

Probabilistic Analysis of Overload-Free Property for Critical Traffic

Zhiyun Tang*, Yahui Li*, Ke Ruan†, Yongqing Zhu†,

Yingjun Ye†, Jilong Wang*, Xia Yin*, Xingang Shi*, Han Zhang*

*Tsinghua University, China †Research Institute, China Telecom Corporation, China

tzy22@mails.tsinghua.edu.cn, liyahui@tsinghua.edu.cn

Abstract—Network structures are sophisticated and hence vulnerable to errors. Link failures and traffic load fluctuations lead to complexity in network states. Different failure scenarios can result in varying network traffic distribution patterns. Meanwhile, the load on links within the same failure scenario dynamically changes with fluctuations in traffic. Network administrators are particularly concerned about whether links along the paths traversed by critical traffic are overload-free guaranteed when link failures occur. Yet, no attention was ever paid to overload-free property analysis for critical traffic. We propose Offaela, an efficient and accurate probabilistic analysis framework that verifies an overload-free property for critical traffic. We prudently formulate the problem and prove its computational hardness, then storm this fortification by proffering a failure scenario merging algorithm and adopting a randomized approximation method. Evaluations on real networks show that Offaela outperforms the state-of-the-art solution by 4.83x and can provide availability analysis assistance such as identifying vulnerable failure scenarios.

I. INTRODUCTION

As the intranet becomes more extensively integrated into various sectors such as business, healthcare, education, etc., and serves tens of billions of customers concurrently, ensuring customer service quality has increasingly become a focal point for Internet Service Providers (ISPs). Providers often enhance customer service quality by improving network reliability and availability and tailoring network solutions for unique customers [1][2]. When customers are deciding from which provider they wish to get service, network availability often becomes a crucial reference point.

For network administrators, certain critical customer traffic is of paramount importance due to commercial requirements. Quite a few network administrators have assigned priorities to traffic demands [3]. Administrators are obligated to ensure that the overload-free property on the links along the paths traversed by critical traffic (CCT) is verified to the fullest extent. It is significant to quantitatively analyze the availability of critical customer traffic for ISPs. In recent years, numerous efforts [4]-[13] have been made with the aim of assessing network availability to assist administrators in operating the network. Data plane verification tools like Anteater [4] and VeriFlow [5] proposed a Boolean variable to assess network traffic reachability, but only single-flow traffic is taken into consideration, and simply classifying the problem as a yes-or-

Corresponding author: Yahui Li (email: liyahui_ego@126.com)




Fig. 1: Traffic load distribution under different failure scenarios

no question is extremely inadequate. Control plane verification tools like Batfish [6] and Minesweeper [7] detect reachability property violations proactively by deriving data plane models from configuration files. These works are neither competent for simulating traffic load distribution nor for probabilistically analyzing failure scenarios. Jingubang [8] analyzes traffic load property using a traffic distribution graph. However, it cannot probabilistically reason about failures. Netdice [12] improved upon these works by evaluating availability under probability failure scenarios; however, Netdice's solution is not practically multi-flow-adaptive. Pita [13] probabilistically analyzes traffic load property, yet its brute-force method for failure scenarios limits its applicability to tiny networks. Besides, no known work is capable of assessing whether links along the paths traversed by critical traffic are overload-free guaranteed when link failures occur. In this paper, we propose a probabilistic analysis framework that verifies an overload-free property for critical traffic in ISP's backbone networks.

Verifying overload-free property for critical traffic is sophisticated due to diverse failure scenarios and dynamic traffic demands. As exemplified in Fig. 1, suppose there are two demands in the network: d_1 is the critical customer traffic demand, and d_2 is a normal demand. Assuming that all link capacities are 10 Gbps. Traffics A→B and E→G do not share any link when no failure happens. Therefore, the volume of d_2 does not impact the overload-free property of traffic A→B. If the volume of d_1 uniformly distributes on [3, 12], the probability of no overload occurring for critical customer traffic A→B is $\frac{7}{9}$. However, when link E-F fails, traffic E→G is rerouted and thus shares a common link with traffic A→B. In this case, determining whether critical customer traffic experiences congestion necessitates the extensive consideration of demands d_1 and d_2 . Assuming the volume of d_2 uniformly distributes on [5, 18], the probability of traffic A→B being overload-free decreases to $\frac{2}{27}$.

