\mathcal{S}(\boldsymbol{x}_{i:i})}\Phi\big(S(\boldsymbol{x}_{1:i})\big)\Big) = \sum_{S(\boldsymbol{x}_{i:i})\in\mathcal{S}(\boldsymbol{x}_{i:i})}\Phi\big(S(\boldsymbol{x}_{1:i})\big)\cdot\big(1+O(\epsilon)\big). \quad (55)$$

As $O(\epsilon)$=$o(1)$, thus,

$$\sum_{S(\boldsymbol{x}_{1:i})\in\mathcal{S}(\boldsymbol{x}_{1:i})}\Phi\big(S(\boldsymbol{x}_{1:i})\big)=\mathbb{P}\big[\mathcal{R}=1|\boldsymbol{x}_{1:i}\big]\big(1+o(1)\big). \quad (56)$$

Hence, Eq. (24) is proved. The proof of Lemma 3 is completed.

L. The Proof of Theorem 9

According to Lemma 3, we have:

$$\hat{\mathbb{P}}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i-1},x_i=1\right]p_i+\hat{\mathbb{P}}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i-1},x_i=0\right](1-p_i) = p_i\left(\mathbb{P}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i-1},x_i=1\right]+o\left(\mathbb{P}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i-1},x_i=1\right]\right)\right) +(1-p_i)\left(\mathbb{P}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i-1},x_i=0\right]+o(\mathbb{P}[\mathcal{R}=1|\boldsymbol{x}_{1:i-1},x_i=0])\right) = \mathbb{P}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i-1}\right](1+o(1)). \quad (57)$$

As there is a normalizing step in Algorithm 3, $q_i(1)$ is normalized by the following:

$$q_i(1)/(q_i(1)+q_i(0)) = \frac{p_i\hat{\mathbb{P}}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i-1},x_i=1\right]}{p_i\hat{\mathbb{P}}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i-1},x_i=1\right]+(1-p_i)\hat{\mathbb{P}}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i-1},x_i=0\right]} = \frac{\mathbb{P}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i-1},x_i=1\right](1+o(1))}{\mathbb{P}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i-1}\right](1+o(1))}p_i = \frac{\mathbb{P}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i-1},x_i=1\right]}{\mathbb{P}\left[\mathcal{R}=1|\boldsymbol{x}_{1:i-1}\right]}p_i(1+o(1)) \quad (58)$$

Similar conclusion holds for $q_i(0)$. Therefore, the sampling distribution satisfies:

$$q(x_i|\boldsymbol{x}_{1:i-1})=\frac{\mathbb{P}(\mathcal{R}=1|\boldsymbol{x}_{1:i})}{\mathbb{P}(\mathcal{R}=1|\boldsymbol{x}_{1:i-1})}p_i(x_i)(1+o(1)), \quad (59)$$

and it matches the definition in Theorem 5. Hence, according to the proof of Theorem 6, SEED-VRE can achieve the VRE pro-perty. The proof of Theorem 9 is completed.

M. The Proof of Theorem 10

By the definition of the one-run variance $\left(\sigma_q^{(k)}\right)^2$, we have:

$$\begin{aligned}\left(\sigma_q^{(k)}\right)^2 &= \int \frac{\left(\mathcal{R}^{(k)}(\boldsymbol{x})p(\boldsymbol{x})\right)^2}{\sum_j w_j q^{(j)}(\boldsymbol{x})}\mathrm{d}\boldsymbol{x} - \left(\mu^{(k)}\right)^2 \\ &\leq \int \frac{\left(\mathcal{R}^{(k)}(\boldsymbol{x})p(\boldsymbol{x})\right)^2}{w_k q^{(k)}(\boldsymbol{x})}\mathrm{d}\boldsymbol{x} - \left(\mu^{(k)}\right)^2 \\ &= \frac{1}{w_k}\left(\left(\sigma_{q^{(k)}}^{(k)}\right)^2 + (\mu^{(k)})^2\right) - \left(\mu^{(k)}\right)^2. \end{aligned} \tag{60}$$

Thus Eq. (27) is proved. As we have discussed that the SEED algorithms can guarantee at least a bounded relative error, i.e., $\sigma_{q^{(k)}}^{(k)}=O\left(\mu^{(k)}\right)$. Thus $\sigma_q^{(k)}=O\left(\mu^{(k)}\right)$. The proof of Theorem 10 is completed.

