Copyright Notice

Institute of Communication Networks and Computer Engineering
University of Stuttgart
Pfaffenwaldring 47, D-70569 Stuttgart, Germany
Phone: ++49-711-685-68026, Fax: ++49-711-685-67983
Email: mail@ikr.uni-stuttgart.de, http://www.ikr.uni-stuttgart.de

Final published version: https://doi.org/10.1364/JOCN.510726

[Invited] Availability Estimation of Optical Network Links Using a Bayesian Model

FILIPPOS CHRISTOU

Institute of Communication Networks and Computer Engineering (IKR), University of Stuttgart, Stuttgart, Germany
filippos.christou@ikr.uni-stuttgart.de

Compiled February 16, 2024

Common daily activities rely on information exchange and processing. Optical fiber links are the predominant way for transmitting vast amounts of data, thus enabling processing in remote data centers. As critical use cases like autonomous driving and biomedical procedures start relying on such infrastructure, the system's availability becomes even more relevant. Assessing the availability of an optical link is a well-known problem, but inconclusive nonetheless. Finding the true link availability requires a perfect understanding of the complete underlying system, which is impossible to capture to such an extent. Hence, different approaches or models arise as we focus on approximating the true value. Here, we develop a hierarchical Bayesian model and compare it to various baselines. We show that the estimation methods present different behavior for separate scenarios. Moreover, a use case is investigated where services with varying availability requirements must be deployed. Using a Bayesian model to estimate the link availabilities produces, on average, the best accuracy among the considered baselines and provides worthy uncertainty estimations. Such estimations increase the network operator's trust and allow more decision-making flexibility.

http://dx.doi.org/10.1364/ao.XX.XXXXXX

1. INTRODUCTION

It is safe to argue that the most essential aspect of an engineered system is that it operates in the first place. For optical network links, this means that data can traverse a fiber cable and reach the destination successfully. As perpetual operation cannot be guaranteed, scientists devised measures to describe the degree to which a system is expected to operate. All constructed measures came under the umbrella of the generic term *availability*. The availability of a network link signifies the time amount for which the particular connection is operational. Optical network links can fail due to several reasons, like fiber cuts (e.g., due to nearby construction activities), device failure (e.g., due to an amplifier or transponder fault), human errors (e.g., misconfigurations), or natural disasters [1].

Upon failure, fault management schemes are activated to reestablish connectivity using restoration or protection techniques [2]. To develop efficient fault management schemes, it is crucial to have a good availability model of the underlying system behavior since different models lead to different solutions. For example, [3] differs from previous research in that it assumes that up to two links can fail simultaneously in a network instead of just one. This modeling difference leads the authors to develop a specialized backup precomputation algorithm. Having a reliable model of the link availabilities and, consequently, the end-to-end connections can have a real impact and significantly help convince network operators to stop overprovisioning resources and decrease power consumption [4] and other costs.

A good underlying availability model is also pivotal when dealing with stringent availability requirements. Service Level Agreements (SLAs) are contractual agreements between service providers and customers that define the Quality of Service (QoS) that shall be provided [5]. Certain metrics are used to evaluate whether the service fulfills the requirements. For example, the agreed availability is assessed by measuring the *empirical interval availability* throughout the contract period. Penalties are incurred when the measured metrics do not comply with the SLA. An accurate underlying availability model can help the network operator minimize these penalties and provide the agreed service to the customers.

Yet, to construct availability metrics, one must first determine a model of the underlying system. An engineering model is a representation or abstraction of a system that is used to understand, analyze, or predict its behavior. Models can be either data-driven, depending solely on measurements, or incorporate expert knowledge regarding the system's dynamics [6]. Traditionally, pure data-driven approaches have been more popular due to their simplicity and overall decent results. However, in cases where measurements are scarce and data might be missing, like failure incidents in a fiber link, such approaches become

unreliable. Bayesian modeling [7] is a technique that combines data observations and expert knowledge to boost estimation accuracy while also having increased interpretability. Given its probabilistic nature, it also provides valuable uncertainty assessments, which can be propagated to newly constructed entities that might be of interest.

In this work, we build a hierarchical Bayesian model, which operators can use to accurately estimate the link availabilities and manage their networks efficiently thereafter. Several positive repercussions can stem from this groundlaying model, impacting a wide sphere of applications and research for reliable optical networks. In the following Section, we mention past work in the field and the contribution of this paper. In Section 3, we introduce the model and an extended version before conducting a prior analysis. In Section 4, we describe the Bayesian inference. Section 5 lists the baselines and their metrics and proceeds on an extended evaluation. Section 6 offers a summary of the results and a higher-level discussion. We close the paper with a conclusion in Section 7.

2. PREVIOUS WORK

Estimating availability measures in Optical Transport Networks (OTNs) has been studied to a considerable extent. Following, we mention a representative subset of the most relevant publications to the current work.

Most research takes link availabilities for granted and proceeds to calculate network-wide availability measures. The motivation is to conduct availability-aware network design and dimensioning or network reconfiguration. For example, [8] develops a theoretical and computational approach for network availability analysis in mesh networks, given that single-span failures are restorable. [9] assumes the Mean Time To Failure (MTTF) and Mean Time To Repair (MTTR) to be independent memoryless processes with known mean values and proposes an availability-aware service provisioning strategy under different protection modes. [10] introduces a framework for OTN design, focusing on end-to-end reliability and cost. An optical link is modeled to consist of several components like Reconfigurable Optical Add-Drop Multiplexers (ROADMs), amplifiers, transponders, and regenerators. The authors argue that the end-to-end availability of a few reference connections cannot represent the overall reliability well. Therefore, they calculate the end-to-end availability for all demand pairs while considering various protection schemes like dedicated protection, shared protection, no protection, or a mixture of those. [11] assumes the same link availabilities as [10] and proceeds to design a communication network with high availability requirements of up to 99.9999 %. The authors are interested in the *all terminal reliability* and strive to redesign an existing network using greedy iterative techniques and Integer Linear Programming (ILP). [12] analyzes the influence of different topologies and their properties under different protection scenarios. The link availability values are still considered known and taken from [13].

[13] proposes an availability model covering various types of IP-optical network equipment, including fibers, and utilizes a trio of scalar values to represent different availability measures. [14] presents a method to estimate link-dependent parameters of OTNs, such as the number of amplifiers and modulation schemes, using limited information, namely the statistical distribution model of the link lengths. This method improves over previous work that used the average link length.

Several publications are focused on estimating the interval availability due to its importance as an SLA metric. [15] develops a Markov model to estimate the interval availability for protected connections assuming Poisson node and link failures and repairs. [16] introduces a way to limit network operators' expenses by increasing the MTTR to reduce the related repair costs as a trade-off with potential SLA penalties. This is achieved by modeling the interval availability and binding it to the expected penalty compensations through a financial model. In subsequent work, [17] uses a stochastic approach to estimate the probability of SLA compliance for a service, such as to decrease the availability overfulfillment in favor of releasing resources reserved for protection. Interestingly, [18] argues that in some cases, it is more appropriate to focus on the continuity of a connection instead of the overall availability.

When anticipating multiple network faults, a probabilistic Shared Risk Link Group (SRLG) framework is proposed in [19] to model correlated failures and develop routing mechanisms. Likewise, [20] proposes a general stochastic model to account for geographically correlated link failures caused by natural disasters, along with an efficient method to calculate the joint failure distribution of specific links.

A limited number of papers have conducted a Bayesian analysis on similar matters. [21] proposes a probabilistic failure localization algorithm based on Bayesian networks for Root Cause Analysis (RCA). [22] also uses a Bayesian network encoded with expert knowledge to propose a fault propagation model for the GPON-FTTH (Gigabit Passive Optical Network-Fiber To The Home) access network.

Our work differs because it focuses on estimating the link availability, which is commonly presupposed. Namely, we invent a Bayesian model that combines expert knowledge and data observations to produce a more accurate one-shot estimate. This work was first presented in [23] and is now being extended in various ways. To begin with, an alternative proof-of-concept version of the model is given, allowing for more granular control over priors on a per-link basis. In addition, we present an investigation of the influence of priors on the posterior model. We offer a more robust evaluation across several topologies and more advanced baselines. Lastly, we deploy the model in a simulated scenario where services require different availability levels and report on the findings.

3. MODELING

This section provides a fundamental overview of Bayesian modeling in subsection A, introduces the link availability model in B, and conducts a prior analysis in C. For the reporting of the model, we adhere to guidelines supplied by [24]. Interested readers are also referred to [23] for a broader perspective.

A. Introduction to Bayesian Modeling

Bayesian modeling is a type of statistical modeling that relies on Bayes' theorem.

$$P(\theta|y) = \frac{P(y|\theta) \cdot P(\theta)}{P(y)} \tag{1}$$

y is the evidence (or data) observed during an experiment, and θ is the parameter we want to estimate. $P(\theta)$ is the *prior probability* expressing the initial belief (or expert knowledge). $P(y|\theta)$ gives the probability of the data given the parameters, also known as *likelihood*. After receiving the data y, we can infer the *posterior probability* $P(\theta|y)$, which is also the final estimate

of the Random Variable (RV) θ of interest. The denominator $P(y)$, or *marginal likelihood*, is frequently hard to calculate and can often be ignored, leading to the posterior being proportional to the likelihood times the prior.

$$P(\theta|y) \propto P(y|\theta) \cdot P(\theta) \tag{2}$$

$P(y)$ can be ignored because it is a constant and serves only as a scaler to standardize the posterior probability density function (PDF) such that it integrates to 1 and adheres to the probability axioms. Hence, we can still draw meaningful conclusions from the unstandardized posterior by relatively comparing the probability outcomes with each other.

There are numerous methods to infer the posterior distribution. As this was initially done analytically, statisticians developed the notion of *conjugate priors*, according to which certain distribution families for the prior and the likelihood give a closed-form solution of the posterior distribution. However, this is restricting as it can only be used for models with the exact distribution families. Due to advancements in computing and *probabilistic programming* [25], inferring the posterior using Markov Chain Monte Carlo (MCMC) sampling techniques has gained significant attention. MCMC techniques do not explicitly calculate the posterior but instead generate a chain of values that are representative samples of the posterior distribution. Consequently, they are generic and can be used more broadly. Certain metrics exist to evaluate the quality of the samples [26], such as the effective sample size (ESS), which measures the number of independent samples in the chain. Another is $\hat{R}$ (*R hat*), which is an indicator of the convergence of the Markov chain to the target distribution. The ESS should be reasonably large, although the number of generated samples clearly caps it. When the interest lies in high-order statistics of the posterior, the ESS should accordingly be high. The Markov chain samples often begin with a high correlation between many consecutive samples, leading to biased sampling. As sampling continues, the Markov chain [27] theoretically converges to the stationary distribution, which is also the posterior distribution. The $\hat{R}$ tests this convergence and should approach the value 1.00 from above.

