IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received December 4, 2019, accepted December 31, 2019, date of publication January 21, 2020, date of current version January 28, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.2968380

Large-Scale Characterization and Segmentation of Internet Path Delays With Infinite HMMs

MAXIME MOUCHET[1], SANDRINE VATON[1], THIERRY CHONAVEL[1], EMILE ABEN[2], AND JASPER DEN HERTOG[2]

[1]Lab-STICC, IMT Atlantique, 29280 Plouzané, France
[2]RIPE NCC, 1012 Amsterdam, The Netherlands

Corresponding author: Maxime Mouchet (maxime.mouchet@imt-atlantique.fr)

The work of Maxime Mouchet was supported by the Futur and Ruptures Program of Institut Mines Télécom for his Ph.D.

ABSTRACT Round-Trip Times are one of the most commonly collected performance metrics in computer networks. Measurement platforms such as RIPE Atlas provide researchers and network operators with an unprecedented amount of historical Internet delay measurements. It would be very useful to process these measurements automatically (statistical characterization of path performance, change detection, recognition of recurring patterns, etc.). Humans are quite good at finding patterns in network measurements, but it can be difficult to automate this and enable many time series to be processed at the same time. In this article we introduce a new model, the HDP-HMM or infinite hidden Markov model, whose performance in trace segmentation is very close to human cognition. We demonstrate, on a labeled dataset and on RIPE Atlas and CAIDA MANIC data, that this model represents measured RTT time series much more accurately than classical mixture or hidden Markov models. This method is implemented in RIPE Atlas and we introduce the publicly accessible Web API. An interactive notebook for exploring the API is available on GitHub.

INDEX TERMS Round-trip times, RIPE Atlas, hidden Markov models, nonparametric Bayesian models, anomaly detection, time series clustering.

I. INTRODUCTION

A. SCOPE OF THE PAPER

Network management has traditionally been entrusted to humans. But this is expensive, error-prone, and slow to adapt to changes. The task of human experts is very complex because of the large number and heterogeneity of equipments, as well as the wide variety of applications.

We believe that the future of network management is in automation, or driverless (autonomous) networks. References [1]–[4]. For self-driving networks to become reality, recent machine learning techniques must be used to extract information from network measurements and automate decision-making. Different needs should be addressed: statistical characterization, prediction, detection of changes or anomalies, classification, etc. The results must be reliable and accurate in order to automate decision-making related to network management or to security and resilience, and the analysis should be scalable and fully automated (no human intervention).

The associate editor coordinating the review of this manuscript and approving it for publication was Xiaofei Wang.

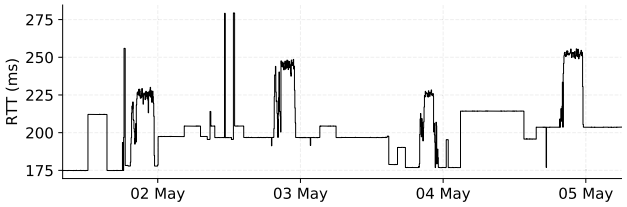

FIGURE 1. RTT between two RIPE Atlas anchors from May 1st to May 5th, 2018.

Delay is an important performance metric. In particular, it is easy to measure Round Trip Time (RTT) and data is readily available from measurement infrastructures at the Internet scale like RIPE Atlas [5]. Humans are quite good at finding patterns in this latency data, as you can see for yourself in Figure 1, but this is difficult to automate. An automated solution would allow many time series to be processed at the same time.

In this article we propose to use a Hierarchical Dirichlet Process Hidden Markov Model (HDP-HMM), also called infinite HMM, or nonparametric Bayesian HMM.

This model mimics human cognition very well (in terms of segmentation of the series, recognition of different states,

etc.). These models are used to segment audio sequences for which they give very good results for speaker recognition [6]. These recent techniques are more complex than standard HMMs but they are worth the effort.

The goal of the article is to recall the major principles of infinite HMMs and apply this theory to network measurement data. Whereas [6]–[8] are written for statisticians, we want to make the theory accessible to a wider audience and show the potential of this model for automating the analysis of a wide variety of delay time series.

The method has been implemented in RIPE Atlas to automate the processing of anchor to anchor RTT measurements, and a Web API is available. The article introduces the API and an accompanying notebook is provided to get started with the API: https://github.com/maxmouchet/atlas-trends-demo. To the best of our knowledge, this is the first time that such a method has been applied to automatically segment large delay measurements databases, which is made possible thanks to the robustness of this model. On the contrary, the validation of segmentation methods based on more standard approaches has been limited to a few subsets of time series in the literature.

B. STATE OF THE ART

Network delay modeling and prediction is a well-studied problem. Some of the simplest models assume independent observations and can be used to detect anomalous delay values. They, however, cannot predict the delay or find recurring patterns in a delay series since they do not account for time dependencies. Such models include the Pareto distribution [9] and mixtures of Weibull [10] or Normal distributions [11].

More complex time series models have been used for short-term delay predictions (from seconds to minutes), with applications such as telerobotics. These include autoregressive models [12], [13] and deep neural networks [14], [15]. Unfortunately, their parameters are more difficult to interpret and they do not provide a segmentation of the data.

HMMs are another kind of time-series model that can model different delay distributions and the dynamics between them. In [16], a discrete-time HMM is used to model packet losses, while in [17], a continuous-time HMM is used to model both packet losses and delays. In [18], a HMM is used to model inter-packet times and packet sizes. HMMs have few parameters and they are easily interpretable (state transition probabilities, means, variances, . . .).

However, standard HMMs require the use of heuristics to determine the number of hidden states. To remedy this problem, we use a nonparametric HMM for which the number of hidden states is inferred from the data. A nonparametric mixture model has been used in the past to model the delay of a set of hosts measured over disjoint time intervals [19]. In comparison, our model is a nonparametric HMM and concerns the delay between two hosts over a large and continuous time interval, from a few hours to a few weeks.

We first introduced the use of the HDP-HMM for RTT time series in [20], [21]. In these papers we explain how the time-dependency of such a Markov model can be used to reduce the frequency of measurements in routing overlays. This article expands on the statistical details behind the model, describes two new applications (CAIDA MANIC measurements, and anomaly detection), and introduces a RIPE Atlas API for time series segmentation as a service. We also demonstrate the genericity of the HDP-HMM model, which can fit any measured RTT time series much more accurately than classical mixture and hidden Markov models.

C. STRUCTURE OF THE ARTICLE

The paper is structured as follows. Section II is a reminder on mixture models (MM) and hidden Markov models (HMM). In section III we describe their nonparametric Bayesian counterparts, the Dirichlet Process MM (DPMM) and the Hierarchical Dirichlet Process HMM (HDP-HMM, or infinite HMM). In the same section, we explain how to automatically calibrate these models, that is how their parameters can be inferred from measurements without human intervention.

In section IV the accuracy of the model is demonstrated on a dataset that has been labeled by humans, as well as on some RIPE Atlas RTT time series where we discuss the matching between routing configurations (from traceroutes) and states learned by the statistical model. We also briefly address the analysis of some CAIDA MANIC measurements. In Section V we introduce a Web API that makes it possible to request the HDP-HMM analysis of anchor to anchor RTT measurements in RIPE Atlas. We also demonstrate that analyzing the frequency of state changes in RTT time series over Atlas allows a very precise detection of the moment of occurrence of events affecting large infrastructures of the Internet (such as IXPs). In Section VI we conclude and present our vision of the research axes to be developed in the future.

Readers who are less interested in the description of the Bayesian nonparametric context might wish to skip most of sections II and III, and read their summaries instead (subsections II-E and III-F).

II. A REMINDER ON MIXTURE MODELS AND HIDDEN MARKOV MODELS

In the next two sections our goal is to help the reader understand the HDP-HMM model, starting from simpler and more popular models such as mixtures or HMMs.

A. A TAXONOMY OF STATISTICAL MODELS

We start by providing a taxonomy of the different models discussed in the article. Our taxonomy takes into account three criteria: (i) whether there is naturally a notion of ''hidden state'' in the statistical model (ii) whether time dependency is taken into account, and (iii) whether the number of states is supposed to be known (and finite) or unknown (and potentially infinite).

The RTT is stable over long periods of time (usually a few hours) before its distribution switches from one probability law to another (see Fig. 1). This can be explained by reference to IP-level routing changes, congestion, and traffic

TABLE 1. Taxonomy of models.

Model	Number of states	Time dependency
Mixture Model	Fixed	No
Hidden Markov Model	Fixed	Yes
Dirichlet Process Mixture	Infinite	No
Hierarchical Dirichlet Process Hidden Markov Model	Infinite	Yes

engineering at lower layers than layer 3 [22]. Propagation delays give a lower bound on the RTTs, and as router queue lengths increase with the traffic, so do the observed RTTs. From a statistical point of view, it is natural to think of models with ''hidden states'' such as MMs or HMMs.

Knowing that the delay is stable over several hours means that, if the path quality is measured at a frequency of one ''ping'' every few minutes, the delay distribution remains stable for tens or hundreds of time slots. In order to have a model that can be used for prediction, it is necessary to account for this temporal dependence. This is made possible by HMMs, while mixture MMs assume independent observations.

But a classical problem in statistics with MMs or HMMs is that the order of the model is assumed to be known (and finite). In practice this hypothesis is unrealistic in most applications considered. This is where models with Dirichlet processes (DP) priors on the number of components of the mixture, or of the HMM, find their interest.

In the Dirichlet Process MM (DPMM) and the Hierarchical Dirichlet Process HMM (HDP-HMM), the number of model states is ''infinite''. And the order of the model can be learned from the measured data, as is the case for the other parameters of the model. This is an important property as it means that the model is flexible enough to adapt to a large number of time series, without manual human intervention (initialization of algorithms, etc.).

In Table 1 we have summarized which properties are satisfied by which models. This justifies the choice of the HDP-HMM to characterize RTTs and to automate their processing.

This flexibility is obtained at the cost of a greater complexity of the model of inference algorithms for parameter estimation. However, we could provide an efficient implementation for it embedded in an operational API (see Section V-B).