N. The Proof of Theorem 11

Consider one particular flow f_k. Similar to the definition in Eq. (18), we use $\mathcal{F}^{(k)}$ to denote the collection of all link statuses $\boldsymbol{x}$ satisfying $\mathcal{R}^{(k)}(\boldsymbol{x})=1$. Assume we are given a set of failure configurations $\mathcal{L}^{(k)}\subset\mathcal{F}^{(k)}$, and there is no inclusion among members of $\mathcal{L}^{(k)}$. When a new sample of $\boldsymbol{x}$ is generated, and $\Psi^{-1}(\boldsymbol{x})\notin\mathcal{L}^{(k)}$: If $\mathcal{R}^{(k)}(\boldsymbol{x})=0$ (i.e., flow succeeds), $\Psi^{-1}(\boldsymbol{x})$ cannot be a SEED; Otherwise, if $\mathcal{R}^{(k)}(\boldsymbol{x})=1$ (i.e., flow fails), we consider the following cases:

1) If $\Psi^{-1}(\boldsymbol{x})\in\mathrm{span}\left(\mathcal{L}^{(k)}\right)$, then $\left\{L|L\subsetneq\Psi^{-1}(\boldsymbol{x}), L\in\mathcal{L}^{(k)}\right\}\neq\emptyset$. As-sume that $\Psi^{-1}(\boldsymbol{x})$ is a SEED, according to the definition of SEED in Definition 6, $\forall L\subsetneq\Psi^{-1}(\boldsymbol{x})$, $L\notin\mathcal{F}^{(k)}\subset\mathcal{L}^{(k)}$. Due to the contradiction, $\Psi^{-1}(\boldsymbol{x})$ cannot be a SEED and we drop it.
2) If $\Psi^{-1}(\boldsymbol{x})\notin\mathrm{span}\left(\mathcal{L}^{(k)}\right)$, to make $\mathrm{span}\left(\mathcal{L}^{(k)}\right)$ provide a better coverage for $\mathcal{F}^{(k)}$, we add $\Psi^{-1}(\boldsymbol{x})$ to the SEED set and so extend $\mathrm{span}\left(\mathcal{L}^{(k)}\right)$. Before extending, $\mathcal{L}^{(k)}$ is pruned by removing $\{L|L\in\mathcal{L}^{(k)}, L\supset\Psi^{-1}(\boldsymbol{x})\}$ from it, which is a subset of $\mathrm{span}\left(\Psi^{-1}(\boldsymbol{x})\right)$.

As $\mathrm{span}\left(\mathcal{L}^{(k)}\right)$ is monotone increasing, the coverage of $\mathcal{L}^{(k)}$ can be improved through the update in Algorithm 4, and eventually $\mathrm{span}\left(\mathcal{L}^{(k)}\right)\rightarrow\mathcal{F}^{(k)}$. Also, once $\Psi^{-1}(\boldsymbol{x})$ is added into $\mathcal{L}^{(k)}$, it can only be replaced by its subset. This means that $\forall L\in\mathcal{L}^{(k)}$, L is monotone decreasing. Hence, according to the definition of SEED in Definition 6, $\mathcal{L}^{(k)}\rightarrow\mathcal{S}^{(k)}$. Hence, the proof of Theorem 11 is completed.

REFERENCES

[1] Cisco. (2013). *Best Practices in Core Network Capacity Planning*. [Online]. Available: https://communities.cisco.com/docs/DOC-35673

[2] Cisco. (2013). *Planning and Designing Networks With the Cisco MATE Portfolio*. [Online]. Available: https://communities.cisco.com/docs/DOC-36973

[3] S. J. Taylor and B. Letham. (2017). *Prophet: Forecasting at Scale*. [Online]. Available: https://research.fb.com/prophet-forecasting-at-scale/

[4] Juniper. Net. (2015). *Forecasting and Capacity Planning Assessment*. [Online]. Available: https://www.juniper.net/assets/cn/zh/local/pdf/datasheets/1000509-en.pdf

[5] A. K. Bangla *et al.*, "Capacity planning for the Google backbone network," in *Proc. ISMP*, 2015.

[6] J. M. Hammersley and D. C. Handscomb, "The general nature of Monte Carlo methods," in *Monte Carlo Methods*. Dordrecht, The Netherlands: Springer, 1964, pp. 1–9.

[7] P. W. Glynn and D. L. Iglehart, "Importance sampling for stochastic simulations," *Manage. Sci.*, vol. 35, no. 11, pp. 1367–1392, Nov. 1989.

[8] Y. Jiang, R. Li, R. Kang, and N. Huang, "The method of network reliability and availability simulation based on Monte Carlo," in *Proc. Int. Conf. Quality, Rel., Risk, Maintenance, Saf. Eng.*, Jun. 2012.

[9] W.-C. Yeh, Y.-C. Lin, Y. Ying Chung, and M. Chih, "A particle swarm optimization approach based on Monte Carlo simulation for solving the complex network reliability problem," *IEEE Trans. Rel.*, vol. 59, no. 1, pp. 212–221, Mar. 2010.