B. Link Availability Bayesian Model

We will lay the foundations in B.1 and later introduce the model in B.2.

B.1. Network Availability

Since *availability* is an overloaded term, we first clarify its use in this paper. Let $G(V, E)$ be the graph of a network with a set of nodes V and a set of undirected edges (i.e., links) E with link length $l_e \ \forall \ e \in E$. We are interested in estimating the *steady-state availability* (from now on, also *true availability*) of every link a_e. The steady-state availability describes the availability of a system in the long run and is given by

$$a_e = \frac{f_e}{f_e + r_e} \tag{3}$$

with f_e being the MTTF and r_e the MTTR of the link. These values can be observed in a network through the time series of uptimes and downtimes. Assuming the measurement period spans from t_o to t_s, let $U_e = \{u_e^1, u_e^2, ...\}$ be the set of experienced uptimes of link e in chronological order, and $D_e = \{d_e^1, d_e^2, ...\}$ be the set of experienced downtimes between them, as shown in

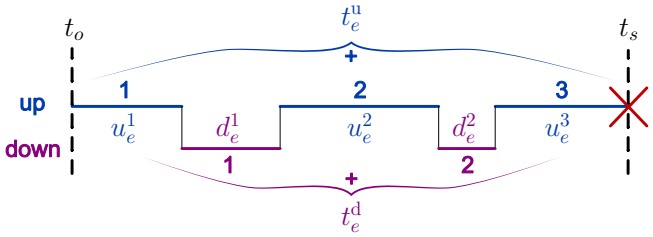

Fig. 1. Uptimes, downtimes, and state representation of link e.

Fig. 1. Let U be the superset containing all $U_e \ \forall \ e \in E$ and D be the superset containing all $D_e \ \forall \ e \in E$.

The *empirical interval availability* (from now on, just *interval availability*) is the availability of a system experienced in a time window and is given by

$$a_e^{\mathrm{int}} = \frac{t_e^{\mathrm{u}}}{t_e^{\mathrm{u}} + t_e^{\mathrm{d}}} \tag{4}$$

t_e^{u} and t_e^{d} are the total uptime and downtime of link e respectively.

$$t_e^{\mathrm{u}} = \sum U_e = \sum_i u_e^i \tag{5}$$

$$t_e^{\mathrm{d}} = \sum D_e = \sum_i d_e^i \tag{6}$$

The interval availability approximates the steady-state availability as $\Delta t = t_s - t_o \to \infty$. Despite that, the interval availability has long been used as a substitute for the steady-state availability [15]. This is a fair estimation only when the measurement period is long enough since link failures are scarce.

Another flavor of empirical availability for a time interval can be calculated using Eq. (3) with the empirical MTTF f_e^{int} and MTTR r_e^{int} instead. However, the last measurement of either U_e or U_d must not be included in the calculation since the measurement was interrupted, and it is not a representative sample of the underlying process. This does not hold for the first measurement, as we assume that our measurements start with the network equipment provisioned in a working (up) state. For example, in Fig. 1, only u_e^3 must be ignored as it is interrupted due to the end of the measurement period. Given U_e' and D_e' notate these meta-processed data sets where the last disrupted measurement has been discarded, we can formulate Eq. (7) and Eq. (8).

$$f_e^{\mathrm{int}} = \frac{\sum U_e'}{|U_e'|} \tag{7}$$

$$r_e^{\mathrm{int}} = \frac{\sum D_e'}{|D_e'|} \tag{8}$$

The $|\cdot|$ denotes the cardinality of the enclosed set. U', D' are similarly the supersets of U_e' and D_e' respectively $\forall \ e \in E$.

The target of this work is to develop an accurate estimate for MTTF $\hat{f}_e$ and MTTR $\hat{r}_e$, such that using Eq. (3), we derive an availability estimate $\hat{a}_e$ for each link.

$$\hat{a}_e = \frac{\hat{f}_e}{\hat{f}_e + \hat{r}_e} \ \forall \ e \in E \tag{9}$$

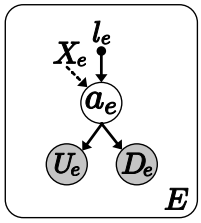

Fig. 2. Bayesian network for links $e \in E$

From link availabilities, we can calculate path availabilities. Given that a path p covers the links E_p, the path availability is

$$a_p = \prod_{e \in E_p} a_e \tag{10}$$

The same formula holds for the estimated path availability $\hat{a}_p$ but not for the empirical interval path availability a_p^{int}. The a_p^{int} is calculated using Eq. (4) but with the overall uptimes t_p^u and downtimes t_p^d across the complete path p instead.

B.2. Bayesian Model

A Bayesian model can be represented using a Directed Acyclic Graph (DAG), known as a Bayesian network, to illustrate the relationships between the RVs [28]. Fig. 2 illustrates the core of the model. The directed links express that the distribution of the target node is conditioned on the source node. We denote observed variables by shading the corresponding nodes; e.g., the availability link node a_e is not directly observed. Deterministic parameters, like link length l_e, are denoted by smaller solid circles. Finally, X_e with the dashed arrow signifies any combination of unknown variables that might influence the overall availability, but we ignore or do not know. We assume such influence is independent per link, and we trust the data to represent them in the model implicitly. In Fig. 2, the edges are drawn such that they expose the causality, i.e., some unknown entities and the link length influence the availability of a link, which in turn influences the data generation of uptimes and downtimes.

Embracing the Bayesian framework means $\hat{f}_e$ and $\hat{r}_e$ are now considered to be RVs, and our estimates of f_e and r_e will have a stochastic nature. Bayesian models are always generative and capable of imitating the data generation process. We begin by addressing the generation of the measurements. We model the uptimes and downtimes to follow an exponential distribution parametrized with $\hat{f}_e$ and $\hat{r}_e$ respectively.

$$u_e^i \sim \text{Exponential}(\hat{f}_e) \; \forall \, i \in |U_e'| \; \forall \, e \in E \tag{11}$$

$$d_e^i \sim \text{Exponential}(\hat{r}_e) \; \forall \, i \in |D_e'| \; \forall \, e \in E \tag{12}$$

$\sim$ signifies that the Left-Hand Side (LHS) follows the distribution on the Right-Hand Side (RHS). The LHS can be either data, like Eq. (11) and (12), or an RV.

We set the repair times to follow the same distribution for all links $\hat{r} = \hat{r}_e$, which is a common assumption among several studies like [9, 15, 17, 29]. A typical prior distribution for an exponential likelihood model is the inverse-gamma distribution, as they form a conjugate prior pair.

$$\hat{f}_e \sim \text{Inverse-Gamma}(\alpha_e^{\text{f}}, \beta_e^{\text{f}}) \; \forall \, e \in E \tag{13}$$

$$\hat{r} \sim \text{Inverse-Gamma}(\alpha^{\text{r}}, \beta^{\text{r}}) \tag{14}$$

To exploit the interrupted last measurement $u_e^{|U_e|}$ or $d_e^{|D_e|}$, we will use the t_e^u and t_e^d, to which they are implicitly included. An RV that counts events following an exponential inter-arrival time distribution is a Poisson distribution with rate parameter λ. The average rate of occurrence λ is equal to the total time of the experiment divided by the mean of the underlying exponential distribution. With that, we model the counts of uptime (i.e., failures) and downtime (i.e., repairs) events.

$$|U_e'| \sim \text{Poisson}\left(\frac{t_e^{\text{u}}}{\hat{f}_e}\right) \; \forall \, e \in E \tag{15}$$

$$\sum_e |D_e'| \sim \text{Poisson}\left(\frac{\sum_e t_e^{\text{d}}}{\hat{r}}\right) \tag{16}$$

Eq. (16) is aggregated across all links because the sum of independent Poisson distributions $X_i \sim \text{Poisson}(\lambda_i)$ is still Poisson $\sum_i X_i \sim \text{Poisson}(\sum_i \lambda_i)$. The time duration of the experiment for the failure events is the total uptime t_e^u, and for the repair events, it is the total downtime t_e^d.

To model the influence of the link length on the availability, we consult [30], which suggests that the relationship of failures over link length follows a reciprocal function, out of which we derive the most generic form.

$$\tilde{f}(l_e; k, s, h) = k + \frac{s}{l_e - h} \tag{17}$$

$\tilde{f}(l_e)$ is a continuous function estimate of f_e. We bind $\tilde{f}$ to the link estimations $\hat{f}_e$ by conditioning the first one to the second, i.e., $\Pr(\hat{f}_e) = \Pr(\alpha_e^{\text{f}}, \beta_e^{\text{f}} | \tilde{f})$, where $\Pr()$ denotes the probability. We achieve this by letting $\tilde{f}$ participate in the calculation of the inverse-gamma parameters in Eq. (13).

$$\beta_e^{\text{f}} = \tilde{f}(l_e; k, s, h) \cdot (\alpha_e^{\text{f}} - 1) \tag{18}$$

Eq. (18) is such that given a shape parameter α_e^{f}, the scale parameter β_e^{f} of the inverse-gamma distribution is calculated so that the mean value of the distribution in Eq. (13) is $\tilde{f}(l_e; k, s, h)$. We note that Eq. (17) will be forced to be the mean of Eq. (13) as a prior, but depending on the data during inference, it might deviate. The technique presented here is known as hierarchical or multilevel modeling, according to which a layer of RVs (e.g., $\hat{f}_e$) depends on another RV (e.g., $\tilde{f}$). Hierarchical modeling is a powerful Bayesian technique that enables grouping observed data and analyzing nested stochastic processes.

Now we specify the parameters k, s, and h of Eq. (17). We set $h = 0$, so that $\lim_{l_e \to 0} \tilde{f}(l_e) = \infty$ and to exclude potential negative values of $\tilde{f}$. We want to keep both k and s positive with a long tail on the right side, so we model them using a gamma distribution.

$$k \sim \text{Gamma}(k_\alpha, k_\theta) \tag{19}$$

$$s \sim \text{Gamma}(s_\alpha, s_\theta) \tag{20}$$

Since k, s are RVs, so are $\tilde{f}$ (through Eq. (17)) and β_e^{f} (through Eq. (18)).