The toy example briefly clarifies our overload-free property verification problem. We face a calculation challenge in this verification process. While verifying the overload-free property on a specific single traffic pattern is relatively simple, for works on network simulations and emulations can transform this problem into a tractable one without effort, verifying the property over all possible traffic patterns, whose union is a high dimensional continuous space, is arduous. To conquer this intractability, we resort to a stochastic approximation method [15] based on Multiphase Markov Chain Monte Carlo (Multiphase MCMC) to acquire an approximate solution without loss of accuracy compared to the state-of-the-art solutions. Although the stochastic method can be successfully applied to theoretically overcome the intractability difficulty by reducing the problem into one that is applicable to this method, it still takes tens to thousands of minutes to quantitatively analyze a network's overload-free property; hence, the computational hardness persistently perplexes us.

Our key idea is based on the observation that the forwarding behavior of some flows remains consistent in many failure scenarios, leading to similarities in the overload-free property across these scenarios. This fact provides an insight into merging analysis processes in different failure scenarios, allowing for a substantial reduction in sampling and computation processes. We recognize the potential to optimize the analysis process of overload-free property based on the specific characteristics of routing protocols, enabling rapid and accurate identification of all optimizable scenarios.

The above gives us a glimpse of a more delicate constructive framework for verifying network availability of critical customer traffic, Offaela. In Section II, we formally formulate the problem. In Section III, we elaborate on our solution, Offaela, and how it effectively solves this problem. In Section IV, we evaluate Offaela's performance.

We make the following contributions.

- To the best of our knowledge, we are the first to propose a probabilistic verification of a critical-traffic-related property, namely CCT-overload-free availability.
- We propose an efficient and accurate probabilistic analysis framework, Offaela, to calculate the value of CCT-overload-free availability. We proffer a new concept *Critical Endnode Matrix* with prudent inference to reduce the computational overhead and adopt a stochastic method based on Multiphase MCMC for approximation.
- We evaluate the efficiency and effectiveness of Offaela by conducting experiments on real networks. Results show that Offaela outperforms the state-of-the-art approach by up to 4.83x and can identify failure scenarios that are susceptible to overloading.

II. PROBLEM DEFINITION

A. Model

Suppose (V, E) is the network topology in which V stands for routers or nodes in the network, and E represents direct IP links. Each link $e \in E$ has a maximum link speed, also called capacity, and a weight number suggests its cost. We

use $C_e \geq 0$ and $W_e \geq 0$ to represent the capacity and weight, respectively. D is the set of traffic flows. For each demand $d \in D$ we have an upper bound u_d , correspondingly a lower bound l_d , and a continuous random variable x_d uniformly distributed in range $[l_d, u_d]$ which is the traffic volume of demand d . Assuming different traffic distributions would still make our approach viable [13]. We use a uniform distribution here for simplicity in our explanation. h is a given Critical-Customer-Traffic demand, $h \in D$. X is the traffic volume set, for each $x_d \in X, d \in D$, x_d is the traffic volume of demand d .

B. Problem Formulation

Now we begin to give some definitions related to Critical-Customer-Traffic, CCT for short.

Definition 1. CCT-related links Υ_f .

Given a failure scenario $f \subseteq E$, which is the set of failed links under this failure scenario. Υ_f is defined as:

$$\Upsilon_f = \{e \in E \mid \text{s.t. } \gamma_{h,e}^f = 1\} \quad (1)$$

$\gamma_{h,e}^f$ is an indicator that equals 1 when traffic demand h traverses the link e in the failure scenario f and 0 otherwise.

Definition 2. CCT-related demands Θ_f .

Given a failure scenario $f \subseteq E$, Θ_f is defined as:

$$\Theta_f = \{d \in D \mid \exists v \in \Upsilon_f, \text{ s.t. } \gamma_{d,v}^f = 1\} \quad (2)$$

Θ_f is the set of demands that at least share one common link with critical demand h under scenario f .

Definition 3. CCT-overload-free Lebesgue function $\varphi_f(x)$.

Given a failure scenario $f \subseteq E$. x is an input vector of demands $x = [x_{d_1}, x_{d_2}, \dots, x_{d_{|\Theta_f|}}]$, where $d_i \in \Theta_f$ for i in $\{1, 2, \dots, |\Theta_f|\}$ and $x_{d_i} \in [l_{d_i}, u_{d_i}]$. $\varphi_f(x)$ is defined as:

$$\varphi_f(x) = \begin{cases} 1, & \sum_{d \in \Theta_f} x_d \cdot \gamma_{d,v}^f \leq C_v, \forall v \in \Upsilon_f \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

Definition 4. CCT-overload-free probability ρ_f .