B. MIXTURE MODELS

Some of the simplest statistical models that include hidden states are mixture models. MMs are a generative statistical model used to describe data produced by different system states. For instance, in a Gaussian Mixture Model (GMM), observations $y_{1:T} = (y_1, y_2, \ldots, y_T)$ are assumed to be independent and a normal distribution is associated to each hidden state. For continuously distributed observations, conditionally to the underlying state $z_t = k \in \{1, 2, \ldots, K\}$, where K denotes the number of states of the model, the observation y_t follows a distribution with probability density function p_{θ_k}, where θ_k is a parameter vector. For example, in a GMM, θ_k consists of mean and variance parameters, so $\theta_k = (\mu_k, \sigma_k^2)$ and $p_{\theta_k}(y) = \mathcal{N}(y; \mu_k, \sigma_k^2) = (2\pi\sigma_k^2)^{-1/2} \exp\left(-\frac{(y-\mu_k)^2}{2\sigma_k^2}\right)$. Finally, the data distribution writes $p(y_t) = \sum_{k=1:K} \pi_k p_{\theta_k}(y_t)$ where π_k denotes the probability that the state of an observation is k, that is, $\pi_k = \mathbb{P}(z_t = k)$.

MM parameters $\boldsymbol{\phi} = \{\pi_k, \theta_k\}_{k=1:K}$ can be estimated from measurements $y_{1:T}$ according to different criteria and algorithms. A common choice is the Maximum Likelihood Estimator (MLE) which supplies the parameters that maximize the likelihood of the observations: $\boldsymbol{\phi}_{MLE} = \arg\max_{\boldsymbol{\phi}} p(y_1, y_2, \ldots, y_T; \boldsymbol{\phi})$. In general, direct maximization of the likelihood $p(y_{1:T}; \boldsymbol{\phi})$ with respect to $\boldsymbol{\phi}$ is infeasible. The Expectation-Maximization (EM) algorithm [23] is a popular iterative two-step algorithm used to compute the MLE for models with incomplete data, in particular mixture models.

C. HIDDEN MARKOV MODELS

Because of the independent observations assumption, the predictive ability of MMs is limited. Knowing model parameters and which state value z_t has generated the last observation y_t does not bring any information about the next state z_{t+1}. HMMs are a generalization of MMs that take into account temporal dependencies among states. These temporal dependencies are expressed through a Markov property assumed for the states, written as $p(z_{t+1}|z_{1:t}) = p(z_{t+1}|z_t)$. Thus, the probability distribution of the next hidden state z_{t+1} depends on the current state z_t only.

More formally the transition probabilities between successive states are defined via a $K \times K$ matrix Π with entries $\Pi_{ij} = P(z_{t+1} = j | z_t = i)$. The model parameters are now $\boldsymbol{\phi} = \{\Pi, \{\theta_k\}_{k=1:K}\}$, the steady state probability vector $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_K)$ being related to Π through the linear system $\boldsymbol{\pi}\Pi = \boldsymbol{\pi}$ and $\boldsymbol{\pi}\boldsymbol{e} = 1$, where $\boldsymbol{e} = (1, \ldots, 1)^T$.

The MLE of HMM parameters can be estimated using a variant of the EM algorithm known as the Baum-Welch (or *forward-backward*) algorithm [24]. While easy to implement and well-studied, this approach is prone to overfitting on noisy data or data with few samples. Furthermore this method requires the number K of hidden states to be known, which is usually not the case when studying RTTs.

D. LIMITATIONS OF VANILLA MMS AND HMMs

Classical mixtures and HMMs are parametric models, meaning that they have a set of parameters with fixed size. This is a major difficulty when estimating HMM parameters as often the number of hidden states is not known in advance.

One could estimate models for different numbers of states, but the maximum of the likelihood would increase with the number of states as a model of order K is a degenerated case of model of order $K + 1$. A classical approach consists of penalizing the MLE optimization criterion by adding a penalty term to the log-likelihood such as the AIC [25] or the

BIC [26] criterions and by selecting the model that minimizes this penalized criterion. Another approach is to use nonparametric models with unbounded number of parameters.

Another limitation of parametric models is that the EM algorithm usually used to tune the parameters of the model is sensitive to the choice of its initialization point. Appropriate initialization strategies must be used otherwise it may converge to a local but non-global maximum of the likelihood.

Because of these limitations, standard MMs or HMMs cannot be used on a large scale to analyze Internet measurements. In the following section we introduce a new approach for RTT measurement analysis, based on nonparametric Bayesian models, and more particularly the HDP-HMM.

E. SECTION SUMMARY

MMs and HMMs are interesting for characterizing time series of RTTs. They are designed to model phenomena that change state from time to time and in which the value of the observations, here the RTTs, noisily depends on the hidden states. One can imagine that different hidden states result from different routing configurations, traffic engineering choices, or link loads. However, these models are too simple to characterize a large variety of RTT series and are not suitable for automatic processing on a large scale.

We propose to use a more generic model, the HDP-HMM. This model does not make assumptions about the number of states of the system, contrary to vanilla mixtures or HMMs, and it is possible to learn the number of states from the data itself. Contrary to DPMMs it also takes into account time dependency and makes it possible to account for the RTT distribution being stable for a long period of time.

III. NONPARAMETRIC BAYESIAN APPROACH

A more formal approach to models with an unknown number of components can be found in Bayesian statistics. The Bayesian framework allows one to specify models with several layers of uncertainty and infer the parameters in a systematic way. We will make better use of this flexibility to estimate HMMs from RTT series where neither the number of states nor the probability distribution in each state is known.

A. BAYESIAN SETTING

In the MLE approach, estimates of the parameters are inferred from data. In contrast, Bayesian approaches make use of prior distributions upon the model parameters, and output a posterior probability distribution over the model parameters. These prior distributions can account for prior knowledge of the parameter distributions.

When the dimension of the model is unknown, as for MMs or HMMs with unknown order K, one can resort to nonparametric Bayesian approaches, where the number of components of the model is inferred from the data itself.

Bayesian inference can be performed from the posterior likelihood which is defined as $p(\boldsymbol{\phi}|y_{1:T}) \propto p(y_{1:T}|\boldsymbol{\phi})p(\boldsymbol{\phi})$ where $p(y_{1:T}|\boldsymbol{\phi})$ is the likelihood of the data $y_{1:T}$, $p(\boldsymbol{\phi})$ is a *prior distribution* and $\propto$ denotes proportionality.

In general, a direct maximization of the posterior likelihood $p(\boldsymbol{\phi}|y_{1:T})$ with respect to $\boldsymbol{\phi}$ is not feasible as $p(\boldsymbol{\phi}|y_{1:T})$ can be quite complex. Note, however, that there are situations where the likelihood and the prior distribution are such that posterior distribution belongs to the same family as the prior. In this case, the prior is said to be conjugate. Using conjugate priors, when possible, often makes inference simpler.

Markov Chain Monte Carlo (MCMC) techniques, and in particular Gibbs sampling, can be used in very general situations for inference [27]. Alternatively, variational Bayesian methods can be considered ([28], chap. 33). The principle of MCMC methods is to use simulations to draw a large number of samples $\boldsymbol{\phi}$ from the posterior distribution $p(\boldsymbol{\phi}|y_{1:T})$.

B. DIRICHLET PROCESSES AND DP MIXTURES

Modelling a HMM with an infinite number of states is generally achieved by means of a Dirichlet process (DP) prior. DPs were introduced by Ferguson [29] in 1973 and were first applied to mixture models with an unknown number of components in [30]. The extension to the modelling of HMMs was first defined in 2002 in [31]. More recently this has been formalized in the framework of hierarchical Dirichlet processes (HDP) in [8] where HDP-HMMs have been introduced. These models are called nonparametric Bayesian, meaning that they are Bayesian and involve parameter spaces of infinite dimension [32].

A Dirichlet Process (DP) is a stochastic process $G \sim DP(\alpha, H)$, the realizations of which are probability distributions. It is parameterized by a concentration parameter α and a base distribution H. It can be seen as a process indexed by partitions $(A_1, \ldots, A_n)$ $(n > 0)$ of the space E on which H is defined, with n-variate Dirichlet random realizations:

$$(G(A_1), \ldots, G(A_n)) \sim \mathbf{Dir}(\alpha H(A_1), \ldots, \alpha H(A_n)). \quad (1)$$

Here $\mathbf{Dir}(\alpha_1, \ldots, \alpha_n)$ denotes the n-variate Dirichlet *distribution* with parameters $\alpha_{1:n} = (\alpha_1, \ldots, \alpha_n)$, that is to say the probability distribution with density function:

$$p(x_{1:n}; \alpha_{1:n}) = \frac{1}{B(\boldsymbol{\alpha})} \mathbb{1}_{\{1\}}\Big(\sum_{i=1:n} x_i\Big) \prod_{k=1:n} x_i^{\alpha_i - 1} \mathbb{1}_{[0,1]}(x_i) \quad (2)$$

where $\mathbb{1}_A(x) = 1$ if $x \in A$ and 0 otherwise, and $B(\boldsymbol{\alpha})$ is a normalization factor.

Alternative definitions of DPs are also useful both for their understanding and simulation. In particular it can be proved that a *Dirichlet Process* $G \sim \mathrm{DP}(\alpha, H)$, can also be defined via the *stick-breaking* constructive approach [33]. The idea is to build a discrete distribution by assigning probabilities π_k to samples θ_k drawn independently from H. As the probabilities π_k must sum to 1, a unit-length stick is divided as displayed in Figure 2. The stick is first broken into two parts, of lengths η_1 and $1 - \eta_1$. Then the second portion, of length $1 - \eta_1$, is broken again into two parts in proportions η_2 and $1 - \eta_2$. The three resulting portions are now of lengths η_1, $\eta_2(1 - \eta_1)$ and $(1 - \eta_2)(1 - \eta_1)$. The process of breaking the stick into smaller parts continues indefinitely.

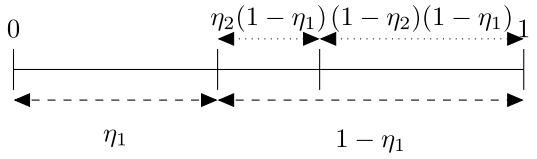

FIGURE 2. The stick-breaking process.