Eq. (11–20) constitute the Bayesian model as also developed in [23] (from here on, just *base model*). In the following subsection, C, we will determine the priors for $\alpha_e^{\text{f}}, \alpha^r, \beta^r, k_\alpha, k_\theta, s_\alpha, s_\theta$ and analyze them. The priors are a way for the operator to convey expert knowledge in the model. Unfortunately, the base model does not support different MTTF priors per link. Such a thing

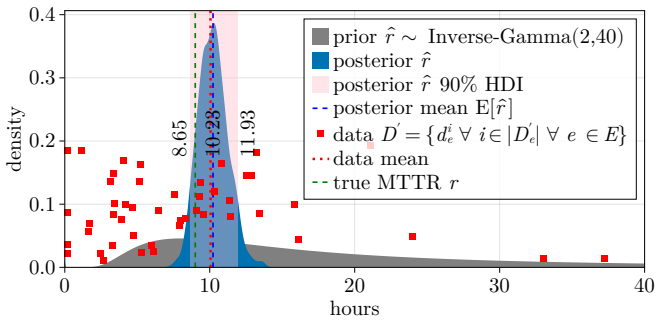

Fig. 3. Prior and posterior distributions for estimated MTTR $\hat{r}$.

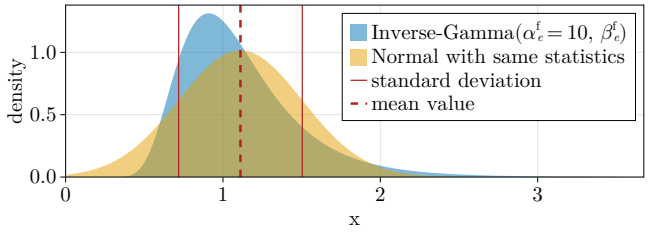

Fig. 4. Comparison of the inverse-gamma scale parameter prior with a normal distribution of similar mean and standard deviation values.

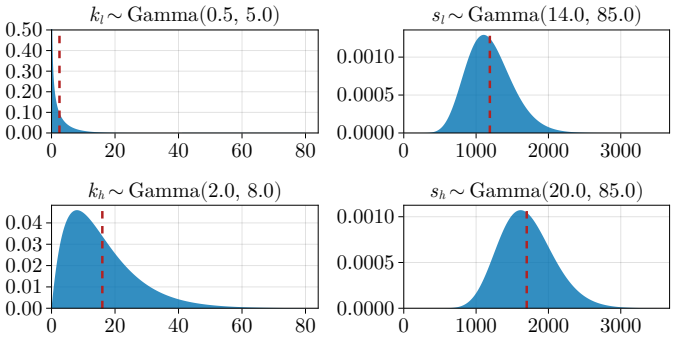

Fig. 5. The PDFs of the different priors for k and s. The dashed red line is the mean value of the distribution.

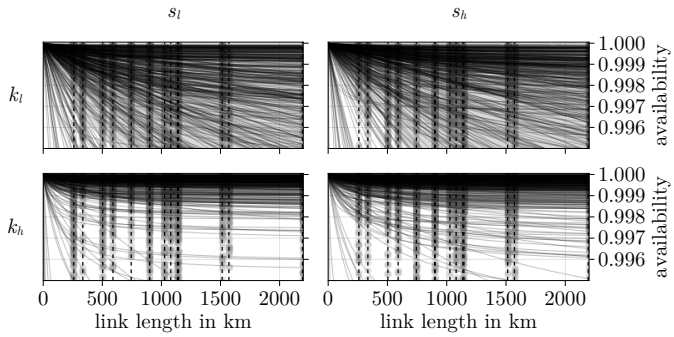

Fig. 6. Prior predictive simulation for different k and s priors.

is common for multilevel models, and it is usually treated as an advantage because it relieves us from the burden of defining priors for every element in a group. However, since network operators might have information regarding certain links behaving as outliers (e.g., due to older equipment), we provide a proof-of-concept version of the base model with such support. First, we define a set of positive RVs that will directly operate on the $\hat{f}_e$.

$$b_e \sim \text{Exponential}(c_e) \;\forall\; e \in E \tag{21}$$

Furthermore, we substitute Eq. (13) with the following.

$$\hat{f}_e \sim \text{Inverse-Gamma}(\alpha^f_e, \beta^f_e) \cdot b_e \;\forall\; e \in E \tag{22}$$

The extended model is given by Eq. (11), (12) and (14–22) (from here on, just *granular model*). The network operator has more granular control over each link prior by accordingly setting the c_e parameter. If such information is missing. c_e is set to 1.

C. Prior Analysis

During this study, we used weakly informative priors, i.e., the priors have been chosen so that they do not bias the inference but still introduce meaningful constraints. The model was designed so that negative or availability values above 1.0 are not supported. Moreover, the priors strongly encourage availability values above 0.99, which we characterize as a broad region. If considerations are different for an application, the priors must change accordingly. Minor modifications of the priors will lead to similar results as we confirm in 5C. Thus, the specific priors proposed are not a necessity to the model and are interchangeable. The priors are tuned for up/downtimes measured in hours and link distances in kilometers (km).

We set the priors $\alpha_r = 2$ and $\beta^r = 40$ for Eq. (14). Fig. 3 shows the prior for $\hat{r}$ in gray. The Inverse-Gamma(2, 40) does not encourage MTTRs close to 0 hours nor excludes very high values because of the long tail.

The shape parameter of Eq. (13) is set to $\alpha^f_e = 10$. The reason for giving α^f_e a scalar prior instead of β^f_e is that it is much easier to find a general prior for the shape rather than the scale. The shape parameter has been chosen such that the inverse-gamma resembles a normal distribution but with some positive skewness (see Fig. 4). Due to hierarchical modeling, we do not need to specify a prior for β^f_e explicitly.

Finally, we must decide the priors for $k_\alpha, k_\theta, s_\alpha, s_\theta$. As this situation is not straightforward, we create two priors for k and s of Eq.(19) and (20), respectively, and analyze their effects. For each case, we define a prior with low $_l$ and high $_h$ values. Fig. 5 shows the resulting gamma distributions k_l and k_h for k and s_l and s_h for s. As it is hard to understand the implications of these RVs, we take advantage of the generative nature of Bayesian modeling and sample the availabilities $\hat{a}_e$ and $\tilde{a}$ using just the priors and no data (commonly known as *prior predictive simulation* [26]). This way, we can see what our prior choices imply about meaningful and well-understood quantities. $\tilde{a}$ is derived from $\tilde{f}$, and it is similarly a continuous function of the link length. Fig. 6 shows the results of the prior predictive simulation. There are 4 total combinations between k_l, k_h and s_l, s_h. The vertical dashed lines are drawn where network links are supposed to be. Each case demonstrates the sampled prior availabilities $\tilde{a}$ and $\hat{a}_e$ with gray lines and scattered gray circles correspondingly. Every gray line is a sample of $\tilde{a}$ representing a plausible relationship. The numerous widely spread lines indicate that the underlying priors hardly exclude any possibilities. The gray circles, which are samples of $\hat{a}_e$, are similarly broad. We note that k_h generally encourages higher availabilities, and k_l may not cover well some very high availabilities near 1.000 00. The effect of s is looser,

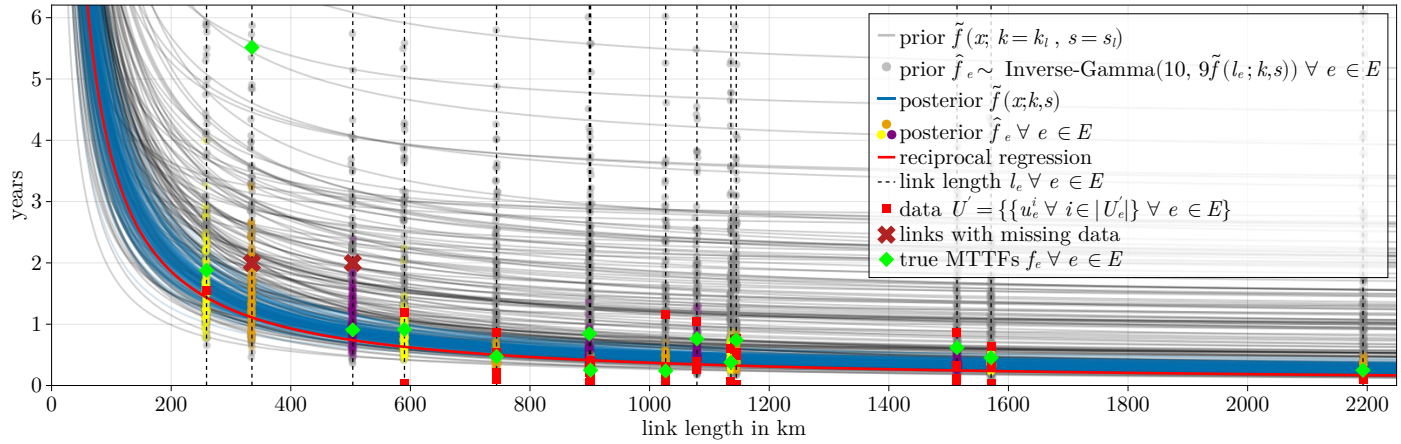

Fig. 7. Prior and posterior distributions for estimated MTTF $\hat{f}_e$ and $\tilde{f}$.

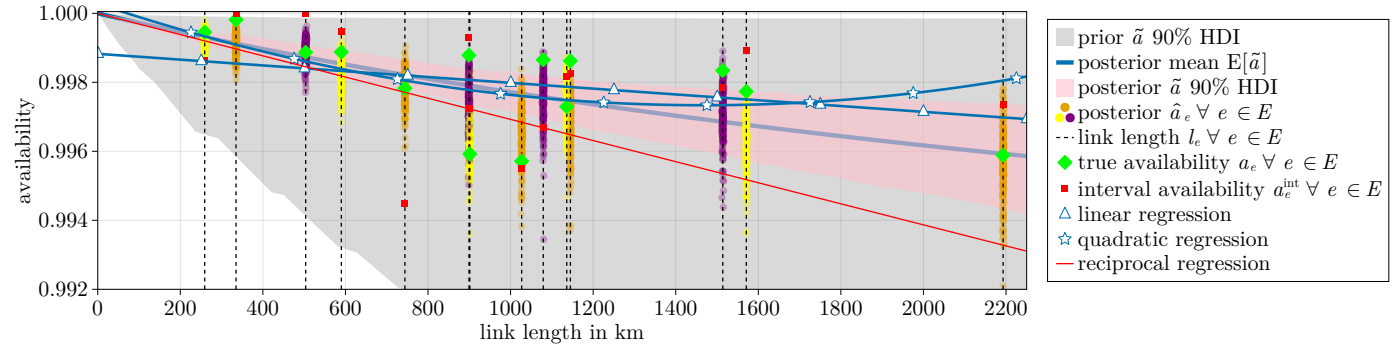

Fig. 8. Estimated availabilities $\hat{a}_e$ and $\tilde{a}$.

but we generally see that s_h generates mildly more sporadic samples, meaning that the estimates might have higher variance. During prior sensitivity analysis in 5C, we find that such reasonably small changes in the priors do not significantly influence the posterior distributions. This tolerance against such prior adjustments renders the model more reliable and stable.