Given a failure scenario $f \subseteq E$, ρ_f is defined as:

$$\rho_f = \frac{\int_{R^{|\Theta_f|}} \varphi_f(x) dx}{\prod_{d \in \Theta_f} (u_d - l_d)} \quad (4)$$

Here $R^{|\Theta_f|}$ is $|\Theta_f|$ -dimensional real number set. ρ_f is the ratio of traffic loads that cause no congestion on CCT-related links to the total variation range of traffic loads.

Definition 5. CCT-overload-free availability Ω .

Given a set of failure scenarios F , where $\sum_{f \in F} Pr(f) = 1$. Ω is defined as follows:

$$\Omega = \sum_{f \in F} \rho_f Pr(f) \quad (5)$$

Failure scenario set F and its corresponding failure probabilities could either be given manually or generated by specifying the maximum number of link failures.

C. Challenges and Solutions

Due to the **complexity and diversity of network states**, the flows sharing links with critical flows vary under different failure scenarios. This leads to a significant amount of intricate

computations. We propose the concept of **Critical Endnode Matrix**, which can assist in optimizing this process in III.B.

Furthermore, the **multitude of failure scenarios** complicates the computation. Estimating reliability in each failure scenario then multiplying it by the probability of the failure scenario and summing over all scenarios is computationally intensive. Although the forwarding behavior consistency lights our way of saving computation costs, the locating of similar failure scenarios is not straight-forward: one needs to compute every link that any traffic demand traverses on all failure scenarios and then compare the forwarding behavior of demands between every pair of failure scenarios. On the contrary, **Critical Endnode Matrix** can spare us the trouble. By following the approach described in Section III.C, we **merge failure scenarios**, reducing the number of reliability estimations required. Experiment results show that this leads to a significant improvement in computation speed, up to 4.83x.

The third challenge falls on the calculation process of Ω value, which is defined in II.B; we face a computational challenge since computing the value of CCT-overload-free probability ρ_f is a **#P-hard problem**. This conclusion is trivial for φ_f is a boolean formula and integrating it over $R^{|\Theta_f|}$ is a #SAT problem. We solve this challenge by **reducing this problem to the hyper-polyhedron volume computation problem**, which has practical polynomial time approximation solutions with less than 10% loss of accuracy [15] in the method of Multiphase MCMC.

III. METHODOLOGY

We provide an overview of our solution in III.A, then introduce two insights that can help accelerate the computation process in III.B. We state our principle of merging failure scenarios in III.C and present the implementation in III.D.

A. Overview

We take network topology (V, E) , link capacity set $\{C_e | e \in E\}$, weight set $\{W_e | e \in E\}$ and a demand set D as input, as well as the lower and upper bounds of each demand $\{l_d | d \in D\}$, $\{u_d | d \in D\}$. Concurrently we have a given set of failure scenarios F with corresponding probability $\{Pr(f) | f \in F\}$.

To calculate the CCT-overload-free availability Ω , we need to determine the distribution of critical customer traffic across links for each failure scenario, as well as the traffic loads and their balancing on those links. In other words, we are required to compute the CCT-related links Υ_f , demands Θ_f , and the indicator function set $\gamma_{h_i, e_j}^f, \forall h_i \in \Upsilon_f, \forall e_j \in \Theta_f$ defined in II.B. Instead of brute-force, we design a more efficient algorithm to compute them by introducing **Critical Endnode Matrix** and two insights in III.B.

Next we approximate the value of CCT-overload-free probability ρ_f in polynomial time by means of Multiphase MCMC. Once we know the value probability ρ_f for all failure scenarios $f \in F$, we can calculate the availability Ω by definition; thus, the ultimate goal is achieved. The calculation process of CCT-overload-free probability ρ_f in some failure scenarios can be merged under certain conditions. Since this process

contains mountains of random points sampling and validating, which contribute to the majority of the runtime, merging failure scenarios can avoid recalculations of CCT-overload-free probabilities, which take up massive computation time. We propose an efficient method by first demonstrating under what conditions scenarios can be merged in III.C, then put this conclusion into practical implementation in III.D.

B. Critical Endnode Matrix

In this section we introduce **Critical Endnode Matrix**, the key idea of this work. We first give the definition of **Critical Endnode Matrix** in Definition 8, then introduce two insights in Theorem 1 and 2 to demonstrate how **Critical Endnode Matrix** efficiently assist us in speeding up our availability analysis.