The weights π_k are defined as $\pi_1 = \eta_1$, $\pi_2 = \eta_2(1-\eta_1)$, $\pi_3 = \eta_3(1-\eta_2)(1-\eta_1)$, and in general $\pi_k = \eta_k \prod_{l=1:k-1}(1-\eta_l)$. The proportions η_k are independent and $\eta_k \sim \textbf{Beta}(1,\alpha)$, where $\textbf{Beta}(a,b)$ is the beta distribution with parameters a and b and probability density function $x^{a-1}(1-x)^{b-1}\mathbb{I}_{[0,1]}(x)$.

The distribution with weights $\boldsymbol{\pi} = [\pi_1, \pi_2, \ldots]$ is called a Griffiths-Engen-McCloskey distribution, denoted by $\boldsymbol{\pi} \sim \textbf{GEM}(\alpha)$. Clearly, $\sum_{k=1:\infty} \pi_k = 1$. We then get the stick-breaking representation of the Dirichlet Process G:

$$G = \sum_{k=1:\infty} \pi_k \delta_{\theta_k}, \quad \text{with } \boldsymbol{\pi} \sim \textbf{GEM}(\alpha) \text{ and } \theta_k \sim H. \quad (3)$$

Note that the π_ks tend to decay to zero at geometric rate. Indeed it can easily be proven that:

$$\mathbb{E}[\pi_k] = \mathbb{E}[\eta_k] \prod_{l=1:k-1} (1-\mathbb{E}[\eta_l]) = \frac{1}{\alpha+1}\left(\frac{\alpha}{\alpha+1}\right)^{k-1}. \quad (4)$$

Now, suppose we want to fit a mixture model to some observations $y_{1:T} = (y_1, y_2, \ldots, y_T)$. Assume that the mixing distributions are in the form $p_\theta(y)$, where θ is a vector of parameters and that the prior distribution over the vector of parameters is $\theta \sim H$. We can build a nonparametric Bayesian generative model of observations in the form of a Dirichlet Process Mixture model (DPMM). In this model the distribution of observations is a mixture:

$$p(y) = \sum_{k=1:\infty} \pi_k \, p_{\theta_k}(y) \quad (5)$$

and the weights π_k and parameters θ_k of the different components of the mixture are defined as a Dirichlet Process:

$$G = \sum_{k=1:\infty} \pi_k \delta_{\theta_k} \sim DP(\alpha, H) \quad (6)$$

C. HIERARCHICAL DIRICHLET PROCESS HMM

The idea of using a DP as a prior in mixture models has been extended to the case of Hidden Markov Models (HMMs). In fact, for some technical reasons that we will explain, the extension of this approach to HMM modelling involves a hierarchy of DPs.

In the Hierarchical Dirichlet Process HMM (HDP-HMM), DPs are used as priors on the rows $\boldsymbol{\pi}_i = (\pi_{i1}, \pi_{i2}, \ldots, \pi_{ik}, \ldots)$ of the transition matrix Π of the hidden Markov chain $(z_t)_t$. This makes it possible to specify that the number of states of the Markov chain is unknown.

But it is also necessary to ensure that the transition probabilities π_{ik}, for all row i, weigh the same emission

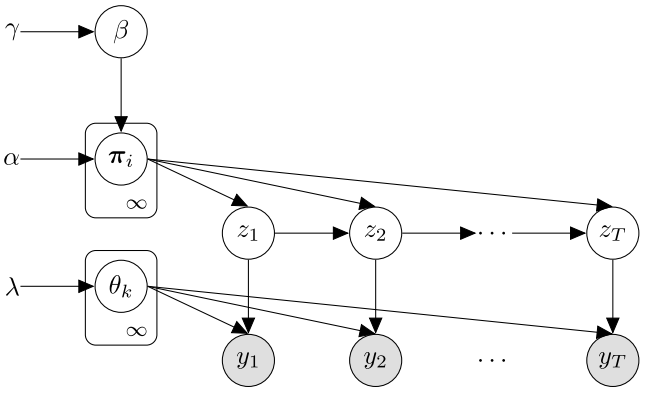

FIGURE 3. The Hierarchical Dirichlet Process - Hidden Markov Model (HDP-HMM).

distribution p_{θ_k}. This is made possible by parameterizing the DPs G_i ($i = 1, 2, \ldots$) by the same *discrete valued* base distribution G_0

$$G_i = \sum_k \pi_{ik} \delta_{\theta_k} \sim DP(\alpha, G_0) \quad (7)$$

where G_0 is modeled by a DP prior with base distribution H_λ:

$$G_0 = \sum_k \beta_k \delta_{\theta_k} \sim DP(\gamma, H_\lambda) \quad (8)$$

This hierarchy of DPs yields the HDP-HMM process [8]. A graphical representation of the HDP-HMM is given in Figure 3, where the arrows represent the dependencies. The HMM itself is represented by states z_t and observations y_t. Its parameters are $(\theta_k)_{k\geq 1}$ and $(\boldsymbol{\pi}_i)_{i\geq 1}$, where $p_{\theta_k}(y_t) = p(y_t|z_t = k)$ and $\boldsymbol{\pi}_i$ denotes the i^{th} row of the transition matrix Π of the HDP-HMM, so $\pi_{ij} = \mathbb{P}(z_{t+1} = j \mid z_t = i)$.

α, γ and λ are hyper-parameters. γ and λ are the parameters of a Dirichlet process $G_0 \sim DP(\gamma, H_\lambda)$ that lies at the top of the HDP hierarchy. These random dependencies and vague priors introduce enough flexibility in the model to let it adapt to many different time series.

D. INFERENCE IN DP MIXTURES

Inference in DPMMs is better addressed via the so-called Polya urn representation of DPs than through stick breaking. Imagine an urn that contains black and colored balls. The "values" of balls are their colors. At initialization the urn only contains α black balls. When drawing a ball from the urn, if the ball drawn is black then a new colored ball is drawn from a base distribution H and the black and colored balls are put back into the urn. When the first ball drawn is not black, it is put back into the urn together with a new one of the same color. The labels of the infinite sequence of draws follow a DP.

We are going to use this formalism in a DPMM, where z_t denotes the hidden state and y_t is the observation. The Polya urn model can be described as follows. Let us introduce $\theta'_t = \theta_{z_t}$ the value of θ associated with z_t. If $z_t = k$ then $\theta'_t = \theta_k$ and y_t is distributed according to $p_{\theta_k}(\bullet)$. Given a sequence of

random variables $(\theta'_t)_{t>0}$ with

$$P(\theta'_1 \in B) = H(B), \quad \text{and}$$
$$P(\theta'_{t+1} \in B | \theta'_{1:t}) = \frac{1}{\alpha + t}\left(\sum_{\tau=1:t} \delta_{\theta'_\tau}(B) + \alpha H(B)\right). \quad (9)$$

it has been shown, in [34], that the distribution of θ'_t converges almost surely to a DP(α,H) (when $t \to \infty$).

The estimation of the parameters and states of a non-parametric mixture model from the posterior distribution $p(z_{1:T}, \theta_{1:K_T} | y_{1:T})$, where $y_{1:T}$ represent the data, can be addressed via Gibbs sampling [7]. The principle of Gibbs sampling [27] is to sequentially update, in turn, the values of $z_t, t = 1, \ldots$ and $\theta_k, k = 1, \ldots$ conditionally to $y_{1:T}$ and to the current value of the other parameters. It requires knowing the distribution of each latent variable conditionally to the observations $y_{1:T}$ and the other latent variables.

Going back to the Polya urn model, let us index by $1, \ldots, K_t$ the distinct colors of the balls present in the urn at time t and let z_t denote the color index of the new ball. As the role of the balls can be exchanged, letting $z_{-t} = \{z_{1:t-1}, z_{t+1:T}\}$ and $n_{-t,k} = \#\{z_\tau \in z_{-t};\ z_\tau = k\}$ be the number of occurrences of the value k among z_{-t}, it can be shown [7] that:

$$P(dz_t | z_{-t}) = \frac{1}{\alpha + T - 1}\left(\sum_{k=1}^{K_{-t}} n_{-t,k}\delta_k(dz_t) + \alpha\delta_{K_{-t}+1}(dz_t)\right) \quad (10)$$

where K_{-t} is the number of distinct elements in z_{-t} with indexing set from 1 to K_{-t}. Equation 10 can be interpreted as follows: knowing the values of $z_{1:t-1}$ and $z_{t+1:T}$, the distribution of z_t is a mixture of the values $k \in z_{-t}$ and of a new index value $(K_{-t} + 1)$. The respective weights of this mixture are $\frac{n_{-t,k}}{\alpha+T-1}$ for any $k \in z_{-t}$ and $\frac{\alpha}{\alpha+T-1}$ for the value $K_{-t} + 1$.

It can be proven [7] that, if observations $y_{1:T}$ and parameters θ'_{-t} are taken into account, then:

$$P(dz_t | z_{-t}, y_{1:T}, \theta'_{-t}) \propto \sum_{k=1}^{K_{-t}} n_{-t,k} p_{\theta_k}(y_t)\delta_k(dz_t) + \alpha\mathcal{I}(y_t)\delta_{K_{-t}+1}(dz_t). \quad (11)$$

where $\mathcal{I}(y_t) = p(y_t \mid z_t = k, \theta'_{-t}) = \int p_\theta(y_t) H_\lambda(d\theta)$.

Note that, provided $\mathcal{I}(y_t)$ is known, the proportionality factor in Eq. (11) can be obtained from the normalization condition $\sum_k P(z_t = k \mid z_{-t}, y_{1:T}, \theta'_{-t}) = 1$. If p_θ and H_λ are conjugate distributions, $\mathcal{I}(y_t)$ can easily be calculated in closed form. In other cases one can resort to Metropolis-Hastings simulation using the prior distribution of z_t in (10) as an auxiliary distribution [7] to calculate $\mathcal{I}(y_t)$.

After sampling $z_t, t = 1 : T$ from Eq. (11), $\theta_k, k = 1, \ldots$ can be sampled from the following distribution [7]:

$$P(d\theta_k | z_{1:T}, y_{1:T}, \theta_{-k}) \propto H_\lambda(d\theta_k) \prod_{\{t; z_t = k\}} p_{\theta_k}(y_t). \quad (12)$$

Here again simulation can be performed directly or via Metropolis-Hastings simulation depending whether p_θ and H_λ have conjugate distributions.

E. INFERENCE IN HDP-HMMs

Inference in HDP-HMM is technically more involved than for mixture models. We briefly summarize it here. Interested readers can find additional information in appendices of [6].