4. INFERENCE

In this section, we go through the posterior inference for both the base and the granular model. Along the way, we introduce alternatives to estimate the link and path availabilities, which we use as baselines in Section 5. We use the No-U-Turn Sampler (NUTS) [31], a state-of-the-art algorithm belonging to the MCMC algorithmic family. We preprocess the data, i.e., normalize all uptimes U by dividing them with their standard deviation. We run chains with 3000 samples and find that the ESS and $\hat{R}$ metrics are within acceptable ranges. This section uses the Abilene topology of Fig. 11. The simulation setup will be described rigorously in 5B.

A. Base Model

Given the red square downtime data points D' in Fig. 3 and the gray prior, we infer the blue posterior estimate $\hat{r}$. The posterior is drawn using a Gaussian kernel estimate from the posterior samples. Since the inference gives a distribution instead of scalar values, we can assess uncertainty measures and calculate several statistics. We can also narrow it down to a scalar value by, e.g., taking the expected value of the posterior samples. A popular measure is the Highest Density Interval (HDI), defined as the narrowest interval containing the specified probability mass [26]. The HDI is the interval that best represents the region more consistent with the data. In Fig. 3, we calculate the 90 % HDI, yielding that the model is 90 % confident that the true MTTR r lies within $(8.65, 11.93)$. In this simulation, the data mean coincides more or less with the posterior mean, and the true MTTR r is slightly lower at 9 hours.

The uptimes U' are shown in red squares in Fig. 7. These, together with the priors (gray lines and circles), lead to the posteriors $\tilde{f}$ and $\hat{f}_e$. The blue lines denote the posterior for $\tilde{f}$ by updating the k and s parameters of the continuous reciprocal relationship Eq. (17). The colorful scattered circles are the posteriors of $\hat{f}_e$, discrete for each link. Some links have never failed and thus have no U_e' data. This is noted with the red cross positioned at the time of measurement completion (e.g., 2 years). The true MTTFs are depicted with the green rhombuses. The posterior distributions have notably narrowed down compared to the priors.

Given the assumption of a reciprocal relationship between the link length and the uptimes, one might be tempted to fit a reciprocal line on the data. The regression can be achieved by finding $k_r, s_r \in R^+$ such that the least squares error between $f_r(x; k_r, s_r) = k_r + s_r/x$ and the data points is minimized. Since this boils down to a non-linear problem, we set $z = 1/x$ and fit the linear $f_r(z; k_r, s_r) = k_r + s_r \cdot z$ instead. The resulting curve is illustrated with the red continuous line in Fig. 7. There are cases where s_r is found to be 0.0 , which produces a straight line.

Working with the complete posterior distributions enables

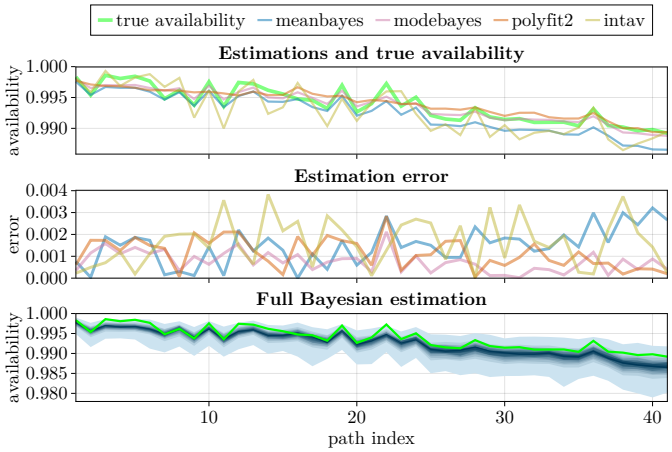

Fig. 9. Path availability estimations.

us to construct new RVs and propagate the uncertainty. We achieve this by doing operations using the posterior samples. For example, we use Eq. (9) for all 3000 samples of the chain to get a representation of the posterior link availability $\hat{a}_e$. We can similarly calculate HDIs and the posterior mean, as shown in Fig. 8. The red squares denote the interval availability, which is 1.0 for links that never failed. The red line is calculated with Eq. (9) using the reciprocal regression for uptimes and the data mean for downtimes. We also use the interval availabilities to conduct a first (linear) and second-order (quadratic) polynomial regression, as shown with the blue lines marked with triangles and stars, respectively. We note that, in the specific simulation, even if the quadratic regression might fit the data well, it misses on representing the actual relationship as it is visibly convex, and thus, the availabilities will grow for longer links, which is unrealistic. Moreover, in both polynomial regressions, the dependent variable must be clipped between 0 and 1.0

We similarly calculate the availability estimates for all shortest paths using Eq. (10) and the end-to-end uptimes and downtimes for the interval availabilities. Fig. 9 shows the estimates for the posterior mean (from now on, also *meanbayes*), the posterior mode (from now on, also *modebayes*), the quadratic regression (from now on, also *polyfit2*), and the interval availabilities (from now on, also *intav*). The posterior mode, also known as Maximum a Posteriori (MAP), is the value with the highest posterior probability. Modebayes is calculated by finding the MAP for each link and then using Eq. (10). On the top of Fig. 9, the estimations are shown, which should follow the green true availability as closely as possible. In the middle, the absolute errors to the true availability are shown. The path indices are sorted with the fewest hops and shortest overall length first. In this instance, modebayes and polyfit2 maintain the lowest error for longer paths, meanbayes error steadily increases as the paths get longer, and intav has high variance. Meanbayes diverges with higher hop counts because uncertainty accumulates, and the long-tailed distributions we used together with the capped availability from above incur significant skewness that shifts the mean value considerably. At the bottom of Fig. 9, we remind the readers that meanbayes and modebayes are only point estimates of a full posterior distribution, which is depicted here from a 90 % HDI with lighter shades to a 10 % HDI with darker shades and increments of 10 %. We notice how the uncertainty and skewness grow for longer paths. This is expected since paths with more hops involve operations with more RVs.

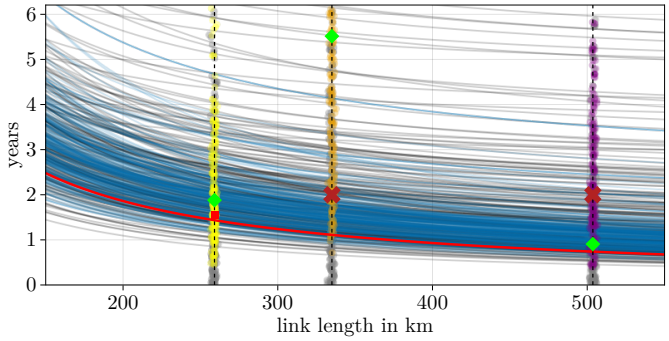

Fig. 10. Granular model $\hat{f}_e$ and $\tilde{f}$. Compare with Fig. 7.

B. Granular Model

Since the inference for the granular model is similar, we only focus on the particular usability of the model. We notice from Fig. 7 that the true MTTF f_e from the second shortest link e_2 with around 330 km is very high, almost 5.5 years. This might be because the link involves new equipment or is well-sealed and protected. The base model could never properly estimate this value with only 2 years of simulation. However, the operator can provide expert knowledge using the granular model and the c_e priors. Following, we supplied the model with $c_e = 1 \; \forall \; e \in E \setminus \{e_2\}$ and $c_{e_2} = 2$, meaning that we expect e_2 to have twice the MTTF than usual. Fig. 10 shows the advancement in estimating MTTF for e_2. Namely, comparing it to Fig. 7, the $\hat{f}_{e_2}$ samples are much closer to the true MTTF f_{e_2}. This leads to a 75 % error reduction in the final availability estimate of e_2 using the posterior mean. However, the current granular model also brings additional noise to the other links. For example, the MTTF estimates for the third link e_3 with almost 500 km are also significantly higher compared to the base model, even if $c_{e_3} = 1$. The reason is that multiplying Eq. (13) with an exponential distribution to get Eq. (22) may amplify the influence of the data (in this case, count of uptimes $|U'_{e_3}|$) or generally spread the estimations.

5. EVALUATION

In this section, we conduct an extended evaluation using the developed model and baselines. Subsection A introduces the metrics considered, B describes the simulation process, C reports the results, and D demonstrates a decision-making scenario. Comments and interpretations of the results are integrated in this section.

A. Methodologies and Baselines

Following, we provide a list of the estimation methodologies, some already introduced in Section 4. The metrics are attained by getting the mean absolute error across all links or shortest paths between the estimation and the true availability.

- *meanbayesgran*: posterior mean of the granular model with uninformative prior $c_e = 1$.
- *meanbayes*: posterior mean of the base model.
- *meanbayes66*: posterior mean of the base model with samples only contained in the 66 % HDI. The objective is to trim the long tail of the distribution such that they do not propagate during path estimation.

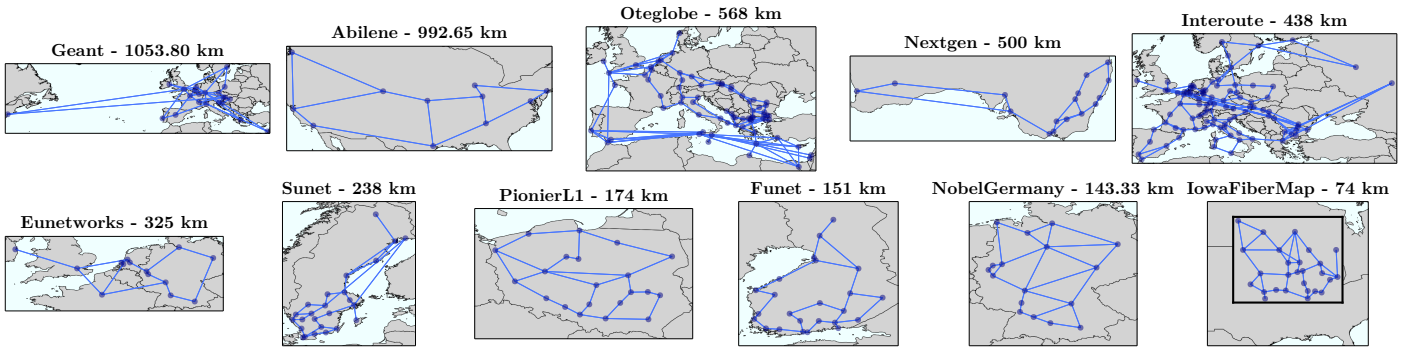

Fig. 11. The considered topologies and their average link length.