Definition 6. node pair hash function ψ

Suppose $V = \{v_1, v_2, \dots, v_{|V|}\}$, $\psi : V \times V \rightarrow \mathbb{N}^+$:

$$\psi(v_i, v_j) = (i - 1) \times |V| + j, \quad \forall i, j \in \{1, 2, \dots, |V|\}$$

Note that ψ is a reversible function and its input is ordered.

Definition 7. demand-CCT hash function $\{\Psi_f^h\}$

Given failure scenario f and a CCT demand h , define function $\Psi_f^h : D \rightarrow \mathbb{N}^+$:

$$\Psi_f^h(d) = \begin{cases} \psi(g_f(d)), & \mathcal{L}_d^f \cap \mathcal{L}_h^f \neq \emptyset \\ -1, & \text{otherwise} \end{cases}, \forall d \in D$$

Here $g_f(d)$ is a function that outputs an ordered pair of nodes which are sequentially the first and last node in $\mathcal{N}_f^h$ that d passes by. $\mathcal{N}_f^h$ is the set of nodes demand h goes through. $\mathcal{L}_d^f$ is the set of links d traverses through under scenario f .

Definition 8. Critical Endnode Matrix ST

Given the conditions in Definition 7, define ST matrix:

$$ST[i][j] = \Psi_f^h(d_{i,j}), \quad \forall v_i, v_j \in V$$

$d_{i,j}$ represents the demand from v_i to v_j . The value $ST[i][j]$ is a positive integer obtained by hashing the node pair (v_i, v_j) .

By definition, Matrix ST logs the first and last node in $\mathcal{N}_f$ that each demand d passes by. We introduce this new concept that can help unravel the mergeable scenarios by analyzing costs on links. Our solution is adaptive to standard ISP backbone networks because most IGP routing protocols such as OSPF, RIP and ISIS, select paths basing on cost.

Now we propose two theorems. Owing to space constraints, the proofs of the theorems introduced in this section are relegated to Appendix A of the full version, which can be downloaded from GitHub¹.

Theorem 1. Let $\mathcal{N}_f$ be the set of nodes that demand h traverses through. Let $\mathcal{L}_{a,b}(a, b \in V)$ be the set of links that demand from a to b traverses through. $\forall v \in V \setminus \mathcal{N}_f$, define function $\phi : V \setminus \mathcal{N}_f \rightarrow \mathcal{N}_f$, $\phi(v)$ is the nearest node in $\mathcal{N}_f$ to v , which means: $\sum_{e \in \mathcal{L}_{v, \phi(v)}} W_e \leq \sum_{e \in \mathcal{L}_{v, u}} W_e, \forall u \in \mathcal{N}_f$. Define a family of sets $\{\mathcal{S}_u\}, u \in \mathcal{N}_f$ as: $\mathcal{S}_u \triangleq \{v \in V \setminus \mathcal{N}_f | \phi(v) = u\}$. Then $\forall a, b \in \mathcal{S}_u, \mathcal{L}_{a,b} \cap \Upsilon_f = \emptyset, \forall u \in \mathcal{N}_f$.

¹<https://github.com/salazar1117/Offaela>

Theorem 1 reveals a potential to accelerate the calculation process of acquiring Θ_f . For two nodes whose nearest neighbor in $\mathcal{N}_f$ are identical, the shortest path between those two nodes shall never share same member with Υ_f .

Theorem 2. $\forall d \in \Theta_f, a_d, b_d \in \mathcal{N}_f$ being the first and last node in $\mathcal{N}_f$ that d passes by, then $\mathcal{L}_d^f \cap \Upsilon_f = L_{a_d, b_d}$. Here $\mathcal{N}_f$ and L_{a_d, b_d} is defined as in Theorem 1. $\mathcal{L}_d^f$ is the set of links d traverses through under scenario f .

Theorem 2 indicates that to calculate matrix Γ , instead of enumerating every link each demand traverses through and verifying whether it belongs to Υ_f , we only need to know the first and last node in $\mathcal{N}_f$ that each demand $d \in D$ passes by, which is recorded in *Critical Endnode Matrix* ST .

C. Failure Scenario Mergence

If $\exists f_1, f_2$ with same relating CCT-related demands and links, and $\gamma_{\theta_j, v_i}^{f_1} = \gamma_{\theta_j, v_i}^{f_2}, \forall v_i \in \Upsilon_f, \forall \theta_j \in \Theta_f$, then the computation phase of CCT-overload-free probability ρ_f for these two failure scenarios can be merged. This observation can lead to a theorem as follow:

Theorem 3. If $\Psi_{f_1}^h = \Psi_{f_2}^h, \forall h \in H$, then $\rho_{f_1} = \rho_{f_2}$.