Letting K denote the current number of states, the Gibbs sampler should sample $z_{1:T}$. Note that $\theta_{1:K}$ can be marginalized out and does not need to be sampled in Gibbs. To make it possible, we will also have to sample the π_j, which in turn requires sampling the weights of the base distribution $G_0 = \sum_{k=1:\infty} \beta_k \delta_{\theta_k}$. As only $(\beta_k)_{k=1:K}$ is concerned for describing the weights of the states of the finite size data set at hand, letting $\beta_{-K} = \sum_{k=K+1:\infty} \beta_k = 1 - \sum_{k=1:K} \beta_k$, we simply sample $(\beta_{1:K}, \beta_{-K})$ that follows a Dirichlet distribution of order $K + 1$. The sampling of $(\beta_{1:K}, \beta_{-K})$ is described in [6].

Note also that we want to implement inference for a sticky HDP-HMM, that is, a modified version of the HDP-HMM that models persistency of the states by biasing the model towards self transitions $(z_{t-1} = j, z_t = j)$. This is ensured by introducing an additional parameter κ and changing the prior upon π_j:

$$\pi_j | \alpha, \beta, \kappa \sim \mathbf{DP}\left(\alpha + \kappa, \frac{\alpha(\sum_k \beta_k \delta_k) + \kappa\delta_j}{\alpha + \kappa}\right). \quad (13)$$

When $\kappa = 0$ we get the standard HDP-HMM, while when $\kappa \to \infty$, π_j tends to only weight state j.

To implement the Gibbs sampler for the states $z_{1:T}$ let $\psi = (\alpha, \beta, \kappa, \lambda)$, and $\boldsymbol{\pi} = (\pi_j)_j$. Then $P(z_t | y_{1:T}, z_{-t}, \psi)$ can be expressed by marginalizing against the $\boldsymbol{\pi_j}$s and θ_ks:

$$P(z_t \mid y_{1:T}, z_{-t}, \psi) \propto P(z_t \mid z_{-t}, \psi) p(y_t \mid y_{-t}, z_{1:T}, \psi) \quad (14)$$

Let us introduce the following notations: $x_{i\bullet} = \sum_j x_{ij}$ and n^{-t}_{jk} denotes the number of transitions from state j to state k, not counting the transitions $z_{t-1} \to z_t$ or $z_t \to z_{t+1}$. Then, the first factor in (14) can be written as:

$$\begin{aligned} &P(dz_t \mid z_{-t}, \psi) \\ &\quad \propto \sum_{k=1}^{K_{-t}} \frac{\alpha\beta_k + n^{-t}_{z_{t-1},k} + \kappa\delta_{z_{t-1},k}}{\alpha + \kappa + n^{-t}_{z_{t-1},\bullet}} \\ &\qquad \times \frac{\alpha\beta_{z_{t+1}} + n^{-t}_{k,z_{t+1}} + \kappa\delta_{z_{t-1},k}\delta_{k,z_{t+1}}}{\alpha + \kappa + n^{-t}_{k,\bullet} + \delta_{z_{t-1},k}} \delta_k(dz_t) \\ &\qquad + \frac{\alpha^2\beta_{-K}\beta_{z_{t+1}}}{(\alpha + \kappa)^2}\delta_{K_{-t}+1}(dz_t). \end{aligned} \quad (15)$$

And the second factor in (14) as:

$$p(y_t \mid y_{-t}, z_{1:T}, \psi) \propto \int_{\theta_{z_t}} p(y_t \mid \theta_{z_t}) H_\lambda(d\theta_{z_t} \mid \{y_\tau; z_\tau = z_t, \tau \neq t\}). \quad (16)$$

As far as discussed earlier, if the θ_ks have conjugate prior distributions, $p(y_t \mid y_{-t}, z_{1:T}, \lambda)$ can be calculated in closed form. Note in addition that to avoid a particular choice of hyperparameters $(\alpha, \gamma, \lambda)$ biasing the solution, they can also be given some prior distribution.

At the end of the process, after the $z_{1:T}$ have been estimated, the θ_ks can be estimated easily, e.g. by maximizing the likelihood $p(\{y_t; z_t = k\} \mid \theta_k)$.

F. SECTION SUMMARY

In this section we have introduced non-parametric Bayesian approaches. In Bayesian statistics some of the parameters on which the data depend are considered random. The term ''non-parametric'' means that there is a large number of parameters that are estimated from the data.

When the number of states of a mixture or a HMM is not known in advance, it is possible to use a non-parametric Bayesian approach using Dirichlet processes (DP) as priors. This is called the Dirichlet Process Mixture Model (DP-MM) or the Hierarchical Dirichlet Process Hidden Markov Model (HDP-HMM). Equivalently, the name infinite (or non-parametric) mixture or HMM can also be used.

Missing data, that is, states and parameters, can be estimated from observations using a Gibbs sampling algorithm which comes up to randomly simulating, in turn, the different components of the model which are not measured directly. These components are simulated according to some conditional distributions which have been specified in this section.

IV. A FIRST LOOK AT RTTS THROUGH THE HDP-HMM

As stated previously, HDP-HMM is a flexible method for inferring HMM parameters and segmenting data when the number of latent states is unknown. This fits the problem of segmenting RTT time series (remember that of Figure 1), where the number of different states is not a priori known. Furthermore, it is not mandatory to make an assumption on the type of RTT distribution in each state (Gaussian, exponential, or other kind of parametric distributions). This distribution can be assumed nonparametric, which introduces even more flexibility and allows a very generic model that adapts to a very large number of traces.

In this section we show that the model produces realistic segmentations from a human point of view, and that the inferred parameters can be interpreted easily, given the application domain. In addition, we provide two validations for the model. We show on a labeled change point dataset that the model performs at least as well as ad-hoc change point detection methods. And we also show that the states inferred from the RTT time series match well with the AS and IP paths seen in RIPE Atlas traceroutes.

A. A NONPARAMETRIC OBSERVATION MODEL

Many parametric models have been proposed in the literature to explain the distribution of the delay in computer networks and on the Internet. For example, in [11] a Gaussian mixture model is proposed, in [10] a Weibull mixture model, and in [9] a Pareto distribution. In practice, however, it seems that the distribution can be very different depending on the network state. For example, in some states the delay can be relatively stable with occasional spikes above a baseline, in which case it might be modeled by an exponential distribution, while in other states the delay can experience large variations caused by a high volume of traffic, and might be better modeled by a normal distribution.

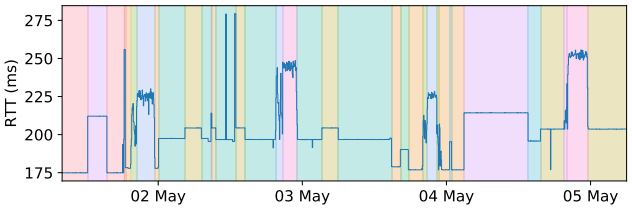

FIGURE 4. Segmentation of RTT observations between `at-vie-as1120` **and** `sg-sin-as59253`**.**

In this work, we choose instead to use nonparametric Dirichlet Process Mixture Models (with a Gaussian as ''base'' distribution) as emission distributions of the HDP-HMM. As such, the delay in each state is modeled by a varying number of Gaussian components. This allows us to model a wide range of distributions, and we avoid choosing a particular parametric emission distribution for each state of the HDP-HMM. For each Gaussian component, we use a Normal-Inverse-χ^2 prior, which is the conjugate prior to the normal distribution with unknown mean and variance. The use of a nonparametric observation model reinforces the need for Bayesian inference methods, since a more traditional MLE approach would require several layers of penalization.

The segmentation of the series from Figure 1 using an HDP-HMM with DP-GMM emissions is shown in Figure 4, in which states are represented by colors.

As a matter of comparison, we provide in Fig. 5 the segmentation obtained with a HDP-HMM with DPMM emission distributions, with that resulting from parametric and non-parametric MMs and HMMs with a Gaussian observations model. In the case of the Gaussian MM and of the HMM, the number of latent states has been chosen by estimating the model for a varying number of components and choosing the number that minimizes the penalized log-likelihood using the BIC criterion. As we can see the HDP-HMM produces a segmentation close to what a human would do, contrary to other models which generate far too many state changes.

B. CHANGE POINT DETECTION

Quantifying the performance of the HDP-HMM on real RTT time series is not easy since there is no ground truth. The ''network state'' is not known or vaguely defined. But it was possible to compare the performance of the model in a change point detection task where the goal is to detect significant changes in the delay. While not the primary purpose of the HDP-HMM, detecting change points is simply a matter of segmenting the data and finding changes in the inferred state sequence. This can be used to partially validate the quality of the segmentation obtained.

We have benchmarked the HDP-HMM against different change point detection methods on a labeled dataset introduced by [35]. This dataset is particularly interesting because change points in RTT timeseries have been manually labeled by human experts. To our knowledge, there are no other RTT time series datasets that are both realistic and labeled.

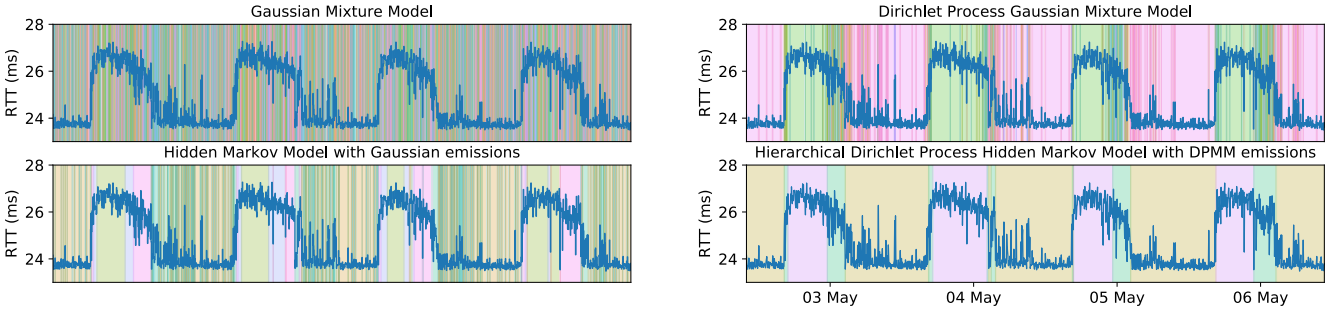

FIGURE 5. Segmentation of a RTT time series with parametric and nonparametric mixture models and HMMs.