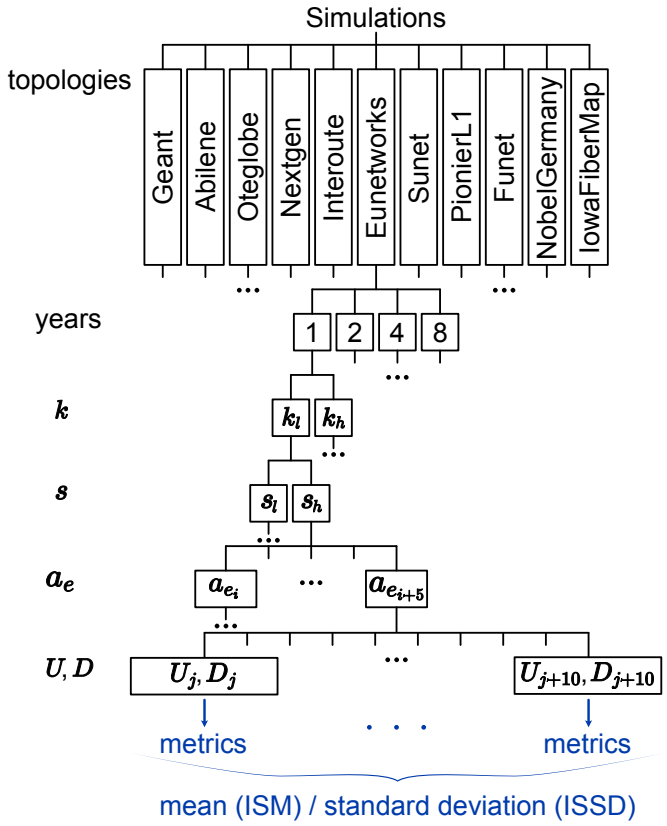

Fig. 12. Simulation parameters and setup.

- *meanbayes33*: posterior mean of the base model with samples only contained in the 33 % HDI.
- *modebayes*: MAP of the base model
- *polyfit1*: linear regression of the interval link availabilities.
- *polyfit2*: quadratic regression of the interval link availabilities.
- *recfit*: Eq. (9) using the reciprocal regression of uptimes for MTTF estimation and the mean downtimes for MTTR.
- *intav*: interval link or path availabilities computed along the measurement period.

All estimates, except intav, first apply to the links, and then we use Eq. (10) to deduce the path availabilities. We will identify meanbayes33 and meanbayes66 together as *meanbayestrimmed* and polyfit1 and polyfit2 as *polyfit*.

B. Simulations

We conducted simulations with different topologies and settings to get robust results. The considered topologies are extracted from [32] and [33] and shown in Fig. 11. We filtered the topologies of [33] down to only metro/core optical fiber meshed topologies. Let d_T be the standard deviation of the link distances in topology T. The overall simulation setup is depicted in Fig. 12. We evaluated each topology across measurement periods of 1, 2, 4, and 8 years. For each case, we analyzed the sensitivity between all four combinations of k and s priors shown in Fig. 6.

For every simulation parameter combination, we generated 5 different true link availabilities sets. These are calculated using the MTTF and MTTR model of [13]. Namely, the MTTR r is set to 9 hours for all links, and the MTTF f_e is calculated using the reciprocal relationship $f(l_e) = 628 \cdot 360 \cdot 24/l_e$. Before we evaluate the reciprocal relationship f to produce the uptimes, we add noise to each link's distance function argument l'_e using a positively truncated normal distribution with a standard deviation of 30 km and a mean value of the real link length l_e. Also, we randomly (uniform distribution) pick 10 % of the links in the network to behave like outliers by adding or subtracting (Bernoulli distribution with probability parameter $p = 0.5$) from the previously noise-inflicted distance the quantity $2 \cdot d_T$. Lastly, the arguments are capped between 30 km and 10 000 km, before passed into the MTTF reciprocal function f. The reason for introducing artificial noise is twofold. First, data generation becomes more realistic. Second, the noise differentiates the Bayesian model from the underlying data generation process. This makes the Bayesian inference and estimation more challenging, but the overall results more general and relevant to situations where the designed model is not identical to the system dynamics.

For each set of different MTTFs f_e and MTTR r derived (yielding a_e), we parametrize exponential distributions to generate uptimes U and downtimes D for the links. We use a distinct seed for 10 different cases to generate different data. For each simulation, the metrics of the subsection A are noted. We also calculate the inter-simulation mean (ISM) and inter-simulation standard deviation (ISSD) of the metrics given by the same true availabilities. A reliable metric should have a low standard deviation, meaning it should not be too sensitive to different data, given that the underlying generative process is the same. In total, for each simulation parameter combination of topologies, years, and k and s priors, we get $5 \cdot 10 = 50$ differently seeded simulations. Altogether, we conducted $11 \cdot 4 \cdot 2 \cdot 2 \cdot 5 \cdot 10 = 8800$

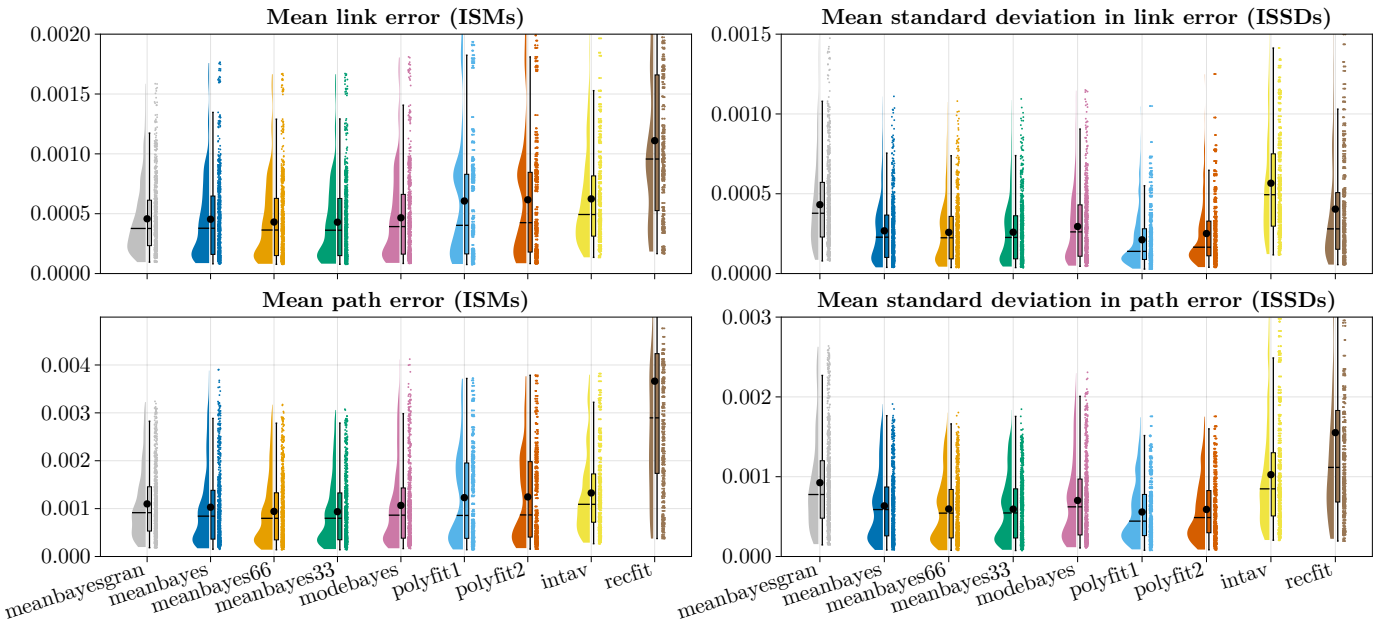

Fig. 13. ISMs and ISSDs across all simulations.

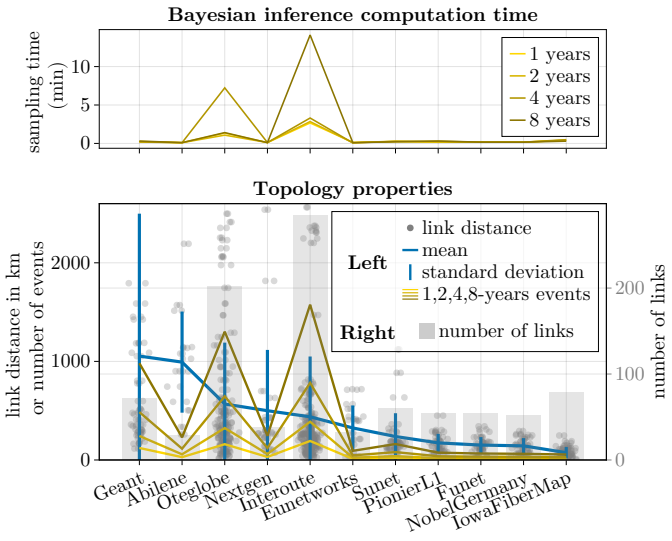

Fig. 14. Top, the average computation time of the Bayesian inference per topology and measurement period. Bottom, an analysis of the considered topologies.

simulations and produced $11 \cdot 4 \cdot 2 \cdot 2 \cdot 5 = 880$ ISMs and ISSDs.

A small analysis of the topology properties is shown at the bottom Fig. 14. In the plot, the topologies are sorted from the highest average link length to the lowest. All link distances l_e for each topology within the y-axis limits are shown with gray circles. The link distances' mean and standard deviation are shown in blue. In addition, the total average of failure events across all simulations is shown per topology for different measurement periods in yellow. The number of events is vital since each event conveys information to the estimators. The number of events is positively associated with the number of links and their length. The number of links per topology is shown on the right axis. The top of Fig. 14 gives the average sampling time of the Bayesian inference for the different topologies and measurement periods. The reported times include the warmup stage of the NUTS algorithm and the further generation of 3000 samples. For longer periods, more data gets produced, which in turn demands more resources. Additional links also lead to additional computations, since the number of RVs increases as well. The longest inference, which involved the Interoute topology and the 8-years timeframe, took, on average, 14 minutes. The experiments were executed using the *AMD EPYC 7282 16-Core* processor with 2.8 GHz clock speed in a single thread. Understandably, MCMC sampling using NUTS is an offline methodology and cannot be used to incrementally process new live data. Given that new data generate at a much slower rate than the inference completion, it is still realistic to trigger the Bayesian inference with the occurrence of every new event. Otherwise, the inference update should be strategically scheduled periodically when the output estimations are needed. The baseline intav is instantaneous to calculate, while polyfit and recfit might take a couple of seconds.