Theorem 3 unveils a fact that, for two given failure scenarios. If both the first and last node in $\mathcal{L}_h$ that $d, \forall d \in D$ passes by in the two failure scenarios are the same, $\mathcal{L}_h$ being the set of nodes h goes by, then the CCT-overload-free availability values of the two failure scenario are equal. Thus, the value is required to be calculated only once, saving a lot of calculation labor from avoidance of recomputing.

D. Implementation

Algorithm 1: Offaela: Computing availability Ω

```

input : network topology ( $V = \{v_1, v_2, \dots, v_{|V|}\}, E$ ) with
        weight and capacity for each link,
        Critical-Customer-traffic demand  $h$ , failure scenarios  $F$ 
        with probability  $Pr(f)$  for  $f \in F$ 
output : Approximation of availability  $\Omega$ 
1  $\Omega \leftarrow 0, P \leftarrow \emptyset;$  //  $P$  is a set of arrays
2 for  $f \in F$  do
3   Initialize  $Q$  as an empty array; //  $Q$  is an array of matrices
4    $flag \leftarrow 0; G \leftarrow$  weighted adjacency matrix of  $(V, E \setminus f);$ 
5   if  $h$  is unreachable in  $(V, E \setminus f)$  then
6      $flag \leftarrow 1; \text{break};$ 
7   Calculate  $\mathcal{L}_h$ , the set of nodes  $h$  passes by;
8    $K \leftarrow \{j | v_j \in \mathcal{L}_h\};$ 
9    $ST \leftarrow \text{Calculate\_ST}(G, h, K);$  ▷ Algorithm 2
10  Add  $ST$  into array  $Q;$ 
11 if  $\{flag = 1\};$  //  $\rho_f = 0$  if critical
12   continue; // demand is unreachable
13 if  $Q \in P$  then
14    $\Omega \leftarrow \Omega + \zeta_Q \cdot Pr(f); \text{continue};$ 
15 Calculate demand set  $\Theta_f$  and matrix  $\Gamma$ ;
16  $\rho_f \leftarrow \text{Multiphase\_MCMC}(\Theta_f, \Gamma);$  // see appendix C of full version1
17  $\Omega \leftarrow \Omega + \rho_f \cdot Pr(f);$ 
18  $\zeta_Q \leftarrow \rho_f;$  Add  $Q$  into  $P;$ 

```

Algorithm 1 illustrates the overall process of Offaela. G stands for the weighted adjacency matrix of the current topology (under failure scenario f). For each f , we compute the *Critical Endnode Matrix* ST . Each element $ST[i][j]$ representing the hash value of the pair of nodes that are the first

and last node in $\mathcal{N}_f$ that demand from node v_i to node v_j passes by as defined in III.B, where $\mathcal{N}_f$ is the set of nodes the critical demand h passes by in f . If ST is equivalent to a *Critical Endnode Matrix* of another failure scenario f' , then $\rho_f = \rho_{f'}$, thus we can skip the random approximation phase for f . Otherwise, we finish the procedure of MCMC for f in line 21 and record the value of *Critical Endnode Matrix* as well as ρ_f .

Algorithm 2: Calculate ST

```

input :  $G, h, K$ 
output :  $ST$ 
1  $NT, DS \leftarrow \text{Dijkstra}(G); m \leftarrow |K|;$ 
2 Generate a zero matrix  $ST$  of size  $n \times n;$ 
3 Generate an array  $M$  of  $m$  empty sets;
4 for  $i = 0$  to  $n - 1$  do
5    $j \leftarrow j_i, \text{ s.t. } DS[i][j_i] = \min\{DS[i][j] | j \in K\};$ 
6   Add  $i$  into set  $M[j];$  // Find  $v_i$ 's closest node in  $\mathcal{N}_f$ 
7 for  $i, j = 0$  to  $n - 1$  do
8   if  $\exists r, \text{ s.t. } i, j \in M[r]$  then
9      $ST[i][j] \leftarrow -1;$  // Demand from  $v_i$  to  $v_j$  is irrelevant
10 for  $i, j = 0$  to  $n - 1$  do
11   if  $ST[i][j] = 0$  then
12     DFS_critical_pair( $K, i, j$ ); ▷ Algorithm 3