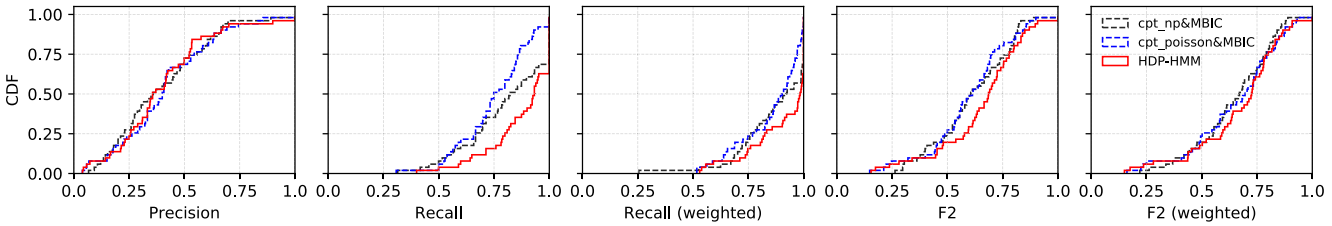

FIGURE 6. Benchmark of the HDP-HMM against classical change point detection methods on a human-labelled change point dataset [35]. The weighted recall gives more importance to large delay changes.

The dataset consists of 50 RTT series of varying length for a total of 34,008 hours of observations. In [35] change point detection is performed by minimizing $\sum_{i=1}^{m+1} C(y_{\tau_{i-1}+1:\tau_i}) + \beta f(m)$. m is the number of changes, C is a cost function that measures the stability of the delay over a range of successive values, and $f(m)$ is a penalty that prevents overfitting. Different cost functions and penalties are considered.

We have compared the performance of the segmentation obtained by HDP-HMM with the best performing changepoint detection methods of [35]. In our approach a HDP-HMM model is learnt on each timeseries, the most likely hidden state sequence is computed, and changepoints are simply defined as changes in the hidden state sequence.

In Figure 6 we show that the HDP-HMM performs similarly to the best performing change point detection methods of [35] in terms of precision ($\frac{\#\text{ True Positive}}{\#\text{ True Positive}+\#\text{ False Positive}}$), while performing better in terms of recall ($\frac{\#\text{ True Positive}}{\#\text{ True Positive}+\#\text{ False Negative}}$). This means that our model is more sensitive to small changes in the delay without generating unnecessary false alarms.

C. RIPE ATLAS MEASUREMENTS

In addition to detecting significant changes in the delay, the HDP-HMM also provides a notion of hidden states. In this section we validate the quality of this clustering both visually and by studying the correlation with AS and IP paths revealed by traceroutes.

1) DATASET

RIPE Atlas offers two types of measurement sources: probes and anchors. Probes are deployed in heterogeneous environments while anchors are restricted to high-availability environments such as data centers, universities, and IXPs (Internet eXchange Points). Anchors tend to be located in well-connected autonomous systems and measurements between anchors represent more stable paths than may be observed from probes located at the edges of the Internet. Moreover, anchors are more powerful and perform so-called *anchoring mesh measurements*, where various measurements are performed regularly between each pair of anchors. This allows us to collect traceroute results both on the outward and on the return path.

Our dataset consists of one week of IPv4 RTT measurements between all Atlas anchors and the `at-vie-as1120` anchor.[1] Delay is measured every four minutes using three ICMP (Internet Control Message Protocol) pings towards the target anchor. The minimum value of the delay is kept for each time step. Considering the subset of anchors that were online over the time period, we collected 301 series of 2520 data points. The associated traceroute measurements were also collected, both on the outward path, and on the return path. Traceroutes are performed every fifteen minutes using three ICMP probe packets for each hop.

2) INFERENCE

Each series were segmented using our Julia implementation of the Gibbs sampler. It takes less than 2 seconds on a single thread of a 2.80GHz Intel Core i7-7600U CPU to process a 2520 point time series (1 week of an Atlas RTT measurement) with 300 iterations of the sampler. The task is highly parallelizable as each time series can be processed independently.

[1]RTT measurement results are available at https://atlas.ripe.net/measurements/1437285. We considered the period between the 2nd and the 9th of May 2018.

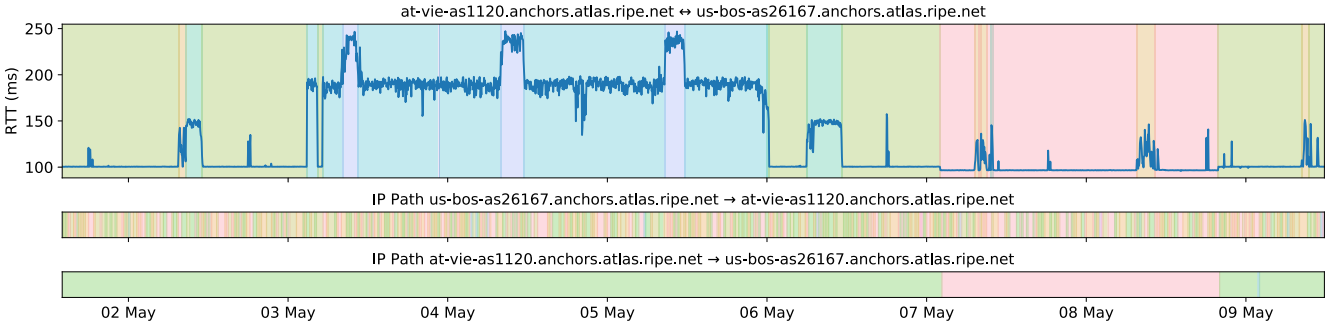

FIGURE 7. Segmentation of RTT observations between `at-vie-as1120` and `us-bos-as26167` using an HDP-HMM with DP-GMM emissions. Each color identifies a state or an IP path observed in the traceroute.

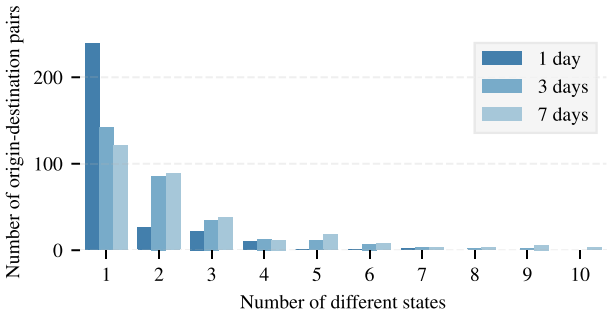

FIGURE 8. Distribution of the number of states learned for different timescales.

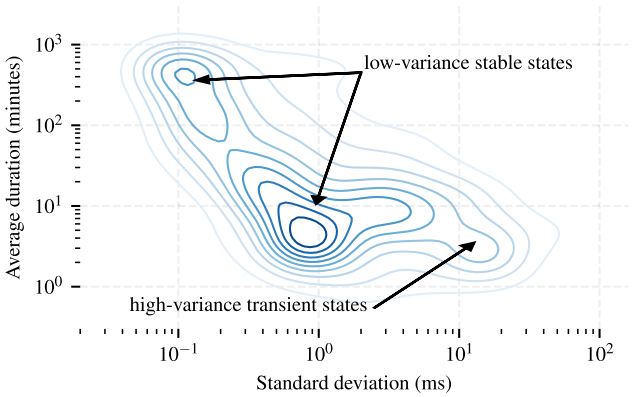

FIGURE 9. Density estimation of the (standard deviation, average duration) couple. Darker colors indicate a higher density.

Using 4 threads, 300 one-week long time series can be processed in 6 minutes.

Figure 8 shows the distribution of the number of states in the resulting HMMs for different measurement timescales. It is clear that the number of states grows with the length of the series. This is not surprising and visual inspection by a human expert would also probably identify more states in longer timeseries. One, three, and seven-day long series have less than 8, 10, and 11 states respectively. This confirms the capability of the HDP-HMM to learn more complex models as the number of RTT observations, and possibly the number of underlying network configurations, grows.

3) STATE DURATIONS VS. DELAY VARIATIONS

An advantage of HMMs over other timeseries models (e.g. autoregressive models or neural networks) is that the parameters can be interpreted easily, given the application domain. In our case, the state transition matrix Π gives us information about the frequency of network configuration changes and the relation between them, while the observation distributions give us in particular the mean value of the delay and its variance (of the delay in each configuration).

On most time series we can distinguish two types of states: those where the delay is relatively constant (such as the green one on Fig. 7), and states where the delay is very variable (such as the purple one). This is reflected by the variance of the delay in the state. And the average duration of a HMM in a state i is given by $1/(1-\pi_{ii})$ where π_{ii} is the probability of self-transition. In the example of Fig. 7 the average duration of the purple state is of 45 timesteps (= 3 hours) and of 149.5 timesteps (= 9 hours 58 minutes) for the green state. The standard deviation of the delay in the purple state is of $\sigma = 10.3$ msec while the standard deviation of the green state if of $\sigma = 4.1$ msec. States with a high variance could possibly be explained by intra-domain load-balancing (since Atlas ping flow ID is not constant), congestion, or in-path devices delaying the processing of ICMP packets. However, assessing the cause of such variations and studying the possibility of detecting them from delay measurements is to be done in future works.

Figure 9 displays the standard deviation of the RTT against the average duration in a state. In the analyzed dataset, the average state duration decreases as the RTT standard deviation increases. This is not surprising as we expect Internet paths to spend more time in stable states.

4) RELATIONSHIP WITH THE AS AND IP PATHS

We hypothesized that the distribution of delay observations is conditioned on the underlying network state, such as the inter and intra-AS routing configuration, as well as the traffic level. As shown in Figure 10, the majority of the states learned over all the paths in our dataset match only one AS path and one IP path. For example there are 595 states which always correspond to the same AS path over the 746 states learned.