The Bayesian model was implemented using Turing [34], a general-purpose probabilistic programming language for Bayesian inference in Julia [35]. All the figures in this paper were generated using Makie [36]. Data processing is made easy using DataFrames [37].

C. Results and Discussion

As the results are numerous, we do not offer a numerical value for all of them. We communicate them using different plots and highlight the most important aspects in the text.

Fig. 13 shows all the ISMs and ISSDs of the metrics in a raincloud plot across all simulations for links and paths. The raincloud plots [38] visualize the raw data, the empirical probability density, a box plot, and key summary statistics such as median and mean. Regarding the boxplots, the crossbars

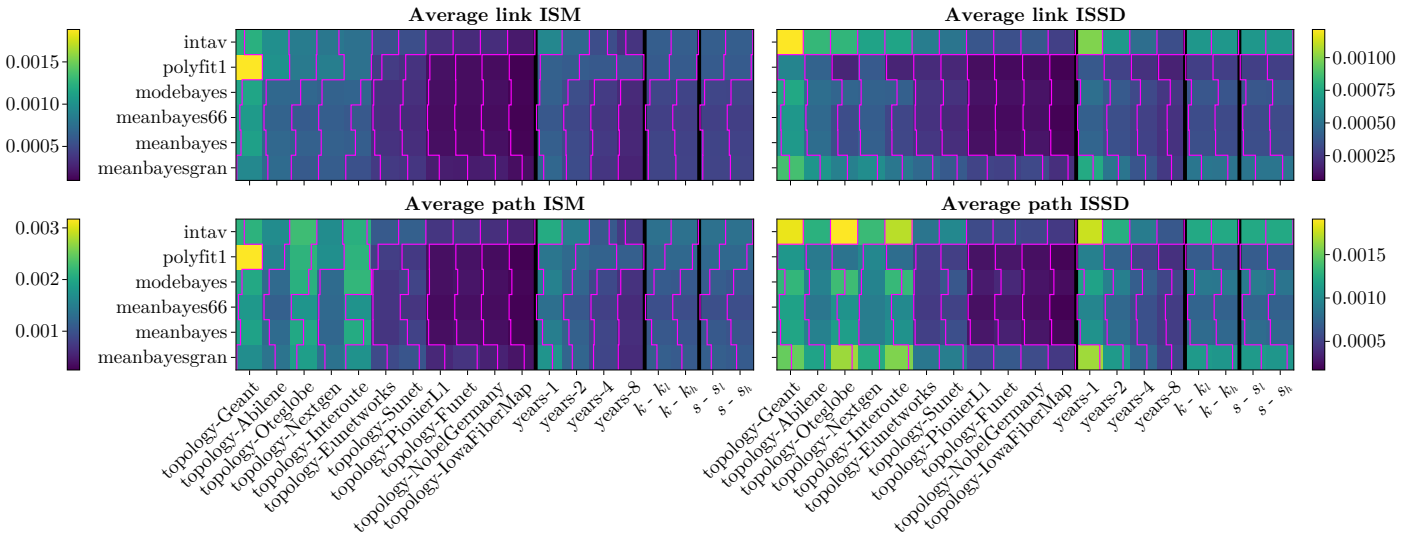

Fig. 15. Average of ISMs and ISSDs across different simulation settings.

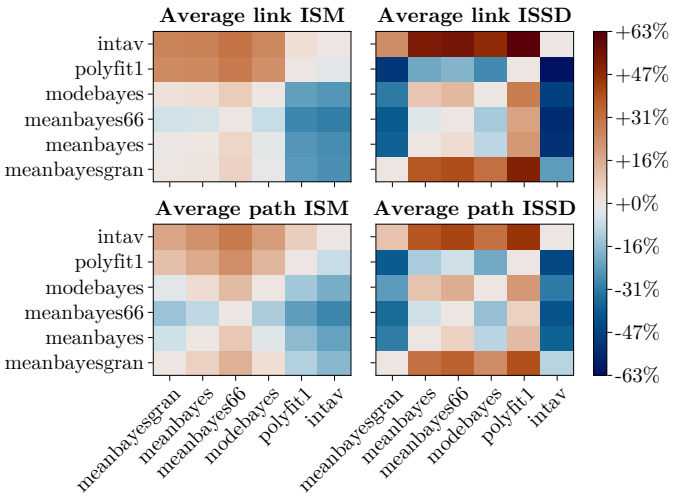

Fig. 16. Average of ISMs and ISSDs across all simulations.

span the Interquartile Range (IQR), with the midline marking the median, and the error bar whiskers continue to span up to 1.5 × IQR. The black circular marker is the average value. Looking at the link ISMs, we notice that the Bayesian methods generally have the lowest errors and the recfit the highest. The lowest error is given by bayesmean33 and beats polyfit1, the best non-Bayesian estimate, by an average of −24%. However, bayesmean66's mean value is very close to bayesmean33 with a negligible difference of +0.6%. Since their probability densities are very similar, we argue that they are indistinguishable. We can also support this with more traditional methods like the two-sample KolmogorovSmirnov test [39], according to which we calculate that the null hypothesis that the samples of bayesmean66 and bayesmean33 follow the same distribution is not rejected. The same holds for polyfit1 and polyfit2.

To better grasp the average behavior (black circles), we plotted their relative values in Fig. 16. Each heatmap is symmetrical along the antidiagonal and shows the relative difference of the y-axis to the x-axis metric. The blue colors signify lower error or standard deviation for the y-axis respective to the x-axis metric, while the red signifies higher. As we move from the link estimates to path estimates, on the bottom left of Fig. 13 and Fig. 16, the relative differences decrease, benefiting polyfit the most. This is because potentially wrong link estimates are averaged and can lead to unexpectedly better path estimates. In addition, Bayesian estimators, except for meanbayestrimmed, are also becoming relatively less successful because of the long tails that propagate and damage the point estimates of path availabilities. Although intav performs similarly to polyfit, it ends up slightly worse in the path estimations. Regarding the ISSDs, the weakness of intav is shown, as they present high variance. On this matter, polyfit is the most consistent method, with meanbayestrimmed coming second. The meanbayesgran metric, unfortunately, generally demonstrates considerably worse performance than meanbayes.

Fig. 13 and Fig. 16 serve to find the best estimation method, agnostic to the scenario. In Fig. 15, we isolate the influence of the different simulation settings and examine the results separately, e.g., for each topology and measurement period. The y-axis shows the metrics omitting recfit because of its poor performance. We also omit meanbayes33 and polyfit2 since they are statistically identical to meanbayes66 and polyfit1, respectively. The values depicted in the heatmap are the averages of ISMs and ISSDs, similar to Fig. 16 or the black circular markers of Fig. 13. However, the averages are not across all simulations but only for those that hold constant the parameter value depicted on the x-axis. For example, for *years-2*, we will get the average across all simulations run with a measurement period of 2 years. Black lines separate the simulation parameter types. The magenta stairs-plot lines are drawn per simulation parameter and evenly scale the heatmap column values from minimum on the left to maximum on the right. This serves so that small differences along a column can be better spotted.

One of the major points in this figure is that the estimates get better for smaller networks (in link length and number). This might appear counterintuitive since smaller links generate fewer failure events, and the models end up with less information. However, the link distances of the smaller networks are much more concentrated, which significantly facilitates the predictions. In addition, bigger networks may provide the models with more

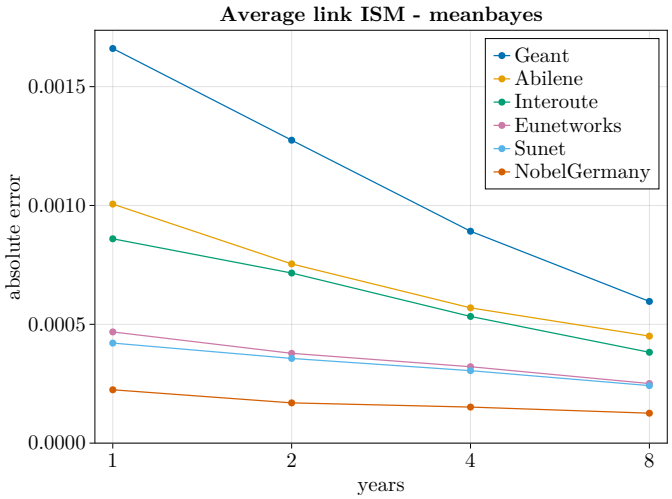

Fig. 17. Effect of growing measurement periods on several topologies.

data, but the stakes are higher since a much greater range of links and outliers must be covered. As a result, small and big networks are why some metrics demonstrate a bimodal distribution of results, which is also visible in Fig. 13 (e.g., polyfit ISMs). Moreover, for small networks, the benefit of Bayesian methodologies diminishes, and meanbayes66 is only marginally better than polyfit1. However, the Bayesian metrics are doing notably better for the big topologies. The main reason polyfit is suffering in big topologies is that it cannot independently adapt to each link's diverse behavior. On the contrary, intav treats each link separately and cannot consolidate the amount of information to identify interlink patterns. The Bayesian estimators achieve both by simultaneously treating each link individually while also respecting the overall trend. To prove that more data indisputably benefit the models, we notice that as simulations run longer (from *year-1* to *year-8*), the metrics improve, i.e., turn darker. An interesting phenomenon is that as data input grows, polyfit's performance remains unchanged, and consequently, even the interval availabilities start to beat it. Polyfit picks up quickly and is on par with meanbayes66 when data is scarce but is rapidly saturated and left behind when information is increased. Finally, the influence of the different priors of Fig. 6 is negligible, meaning that the model is robust, and small changes in priors will not matter.

To consider the different measurement periods and the topologies jointly, Fig. 17 is provided. The plot shows the effect only for the average link ISMs, the meanbayes estimator, and limited topologies to avoid clutter. Fig. 17 instructs us that as the measurement period grows, we can expect greater improvement in scenarios of poor performance.

D. Model Deployment and Decision-Making

In this part of the evaluation, we employ all the different metrics as decision variables for selecting routing and protection paths for various services. During the previous evaluations, we have only used point estimates (like mean and mode) to describe the full posterior. This section will showcase a scenario that uses the full posterior distribution for sophisticated decision-making. The simulations here were run with the same structure described in Fig. 12, but we will only showcase the average values across all simulations. Given a network topology, we generate 5 connection requests per node pair, and each is randomly assigned an availability requirement. Let a_{std} be the standard deviation of the true availabilities from all the shortest paths in a network. The availability requirements are generated with a normal distribution around the true shortest path availability with a standard deviation of $a_{std}/2$.