```

Algorithm 2 shows the calculation process of *Critical Endnode Matrix* ST . Matrix NT is the output matrix of Dijkstra algorithm. $NT[i][j]$ stands for the next hop of node traffic demand from v_i to v_j passes by. Matrix DS is another output matrix of Dijkstra algorithm with each value meaning the total cost between nodes, equivalent to ∞ if the demand is unreachable. For each node v_i , we find node v_{j_i} , such that v_{j_i} is the node closest to v_i in $\mathcal{N}_f$. Then for each node v in $\mathcal{N}_f$, we have a set of nodes whose nearest node in $\mathcal{N}_f$ is v . With Theorem 1, we can infer the demand between those nodes in the same set would never share common links with critical demand h . Hence we give value -1 to the corresponding ST element, marking this demand as an insignificant demand since its load would never affect a critical demand in this case. This saves at least $\frac{(n-m)(n-2m)}{mn(n-1)}$ of calculation burden, m being the total number of nodes critical demand h passes by, reducing the computational complexity of this phase from $O(n^2)$ to $O(\frac{m-1}{m}n^2)$. Thus in line 20 of Algorithm 1 Θ_f is the set of demands whose corresponding ST value does not equal to -1, and Γ can be deduced by the ST matrices as explained in section III.B.

Algorithm 3: DFS_critical_pair

```

input :  $K, a, b$  // Suppose  $d$  is the demand from  $v_a$  to  $v_b$ 
output :  $ST[a][b]$  // The first and last node in  $\mathcal{N}_f$   $d$  passes by
1 if  $ST[a][b] \neq 0$  then
2   return  $ST[a][b];$  // Base case:  $ST$  value for  $d$  is known
3 if  $a \in K$  and  $b \in K$  then
4    $ST[a][b] \leftarrow \psi(v_a, v_b);$  // Base case:  $v_a, v_b \in \mathcal{N}_f$ 
5   return  $ST[a][b]$ 
6  $a_1 \leftarrow NT[a][b]$  if  $a \notin K$ , otherwise  $a;$ 
7  $b_1 \leftarrow NT[b][a]$  if  $b \notin K$ , otherwise  $b;$ 
8  $tmp \leftarrow \text{DFS\_critical\_pair}(K, a_1, b_1);$  //  $ST[a][b] = ST[a_1][b_1]$ 
9  $ST[a][b] \leftarrow ST[a_1][b] \leftarrow ST[a][b_1] \leftarrow tmp;$ 
10  $ST[a][b] \leftarrow \text{ACLcheck}(a, b, ST[a][b]);$ 
11 return  $ST[a][b]$ 

```

Algorithm 3 is a depth-first search (DFS) for efficiently computing the value of ST . Given node pair v_a, v_b , if both




Fig. 2: Average runtime of Offaela and Pita at fixed demand scale upper bound. The x-axis represents the demand scale.




Fig. 3: Percentage of single Critical-Customer-Traffic demand that Offaela runs without MCMC process at fixed demand scale upper bound and k value




Fig. 4: Percentage of failure scenarios merged and pruned on average at fixed demand scale upper bound and k value

Network	Nodes	Edges	Demands	# Failure scenarios
Abilene	12	15	132	121
Cernet	40	58	1560	1712
Geant	27	38	702	742
FITI	40	44	1560	991

TABLE I: Topologies for evaluation

Demand Scale Failure Scenario	0.8		1.0		1.2		1.4		1.6		1.8	
	k=1	k=2	k=1	k=2	k=1	k=2	k=1	k=2	k=1	k=2	k=1	k=2
Average Availability (%)												
Abilene	99.854	99.171	98.607	82.430	98.499	80.132	79.724	70.716	64.996	51.019	62.793	55.821
Cernet	99.776	98.779	99.732	95.948	96.763	95.688	96.718	95.555	96.701	95.614	96.518	95.595
Geant	99.513	96.144	96.965	93.537	95.854	93.211	92.842	88.715	89.492	88.320	89.518	89.561
FITI	99.533	98.484	99.912	97.519	95.380	93.912	94.191	93.514	94.031	92.921	93.910	92.732

TABLE II: Average availability at fixed demand scale upper bound and k value

nodes belong to $\mathcal{N}_f$, we already have $\Psi(d_{a,b})$ as the ST value. Otherwise, suppose $v_a \notin \mathcal{N}_f$, we define v_{a_1} as the next hop of demand from v_a to v_b , then it is trivial to infer that the first and last node in $\mathcal{N}_f$ that demand from node v_a to node v_b passes by is the same for demand from node v_{a_1} to node v_b . Before returning the value, we perform an access control list (ACL) check on the traffic. If there are corresponding ACL entries on the transit nodes that prevent the traffic from being delivered successfully, we will modify the value of ST based on the actual behavior of the traffic forwarding. For each demand, our algorithm only needs to compute its next hop once; thus, the amortized complexity to compute Γ for each demand d goes down from $O(|\mathcal{L}_d^f|)$ to $O(1)$. $\mathcal{L}_d^f$ is the set of links d travels through. In summary, the computational complexity for each failure scenario is $O(\frac{m-1}{m}n^2)$, while brute-force executes $\sum_{d \in D} \sum_{e \in E} \gamma_{d,e}^f$ times of operations, making the complexity of the latter larger than $O(n^2)$ since at least one link is passed through by each traffic.