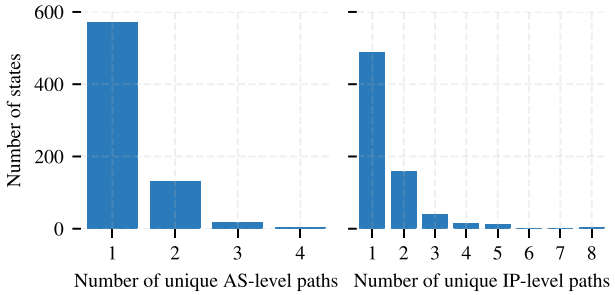

FIGURE 10. Distribution of the number of states associated with a given number of unique paths.

In other words, only 16% of the states learned can match two AS paths or more. States associated with more than one AS path can be explained by delay differences too small to be separated into two clusters.

Conversely, one AS or IP path can be mapped to several states. For example in Figure 7 we only observe the AS path `ASN MARKLEY` → `GTT BACKBONE` → `NTT COMMUNICATIONS` → `ACONET SERVICES` in the traceroutes from `us-bos-as26167` to `at-vie-as1120` and `ACONET SERVICES` → `ACONET` → `NEXTLAYER AS` → `NTT COMMUNICATIONS` → `ASN MARKLEY` in the reverse traceroutes (as resolved using the RIPEstat API). In the forward traceroutes we observe IP path changes every 15 minutes, in the GTT and NTT ASes, probably due to intra-AS load-balancing, while in the reverse traceroutes we only observe two different IP paths in NTT AS that are perfectly correlated to state changes in the model.

D. CAIDA MANIC AND OTHER MEASUREMENTS

In addition to RIPE Atlas delay measurements, the HDP-HMM fits other kinds of network measurements as well. In this section we show the results obtained on delay measurements from the CAIDA MANIC project [36]. The CAIDA MANIC project uses Time Series Latency Probes (TSLP) to measure inter-domain congestion. Once a peering link between two ASes has been identified, ICMP probes are sent to the near-end (i.e. the last router in the first AS) and the far-end (i.e. the first router in the second AS) of the link. The intuition is that if there is congestion the router queues will fill up, and the delay between the near-end and the far-end will increase. Using the same model as for the RIPE Atlas RTT series, we segment the delay difference time series (far-end - near-end) from publicly available measurements.

In Figure 11 we show the resulting segmentation for a peering link experiencing periodic congestion. Three states are learned. The green state, corresponding to a non-saturated link, has a standard deviation of 0.1 ms, while the standard deviations for the red and blue states are 7 ms and 11 ms, respectively. The blue state seems to correspond to a state of increased traffic, while the red state seems to correspond to a saturated link. Because the model accounts for temporal dependencies, it is able to clearly separate those two states even though their distributions overleap.

In addition, we have tested the use of this segmentation method with application-level delay measurements such as web Page Load Times (PLTs) and obtained promising results. Figure 11 shows the PLT for the baidu.com website measured from a probe located in Lannion, France. Measurements have been performed with the WebView platform.[2]

V. LARGE-SCALE MEASUREMENT ANALYSIS

Internet monitoring projects such as RIPE Atlas provide a large amount of latency information. Due to its scale, RIPE Atlas has a good chance of providing enough information to allow the detection of anomalous latency patterns in important network components, such as IXPs or large transit providers. However, detecting and characterizing these anomalies has proven challenging (e.g. the analysis in [37] took weeks). In this section, we start by showing the validity of the model on a very large scale considering all the measurements of the Atlas anchoring mesh. We show that real data are as likely under the model considered as simulated data would be. Then we introduce the trends API that we have developed and that publicly provides an on-demand segmentation service for Atlas RTT measurements. Finally, we show how aggregating change points learned with the HDP-HMM from a large number of origin-destination pairs is a simple and elegant method for detecting and characterising anomalies in key Internet infrastructures.

A. VALIDATION OF THE HDP-HMM MODEL AT LARGE SCALE ON RIPE ATLAS

In Section IV-B, we have shown that the HDP-HMM model is at-least as good as classical change point detection methods on a labelled RTT change points dataset. This however, does not tell us whether the model fits RTT data well from a statistical point of view. In this section, we propose to compare the likelihood of the time series (with respect to their inferred model) with the likelihood of time series simulated according to an HDP-HMM model. If the models fit the data well, we can expect the likelihood of the data with respect to the model to follow the same distribution as the likelihood of synthetic data generated by the model.

To perform a comparison, 100k time series of one week duration (2520 data points) from the anchoring mesh measurements were considered. For each time series, the model was learned, and their likelihood $p(\boldsymbol{y}|\boldsymbol{\pi}, \boldsymbol{\theta})$ computed with respect to the model. In addition, for each HMM with parameters $(\boldsymbol{\pi}, \boldsymbol{\theta})$, a time series $\boldsymbol{y}'$ was sampled and its likelihood $p(\boldsymbol{y}'|\boldsymbol{\pi}, \boldsymbol{\theta})$ computed.

The distributions of the likelihood on observed and synthetic time series are compared in Figures 13 (Q-Q plot) and 14 (histograms). It can be seen that both distributions are similar, with the simulated time series being slightly more likely. This demonstrates that the HDP-HMM explains well the diversity of observed trajectories in RIPE Atlas measurements.

[2]https://webview.orange.com

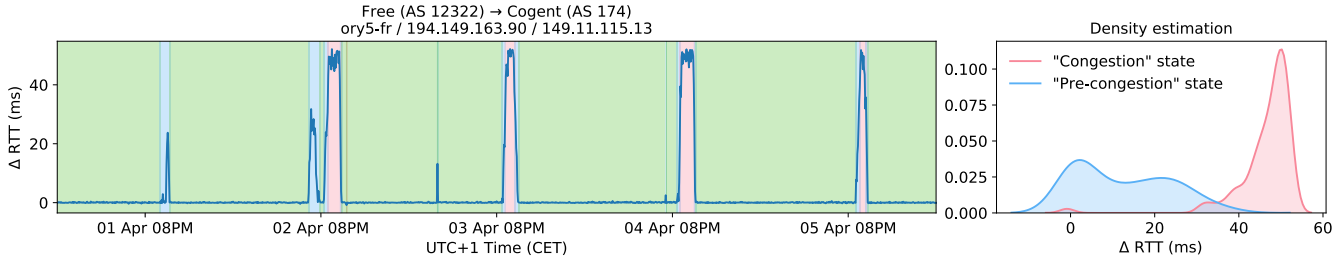

FIGURE 11. Segmentation of a RTT difference (far - near) time series obtained with TSLP probes from the CAIDA MANIC project. Each color identifies a state.

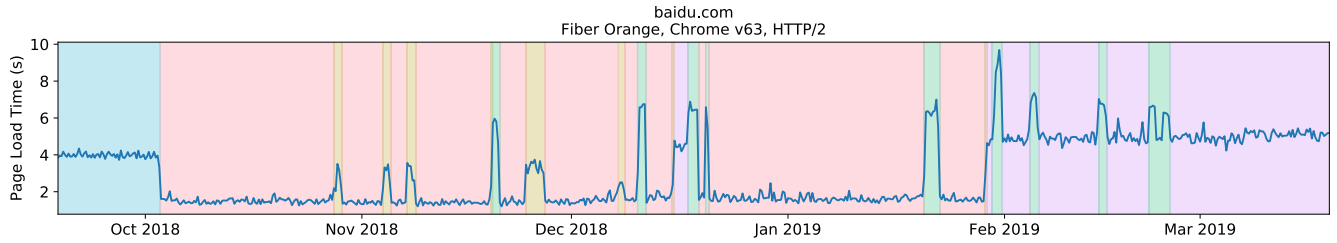

FIGURE 12. Segmentation of a PLT time series obtained from a WebView probe. Each color identifies a state.

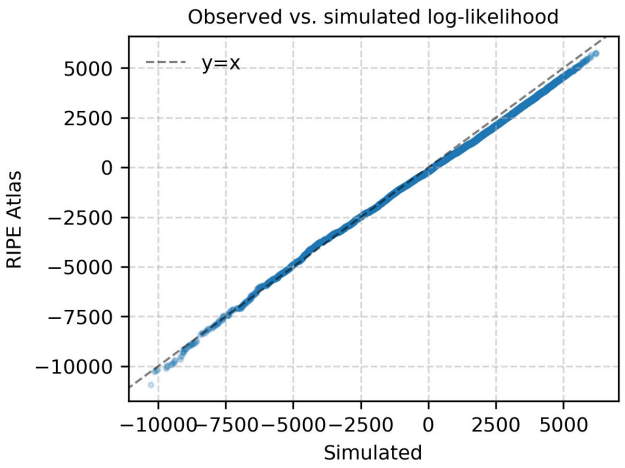

FIGURE 13. Q-Q plot of observed vs. simulated log-likelihood on 100k time series.

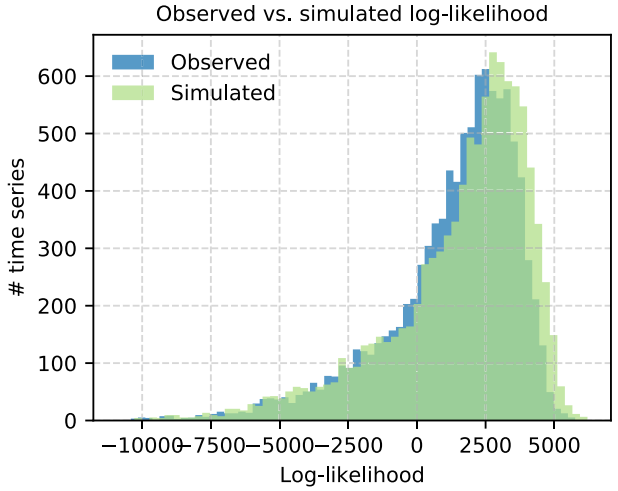

FIGURE 14. Distribution of observed and simulated log-likelihood on 100k time series.

In addition, a Neyman-Pearson (NP) test was performed between two simple hypothesis: H_0 : the time series is distributed according to an HMM, H_1 : the time series is distributed according to an HDP-HMM. The false alarm rate (FAR) was set to $\alpha = 5\%$ and $\alpha = 10\%$. The FAR is the probability of deciding in favor of the HDP-HMM (H_1) when the series is distributed according to an HMM (H_0 is true). When $\alpha = 5\%$ (resp. 10%) the NP test result was that the time series was distributed according to H_1 (HDP-HMM) in 96% (resp. 97%) of the cases. The same experiment was performed to compare the HDP-HMM hypothesis (H_1) against the DPMM hypothesis (H_0). In that case, for $\alpha = 5\%$ and $\alpha = 10\%$, the test result supported the HDP-HMM hypothesis in 99% of the cases.