For the deployment, we use the algorithm developed in [17]. The algorithm receives as input the connection request with its availability requirement and a minimum compliance target. The compliance target is the desired availability compliance probability, which indicates how probable it is that a given connection implementation covers the service availability. Given its probabilistic nature, it can only be fully supported by probabilistic estimations like the Bayesian model. Non-probabilistic estimations will default to a compliance target of 100 %. Using one of the methods in subsection A, the algorithm estimates the availability of the shortest path and calculates the probability that the compliance target will be covered. In non-probabilistic estimates, compliance is achieved if the availability estimation is higher or equal to the availability requirement. For the full Bayesian estimates, the posterior distribution is examined to settle whether the compliance target is reached. If the compliance target is not covered, disjoint protection is added, and the compliance target is tested again. If it fails, the second shortest path is examined, and so on, until the compliance target is achieved or the algorithm quits after some iterations. In the case of paths with equal lengths, the selection order is random. We assume unlimited resources.

Fig. 18 depicts the outcomes of this algorithm. The bar plots show the averaged ISMs and the red error bars show the averaged ISSDs across all simulations. The *fullbayesgran* method utilizes the full posterior of the granular model, and *fullbayes* utilizes the full posterior of the base model. The *fullbayes66* only considers the trimmed posterior distribution enclosed in the 66 % HDI. The meaning of the different bar plots is explained.

- *failed*: the percentage of connections that failed.
- *false-fail*: the percentage of failed connections that should not fail because the true availabilities could handle the requirement.
- *protected*: the percentage of protected connections
- *false-protected*: the percentage of protected connections that should not be protected because the true availability of the selected path alone could handle the requirement.
- *false-coverage*: the percentage of connections wrongly deployed because the true connection availability does not cover the requirement.

The 100 % compliance probability case shows the Bayesian estimates have the lowest false-coverage. Remarkably, when the algorithm uses the full Bayesian posterior, there is almost 0 % false-coverage. This comes at the cost of the false-protected connections being increased. The reasoning is that the model cannot be 100 % sure that some connections can cover the requirement, so it enforces a protection path. Notably, the standard deviations for the full Bayesian posterior are also very low, which testify to outstanding reliability and stable performance. The modebayes, being a point estimate, can only be used for the 100 % compliance target. Although it introduces some false-coverage, it is lower than the polyfit and the interval availabilities. The polyfit and interval availabilities look very similar and exhibit among

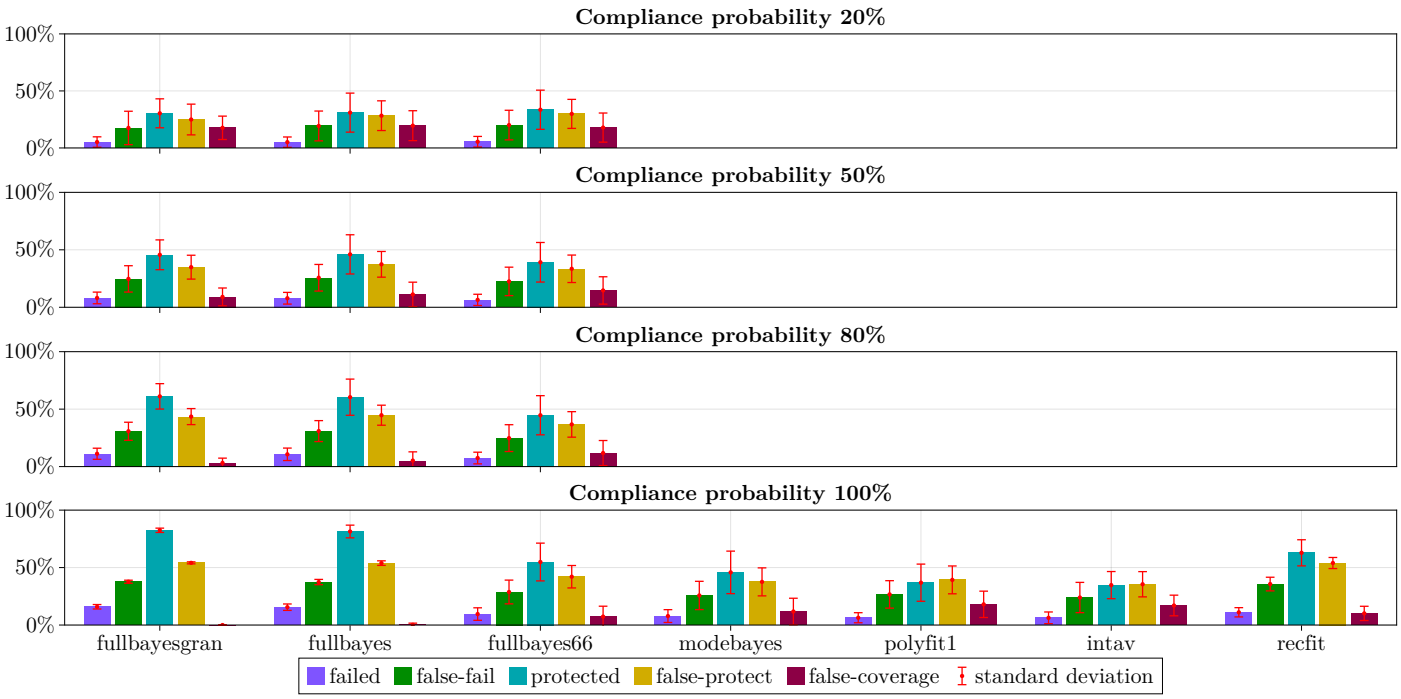

Fig. 18. Evaluation of the models deployed in a decision-making scenario.

the lowest false-fail and false-protect, similar to modebayes. We emphasize the difference between theoretical evaluations like subsection C, where polyfit and intav differ, and practical evaluations, where it might be hard to tell them apart. This gap may occur because the relative estimation difference between polyfit and intav is too low to trigger different routing decisions for the given availability requirements.

Interestingly, as the compliance target probability decreases, the false-fail and the false-protect also decrease, but the false-coverage increases. Ideally, false-fail, false-protect, and false-coverage should all be 0 %. Nevertheless, this is unrealistic to expect from any model. A model without the true estimations will inevitably end up with at least some of these quantities being positive. Even worse, non-probabilistic models will get trapped in a specific instance of all the trade-off possibilities. In contrast, probabilistic models enable the operator to choose among different situations conveniently. Finding the best trade-off depends on the priorities, and our Bayesian model can be used to navigate between different solutions. We also learn that the effects of fullbayes66 can be reproduced using fullbayes with a lower compliance probability of around 50-80 %. This is a sensible finding since both approaches rely on a similar portion of the posterior distribution.

6. DISCUSSION

This subsection recaps some of the results and remarks of the previous sections and offers a higher-level perspective on modeling.

A. Results Summary

The developed Bayesian model was found to attain the overall best performance when evaluating it among different topologies and measurement periods. Its performance beats the baselines like empirical interval availabilities, their polynomial fit, and reciprocal regression of uptimes. The reciprocal regression provides particular interest as a baseline since it uses the same information as the Bayesian model: the downtimes and the uptimes with the assumption that they follow a reciprocal relationship against the link length. More specifically, we developed several point estimates out of the Bayesian model, which we compared. These included the mean of the posterior distribution, the mode of the posterior distribution, and the mean of the posterior distribution after trimming it to an HDI. The posterior mean for the 66 % and 33 % HDIs behave similarly, pointing out that the posterior should be trimmed just enough to get rid of the long tail, which severely influences the mean point estimate of path availabilities. The mean of the posterior trimmed distributions achieved on average 24 % higher accuracy in estimating the true link availabilities compared to the best non-Bayesian baseline. This win is higher for bigger topologies because of the high link length variability and the numerous outliers and decreases for smaller topologies. Allowing longer measurement periods improves the model as more information becomes available. By conducting a prior sensitivity analysis, we confirm that the Bayesian model is robust and that small changes in the priors will not matter. The proof-of-concept granular Bayesian model is handy when the network operator possesses expert knowledge on a per-link basis, which can be shared through priors. A small isolated experiment led to a 75 % error reduction in a link availability estimation. However, the overall granular model performs worse when using uninformative priors than the base Bayesian model. Future efforts can focus on mitigating this effect of the granular model.

Limiting the distributions to an HDI has been a successful approach. However, when dealing with realistic scenarios and some stochasticity is involved, the full posterior distribution should be used. In particular, we found out that the trimmed posterior approach can be considered as a special case of dealing with the full posterior distribution. The probabilistic nature of

the Bayesian model can not only be used for building stochastic policies, where consecutive choices must be made but also for one-off decisions. For example, we showcased that our model can flexibly tune a decision-making algorithm, making routing and protection choices. Moreover, when needed, the developed model achieved remarkable stability and performance in satisfying the service requirements.

B. Modeling Remarks

For the development of the Bayesian model, several assumptions were made. For example, the downtimes were modeled using an exponential distribution. Although this is a common assumption, it is implausible that the underlying process indeed follows an exponential distribution. When modeling using an exponential distribution, we are actually interested in finding a continuous, non-negative distribution with a specified mean value. Searching through different distribution families, it turns out that the exponential distribution provides the maximum entropy on this matter [40], i.e., it is the widest and least informative distribution that holds onto the requirements. This constitutes an excellent argument for why we choose the exponential distribution. It is not because we think that the underlying process is an exponential distribution but because we are settling with focusing only on the mean value of the underlying process. Using Bayesian inference, we can calculate the posterior distribution, which can also be seen as entropy maximization [26, 41], meaning that the posterior has the least possible divergence from the priors such that it is consistent with the data received. With this in mind, the statistical choices in the model, which at times might seem overly simplified, can be excused.

A statistical model, being only an approximation to the truth, should be shaped to focus on addressing specific aspects of interest rather than replicating reality since that might lead to overfitting with highly complex models [42]. The current work provides a general equipment-agnostic model that only consumes the overall link uptimes and downtimes. The model can successfully personalize the behavior of each link so that, for example, unknown effects (like X_e from Fig. 2) can be accommodated through the data while respecting the overall reciprocal behavior between the uptimes and the link length. If more data is available, the current model can be modified to include several networking components in the form of nodes. If the equipment lifetime is also available, the time dimension may be integrated, e.g., using a bathtub failure rate curve [43]. As these require separate analyses and might lead to disruptive changes, such efforts are encouraged for future work.