IV. EVALUATION

In this section, we evaluate the efficiency of our framework, Offaela. Our evaluations focus on the following questions:

- How efficient is Offaela? (Section IV.B)
- Does Offaela demonstrate scalability? (Section IV.B)
- How does our optimization enhance Offaela's performance? (Section IV.C)
- How can Offaela be utilized to identify vulnerable failure scenarios? (Section IV.D)

A. Experimental Setting

Setup. We implement Offaela in Python with the open-source implementation of VolEsti [15] as the basis for our hyperpolyhedron volume estimation algorithm. All experiments run on a computer with 48 CPU cores at 3.2GHz and 64GB RAM.

Topologies and demand scaling. We use three network topologies from Topology Zoo [14]: Abilene, Cernet, Geant,

and FITI, as shown in TABLE 1. We utilize gravity model [16] to generate traffic matrices such that without any link failure, the resulting link utilization is between 0.6 and 1.8.

Deriving failure scenarios. To evaluate the robustness and tractability of our model, we consider scenarios in which at most 2 links fail simultaneously ($k = 1, 2$), with link failures independent and identically distributed. Presume the failure probability of any link is 0.001 and scale up the probability of each failure scenario until $\sum_{f \in F} Pr(f) = 1$ is satisfied.

B. Verification Performance

We now evaluate the performance of our verification solution in analyzing the overload-free property for CCT demands. For comparison, we implement the state-of-the-art solution, Pita [13], based on the algorithm outlined in Appendix A of [13]. Apart from Pita, no other work can be applied to probabilistically verifying the overload-free property of networks under dynamic traffic demands. Our solution produces the same probability outcomes as Pita.

Running time. We compute CCT-overload-free availability Ω for each single critical traffic at demand scale ranging from 0.6 to 1.8. We report the average running time of Offaela and Pita at fixed demand scale upper bound in Fig. 2. We executed 20, 533, 200 and 705 single CCT demand scenarios for Abilene, Cernet, Geant and FITI, respectively. Offaela outperforms Pita on all four topologies. On average, with a demand scale upper bound of 1.8, the runtime of Offaela compared to Pita improved by 1.15x, 4.83x, 2.29x, and 2.57x on Abilene, Cernet, Geant, and FITI, respectively, showing the superiority of Offaela over Pita.

In summary, the answer to the **first question** is that **Offaela outperforms state-of-the-art solution** in analyzing overload-free property for critical traffic in the aspect of runtime.

Scalability. To examine the scalability of Offaela, we scale the demand up and record the average running time. Fig. 2 elucidates average runtime with respect to the demand




Fig. 5: Overload-free probabilities under each single-link failure in Abilene

scale at $k = 1$. Despite the positive correlation between demand scale and running time, the growth rate of average runtime slows down as the demand scale increases. This phenomenon arises from the fact that most of the evaluation time is allocated to the MCMC process, whose speed is contingent upon the number of CCT-related demands $|\Omega_f|$. As the demand scale fluctuates, $|\Omega_f|$ tends to remain constant in each failure scenario. We further analyzed the proportion of single CCT demands requiring at least one MCMC process, as illustrated in Fig. 3. The rate of growth in the number of failure scenarios susceptible to overload, leading to the MCMC process, slows down. The average runtime of Offaela remains within a relatively low range for the total count of MCMC processes Offaela conducted does not significantly increase. Consequently, our approach exhibits a gradual increase in runtime with demand scale expansion, rendering it suitable for a wide range of usage and demonstrating high scalability.

In summary, the answer to **the second question** is that Offaela is **adaptive to the scaling up of demand scale**.

C. Offaela Optimization

In this subsection we inspect the main optimization applied in Offaela and evaluate its effectiveness.