Thus, we have not only visually verified on a large number of series that the segmentation obtained with the model is consistent with what a human expert would produce (Section IV). In addition, we have checked on a very large scale (about 100k randomly chosen series among the Atlas mesh measurements) that all these series are well modeled by the HDP-HMM.

B. RIPE ATLAS TRENDS API

In order to make our method widely accessible, we have developed a publicly exposed Web API into RIPE Atlas. Given an origin-destination pair (measurement and probe ID) and a time frame (start and stop time), the *trends* API provides the segmentation of a RIPE Atlas delay measurement. The API offers three endpoints, described in Table 2.

The `/ticks` endpoint returns the minimum RTT for a given pair with a constant time interval (duplicated results due to probe connectivity problems are sup-

TABLE 2. Endpoints of the Atlas Trends API.

Method	Path	Parameters
GET	/ticks/:msm_id/:prb_id	start, stop
GET	/trends/:msm_id/:prb_id	start, stop
GET	/trends/:msm_id/:prb_id/summary	start, stop


```
{
 "n_states": 5,
 "states": {
  "1": {
   "rtt": {
    "max": 224.688,
    "median": 206.499,
    "iqr": 0.133,
    "min": 199.882
   },
   "duration": {
    "total_time": 145680,
    "avg_time": 72840.0
   }
  }, // ...
 },
 "segments": [
  {
   "start_time": 1550448206,
   "stop_time": 1550463568,
   "state": 1
  },
  {
   "start_time": 1550463568,
   "stop_time": 1550546126,
   "state": 3
  }, // ...
 ]
}
```


FIGURE 15. RIPE Atlas Trends API sample JSON output.

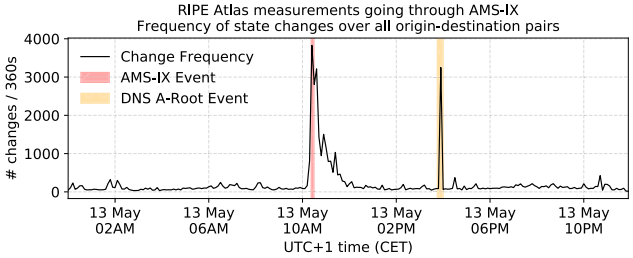

FIGURE 16. Change frequency on the 13th of May 2015 for the 20k pairs that saw AMS-IX in their traceroutes the day before.

pressed, and missing results are explicitly inserted). The `/trends` endpoint returns the minimum RTT and the associated segmentation. For example, the URL https://trends.atlas.ripe.net/api/v1/trends/1437285/6222/?start=2018-05-02&stop=2018-05-10 gives the segmentation in Figure 7 (it should take less than 10 seconds to segment one week of data). A summary of the time series, as shown in Figure 15, can also be requested by appending `/summary` to the path. Start and stop time are specified as `YYYY-MM-DDTHH:MM` where `THH:MM` is optional and defaults to the start of the day.

In addition to this article, we provide interactive notebooks to document and demonstrate the API, and compare various statistical models. Links to interactive Google Colab sessions, as well as the notebooks source and code to facilitate the usage of the API are provided on GitHub [38].

C. MONITORING LARGE INTERNET INFRASTRUCTURES

As shown in [37], a significant number of Atlas origin-destination pairs reliably go through large Internet infrastructures, such as IXPs (AMS-IX, DE-CIX, etc.) and transit providers (Level 3). By *reliably*, we mean that the relevant infrastructure was consistently seen in traceroutes for the relevant pairs over a given time frame. Furthermore, Atlas provides measurements towards the 13 DNS root servers from every probe (more than 10k probes), although such

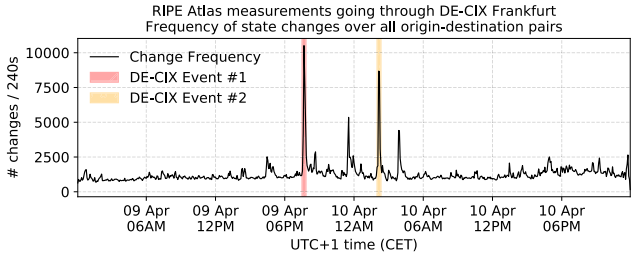

FIGURE 17. Change frequency between the 9th and the 10th of April 2018 for the 60k pairs that saw DE-CIX Frankfurt in their traceroutes the day before.

measurements are more difficult to exploit due to the anycast nature of DNS root servers.

In order to detect anomalous events in those infrastructures, we propose to aggregate the change points learned from each time series individually, so as to obtain a *state-change frequency*, which represents the number of state changes in a given time frame over all the origin-destination pairs considered. One problem is the selection of those origin-destination pairs. One could imagine learning the model for all the origin-destination pairs available in Atlas, or a large subset, such as anchoring mesh measurements (160k origin-destination pairs), and then looking for events in state-change frequency. However, preliminary experimentation shows that the signal obtained is too noisy and requires a lot of manual processing to find relevant events. Instead, we propose to monitor each infrastructure individually, by considering only the origin-destination pairs for which the infrastructure has been seen in recent traceroute measurements.

To validate the ability of our method to detect events, we analyzed two events which have been discussed in the literature (as this provides some groundtruth against which to compare our results): AMS-IX outage in May 2015 [37], [39], [40], and DE-CIX Frankfurt outage in April 2018 [41].

1) AMS-IX MAY 2015 OUTAGE

According to [37], on the 13th of May 2015, AMS-IX experienced a partial outage due to a switch interface generating looped traffic on the peering LAN. The event lasted for seven minutes and two seconds, from 10:22:12 to 10:29:14 (UTC time) before the switch interface was disconnected. This event caused some peers located at AMS-IX to lose their BGP session. In [37], the event was studied using traceroutes, by looking at the percentage of paths seeing AMS-IX peering LAN in their traceroute over time. However changes in the IP paths often result in changes in the delay. The models were learned, and the changepoints extracted, using the ping measurements corresponding to the same origin-destination pairs, provided by RIPE NCC.

By default RIPE Atlas ping measurements are performed every 4 minutes, with a jitter of 2 minutes to maximize the temporal coverage over all the probes participating in a measurement. We therefore counted the number of change-points in buckets of 6 minutes. We show the resulting *state*

change frequency in Figure 16. The real event durations are highlighted in red. The event corresponds to a clearly visible increase in the number of changes. The frequency stays high for a few hours as first of all many peers switch to alternative paths, and then some of them come back to AMS-IX.

We also see a spike between 14:45 and 15:00 (UTC). Further investigation showed that almost all the changepoints that occured during this period are related to measurements targeting the DNS Root-A server. We repeated a similar procedure for all the origin-destination pairs in the Atlas built-in measurement to this DNS server and we saw a similar spike, but all source ASes seem to be affected equally, leading us to believe that the spike was caused by an event close to one of the DNS Root-A instances.

2) DE-CIX APRIL 2018 OUTAGE

According to [41], between April 9th and April 20th 2018, some networks located at DE-CIX Frankfurt lost their connectivity to the route servers, and as a result rerouted their traffic to other interconnections, or experienced an interruption of traffic. An analysis of the rates of BGP updates received by route collectors located at DE-CIX showed that the rates of updates dropped close to zero between 19:43 and 23:28 on the 9th of April, and between 02:02 and 03:51 on the 10th of April. Applying the same methodology as for the AMS-IX event, the state changes frequency for this time frame are shown in Figure 17. The two largest spikes match exactly the two times where the rates of BGP updates dropped to zero. The two smaller spikes match with the two times when the collectors started receiving BGP updates again.

VI. CONCLUSION

In this paper we have shown that the HDP-HMM model, a hidden Markov chain model with a potentially infinite number of states, is a very promising method for analyzing RTT time series over the Internet on long time scales (hours to weeks). We have reviewed the principles of this model that produce accurate segmentation of time series and identification of hidden states. Unlike black box approaches, the HDP-HMM provides some explainable parameters that can be used as input in different network management tasks such as the choice of routes, QoS prediction, or optimization of the measurement strategy.

As shown in Sections IV and V, segmentation results are very close to what a human expert would provide, and any measured RTT time series is very accurately represented by the model. Moreover the analysis method is fully automated with no human intervention, even in the initialization phase, and it is scalable. As proof, it has been implemented on an Internet-wide operational measurement infrastructure, RIPE Atlas, with a publicly available Web API.

We have shown that this method can accurately detect moments when abnormal events occur on the Internet. In the future we would like to automate this detection, and in particular to locate anomalies (infrastructure failures, etc.) precisely. This will require the use of other methods exploiting the diversity of the measured paths and tomographic approaches, or using a preliminary time series filtering strategy. We will also work on real-time processing of measured data to detect novelties in RTT series with HDP-HMMs in an almost instantaneous way, based on some recent sequential approaches to inference in HDP-HMMs [42], [43]. Moreover, having a precise RTT time series segmentation tool will allow us to perform a large scale analysis of the statistical characteristics of Internet paths performance and to revisit previous results [44], [45].

REFERENCES

[1] R. Boutaba, M. A. Salahuddin, N. Limam, S. Ayoubi, N. Shahriar, F. Estrada-Solano, and O. M. Caicedo, "A comprehensive survey on machine learning for networking: Evolution, applications and research opportunities," *J. Internet Services Appl.*, vol. 9, no. 1, p. 16, 2018.

[2] B. Koley, "The zero touch network," in *Proc. IEEE CNSM*, Nov. 2016, pp. 1–39.

[3] M. Boucadair and C. Jacquenet, *Emerging Automation Techniques for the Future Internet*. Philadelphia, PA, USA: IGI Global, 2018.

[4] *ONAP Platform*. Accessed: May 9, 2019. [Online]. Available: https://www.onap.org/platform-2

[5] R. N. Staff, "RIPE Atlas: A global Internet measurement network," *Internet Protocol J.*, vol. 18, no. 3, pp. 2–26, 2015.

[6] E. B. Fox, E. B. Sudderth, M. I. Jordan, and A. S. Willsky, "A sticky HDP-HMM with application to speaker diarization," *Ann. Appl. Stat.*, vol. 5, no. 2A, pp. 1020–1056, Jun. 2011.