[10] P. L'ecuyer, G. Rubino, S. Saggadi, and B. Tuffin, "Approximate zero-variance importance sampling for static network reliability estimation," *IEEE Trans. Rel.*, vol. 60, no. 3, pp. 590–604, Sep. 2011.

[11] K. Lee, H.-W. Lee, and E. Modiano, "Reliability in layered networks with random link failures," in *Proc. IEEE INFOCOM*, Dec. 2010.

[12] Z. I. Botev, S. Vaisman, R. Y. Rubinstein, and P. L'ecuyer, "Reliability of stochastic flow networks with continuous link capacities," in *Proc. Winter Simulation Conf.*, Dec. 2014.

[13] I. Gertsbakh, R. Rubinstein, Y. Shpungin, and R. Vaisman, "Permutational methods for performance analysis of stochastic flow networks," *Prob. Eng. Inf. Sci.*, vol. 28, no. 1, pp. 21–38, Jan. 2014.

[14] E. Datta and N. K. Goyal, "Reliability estimation of stochastic flow networks using pre-ordered minimal cuts," in *Proc. Int. Conf. Microelectron., Comput. Commun. (MicroCom)*, Jan. 2016.

[15] *Abilene*. Accessed: Feb. 4, 2004. [Online]. Available: https://www.internet2.edu/

[16] J. Xiong and W. Gong, "A novel algorithm on network reliability estimation," *Math. Comput. Model.*, vol. 41, nos. 2–3, pp. 119–133, Jan. 2005.

[17] M. O. Ball, "Computational complexity of network reliability analysis: An overview," *IEEE Trans. Rel.*, vol. 35, no. 3, pp. 230–239, Aug. 1986.

[18] A. M. Shooman, "Algorithms for network reliability and connection availability analysis," in *Proc. Electro/Int.*, Jun. 1995.

[19] A. E. Conway, "Fast simulation of service availability in mesh networks with dynamic path restoration," *IEEE/ACM Trans. Netw.*, vol. 19, no. 1, pp. 92–101, Feb. 2011.

[20] W. Jiang, R. Zhang-Shen, J. Rexford, and M. Chiang, "Cooperative content distribution and traffic engineering in an ISP network," in *Proc. ACM SIGMETRICS*, Jun. 2009, pp. 239–250.

[21] Y. Zhang. (2014). *Six Months of Abilene Traffic Matrices*. [Online]. Available: http://www.cs.utexas.edu/~yzhang/

[22] I. Raicu, I. Foster, and A. Szalay, "Harnessing grid resources to enable the dynamic analysis of large astronomy datasets," in *Proc. IEEE Supercomput.*, Feb. 2006.

[23] S. Jain *et al.*, "B4: Experience with a globally-deployed software defined WAN," in *Proc. ACM SIGCOMM*, Aug. 2013.

[24] A. Müller, "Improved variance reduced Monte-Carlo simulation of in-the-money options," *J. Math. Finance*, vol. 6, pp. 361–367, Aug. 2016.

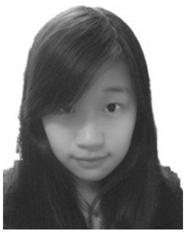

Tingwei Liu received the B.E. and M.E. degrees in electronics and information engineering from the Huazhong University of Science and Technology, China, in 2013 and 2016, respectively. She is currently pursuing the Ph.D. degree with the Department of Computer Science and Engineering, The Chinese University of Hong Kong (CUHK). Her research interests include networking algorithm and protocol design, network availability measurement, and online learning algorithms.

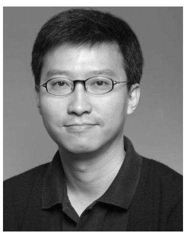

John C. S. Lui (Fellow, IEEE) was born in Hong Kong. He received the Ph.D. degree in computer science from the University of California at Los Angeles, Los Angeles, CA, USA, in 1992.

He was the Chairman of the Department from 2005 to 2011. He is currently a Professor with the Department of Computer Science and Engineering, The Chinese University of Hong Kong (CUHK), Hong Kong. His current research interests are in communication networks, network/system security (e.g., cloud security and mobile security), network economics, network sciences (e.g., online social networks and information spreading), cloud computing, large-scale distributed systems, and performance evaluation theory. He is a Fellow of the Association for Computing Machinery (ACM), a Croucher Senior Research Fellow, and an elected member of the IFIP WG 7.3. He received various departmental teaching awards and the CUHK Vice-Chancellor's Exemplary Teaching Award. He is also a co-recipient of the IFIP WG 7.3 Performance 2005 and the IEEE/IFIP NOMS 2006 best student paper awards. He serves on the Editorial Board of the IEEE/ACM TRANSACTIONS ON NETWORKING, the IEEE TRANSACTIONS ON COMPUTERS, the IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS, the *Journal of Performance Evaluation*, and the *International Journal of Network Security*.