Merge and prune failure scenarios. We illustrate the average percentage of failure scenarios merged or pruned in Fig. 4. Merging indicates the scenarios where the calculation of CCT-overload-free probability ρ_f is bypassed, as elucidated in III.C. Pruning refers to scenarios where maximum traffic load does not cause congestion or minimum traffic does. Both merging and pruning can optimize Offaela's analysis by eliminating the need to recalculate ρ_f . The percentages of scenarios merged are 31.86%, 70.45%, 46.57% and 21.67% for the four topologies respectively at $k = 2$. Besides, it can be noticed that as k value and demand scale grows larger, the percentage of scenarios pruned goes down on all topologies except Cernet. Also, the total percentage of pruned and merged failure scenarios exceeds 80% in all traffic schemes and topologies except Abilene. This indicates that many scenarios are spared from the MCMC process, thus saving time.

In conclusion, the answer to **the third question** is that **failure scenario merging and pruning** are significantly improving the performance of Offaela.

D. Analyzing Real Networks

The results of the analysis on real networks are presented in this subsection. We demonstrate how Offaela can be applied to identify vulnerable failure scenarios.

Table II records the average availability at fixed demand scale upper bound and k value. When the demand scale is lower than 1.2, the average availabilities of all four topologies are above 80%. We also analyze and illustrate the overload-

free probability of Abilene under each single-link failure with an upper demand scale of 1.8 in Fig. 5, considering traffic demand from node 1 to node 5 as the CCT demand. When link 2-6 goes down, the availability becomes 84.747%. When link 1-2 fails, the demand from node 1 to node 5 becomes unreachable. The failure of link 2-5 (6-7) has a devastating impact on the network; the availability decreased to near zero because the failure of either would make the other link a cut in the topology graph.

Thus, the answer to **the fourth question** is that **Offaela is capable of identifying failure scenarios that are vulnerable to congestion**.

V. CONCLUSION

We presented Offaela, an efficient and accurate probabilistic analysis tool for verifying the overload-free property of critical traffic in ISPs. We established a refined model and formulated a property verification problem addressed by Offaela. Utilizing a randomized approximation method and a novel inference algorithm, Offaela significantly reduces calculation overhead. On real networks we evaluate its effectiveness in assisting network administrators proactively assess the potential impact link failures can make on critical traffic. This tool proves invaluable for network management and is currently being deployed in the network operation center of China Telecom.

ACKNOWLEDGMENT

We thank the anonymous reviewers for their invaluable feedback. This work was supported by the National Natural Science Foundation of China under Grant No.62394322.

REFERENCES

- [1] Llopis, Joan Meseguer, et al. Minimizing latency of critical traffic through SDN. In NAS, 2016.
- [2] Saha N, Bera S, et al: Traffic-aware QoS routing in software-defined IoT[J]. IEEE Transactions on Emerging Topics in Computing, 2018.
- [3] Hong, Chi-Yao, et al. Achieving high utilization with software-driven WAN. In SIGCOMM, 2013.
- [4] Mai, H., Khurshid, A., et al. Debugging the data plane with anteaer. ACM SIGCOMM Computer Communication Review, 2011.
- [5] A. Khurshid, X. Zou, et al. Veriflow: Verifying network-wide invariants in real time. In NSDI, 2013.
- [6] Fogel, Ari, et al. A general approach to network configuration analysis. In NSDI, 2015.
- [7] Beckett, Ryan, et al. A general approach to network configuration verification. In SIGCOMM, 2017.
- [8] Li, Ruihan, et al. Reasoning about Network Traffic Load Property at Production Scale. In NSDI, 2024.
- [9] Subramanian, Kausik, et al. Detecting network load violations for distributed control planes. In PLDI, 2020.
- [10] Gember-Jacobson, Aaron, et al. Fast control plane analysis using an abstract representation. In SIGCOMM, 2016.
- [11] Kazemian, Peyman, George Varghese, and Nick McKeown. "Header space analysis: Static checking for networks." 9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12). 2012.
- [12] Steffen, Samuel, et al. Probabilistic verification of network configurations. In SIGCOMM, 2020.
- [13] Zhang, Y., Xu, H., et al. Probabilistic Analysis of Network Availability. In ICNP, 2022.
- [14] "Topology zoo." <http://www.topology-zoo.org/>.
- [15] Chalkis A, Emiris I Z, et al. A practical algorithm for volume estimation based on billiard trajectories and simulated annealing[J]. ACM Journal of Experimental Algorithmics, 2023, 28: 1-34.
- [16] M. Roughan, "Simplifying the Synthesis of Internet Traffic Matrices," ACM SIGCOMM Computer Communication Review, 2005.