[7] R. M. Neal, "Markov chain sampling methods for Dirichlet process mixture models," *J. Comput. Graph. Statist.*, vol. 9, no. 2, pp. 249–265, Jun. 2000.

[8] Y. W. Teh, M. I. Jordan, M. J. Beal, and D. M. Blei, "Hierarchical Dirichlet processes," *J. Amer. Statist. Assoc.*, vol. 101, no. 476, pp. 1566–1581, Dec. 2006.

[9] W. Zhang and J. He, "Modeling end-to-end delay using Pareto distribution," in *Proc. 2nd Int. Conf. Internet Monitor. Protection (ICIMP)*, Jul. 2007, p. 21.

[10] J. Hernandez and I. Phillips, "Weibull mixture model to characterise end-to-end Internet delay at coarse time-scales," *IEE Proc., Commun.*, vol. 153, no. 2, pp. 295–304, 2006.

[11] Y. Sato, S. Ata, I. Oka, and C. Fujiwara, "Using mixed distribution for modeling end-to-end delay characteristics," Graduate School Eng., Osaka City Univ., Fac. Mod. Manage. Inf., Osaka Seikei Univ., Osaka, Japan, Tech. Rep., Jan. 2005.

[12] M. Yang, J. Ru, X. R. Li, H. Chen, and A. Bashi, "Predicting Internet end-to-end delay: A multiple-model approach," in *Proc. IEEE 24th Annu. Joint Conf. IEEE Comput. Commun. Societies*, vol. 4, Mar. 2005, pp. 2815–2819.

[13] E. Kamrani, H. R. Momeni, and A. R. Sharafat, "Modeling Internet delay dynamics for teleoperation," in *Proc. IEEE Conf. Control Appl. (CCA)*, pp. 1528–1533, Aug. 2005.

[14] S. Belhaj and M. Tagina, "Modeling and prediction of the Internet end-to-end delay using recurrent neural networks," *J. Netw.*, vol. 4, no. 6, pp. 528–535, 2009.

[15] A. Dong, Z. Du, and Z. Yan, "Round trip time prediction using recurrent neural networks with minimal gated unit," *IEEE Commun. Lett.*, vol. 23, no. 4, pp. 584–587, Apr. 2019.

[16] K. Salamatian and S. Vaton, "Hidden Markov modeling for network communication channels," *SIGMETRICS Perform. Eval. Rev.*, vol. 29, no. 1, pp. 92–101, Jun. 2001.

[17] W. Wei, B. Wang, and D. Towsley, "Continuous-time hidden Markov models for network performance evaluation," *Perform. Eval.*, vol. 49, nos. 1–4, pp. 129–146, Sep. 2002.

[18] A. Dainotti, A. Pescapé, P. S. Rossi, F. Palmieri, and G. Ventre, "Internet traffic modeling by means of Hidden Markov Models," *Comput. Netw.*, vol. 52, no. 14, pp. 2645–2662, Oct. 2008.

[19] R. Fontugne, J. Mazel, and K. Fukuda, "An empirical mixture model for large-scale RTT measurements," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, Apr. 2015, pp. 2470–2478.

[20] M. Mouchet, S. Vaton, and T. Chonavel, "Statistical characterization of round-trip times with nonparametric hidden Markov models," in *Proc. IFIP/IEEE Symp. Integr. Netw. Service Manage. (IM)*, Apr. 2019, pp. 43–48.

[21] S. Vaton, O. Brun, M. Mouchet, P. Belzarena, I. Amigo, B. J. Prabhu, and T. Chonavel, "Joint minimization of monitoring cost and delay in overlay networks: Optimal policies with a Markovian approach," *J. Netw. Syst. Manage.*, vol. 27, no. 1, pp. 188–232, Jan. 2019.
[22] H. Pucha, Y. Zhang, Z. M. Mao, and Y. C. Hu, "Understanding network delay changes caused by routing events," in *Proc. ACM SIGMETRICS Int. Conf. Meas. Model. Comput. Syst.*, New York, NY, USA, 2007, pp. 73–84.
[23] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the em algorithm," *J. Roy. Stat. Soc. B, Methodol.*, vol. 39, no. 1, pp. 1–38, 1977.
[24] L. R. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Proc. IEEE*, vol. 77, no. 2, pp. 257–286, 1989.
[25] H. Akaike, "A new look at the statistical model identification," *IEEE Trans. Autom. Control.*, vol. 19, no. 6, pp. 716–723, Dec. 1974.
[26] G. Schwarz, "Estimating the Dimension of a Model," *Ann. Statist.*, vol. 6, no. 2, pp. 461–464, Mar. 1978.
[27] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*. New York, NY, USA: Springer-Verlag, 1998.
[28] D. J. C. MacKay, *Information Theory, Inference and Learning Algorithms*. New York, NY, USA: Cambridge Univ. Press, 2002.
[29] T. S. Ferguson, "A Bayesian analysis of some nonparametric problems," *Ann. Statist.*, vol. 1, no. 2, pp. 209–230, Mar. 1973.
[30] C. E. Antoniak, "Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems," *Ann. Statist.*, vol. 2, no. 6, pp. 1152–1174, Nov. 1974.
[31] M. J. Beal, Z. Ghahramani, and C. E. Rasmussen, "The infinite hidden Markov model," in *Proc. Adv. Neural Inf. Process. Syst.*, 2002, pp. 577–584.
[32] S. J. Gershman and D. M. Blei, "A tutorial on Bayesian nonparametric models," *J. Math. Psychol.*, vol. 56, no. 1, pp. 1–12, Feb. 2012.
[33] J. Sethuraman, "A constructive definition of Dirichlet priors," *Statistica Sinica*, vol. 4, pp. 639–650, Jul. 1994.
[34] D. Blackwell and J. B. Macqueen, "Ferguson distributions via Pólya urn schemes," *Ann. Statist.*, vol. 1, no. 2, pp. 353–355, Mar. 1973.
[35] W. Shao, J. L. Rougier, A. Paris, F. Devienne, and M. Viste, "One-to-one matching of RTT and path changes," in *Proc. 29th Int. Teletraffic Congr. (ITC)*, vol. 1, Sep. 2017, pp. 196–204.
[36] *MANIC: Measurement and ANalysis of Internet Congestion*. Accessed: Mar. 29, 2019. [Online]. Available: https://manic.caida.org/
[37] E. Aben, "Does the Internet route around damage? A case study using RIPE atlas," Tech. Rep., 2015. [Online]. Available: https://labs.ripe.net/Members/emileaben/does-the-internet-route-around-damage
[38] M. Mouchet, "Demonstration of the RIPE atlas trends API," Tech. Rep., 2019. [Online]. Available: https://github.com/maxmouchet/atlas-trends-demo
[39] R. Fontugne, E. Aben, C. Pelsser, and R. Bush, "Pinpointing delay and forwarding anomalies using large-scale traceroute measurement," in *Proc. Internet Meas. Conf.*, 2017, pp. 15–28.
[40] V. Giotsas, C. Dietzel, G. Smaragdakis, A. Feldmann, A. Berger, and E. Aben, "Detecting peering infrastructure outages in the wild," in *Proc. ACM SIGCOMM*, Aug. 2017, pp. 446–459.
[41] E. Aben and S. Strowes, "Does the Internet route around damage in 2018?" Tech. Rep., 2018. [Online]. Available: https://labs.ripe.net/Members/emileaben/does-the-internet-route-around-damage-in-2018
[42] A. Bargi, R. Y. D. Xu, and M. Piccardi, "AdOn HDP-HMM: An adaptive online model for segmentation and classification of sequential data," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 29, no. 9, pp. 3953–3968, Sep. 2018.
[43] N. Tripuraneni, S. Gu, H. Ge, and Z. Ghahramani, "Particle gibbs for infinite hidden Markov models," in *Advances in Neural Information Processing Systems*, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, Eds. Red Hook, NY, USA: Curran Associates, Inc., 2015, pp. 2395–2403.
[44] Y. Zhang and N. Duffield, "On the constancy of Internet path properties," in *Proc. 1st ACM SIGCOMM Workshop Internet Meas.*, 2001, pp. 197–211.
[45] M. Iodice, M. Candela, and G. Di Battista, "Periodic path changes in RIPE atlas," *IEEE Access*, vol. 7, pp. 65518–65526, 2019.

MAXIME MOUCHET received the Engineering degree in telecommunications from IMT Atlantique, where he is currently pursuing the Ph.D. degree. His work concerns the optimization of active monitoring in computer networks through statistical modeling and prediction of the QoS.

SANDRINE VATON received the Engineering degree from Télécom Paris, the master's degree in probabilities and finance from Université Pierre et Marie Curie (UPMC), Paris, France, the Ph.D. degree in signal processing from Télécom Paris, and the Accreditation to Supervise Research (HDR) in computer science from Université Rennes 1, France. She is currently a Full Professor with IMT Atlantique, Brest, France. Her research interests are statistical modeling of telecommunications networks and traffic, performance evaluation, network monitoring, and anomaly detection.

THIERRY CHONAVEL received the Ph.D. degree from Télécom Paris, in 1992. Since 1993, he has been a Professor with IMT Atlantique (formerly Télécom Bretagne). His research is related to statistical signal processing methods with applications to several fields (transmissions, speech, sonar, radar, and networks). In the area of hidden Markov modeling, he contributed to techniques for tracking states with unknown and varying dimension in dynamical systems observed from sensor arrays.

EMILE ABEN received the M.Sc. degree in chemistry from the University of Nijmegen, Nijmegen, The Netherlands. He has been a Research Engineer with the RIPE NCC Science Group, since 2009, where he is currently a Senior Research Engineer of the Research and Development Department. Before ten years, he worked as a Web Developer, a Sysadmin, a Security Consultant, and a Researcher. He is interested in the Internet measurement and technology changes, such as the transition to IPv6.

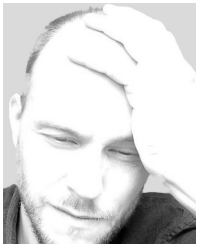

JASPER DEN HERTOG is currently a Software Engineer with the RIPE NCC Research and Development Department. He is also a Core Developer of RIPE Atlas and RIPE IPMap. He is also specialized in cleaning, aggregating and visualizing networking related datasets, and turning them into APIs.