The findings in this paper present convincing arguments for the use of Bayesian modeling in reliable optical networks. The study provides theoretical value backed up with diverse simulation setups. Future efforts should also focus on the deployment of the model in real-life scenarios. Since the data requirements (i.e., the link lengths, uptimes, and downtimes) are easily calculated, the setup can be tested across a wide range of networks. Although the simplicity of the model grants it generality, it can break if the data distributions are very different from the considered exponential. For example, if the true downtime distribution is bimodal, the Bayesian inference can fail miserably as it will continue to assume unimodality. A natural extension would be to stop modeling all the downtimes with the same distribution. It could be especially helpful to divide the link downtimes into terrestrial and submarine groups, leveraging hierarchical modeling similar to the uptimes. Finally, special care was taken to provide broad priors applicable to a wide range of use cases. Network operators can either adopt the values stated in this paper or adapt them (e.g., by narrowing them down) given their expertise.

7. CONCLUSION

This work breaks new ground by developing a Bayesian approach to model link availabilities in optical transport networks by consuming their overall uptimes and downtimes. The developed model delivers higher accuracy compared to baselines like empirical interval availability or polynomial and reciprocal regressions. The uncertainty of the model can be propagated to construct new estimators, e.g., for paths or end-to-end connections. We provided a proof-of-concept alternative to the developed model that enables the operator to convey more expert knowledge on a per-link basis. Although the alternative model can be convenient in certain situations, it generally performs poorer than the base model, which should be preferred if no link-based prior knowledge is available. We conducted a thorough evaluation of many different metrics and scenarios. We found that using an HDI trimmed posterior mean offers the best point estimate with the lowest error against the true availabilities. However, since Bayesian modeling outputs a probabilistic estimation, practical applications should endorse their stochastic nature and employ the full posterior distribution. Embracing the probabilistic nature enables flexibly tuning the underlying decision-making procedure and assessing uncertainty, leading to reliable algorithms and increased stability. The results suggest that this approach is a valuable tool for attaining accurate availability estimations and is appropriate for critical situations. The output availability estimations generated by this model can be utilized in conducting independent studies on network availability.

8. FUNDING

German BMBF (16KIS1312) in the framework of the CELTIC-NEXT EUREKA project AI-NET-ANTILLAS (C2019/3-3)

REFERENCES

1. S. S. Ahuja, S. Ramasubramanian, and M. Krunz, "Single-link failure detection in all-optical networks using monitoring cycles and paths," IEEE/ACM Transactions on Netw. **17**, 1080–1093 (2009).
2. A. Pattavina, S. D. Patre, and M. Martinelli, "Protection and restoration schemes in optical networks : A comprehensive survey," (2013).
3. H. Choi, S. Subramaniam, and H.-A. Choi, "On double-link failure recovery in wdm optical networks," in *Proceedings.Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies*, vol. 2 (2002), pp. 808–816 vol.2.
4. J. López, Y. Ye, V. López, F. Jimenez, R. Duque, P. M. Krummrich, F. Musumeci, M. Tornatore, and A. Pattavina, "Traffic and power-aware protection scheme in elastic optical networks," in *2012 15th International Telecommunications Network Strategy and Planning Symposium (NETWORKS)*, (2012), pp. 1–6.
5. E. Bouillet, D. Mitra, and K. Ramakrishnan, "The structure and management of service level agreements in networks," IEEE J. on Sel. Areas Commun. **20**, 691–699 (2002).
6. L. Breiman, "Statistical modeling: The two cultures (with comments and a rejoinder by the author)," Stat. Sci. **16**, 199 – 231 (2001).
7. R. van de Schoot, S. Depaoli, R. King, B. Kramer, K. Märtens, M. G. Tadesse, M. Vannucci, A. Gelman, D. Veen, J. Willemsen, and C. Yau, "Bayesian statistics and modelling," Nat. Rev. Methods Primers **1** (2021).
8. M. Clouqueur and W. Grover, "Availability analysis of span-restorable mesh networks," IEEE J. on Sel. Areas Commun. **20**, 810–821 (2002).

9. J. Zhang, K. Zhu, H. Zang, N. S. Matloff, and B. Mukherjee, "Availability-aware provisioning strategies for differentiated protection services in wavelength-convertible wdm mesh networks," IEEE/ACM Transactions on Netw. **15**, 1177–1190 (2007).
10. M. Mezhoudi and C.-H. K. Chu, "Integrating optical transport quality, availability, and cost through reliability-based optical network design," Bell Labs Tech. J. **11**, 91–104 (2006).
11. V. Kounev, M. Lévesque, D. Tipper, and T. Gomes, "Reliable communication networks for smart grid transmission systems," J. Netw. Syst. Manag. **24**, 629–652 (2016).
12. J. Segovia, E. Calle, and P. Vila, "Availability analysis of gmpls connections based on physical network topology," in *2008 International Conference on Optical Network Design and Modeling,* (2008), pp. 1–6.
13. S. Verbrugge, D. Colle, P. Demeester, R. Huelsermann, and M. Jaeger, "General availability model for multilayer transport networks," in *DRCN 2005). Proceedings.5th International Workshop on Design of Reliable Communication Networks, 2005.*, (2005), pp. 8 pp.–.
14. S. K. Routray, G. Sahin, J. R. F. da Rocha, and A. N. Pinto, "Estimation of link-dependent parameters in optical transport networks from statistical models," J. Opt. Commun. Netw. **6**, 601–609 (2014).
15. D. A. Mello, H. Waldman, and G. S. Quitério, "Interval availability estimation for protected connections in optical networks," Comput. Networks **55**, 193–204 (2011).
16. T. Enderle and A. Kirstädter, "Reducing network operators' expenses by adjusting the mttr," in *2021 31st International Telecommunication Networks and Applications Conference (ITNAC),* (2021), pp. 134–139.
17. T. Enderle, "Protected connection provisioning with low availability overfulfillment in meshed core networks," in *Photonic Networks; 23th ITG-Symposium,* (2022), pp. 1–8.
18. P. Choda, A. Mykkeltveit, B. E. Helvik, and A. Jajszczyk, "Continuity-based resilient communication," in *2009 7th International Workshop on Design of Reliable Communication Networks,* (2009), pp. 335–342.
19. H.-W. Lee, E. Modiano, and K. Lee, "Diverse routing in networks with probabilistic failures," IEEE/ACM Transactions on Netw. **18**, 1895–1907 (2010).
20. J. Tapolcai, B. Vass, Z. Heszberger, J. Bíró, D. Hay, F. A. Kuipers, and L. Rónyai, "A tractable stochastic model of correlated link failures caused by disasters," in *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications,* (2018), pp. 2105–2113.
21. M. Ruiz, F. Fresi, A. P. Vela, G. Meloni, N. Sambo, F. Cugini, L. Poti, L. Velasco, and P. Castoldi, "Service-triggered failure identification/localization through monitoring of multiple parameters," in *ECOC 2016; 42nd European Conference on Optical Communication,* (2016), pp. 1–3.
22. S. R. Tembo, S. Vaton, J.-L. Courant, and S. Gosselin, "A tutorial on the em algorithm for bayesian networks: Application to self-diagnosis of gpon-ftth networks," in *2016 International Wireless Communications and Mobile Computing Conference (IWCMC),* (2016), pp. 369–376.
23. F. Christou, "Availability estimation of optical network links using multilevel bayesian modeling," in *2023 International Conference on Optical Network Design and Modeling (ONDM),* (2023), pp. 1–6.
24. J. K. Kruschke, "Bayesian analysis reporting guidelines," Nat. Hum. Behav. **5**, 1282–1291 (2021).
25. A. D. Gordon, T. A. Henzinger, A. V. Nori, and S. K. Rajamani, "Probabilistic programming," in *Future of Software Engineering Proceedings,* (Association for Computing Machinery, New York, NY, USA, 2014), FOSE 2014, p. 167181.
26. M. Richard, *Statistical Rethinking A Bayesian Course with Examples in R and Stan* (CRC Press, 2020).
27. B. Steve, G. Andrew, L. J. Galin, and M. Xiao-Li, *Handbook of Markov Chain Monte Carlo* (Chapman & Hall/CRC., 2011).
28. B. Christopher M., *Pattern Recognition and Machine Learning* (Springer New York, 2006).
29. A. Alashaikh, T. Gomes, and D. Tipper, "The spine concept for improving network availability," Comput. Networks **82**, 4–19 (2015). Robust and Fault-Tolerant Communication Networks.
30. J.-P. Vasseur, M. Pickavet, and P. Demeester, *Network Recovery: Protection and Restoration of Optical, SONET-SDH, IP, and MPLS* (Morgan Kaufmann, 2004).
31. M. D. Hoffman and A. Gelman, "The no-u-turn sampler: Adaptively setting path lengths in hamiltonian monte carlo," (2011).
32. S. Orlowski, R. Wessäly, M. Pioro, and A. Tomaszewski, "Sndlib 1.0survivable network design library," Networks **55**, 276 – 286 (2009).
33. S. Knight, H. Nguyen, N. Falkner, R. Bowden, and M. Roughan, "The internet topology zoo," Sel. Areas Commun. IEEE J. on **29**, 1765 –1775 (2011).
34. H. Ge, K. Xu, and Z. Ghahramani, "Turing: a language for flexible probabilistic inference," in *International Conference on Artificial Intelligence and Statistics, AISTATS 2018, 9-11 April 2018, Playa Blanca, Lanzarote, Canary Islands, Spain,* (2018), pp. 1682–1690.
35. J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, "Julia: A fresh approach to numerical computing," SIAM Rev. **59**, 65–98 (2017).
36. S. Danisch and J. Krumbiegel, "Makie.jl: Flexible high-performance data visualization for Julia," J. Open Source Softw. **6**, 3349 (2021).
37. M. Bouchet-Valat and B. Kamiski, "Dataframes.jl: Flexible and fast tabular data in julia," J. Stat. Softw. **107**, 1–32 (2023).
38. M. Allen, D. Poggiali, K. Whitaker, T. R. Marshall, J. van Langen, and R. A. Kievit, "Raincloud plots: a multi-platform tool for robust data visualization," Wellcome Open Res **4**, 63 (2021).
39. V. W. Berger and Y. Zhou, *KolmogorovSmirnov Test: Overview* (John Wiley & Sons, Ltd, 2014).
40. J. van Campenhout and T. Cover, "Maximum entropy and conditional probability," IEEE Transactions on Inf. Theory **27**, 483–489 (1981).
41. A. Giffin and A. Caticha, "Updating probabilities with data and moments," in *AIP Conference Proceedings,* (AIP, 2007).
42. S. Nate, *The Signal and the Noise: Why So Many Predictions FailBut Some Don't* (Penguin Group, 2012).
43. A. G. Varuvel and R. Prasath, *Reliability in Optical Networks* (John Wiley & Sons, Ltd, 2022), chap. 15, pp. 277–316.